Gartner Group presents

Competition

Change

Dataquest Semiconductors '97

October 22-23, San Diego, California
Loews Coronado Bay Hotel

Massive

Structural

Changes

Are Coming!

Are You

Ready?

Innovation

Dataquest
A Gartner Group Company

EXECUTIVE SUMMARY REPORT

Table of Contents

Panelists: 17-69

Bill Howe
Vice President
Intel Corporation

Miin Wu
Asia Pacific
Founder and President
Macronix., International Co., Ltd.

Dan Auclair
Senior Vice President
Operations and Technology
SanDisk

Bruno Beverina
Vice President, FLASH Memory Operation
SGS-Thomson

Wally Maghribi
Group Vice President
AMD

Dr. Hyung-Kyu Lim
Senior Vice President and General Manager
Memory Division
Samsung Electronics Co., Ltd.

Chapter 1: Technology and the Global Economics

Donald H. Straszheim
President
Milken Institute

Agenda: Introduction; How Technology Changes the Economy; Consolidation; Observations; Globalization; Conclusion

INTRODUCTION

Technology is changing the way we live and work and some associated issues.

What is striking about this list is how attitudes change among investors. Just like in the markets now people are picking out those companies that we all know and love that are going to be the leaders for the foreseeable future.

There are times in which attitudes change a great deal and it's often hard to know really what's going to be right around the corner. Let's talk about technology in particular, not just technology in the markets. In 1973, cars were three times as important as computers in terms of industrial output. Now computers, are 50% more important than cars. It is in this context that your industry is utterly changing the way we live and work in the business landscape. When I used to get together the analysts at Merrill Lynch and talk about the companies, I used to say, "The single most important thing that distinguishes a winning company from a losing company is technology." The winning companies in every single industry are those who understand the power of these new technologies. Those that know how to adapt those technologies to their own particular corporate purposes and those who proliferate the technologies the fastest, cut their costs, and beat up on the competition. Nothing is nearly as important as the development of these new technologies. To understand, to adapt, and to proliferate them.

HOW TECHNOLOGY CHANGES THE ECONOMY

Let's talk about another aspect of this technology how it's changing and the companies, the factors in our economy. Here is the top 100, by market cap, firms in the country by 1955, 1975, and 1997. Oil represents heavy industry. Pharmaceuticals driven by technology etc., finance, the whole computer industry.

Let's talk about the pharmaceutical business. What you see here is the number of companies in the top 50 market cap has gone from none to four to seven over the last 40 years and we show the total market cap of these companies and how large they are in all of industry. It's our view that the 20th century was the century of physics and you all know that in your business, and the 21st century is going to be the century of biology driven again by the technologies that you people generate. There are important regulatory issues that are in the news in the technology business.

I believe there's a potential for changes like this in the technology business. Look at how-high tech dominates capital spending. In 1970, high-tech was 7% of overall capital spending, now it's 46% by this chart and this is a great understatement because there's a lot of high tech content in that low tech capital spending. The other 54% is a lot of those machine tools considered low tech that are digital controlled and all the rest. By the way, there's no software in here anywhere so this just shows you how your industries are dominating and revolutionizing our economy.

CONSOLIDATION

Let's talk about a different dimension. I call it "consolidation in banking." There were 14,500 banks as of 1985, now there's about 10,000. What's going on? The staff and computers of one bank can handle the transactions of two so you make two into one, cut the cost, cut the people, and you go do it again. It's gone from 14,500 to 10,000. Canada's got seven banks. What do we need all these banks? The answer is we don't and you will see a further enormous consolidation in banking. The insurance business they are about five years farther behind. If you put one up of the big four accounting firms, used to be the big eight. If you put up securities business or the advertising business, any of these others, enormous economies associated, it's all part of the flattening of organizations as data is made available throughout an organization. You don't need all these layers of management so you're getting these pyramids that are becoming much flatter eliminating all these people and these companies have a much broader reach and are able to accomplish things in their own business that they couldn't before, a lot cheaper than before, a lot more efficiently than before, with less errors than before, and it's going on everywhere.

OBSERVATIONS

The rich are getting richer. A big increase in the number of people at that highest end $75,000 and over. This is being driven again by technology. Those who've got the requisite job skills are finding their wage being bit up and those who don't are falling off the bottom. This is inherently dangerous. Our country's most important problem is the increasing mismatch between the skill mix of people that are out there and the skill mix that companies want to hire and we fail to address this problem at our peril. It's going to take the public sector, federal, state, local. It's going to take private companies, the non-profit sector, the academic institutions all together to solve this problem. You go to a ghetto in any city and you run into this 13-year-old mother carrying her baby around pregnant with her second and you ask yourself, "What chance does this mother, this girl, have of ever being integrated into the mainstream of our society and really contributing to our economy?" The answer is, "What chance? Low." Then you ask, "How about her offspring?" The answer is, "Less." Technology provides the opportunity to solve some of these problems just as it is in the short run, aggravating some of these problems but too many people are being left behind so we've got to figure out how to resolve this. Related to this is education and what you see here is the school year for 13-year-old kids. Taiwan, Korea at the top and the U.S. is at the bottom. A U.S. college graduate has spent less hours in a classroom than a Korean or Taiwan high school graduate. Maybe that explains in part contributes to some of these issues.

GLOBALIZATION

Globalization. This chart simply shows the share of trade in our overall economy back to 1929. For the last 40 years or so, from 1929 to 1973, it averaged 9.5%. Now it's up to 24%. Increasingly it doesn't matter where you headquarters is domicile. Companies' source, raw materials, intermediate products, labor all over the world. They sell all over the world. That's increasingly going to be the case in the future.

CONCLUSION

We're a not-for-profit economic think tank that studies issues we think are important to the future of the economy. Education, labor markets and jobs, financial institutions, capital markets, corporate finance, how technology is effecting the economy and society, globalization and trade, demographics, all these sorts of issues. If we can get our views out on these topics we think we'll have a more informed public, better public policies, and hopefully improved economic outcomes. Technology is one of the areas that we think is more important than any other in the economy in the future.

Chapter 2: Organizational Opportunism — Utilizing Change as a Competitive Advantage

John McCartney
President and Chief Operating Officer
3Com (U.S. Robotics)

Agenda: Introduction; Change in Society and Technology; Structure; Leadership; Conclusion

INTRODUCTION

Today we buy a lot more devices that are much higher quality and much more powerful at much lower prices than we did 14 years ago. Over the next decade we'll buy a lot more than we buy now at much higher quality and much lower prices than we do now. I'm not going to tell you which companies we'll be buying them from or how we're going to survive and what the devices will be, but in terms of growth, there's no question that the engine that you are fueling with technological innovation is going to continue. It's that growth and what it means to our organizations that I'm going to talk about today.

CHANGE IN SOCIETY AND TECHNOLOGY

There's no question that the pace of technological change is accelerating with great impact on the society around us. Even high school students, while they may not be in school for very long each day, know Moore's Law and they understand in some fashion the impact that it has on them. Micro chips are every place from computers to cars to microwaves to smart cards, but the real story is not just the advances in the underlined technology but how they're being used, particularly in communications. Personal computers in the last five years are up 100%. Cellular telephone subscriptions in the U.S. is up 227%. Internet hosts up 1900% in the last five years. Communications and how it effects people's lives is the most dramatic form of a change that technology is bringing. This is a dramatic one for me. 2.7 trillion. That's the number of e-mail messages that will be sent this year. That's 15 times the number of messages handled by the U.S. Postal service. I can remember that if you were active in international business, on your business card you had your telex number and on your company stationery you had your cable gram address. Even five years ago it was extremely unusual to find e-mail, Web site hosts, pagers, cellular telephone numbers on your business cards.

Now I'd wager that many of you have a business card that's so crowded with contact numbers, people have trouble finding your name.

Companies like IBM, GTE, and Compaq are investing 10s of millions of dollars in corporate image campaigns related to this icon or its cousin the "e" for electronic commerce. Tens of millions of dollars, not directed toward the sales of a single product, but toward image and solution. It's easy to envision a time when the I Internet access symbol will be more ubiquitous than McDonald's golden arches or even Coke's contoured bottle. The changes in PCs are even more amazing. The dramatic increase in power and performance and communications capability is driving underlying fabric of change. T he price performance ratio for computers over the last 30 years has changed by a 100,000%. Think about what that means and if the next 30 years changes much how instantaneous communications will be around the world. This technological change shouldn't blind us to another important issue and that is how fast are our organizations changing to adapt to do this market environment. We individually seem able to change our behaviors to take advantage of these advances and products but our companies changing as quickly.

W e can see the impact of speed if we look at sources of innovation in our industry. Why is it that David can beat Goliath? It's because David got in the first shot. He was there while Goliath was still winding up. It's why so much innovation comes from start ups.

STRUCTURE

In the semiconductor industry, a lot of the foundation in early growth came from supporting defense and government customers. If your customers are organized in a hierarchical fashion there's a tremendous incentive for you to organize in the same way. What does this represent? What does this structure, whether it's a network or organization, represent? It represents mainframes and dumb terminals. Data collected at a huge number of points, tremendous reach for data collection funneled up to the top where the data is collated, analyzed and decisions are made and those decisions or pieces of them are funneled back down. There's no empowerment at the bottom. There's no decision making at the bottom. There are vast resources that are not applied to the marketplace but only applied in data collection. As their structures get bigger and bigger and bigger, they get slower and slower and slower and it creates a culture.

What have PCs done to us? PCs made communication across any point in the network possible and it's created an environment where if allowed decisions can be made down at the node. Information can be shared. Tremendous amounts of power, tremendous amounts of velocity can be generated throughout organizations and not just at the top if organizational structures are adapted to use it. The greatest change that PCs have brought about is just beginning.

With today's tools in front of us we could change the "v" in our organizations dramatically. There are some companies that are doing this. Some companies have dramatically changed their focus in a very short period of time. Texas Instruments is a completely different company today than it was 18 months ago. I have no idea what the internal organizations and hierarchies have done but certainly the speed at which a corporate transformation took place was dramatic. I can't believe that there's one of you in this room that competes against Texas Instruments in their various parts of business that isn't more worried about them today than you were 18 months or two years ago. That opportunity is there in front of all of us.

When you use cross-functional teams, are they quickly evaluated and is the evaluation fed back to the next group on a cross-functional project for improvement or is there just check mark so that they can go back to McKenzie who helped with the last reorganization and says that we use cross-functional teams. Do you use 360 degree feedback at every level of your organization including senior executive staff? Do you spend more time focused on internal organizational and development issues or more time focused on the marketplace and your competitors? Those are indicators I would say of structures that indicate speed is there or can be there within your organization.

LEADERSHIP

Besides changing structure the next big challenge is leadership. We have moved because of the power of technology from a phase where management is the most important thing to leadership becoming the most important thing. It is impossible for a manager, a director, an executive V ice President, or the P resident of a large business unit to be involved in every decision and it is not difficult to find qualified managerial talent or to develop it.

What is difficult is leadership because all of these resources even focused at high velocity don't do anything but generate energy, not force, but energy unless they're well directed. A few key attributes of leadership of business organizations relate to vision. The senior leadership creates the vision. If we don't know that we're in bad shape. It articulates the vision in terms that people within the organization can understand and it energizes employees to achieve the vision. It's just one aspect, but it's an incredibly important aspect and you know you can do a lot of this with e-mail. You do it through communication. You do it through constant interaction with employees at all levels, across, up and down, and sliced because that's what generates change. George Bernard once said that progress is only possible with change and people who cannot change their minds cannot change anything. T he ability to articulate this vision and move an organization forward is one of the key aspects of leadership today.

The networking environment of the future will be much different than the networking environment today. We believe that the term "pervasive networking" will apply. We are basically overwhelmingly 90% plus network enabled. Less than half of the people in the world have ever made a telephone call. The world will become network enabled. Communication as we know it today will be extended everywhere over the next few decades and Bandwidth will increase dramatically. If you're everywhere around the edge, that if you drive a network vision that's simple to use, incredibly complex in the backbone, but simple for people to use, whether that's an individual consumer in Malaysia or whether it's a network administrator at the largest healthcare organization in the United States, if it's simple for them to use and manage it will grow. We believe that multi-media applications are right around the corner. We know that this business will be global and we hope that it will be deregulated. That's the vision that we see of the network of the future.

Characteristics that fit extremely successful companies.There are three that I identify with and which we use internally. The first is ingenious. What we do is all about innovation and invention, the creation of intellectual property. If we don't create, if we don't innovate, then we will never be successful. I would suggest that successful organizations are agile. You need a strategic vision but you need to be able to adapt it and change it because the market will change for sure. From book distribution to record clubs to CD ROM to a leader in internet access and content and now to interactive television. A dramatic transformation of a strategic vision over a short period of time.

The biggest business success story in technology, you might argue about technology, but the biggest success story in business over the last 15 years is Microsoft. Is Microsoft ingenious? Gaining control of an operating system they didn't invent? Convincing IBM to use that operating system as a platform, immediately propagating it across an entire industry and surrounding Apple, and then leveraging that in every conceivable fashion into marketing power. That's ingenious. You can argue about whether the products are any good or not but it's ingenious. Agile. From operating systems to application software back to operating systems with graphical interfaces then operating systems for servers and now, I lose track, are we on the fourth or fifth internet strategy. No remorse that the first one was wrong. No remorse at all, we're just on to the next one and we're going. Relentless, I don't think I need to provide any Microsoft examples about relentless. I would suggest that these are three attributes of successful companies and the companies with hose attributes are able to focus on velocity and use their resources to generate competitive force in the marketplace.

CONCLUSION

Apply your resources—whatever they are—with great velocity. If you have leadership that has strategic vision; and if you are ingenious, agile and relentless I congratulate you.

Chapter 3: Tracking the "Food Chain:" Dataquest's Worldwide Outlook for Key Technologies

Moderator:

Gene Norrett
Corporate Vice President and Director
Semiconductors Group
Dataquest

Panelists

Gregory S. Sheppard
Chief Analyst
Semiconductor Application Markets Program
Semiconductors Group
Dataquest

Joseph Grenier
Vice President and Director
Semiconductor Device Group
Dataquest

Clark J. Fuhs
Director and Principal Analyst
Semiconductor Equipment, Manufacturing, and Materials Program
Semiconductor Contract Manufacturing Services Program
Semiconductors Group
Dataquest

GREGORY SHEPPARD: What is going to drive the chip market from an in-use or demand perspective, and trying to characterize how the semiconductor industry can position itself to best take advantage of that profitably. That's the hardest part, to make sure the profits are flowing.

The most prominent one that comes to mind is the information or IT technology upgrade cycle. This is the idea that, back in the 60s you had an IBM 360 and that got upgraded by a 370, etc. Now we've moved into the client/server model of computing and a lot has not changed, although certainly the processing power has proliferated. We are assuming that we'll continue to see cycles of IT upgrade.

The Internet is going to increase five-times itself in terms of usage over the next five years. If you look at it one way, almost all of you probably have two Internet addresses, one at home and one at work. Something along a quarter billion users is not unfathomable in 2001. It's still very much a U.S. centric phenomenon to date.

Next is wireless pervasion. It's exploded. We're projecting that we'll see over 300 million users of wireless services of all kinds, cellular telephony, paging, the wireless data satellite systems, etc. This is certainly going to have an impact on the companies that are positioned in supplying DSP processors as well as RF IF circuitry in this area. Wireless will grow to be about 10-12% of the market by 2001, basically up from 1%-2%.

We also have the concept of Digital Consumer. This is really the concept of all the purchasing that goes on in the home. If you include the greater consumer spending and throw in PCs and communications technology along with the audio visual and interactive entertainment equipment, we're seeing spending that's starting to rival what's spent on food and clothes, the family car, etc. Likewise, the huge developing middle classes of Asia /Pacific, Latin America, and Africa remain greatly untapped.

On a global basis, we incorporate the growth rate we are projecting into the future. The only possible problem area is the impact of the currency crisis we've been seeing in Asia/Pacific. Those countries are production centers for export to the rest of the world, not necessarily having a huge impact on consumption.

Data processing remains the largest sector and we're anticipating it will continue to grow. The second largest area is communications. The third largest area is industrial, which is applications that measure and control. The medical electronics area is also included here. This is followed by consumer, then we have military/civil aerospace. The next area is automotive, which is quite active.

If we map it into specific applications, we see that the PC is definitely *number one* as far as being the largest application out there. This is essentially all the chips that would go into the mother board, after market memory, graphics, audio, those sorts of things that would generally apply to a PC. The other applications are pretty evenly spread. Most notable is cellular and cordless and other computers.

What's really hot out there today? The sub-thousand dollar PC is catching the world by storm with a lot of recent articles about success in this area. Likewise, we're seeing the emergence of the Net PC. This is targeted at companies. Essentially, it's a stripped down PC with the disc drive removed. Pentium II systems are rolling out with abandon now.

Moving into the Internet area we have the DSL modems, and we have the remote access systems that would in some cases incorporate these modems inside. ADSL takes a twisted pair of lines and makes them run ten megabytes per second or greater. Symmetrical DSL provides a megabyte per second capability still using the same twisted pair line. Of course, the ISDN and cable modem are the other hot areas in this space.

JOSEPH GRENIER: Worldwide electronic production equipment is on top. Currently, U.S. equipment production accounts for about 4% of our gross GDP.

The main driver of semiconductors is PC. PCs drive the DRAM market and here's the impact of PCs on multi-megabytes per system. 1997 is going to grow from 32 megabytes to 152 megabytes per system.

PCs drive the microprocessor market. PCs drive the multi-media market. PCs have been evolving from a productivity tool for spreadsheets, word processing, and other data processing tasks, into an information, communications and entertainment tool. The PCs now have MPEG2, AC3, audio, telephony, fax, modem, and all these features. To do all this requires a lot of processing power, and that increase in processing power, and complexity of chips, is going to drive the semiconductor industry.

It's really the special ICs that are allowing advances to happen in the communications industry. First, let's look at the total impact of the communication market on the semiconductor market. Here I've added to the forecast the chips going into the communication equipment. As recently as 1990, communication chips accounted for about $8 billion and that's going to grow to nearly $60 billion by the year 2001.

In the baseband, which is part of the digital cellular phone which does all the data converting, there are two parts of digital phone: the baseband and the RF or radio section. It's really cellular phones that have spurred the DSP manufacturers as cellular phones are the major DSP application right now.

Digital consumer electronics is going to be coming our way and these are some of the markets: advanced video games, top-set box, digital cameras, digital camcorders. Collectively, they will make a big impact.

Let's look at some of the hot areas in the chip arena. First, a not so hot area: DRAM price slide in 1995 reaches it's maximum at $3.47 per megabit. Now for some of the hot areas: Flash is going to go to $6.3 billion. System integration is another hot area. You can see that it will be the dominant ASIC solution by the year 2000. DSP is also taking off as well.

In summary, the message is that the U.S. semiconductor industry is poised very nicely for the emerging areas. We expect good growth in all areas of electronic equipment, PC units, and semiconductor market. There were a lot of questions about what's happened to that. The industry looks pretty good for the next four years.

CLARK J. FUHS: Our group studies most aspects of how chips are manufactured. We study the fundamentals of capacity silicon consumption, process infrastructure and how that relationship ties to the chip industry. Supply side fundamentals have been a major part of the semiconductor industry of the last couple of years and we expect that to continue going forward. We are calling for a cyclical downturn in 2002 and this is consistent with a chip downturn in 2001.

In 1996, we had a transition year from growth to decline. Backlogs were taken down during the year and that set up the two pause years of 97-98. We expected acceleration of growth and the capacity buys in 1999 and technology buying modes to take us through most of this and next year. Our general outlook has remained unchanged over the last two years. We've always experienced a slowdown in 96 leading to a decline in 97 with single digit growth in 98. We've been able to maintain this record by focussing on the fundamental issues associated with the supply side of our industry.

Here is our range for the possibilities for the Wafer Fab Equipment forecast over the next few years. In 1998, there is actually more downside risk than upside potential. Two things contribute to that. First, the over capacity fundamentals are still weighing on the market and 1997 is coming in a little better in capital spending so there is a somewhat more capacity added than we had originally thought. The second thing is that we suspect the issues going on in Southeast Asia will put a throttle on the capital availability in Asia.

The question will be raised by many chip companies, "What plant do I build to support production in the year 2001?" Our guess is that it will continue to be 200mm wafer fabs, so the race to build the last 200mm wafer fab will have to wait until the year 2002 or 2003. Which companies are likely to lead in 300mm? We've split it into three waves of companies: first wave-second wave, and others.

One comment about the fabless/foundry models: they work. The foundry industry will be a driving force going forward because of concentration of capital and concentration of capacity bringing economies in scale to bear.

In summary, we see the end use demand picture for semiconductors remaining strong. However, the large over capacity today means a significant lag for the equipment market into 1999.

Chapter 4: Emerging Technologies in Equipment Processes

Moderator

Clark J. Fuhs
Director and Principal Analyst
Semiconductor Equipment, Manufacturing, and Materials Program
Semiconductor Contract Manufacturing Services Program
Semiconductors Group
Dataquest

Panelists

Mr. Peter Chang
President
United Semiconductor Corporation

Dr. Michael R. Polcari
Director, Silicon Technology & Advanced Semiconductor Lab
IBM Microelectronics

Dr. Dale Harbison
VP, Semiconductor Group
Manager, Mfg Science & Technology
Texas Instruments Inc.

Dr. Inseok S. Hwang, Ph.D.
Senior Vice President, General Manager
Semiconductor R&D Division
Hyundai Electronics Industries Company, Ltd.

Dr. Eiji Takeda
Dept Manager, System LSI Development Office
Semiconductor & IC Division
Hitachi Ltd.

PETER CHANG: In between, we have two approaches to the embedded DRAM situation. One we call the 0.35 embedded DRAM high performance. The other one we call 0.35 micron embedded DRAM low cost. If you look at the cells, low cost cell is closer to DRAM process cell. For example, the metal layer is only 3 layers on a low cost cell. Basically, it's only one gate outside for all the process cell. Then, most of the technology on the peripheral is very close to the DRAM. Even the peripheral is on the poly side. You look around this thing, that's why they reduce the mask set to close to 20 layers. That will reduce the overall cost.

That application, in some customer's cells, we found are very cost sensitive in this particular. Their application may be concentrated not on the particular performance but on the power savings. That is application they were using for our foundry customers. The other sets for high performance are mainly applications in the graphics area. They want the logic to be super fast. As well, there is a very large bandwidth along with everything else. They want the performance. That is the application for overall. But you have to pay the price. The price is 25 layers. You have to put on all the layers. The cost is relatively high. That's our strategy at this moment. This technology right now is in pilot production. This is the 0.35 embedded DRAM. The complication in the fab itself, we have several different kinds of technology. We have logic technology, DRAM technology, embedded DRAM technology as well as SRAM and volatile memory. Those are the requirements for a foundry business if they wish to survive for the next 10 years. It's very obvious that we are using logic as the technology driver to continue to drive this technology down. This is the full planarization of 0.25 micron of 60 transistor, six layers of metal of processed cell. This is six layer of metal processed protection, and SRAM and DRAM.

We have made plans to invest heavily in the southern part of Taiwan which is Tainin Science Park. We plan to break ground on our first 12" fab in 1999 and plan to pilot as close to the beginning of 2001. That's our current plan and it looks like it's conservative. We believe that there's still a lot of technical difficulty and cost problems related to 12" wafers. This is the total UMC group wafer output forecast. By the end of fourth quarter, 1999, we'll put out close to 492.

In summary of my presentation, in the embedded DRAM situation, we have a deal with our customers. Right now, it's 0.35 microns in pilot production. In our point of view, in the future, the trend is going to be go that direction, but it depends on how big the market is and the question of whether the cost could go down. The big question is whether or not it will become a major force. Second, for 300mm, in our minds, it's not mature yet. We'd like to see the wafer cost go down, the raw material cost down and also the equipment should be expediting their fee to get up to the progress as in other areas.

DR. MICHAEL R. POLCARI: I'd like to talk about our view of system-level integration and what we think are some of the issues and some of the important things that need to be focused on. Cost is a primary concern in all this. There are a number of packaging solutions that may, in the interim, serve to relieve some of those constraints. In the long run, we feel that system-level integration is something that's coming and will be here. If you look at what's been driving it, our ability to go to smaller dimensions and increasing the performance and power reduction is allowing us to integrate everything into a single chip. Process technology itself is not going to allow us to get to where we want to be. We need to combine the design skills necessary for system-level integration as well as some of the design tools and methodology to be able to get to where we want to go.

IP is a major issue in terms of having access to the cores that one would like to invent. One would like to have the ability to have cores from a number of different users which leads one to be concerned about being able to license cores. Also the interoperability of these cores so that one can mix and match cores and come up with a correct type of system-level project that one is working on. We feel that alliances, partnerships and licensing is going to be important in terms of doing this.

As you put more functions on a chip, the tendency of the systems people is to make the chip larger. We are continually talking to our systems-level people about trying to do things that make sense.

Putting all these cores together raises the issue of tests and how one is going to be able to test the kinds of system-level chips. Mixing all of these technologies together to bring a system-level design out is a challenge both to the development and in manufacturing. In manufacturing, introducing different technologies into a fab brings in variability which is always a cost issue. We've been manufacturing memory and logic in the same fab together and we need to be concerned about the variability to make sure the costs are consistent with either technology. Mixing these two together in an integrated process flow brings together another level of cost issues that you need to address.

By having a vertically integrated type of system, you're able to optimize across the boundaries much better than if you have the individual horizontal type model. You still need to worry about things like licensing but that's probably the best solution to being able to bring a solution to the customer.

INSEOK S. HWANG: It looks like the embedded DRAM approach improves performance, power, die-size and cost. However, in order to really meet the cost reduction, we have some issues to be resolved. We have experience that an application requiring a small size of DRAM is not cost effective with the embedded DRAM approach. It may have to use very expensive strict combination of DRAM and logic process. One of the challenges for embedded DRAM is how to reduce increasing mask and process steps.

Should the process be compatible with DRAM or logic. The EML process should be compatible with DRAM especially DRAM with a stacked capacitor. This is because of the management of thermal bodies in the process.

The next area I'm addressing is the process issues in embedded DRAM. There are several issues which we have to solve in an enduring manner like resolution. Especially in embedded DRAM, there is a challenge to implement in the regions between cell and random logic patterns. Planarization is not easy because of high topology difference between cell and peripheral area. Metallization in embedded DRAM process has a very small lithography margin and this is especially true with embedded DRAM with a stacked capacitor. However, we can resolve these issues. When integrating EML process, one of the top issues is the capacitor process because the capacitor process usually requires a formal process.

Next are the 300mm issues. The timing for industry conversion to 300mm is not certain. This is because the initial pilot line start-up cost is very expensive due to poor utilization. All 300mm equipment will not reach the same delivery of maturity. The economics of product revenue and investment may not match. Initially, the wafer cost is very high and secondly, there is a question if global standardization can be really successful. How fast will global standardization be complete? This depends on industry cooperation and I300I and J300I construction cooperation.

DR. DALE HARBISON: The real issue is the cost and the economics. The economics in this situation is what gets in the way. The barrier is certainly not the ability to integrate the two technologies. You have to find a way to justify the additional cost. You need some application with the performance need that will justify that and also that will have the volume to be able to justify the cost that you'll have to do that.

As an industry, we're facing a major cost challenge. Maybe, when they get into volume production in the 2000 to 2002 time frame, they may get down to $600 or $700 per wafer. That is still going to be a major increase in the cost per square centimeter and then you have the cost to build fab's at that technology.

People are complaining about factory costs being $1 billion or $1 1/2 billion for a 20K to 25K wafer factory per month. With 300mm tools, that cost will continue to rise. For our industry to grow, we need to stay on that productivity learning curve. If we can't stay on that green line and have that green line continue to move down at that rate with the transistor cost, then the question is whether or not the industry can continue to grow and how fast it can grow. We have serious trouble if we begin to get off of that line like some of the red-dashed lines show there with the question marks.

How much can we save with 300mm? We expect the cost per square centimeter to improve by 25% to 30% because you get a significant improvement in the usable area, especially with large chips. Depending on the chip size, the cost per chip may be in the neighborhood of 30% to 40% improvement. With labor costs about the same. We've given the suppliers a goal of no more than 1.3 times the cost going from 200mm to 300mm with the same footprint and the same throughput. If that cost increases significantly, when we will convert to 300mm and whether or not we will be able to convert to 300mm becomes a serious issue.

Other things are going to drive up the factory costs. The weight of a 300mm carrier or front-opening pod with 25 wafers is about 17 pounds. This is heavier than is typical and is heavier than ergonomically feasible for someone to carry on. Instead of having bay-to-bay automation like we currently have, we will need a lot of intrabay automation which will require a very sophisticated CIM system and mini environment.

Some general observations are that it seems like the transition to 300mm certainly can be a win-win situation if the IC makers and their suppliers can get the cost to be able to sustain the continued growth of the IC market. The use of larger wafers will continue to drive the progress of semiconductor manufacturing just as it has in the past in going from 4" to 6" to 8." The 300mm will cause significant requirements for improved CIM systems, single wafer processing, process control, automated material handling systems and other things that we can sometimes get along without at 200mm. New levels of collaboration are being fostered in the global 300mm effort as Dr. Hwang said. These are all critical things that are needed and we need a very good global cooperation force to be able to stay on this productivity learning curve.

DR. EIJI TAKEDA: Today I would like to talk to you about the challenges for process and device technology in embedded DRAMs.

The new paradigm shift is occurring in system LSIs as follows. For example, megatrends of nomadic computing ending up PDA, and multimedia computing focusing on the 3 dimensional computer graphics and so forth. These trends of new technology are strongly required, for example, embedded memories, 300mm wafers, production lines and so forth.

These are market fields. There are 3 market fields in every DRAM. One is the embedded system such as a DVD and inkjet printer. Second is the HPC and PDA aiming at the low power. The third application is 3D graphics and games in which we are aiming at the higher performance.

We are now developing 0.35 to 0.25 micron technology. In the near future, FLASH is also embedded in the single chip. In every DRAM, logic performance must be at the same level as the conventional logic performance. In the standard DRAM, transistor performance is a little bit behind the conventional logic.

In higher performance, the RAM on chip approach can provide the shorter time-to-market because the memory architecture, bus design and card design cannot be neglected.

Concerning the modular layout in embedded DRAM, we have to provide the architectural DRAM capacity like a basic method. We are using the multi- bank architecture. One bank consists of 256K bit and we use the CRAM, DRAM technologies. In the near future, we have to fabricate DRAM, logic and embedded DRAM in the same line. DRAM-oriented and logic-oriented processes must be added to the core process.

Chapter 5: Semiconductor Interconnect: The Market Differentiator

Moderator:

Mohan Warrior
Director, Final Global Manufacturing
Motorola

Panelists

Ed Fulcher
Director Advanced Packaging Development
LSI Logic

Steve Anderson
Senior Vice President, Corporate Marketing
AMKOR Electronics, Inc.

Thomas DiStefano
Founder and Vice President, Marketing
Tessera, Inc.

Dr. Rama Shukla
Manager, Components Technology Development
INTEL

John Novitsky
Vice President, Marketing
MicroModule Systems

ED FULCHER: What is a major packaging and interconnect differentiator? It is reducing the system cost. That is what drives the businesses, that is what drives the customers. It is to get more in a system for less money. One way packaging can do this is by reducing the buswidths. When you reduce buswidths you can do it with higher speed, so that you can transfer more data on fewer lines. Packaging plays a role in virtually all those aspects. A third way that the interconnect in packaging can reduce system costs is especially for the ASIC business that we are in. We can't afford a custom package, for every design, otherwise we'd be in the application specific package business, the ASP business, so you must have standard predesigned packages. We find that over 97% of all our designs go into a standard package.

This helps to reduce our cycle time also and get you to market faster.

Package cost itself has been reduced significantly over the past five generations. If you look seven to 10 years ago we were in ceramic packages that the first bar, it's a relative scale of cost of the assembled package. Then we went to pin grade arrays, then we went to a plastic pin grade array which was a printed circuit board laminate type and reduced the cost there about six years ago. Then we moved to a ball grid array from a pin grade array allowing us to be higher density, smaller package and reduced the cost again. We are all working on that and it has not emerged yet but it will soon. Chip scale packages would fit in, especially if they are also low cost flip chip.

What is the product differentiator for packages. Reducing size anywhere increases performance not just in the package. It can also increase sales and revenues. In many products people want it to be smaller. The first people with the smallest application will get the highest number of dollars and the most customers. It applies to all products. All our customers from the highest end to the very lowest want less size in their next generation.

The leading challenge we face is "greater density at lower cost" because you can always have greater density if you are willing to pay for it. Greater density is the challenge. As we go to smaller systems, denser silicon, more I\Os even though we may be reducing the number of busses and shrinking those busses a form of Rents rule still applies and the customer still want more I\Os than they had before in the previous system. So more I\Os, smaller packages, denser die, they all ask for the same thing: greater interconnect density for less money.

STEVE ANDERSON: Our company and others in the packaging industry are being pulled or pushed or thrusted in these two new areas where you see the overlap between the purple and the yellow circles. Where we are talking about purple or semiconductor we are starting to see alliances forming with wafer foundries with packaging people. What are those interconnect technologies trends going on in the wafer, how do we help semiconductor companies address those with packaging.

Our imperatives for interconnect deal with four areas and they are not necessarily exclusive. Number one would be design for cost; we call it chip array. Using standard wafer saw technology and sawing so you get the smallest possible package or waste of laminate. The next thing that would come up is design for test. We see that falling in two areas for packaging. One where you are dealing with gravity fed testers and the other where you deal with pick in place. In the gravity fed where you are typically dealing with small lead count say eight through sixty leads we see the industry moving smaller or moving slower possibly to chip scale unless we have a good cost effective test solution.

We've got to try to move these new packages we are bringing up into existing infrastructure until new infrastructures can be developed. You are going to find us doing panel or strips that match most of the equipment built today in the packaging industry.

The interconnects get very complicated. We see packaging broken down into three areas and this is driven by cost. We get a lot of debate with our customers as to give me best interconnect you can to get the best speed but at the lowest cost. Many times they will take a step backward at technology in order to keep the cost low.

What does that mean? Our customer would like to see it at the same cost or cheaper than they are now buying the lead frame or the standard technology. That is not always possible because the infrastructure and volumes are not there today.

In summary, we see these major trends: CSP packages from our vantage point, we are looking at about 200 semiconductor companies, will certainly fill the void for the low lead products until the direct chip attach solution are more cost effective. Front end and back end manufacturing whether it is one company or not must be integrated and there must be better design tools. We see the substrate technology development must be nurtured.

The last thing that we see moving up in packaging is we are involved now in micro machine optical and sensor packaging. A lot of it off the BGA trend but it is a thing that needs to be understood better and nurtured. We think that's going to be very exciting as optical packaging grows.

DR. DI STEFANO: It's been three years since the CSP first surfaced. That's a very short time in this industry and judging by how far CSP have come in those three short years it's difficult to look ahead and see where this field is going.

In 1997 you see small lead count packages being used for flash memory and some of these personal electronics products. The next step in the propagation of CSPs is to standard DRAM whether it is a high speed bus like a RAM bus or synchronous DRAM where standards are being set for chip size packages of high pin counts. From there we expect the CSPs to migrate to higher pin counts really being gated by the availability of high density substrates to mount these high density packages to. Then for small dies you see some fan out configurations or fan in\out configurations where there are more I\O than will fit under the shadow of the die. This rounds out. the evolution of CSPs into the higher pin counts and the higher performance in the standard products.

Standardization is really an important part, not just for the package but for the infrastructure. A very important part of this is that the package must fit the existing infrastructure in order to ramp up manufacturing very quickly. That's supported by standardization of the equipment from wire bonders, dispensers, die placers to support this growing infrastructure. Looking ahead a little further into the future you see that the high pin count end, area array pads, the drive toward higher I\O especially for processors and ASICs is moving us toward area array contacts the chip itself.

At this point there are a lot of benefits to assembling this chip size package on the wafer itself. Since the package is chip size you can build the package up like a high rise building on top of the chip in the wafer form and simply extend the move toward batch processing instead of 30 parts in a strip or 50 or 100 the natural evolution is lets put these on a wafer and process them on the wafer.

You don't have an integrated circuit until the additional layers are added to that chip and that functionality will give you performance, it will give you a reduction of the pin count coming off the package because you do some of the power/ground distribution right in the interconnect layers themselves.

Looking ahead all this leads you make a few observations and they are really logical observations starting with the premise that the chip size package. First of all, CSPs will replace permanently the parts. The drive is to higher I\O, area array small packages. These CSPs are based on the existing infrastructure. There is no new equipment, there are some modifications that the materials sets require but fundamentally it operates on the standard assembly infrastructure in place today. Upgrading the equipment, modifying the materials somewhat to make CSPs. The result of that is that we expect the adoption of CSPs to be extremely rapid.

DR. RAMA SHUKLA: If you take a look at the computer industry from the 80's to today there has been a phenomenal change in terms of market segmentation. In the early 80's we had three universes, the PCs, the main frames and the supercomputers. They are three different domains, very different technologies. PCs use CMOS or NMOS, main frames use ECOs and supercomputers use gallium arsenide. They are really three different universes and they did not talk to each other very much. They were fragmentations as opposed to segmentations. The fact that these are segmented markets, there is a challenge to keep standardization in terms of technology, in terms of the infrastructure because if you don't do that we'll go right back to the fragmented market. You will have on one end, personal computers which will use completely different technologies versus servers or work stations and that will take the cost right back up. That is a very key challenge in keeping those markets growing without losing the standardization. The last point is having a package which provides some level of flexibility and head room for growth in terms of performance and features. That is exemplified by Pentium II processors.

If you take a look at the interconnection densities for the chip level which is illustrated as metal one through metal four to bond pads package board and connectors. If you look at the distance the electrical signal has to travel and the pitch, the kind of densities. It is a huge range. It is no wonder that the technologies used on silicon are completely different than what is used on the package versus what is used on the board. The board technology is very different from semiconductor fab and they are driven by these different geometries and requirements. They are in the technological sense different universes.

In the past assuming that they are different universes, there was local integration and local optimization. Silicon people optimized interconnects based upon multilevel aluminum, silicon dioxide dielectric structures but what is happening is now partitioning is moving more toward market segments.

The future is going to be even more different where the partitioning rather than applications specific will be more based upon economic partitioning and distribution of the logic meaning globally optimizing all the interconnects.

In summary, the interconnects are at a very interesting threshold. Interconnects play a very significant role and they are the foundation for this market segmentation which is created by the CPU revolution. In the near term focus will be to take the existing technologies and evolve them to maximize the local integration between CPU and cache subsystems. In the long term as the area array technology become more pervasive and as the high performance substrates become available then other stuff such as partitioning functions from the chip to the package or the cartridge becomes possible.

JOHN NOVITSKY: It's basically a horse race. The way the industry is shaped right now there is a class of new companies like MicroModule Systems, that are pursuing this very fine pitch set of materials and we are trying to build an infrastructure fast enough to service an industry with companies that are represented here as well. On the other hand the PCB guys are trying to figure out how to develop some very different techniques and it is a horse race between those two. For those who have historically made ceramics the horse race there is try to get a different set of materials, different than ceramics for some of the dielectric and routing pitch and other historical concerns as well. It is too soon to call again as an industry landscape this is kind of the rough shape of the industry that is trying to service the needs that have been pointed out.

In our business we don't care whether the people use the dense substrate for single chip or multichip in the next few years there will emerge such an infrastructure of a very high volume probably pretty low cost dense capacity to do these types of substrates. If and when it exists, it allows us all sorts of opportunities and particularly to pursue few chip modules.

If it makes sense over time, there is a path to forward integration where you do combine those technologies and bring them back together a few years later. If on the other hand it does not make sense either technically or economically you can still allow yourself to keep that footprint down and allow them to still change independently.

Chapter 6: Surviving and Thriving in the Chipless Business

Gary Smith
Director and Principal Analyst
Electronic Design Automation Program
Online, Multimedia, and Software Group
Dataquest

Agenda: SLAMS, List Know Your Methodology, Electronic Systems Design Automation, Virtual Prototypes, System Design Problems, Design Yields, The Winners

SLAMS

We have been trying to address the silicon problem by changing methodologies. Every time we change methodologies, we get a 10 times increase in productivity and that just isn't working anymore.

The answer is System Level Integration, which is the obvious way we're going to address this problem. SLMs are the blocks we feel will be filling the void of silicon, increasing our productivity enough so we can actually use 5 million, 10 million, 15 million gates — whatever the semiconductor industry tries to throw at us. The definition of system level macro is 2500 gates is a UART. That's the simplest slam.

Think of a world where 40% of all microprocessors and DRAMs are SLMs, no longer pinned out. Virtual components. 60% of all DSP microcontrollers, SRAMs are SLMs, and 80% of all ASPICs are now SLMs.

This cyle started in the early 80s with disk drive controllers. They were all ASICs. The ASSP guys made them standard product and got their price down. Those standard products became SLMs, which would fit into another ASIC design, which was multi-SLM design and became an ASIC again. Every jump in technology is made by ASIC implementation and is then consolidated into an ASSP.

We're shortcutting the pinned out silicon on 80% of most of the ASSPs and a lot of the other designs, so suddenly the silicon business does not become an IC business. They start separating, and there will be an entire industry based around producing designs that never get to silicon as a stand-alone pinned out IC.

KNOW YOUR METHODOLOGY

Very few people can tell me what methodology they are using within the engineering department of their company. Do they have standard methodologies for SLMs? Very few companies know how to do this. Of course, those that do are doing very well. What are your engineering problems? Late design, yield issues, keeping your engineering talent: these are all indications of obsolete design methodology.

Methodology is the driving factor in design today. It was yesterday, and it will be tomorrow. If you're in upper management in any semiconductor company and you don't understand design methodology, you should start reading up on it because that's going to be the difference between whether your new fabs are going to be full or not.

We're seeing is a newly emerging theme in systems design automations. There are now six system design automation study programs or projects going on in the United States and around the world: one in Japan, one in Hong Kong, two in Europe and two in the United States. Four are automotive programs and two are in the aircraft industry. That's something to look for and the direction we're going.

ELECTRONIC SYSTEMS DESIGN AUTOMATION

You've probably heard of ESDA. We are now at the electronic system level. At that level of complexity, we can knock out about a million gate design a year. Gate level is pretty much going away as a usable methodology and then you go down to CAD.

A little about the RTL methodology: This is the new ES level that came out in 1994, designing at the behavioral and architectural level. You're doing hardware, software, code development, code design, partitioning. This is the hot area today and has a possibility of solving many software problems, because the main problem with software today is we don't know how to partition our designs into hardware and software. Hardware gets everything you know, and software gets everything that's left over. The software is always late.

Unfortunately, this whole thing started coming apart a couple of years ago as we got down to 0.5 micron. We found that the RTL design was becoming extremely silicon sensitive. After you're over 80 Mhrtz, there is no such thing as a digital transistor. We've gone all this way with the concept of a transistor being Off or On. Other than that, we don't care; but at 80 Mhrtz, you start to care.

VIRTUAL PROTOTYPES

The other issue is that the systems level was really separated from the RT level, so passing designs from one to another was an inefficient mapping operation. The current idea is using virtual prototypes. Virtual prototype at the RT level is an RT level floor planner used as a cockpit for the design, with block place, SLM place, macro place. You're creating a virtual route.

You also need to lock at clock trees, power nets, scan chains and by metal layer. You also have to have an accurate delay calculator and simultaneous analysis estimates. There's not a tool on the market today that can do that.

The new RTL methodology and its entire tool set is extremely complex. It's all feeding into a virtual prototype which is getting simultaneous estimates of timing, power, Sig-Int, EMI, metal migration and thermal. The hardware / software virtual prototype allows software to run on your design prior to ever making silicon. This is getting really complicated, and we can't do it yet.

SYSTEM DESIGN PROBLEMS

One of the big problems we're seeing in the semiconductor industry is a dramatic change in the challenge set every time you give us a new process. It used to be area. Now it is "How fast can it go?" At 0.8 micron, the power user is saying "Speed first, area second." At 0.7-0.6 micron, we have to look at the power side of it which is getting a little scary. At 0.5 micron, we start burning holes in our designs. At 0.35 micron, we saw the first metal migration. Metal migration is the older term for electro-migration.

We ran across metal migration in RF back in the early 70s, when you started seeing a lot of migration problems in 3 Ghrtz transistors. What happens is your metal line moves, and it tends to short out your design.

You catch most of this in burn in. I'm not sure how many people do life test anymore. That's scary because this year we've had multiple reports of metal migration. If there are people doing 0.25 micron designs and they don't do life tests. This is an extremely sensitive reliability issue.

DESIGN YIELDS

We believe the low yields we're seeing in 0.25 micron, is signal integrity. We are just getting the tools on the market now which can extract the information to give to an analysis tool so we can see if it is a signal integrity problem. We're a long way from solving this and we may have to live with low yield for another six to nine months before we start solving some of these issues.

At 0.18 micron, we have to look at inductives. Right now we extract existence and capacitance, and we are going to have to figure out how to extract inductence figures. Unless we put out some tools in the next nine to 18 months, we are really in trouble at 0.18. We could slow down the world on this one.

If the new physical verification technology, you don't have the problems you were trying to design around at the RT level. Everything is complex today.

WINNERS

The winners will use the latest methodology. They'll continue to invest in the latest tools and develop internal tools to fill the voids. Power users spend about 40% of their design dollars developing their own tools to fill the voids in the commercially made tools available.

Lucent has done the best job in making this switch. A lot of the big companies spend a lot of money in maintaining the latest and greatest schematic capture. Almost nobody was using schematic capture anymore and Lucent decided that every tool not 10 times better than something available outside was dead. They then shifted all those resources into developing the tools they needed to complete the tool set. Today Lucent is a power in SLI. I'd put them at number two; they are really doing a great job.

Have a solid SLM development methodology driven by the President or COO. The problem is always with your best designers; they always know how to do it a little bit better. You have to explain you don't care about performance or size optimization, don't touch that transistor. I need a reusable core. That's what you have to drive from the top.

You want engineers working on new designs, not reinventing the wheel.

That's what's going on in the wonderful world of design. It's sometimes scary but always thrilling. The world is changing and a lot of you guys won't survive.

Chapter 7: DSP: Be All That You Can Be

John Scarisbrick
Sr. Vice President, Semiconductor Group
Worldwide Manager, Application Specific Products
Texas Instruments

Agenda: Introduction, DSP Solutions, Tech Support, Progress

INTRODUCTION

Looking back five years to 1992, you can see that TI had a portfolio of businesses, the largest of which was semiconductors, but we were in notebook computers, computers, defense systems, software, etc. even though semiconductor was the largest business at that time, it was less than half of TI's revenues.

Today, we have focused our efforts and resource in one area, Digital Signal Processing solutions. Semiconductor revenues now represent nearly 85% of our total revenues at TI, and within Semiconductor, Digital Signal Processing solutions represents about 40% of that and growing. It's a dramatic change for TI.

As we look at the DSP market, I think you'll understand why we've chosen this strategic position. The DSP segment is the fastest growing semiconductor market.

We've seen tremendous growth in DSPs, which follows the same pattern that microprocessors experienced early in their life cycle. The first single chip microprocessor was commercially available in the early 70s. In 16 years, the microprocessor market reached the $1 billion milestone, and three years after that, the market reached $2 billion.

The first programmable Digital Signal Processors became available in the early 80s, and in just 12 years, in 1994, the market for DSPs reached $1 billion and just two years after that, in 1996, the DSP market reached the $2 billion milestone. The DSP market has actually built up more momentum in its first 15 years than the microprocessor market did in the same time period in its life cycle.

While the growth trend is very similar to microprocessors, the DSP market is not dependent on a single end equipment like a personal computer. DSP growth is driven by very different trends, and DSP applications are widespread and diverse. They include wireless communications, server control in hard disc drives, modems and other data communications devices, set-top boxes, imaging and multi-media, industrial and automotive control, and literally hundreds of other emerging applications. It really has become a DSP centric world.

We see double digit growth in every area. DSPs are not dependent solely on any single end equipment market. No one segment today represents more than 25% of DSPs total usage, and this provides multiple opportunities for growth and a relatively stable foundation for the DSP marketplace in the future.

With such broad base support, we believe the market demand for DSP solutions will explode. This includes both the DSP core and the mixed signal and analog components that make up Digital Signal Processing solutions. Since 1988, the market for Digital Signal Processes has grown by more than 40% per year, up to around $3 billion this year. We expect continued growth, well above that of the semiconductor market, for the next 10 years. We believe that the Digital Signal Processing solutions market, together with related mixed signal and analog devices, will reach $50 billion over the next 10 years.

The growth in Digital Signal Processing solutions means a corresponding growth in the mixed signal and analog components that go along with the DSPs. These provide the interface between the analog world and the digital world, and are critical to delivering a DSP total solution. You need to deliver both the DSP and the analog technology, and you need to do it well, and offer customers a total solution.

DSP SOLUTIONS

In 1997, TI's DSP solutions can be found in a number of areas. One out of every two digital cellular phones in 1996 used TI DSPs. We outsold all TI's DSP competitors in the digital cellular marketplace, working with creative engineering and systems companies like Ericsson, and we're now shipping around a million DSPs a week to the digital cellular market in what is a very high growth market for us.

Our DSP solutions can also be found in nine out of 10 high performance disk drives doing the server mechanism control and 1 out of 3 high speed modems. The high speed modem market we are very confident we can now declare victory for the U.S. Robotics and X2 technology standard. Last month, our data shows that X2 56 capable modems captured more than half the U.S. resale retail market in the 56k modem arena.

It is TI's DSP based platform and high speed modems, and our collaboration with U.S. Robotics, that has been, and continues to be, a very winning partnership and proposition.

As the world leader in DSP solutions, it's vital to us that the best minds in the industry are pushing the envelope on software technology on DSP applications, and we want to support the development of those applications with our DSP technology. We also want to put our investment money where our future is, in DSPS. We want to grow the entire spectrum of our applications using DSPs.

TECH SUPPORT

Building and retaining leadership in the DSP market requires much more than just developing great technologies and architectures. Perhaps even more important than the technology is the web of support you need to build around your products — software programmers, universities, hardware and software companies, etc., all adding value to the architectural proposition we make to our customers.

One measure of a truly mainstream technology is the high level of support available from third parties. By this measure, DSP is certainly mainstream. We count more than 30,000 programmers writing billions of lines of DSP code, with more than 300 third parties adding value through software and hardware development.

Beyond third parties, there are more than 900 universities teaching engineering students DSP design methodologies on TI architectures. We're fortunate to have a disproportionate share of these DSP industry resources focussed on TI architectures.

With the efforts of people and companies such as these supporting the DSP market and TI's DSP architectures, the future looks more promising than ever. In addition to this value web of third parties and university partners, TI also offers the highest performing DSPs on the market.

We have the broadest range of DSP cores of anyone in the industry. Rather than forcing designers to adopt a single architecture, we offer them a choice, not a compromise for their design.

PROGRESS

This progress is reflected in the broad market acceptance of TI's newest digital processor. Leading Internet access providers, ADSL product developers, telecom switch vendors, and base station manufacturers have already chosen the C6X for their advance designs, and the total design ends for the C6X DSP are going faster than any of our previous generation of Digital Signal Processors.

Judging by the number of tool sets we've shipped to date, we have hundreds of developers working on the C6X already. We estimate around 100 new programmers a week are signing up to develop software for TI's Digital Signal Processors, and several hundred million dollars worth of business has already been identified in multiple market segments for that same C6X product, including communications and the mass market.

We'll see hand-held TV become a video conference forum, allowing parents to check in on their children in daycare, all due to DSP technology. We'll see voice recognition and identification technologies allowing us to control access to long distance calls from our cell phones by recognizing our voices. Things like retinal scan identification at an ATM machine will also make for better security to the electronic money network of the future — all again enabled by DSP technology.

In summary, the next wave in the semiconductor business is digital, and DSPs will be making most of the waves. The winner in the exploding DSP market will be the company that has the architecture, the process technologies, the installed base of knowledge and the product leadership. TI is that company.

Our strategy is to increase our world leadership in DSP solutions, and to continue to widen the gap between us and our competitors. We'll do this by sharpening our DSP focus, and by putting our resources where our strategy is.

Chapter 8: Surviving and Thriving in the Chipless Business

Gary Smith
Director and Principal Analyst
Electronic Design Automation Program
Online, Multimedia, and Software Group
Dataquest

Agenda: SLAMS, Know Your Methodology, Electronic Systems Design Automation, System Design Problems, Design Yields, Winners

SLAMS

We have been trying to address the silicon problem by changing methodologies. Every time we change methodologies, we get a 10 times increase in productivity and that just isn't working.

The answer is System Level Integration, which is the obvious way we're going to address this problem. SLMs are the blocks we feel will be filling the void of silicon, increasing our productivity enough so we can actually use 5 million, 10 million, 15 million gates, whatever the semiconductor industry tries to throw at us. The definition of system level macro is 2500 gates is a UART. That's the simplest slam.

KNOW YOUR METHODOLOGY

Very few people can tell me what methodology they are using within the engineering department of their company. Do they have standard methodologies for SLMs? Very few companies know how to do this. Of course, those that do are doing very well. What are your engineering problems? Late design, yield issues, keeping your engineering talent: these are all indications of obsolete design methodology.

Methodology is the driving factor in design today. It was yesterday, and it will be tomorrow. If you're in upper management in any semiconductor company and you don't understand design methodology, you should start reading up on it because that's going to be the difference between whether your new fabs are going to be full or not.

What we're seeing today is a newly emerging theme in systems design automations. There are now six system design automation study programs or projects going on in the United States and around the world: one in Japan, one in Hong Kong, two in Europe and two in the United States. Four are automotive programs and two are in the aircraft industry. That's something to look for and the direction we're going.

ELECTRONIC SYSTEMS DESIGN AUTOMATION

You've probably heard of ESDA. We are now at the electronic system level. At that level of complexity, we can knock out about a million gate design a year. Gate level is pretty much going away as a usable methodology and then you go down to CAD.

A little about the RTL methodology: This is the new ES level that came out in 1994, designing at the behavioral and architectural level. You're doing hardware, software, code development, code design, partitioning. This is the hot area today and has a possibility of solving many software problems, because the main problem with software today is we don't know how to partition our designs into hardware and software. Hardware gets everything and software gets everything that's left over. The software is always late.

SYSTEM DESIGN PROBLEMS

One of the big problems we're seeing in the semiconductor industry is a dramatic change in the challenge set everytime you give us a new process. It used to be area. Now it is "How fast can it go?" At 0.8 micron, the power user is saying "Speed first, area second." At 0.7-0.6 micron, we have to look at the power side of it which is getting a little scary. At 0.5 micron, we start burning holes in our designs. At 0.35 micron, we saw the first metal migration. Metal migration is the older term for electro-migration.

We ran across metal migration in RF back in the early 70s, when you started seeing a lot of migration problems in 3 Ghrtz transistors. What happens is your metal line moves, and it tends to short out your design.

You catch most of this in burn in. I'm not sure how many people do life test anymore. That's scary because this year we've had multiple reports of metal migration. If there are people doing 0.25

micron designs and they don't do life tests. This is an extremely sensitive reliability issue.

DESIGN YIELDS

We believe the low yields we're seeing in 0.25 micron, is signal integrity. We are just getting the tools on the market now which can extract the information to give to an analysis tool so we can see if it is a signal integrity problem. We're a long way from solving this and we may have to live with low yield for another six to nine months before we start solving some of these issues.

At 0.18 micron, we have to look at inductives. Right now we extract existence and capacitance, and we are going to have to figure out how to extract inductence figures. Unless we put out some tools in the next nine to 18 months, we are really in trouble at 0.18. We could slow down the world on this one.

If the new physical verification technology, you don't have the problems you were trying to design around at the RT level. Everything is complex today.

WINNERS

The winners will use the latest methodology. They'll continue to invest in the latest tools and develop internal tools to fill the voids. Power users spend about 40% of their design dollars developing their own tools to fill the voids in the commercially made tools available.

Lucent has done the best job in making this switch. A lot of the big companies spend a lot of money in maintaining the latest and greatest schematic capture. Almost nobody was using schematic capture anymore and Lucent decided that every tool not 10 times better than something available outside was dead. They then shifted all those resources into developing the tools they needed to complete the tool set. Today Lucent is a power in SLI. I'd put them at number two; they are really doing a great job.

Have a solid SLM development methodology driven by the President or COO. The problem is always with your best designers; they always know how to do it a little bit better. You have to explain you don't care about performance or size optimization, don't touch that transistor. I need a reusable core. That's what you have to drive from the top.

You want engineers working on new designs, not reinventing the wheel.

That's what's going on in the wonderful world of design. It's sometimes scary but always thrilling. The world is changing and a lot of you guys won't survive.

Chapter 9: IC Manufacturing 2000-2010: A Decade of Momentous Change

C.M.(Mark) Mellier-Smith
President and COO
SEMATECH

Agenda: Introduction; Who is SEMATECH

INTRODUCTION

I took the liberty of choosing a some what hypobalic title. Namely we face a decade of momentous change in the semiconductor industry. That really what I want to chat about today. Momentous versus incremental depends on your perspective. If your outside the fab, the changes may look like they are incremental, believe me, if you inside the fab over the next 10 years they look pretty momentous to me.

WHO IS SEMATECH

SEMATECH is a consortium of about 10 U.S. semiconductor companies that's engaged in manufacturing R&D. It started about 10 years ago, it is located In Austin, Texas, we have about 600 employees, of whom about 160 engineers on loan from member companies. It's a very valuable, important way we transfer technology back to our members. We run something like \$120-\$150 million a year budget. Our focus is very much on equipment and uniprocessors, a fully integrated process we leave to our member companies, they usually view that as being proprietary and strategic. We tend to concentrate on individual tools and particular problems at one time.

This is a curve that I believe is fairly familiar to everyone in the room. This is a concentrated business. If nothing else, this curve and this tremendous exponential growth for 30 years is a remarkable achievement in an industry. Since almost everyone in this room is dependent on this curve for their continued employment, it is worth looking at what could possibly cause this turn over on us. What might cause us not to stay on our present growth of about 15%-20% per year. The first is, we are blessed with what appears to be a completely elastic market, you lower the price and the volume goes up even more. If we run out of new markets—you know that's not going to continue. The second problem we may find, we run in to super technical problems. We have had a history of being able to solve anything physics throws at us, but as we get down to 50 nano-meters of less design rules, there are some very significant physics problems, that we are going to have to solve. Another problem we have is managing complexity; whether it is complexity in design, in software, we have also heard a lot about systems on a chip. Anyone involved in systems already knows the biggest problem is not with hard ware, it is with software. That's something that if this industry aspired to doing that, it's going to have to change into a very software centric industry.

Last but not least, we fought off the manufacturing productivity curve. Now what do I mean by that? If you look at this industry for the past 30 years, we've been able to deliver to our customers an improvement in functionality; that functionality could be the number of transistors, it could be the number of bits and memory, it could be millions of instructions per second, if you selling DSPs or microprocessors. All of those have approved approximately 25%-30% compounded annually per dollar of price, and that is truly a remarkable productivity improvement. It is based both on lowering the cost of the transistors by changing the design rules and doing other things. It is also based on performance. Running the clock speeds up, making the transistors do more work for us. As I look at this curve, this is what concerns me. If we start to fall off of this curve, then the things that have made our industry so valuable to our customers and so important to us will begin to change.

Chapter 10: Intellectual Property: The Enabler of System-Level Integration

Dr. Walden C. Rhines
President and CEO
Mentor Graphics Corporation

Agenda: Introduction; Reuse of Prior Designs; Defining Standards; EDA Suppliers, Conclusion

INTRODUCTION

Reusable design information is referred to as intellectual property. It's a name that has stuck but it has the connotation of lawyers and, in fact, there are a lot of lawyers involved. The reality of reusable design information is the one element that will make possible the utilization of the silicon.

REUSE OF PRIOR DESIGNS

We've talked about reuse for years, and it's not a new topic. Everyone has thought for a long time that we should reuse what we've already designed, but it doesn't happen very much.

It's going to happen for some very fundamental reasons. The basic driving force this time is simply the number of transistors or gates or bits available, and over the six-year period, mid '90s to early next decade, that's going to increase by 10x, by virtue of the capacity that's being built. The actual design cycle time required for new products is not increasing. In industries like telecommunications that used to take years to develop a new system, that now they're being required to do wireless cell phones in six to nine months just to keep up with the product cycles.

The number of designers per chip is not increasing appreciably just because of the difficulty of growing design teams in size and coordinating the activity. You might argue the EDA vendors ought to get some tools that can fix this problem, but they've got their own problems with how design tools have to handle all these deep submicron effects and do power analysis and reliability and things. There's really only one place to easily get it, and that's from design reuse, using what's already been designed before.

There's some other reasons why you would want to reuse what you've designed over again. A lot of the design today adds no potential value. It's the same companies designing the same functions again and again throughout the world. Standard protocol blocks, other pieces of design information. They don't add any value.

There's another reason. There's the deep submicron physical effects. When you get a block of physical design put in place and it performs to a given level and you've analyzed it for power and noise and reliability, in the future generations the rule is going to be don't fool with it; reuse the block. There's too much invested in getting it to work the way it works today.

Two things have to occur. Design reuse has to happen in a big way within companies and across companies. We can design more than we can verify, so system verification tools become the key. A lot of said I've got to be able to design everything and reuse it across the company because I want to be self-sufficient and design my own chips. That's not enough. We had years at Texas Instruments where we thought we could do everything better than everyone else. We continued to enter every possible market segment, developing products, memories to microprocessors and IO and so forth.

We always hoped that some day we would see one of these customers that would bring us a printed circuit board that all came from a single vendor. It can't happen. It couldn't happen. It never could have happened. No one company can be best at everything. If you're not going to be best at everything, then you need to use the things from people who are best at what they do. That is what has to occur in the reusable IP business.

The chips of the future will increasingly look like the printed circuit boards of the past. You're already seeing a lot of embedded core DSPs and embedded micro controllers and microprocessors, but the steps ahead are even more dramatic in the reuse of functional blocks.

DEFINING STANDARDS

One of the groups that's defining standards associated with the reuse of intellectual properties is the VSI Alliance, Virtual Socket Interface Alliance. Mentor and other EDA companies were founders of this group which now has over 150 members that try to drive general standards. Rapid is an organization that is made up of independent suppliers of intellectual property who principally are worried about the common rules for marketing and selling their intellectual property. One of the most recent ones was one where Synopsis and Mentor Graphics joined forces because it became clear that VSI was going to give high level specs; but people wanted a process, a methodology, a how-to, and they needed a cookbook, a reuse methodology manual. We developed that cookbook in the consulting services to help people reuse or develop reusable IP.

What's happening that's an equivalent in the semiconductor industry? The model is shifting. Traditionally, vertically-integrated companies defined the product; they designed it, developed it, manufactured it, assembled it, shipped it, supported it. That will always be the case for differentiated processes, like linear and like power, even like DRAMs. This model will stay around probably forever, but an increasing percent of the market will go to the fabless model where you drop out that line that says manufacturer and you define the product and you ship it to a customer and support it, driven by companies like Chips & Technologies.

What's the next wave? It's very clear. It's the chipless company. It's the company like ARM or DSP Group or Chromatics that provide design information, and that type of company is becoming the fastest growing segment of the new product business world today as new companies are formed every week to attack that chipless market.

It's a big opportunity for the vertically-integrated companies, but you need to be sure that you're not subsidizing inefficient manufacturing with good intellectual property. If your manufacturing costs are bad, form a relationship with a foundry designed to the same design rules and shift your capacity back and forth as needed. Don't subsidize one with the other, and the same thing on the other side. If you've got great manufacturing costs, go acquire designs, acquire intellectual property. Every company needs specialized IP that is shared within the company and reused in the company but not made available elsewhere.

Where is all this intellectual property going to come from? Mostly from where it's always come from, from semiconductor companies and from systems companies. The new entry into this age is the independent provider of IP, and the new support challenge is for the EDA vendors to provide the tools and the infrastructure to support reusability for designs created by semiconductor and systems companies as well as by third parties.

The semiconductor companies and systems companies will continue to be the dominant creators of IP, but increasingly there will be third parties that are creating intellectual property that's used. They will own it; they can market it where it's appropriate, but that won't be the primary thing. The primary thing is using it and doing system integration to develop system chips. That's where the greatest value is added, by semiconductor companies and by systems companies, and that will be their primary role. The semiconductor companies will shift the production units and so forth. There are secondary roles for all of these, but the third parties mostly will concentrate on creation.

EDA SUPPLIERS

What about the EDA suppliers? There are actually four models that have been discussed at EDA industry meetings. Most of them support the EDA company providing the tools to make reuse possible, and the consulting services and the things needed so that you could verify and use IP. There are several different approaches in the industry for the other roles.

The one Mentor pursued is the distribution and support role because support is such a key challenge, and that is working directly with the designer of chips and providing whatever standard building blocks that designer needs.

The system integration model, is one that Cadence has popularized where you provide the actual design services and do the actual chip design, or in some cases handle outsourcing of the design process, so the focus is on use and system integration.

The third one, Synopsis has driven—portability tools, their cell-based array architecture, etc.

Lastly, one represented by a lot of companies, Aspect and Mentor and Compass and others, just doing the physical libraries. Most of the ASIC libraries in the world today are designed by EDA companies just because there's such a great economy of scale in reuse of libraries.

CONCLUSION

Reusable IP and the ability to separate design information from the silicon will cause an acceleration in innovation in the industry because it will involve thousands of designers who were not part of the process before, able to independently create intellectual property. To do that they'll need that IP to be supported to customers all over the world. They'll need people who can add all of the design views they need, to put in testability or take out testability or support a particular new tool, and that's what the EDA industry will do. As we move to a greater mix of designers, you'll see a greater leverage of design innovation. Spending more time adding value, providing more specialization, all this adds up to a much more productive industry, an industry that can create system chips that have the best of everything available and achieve the end result in the most efficient way. It's going to be a major discontinuity and will have a large impact on our total industry.

Chapter 11: 2010: A Vision of the Future

Gordon Bell
Senior Researcher
Telepresence Research Group
Microsoft

Agenda: Importance of Computers, Cyberization, Laws at Work, Economics, Distribution, Hardware Technology, Predictions

IMPORTANCE OF COMPUTERS

The reason why computers are so important is that they fundamentally are capable of supplementing, and in many cases, substituting or supplanting other information processors including humans. There is a sub-theory that computers are so powerful in themselves that one builds computers on top of other computers all the time. That's why Sun wants to build Java on top of our PC's. Then there will be a JavaPlus or a HotJava or a SuperJava on top of that so that we have these layers and layers of architecture taking all of these processing cycles and interpreting everything.

Computers are built in a well-defined way, on a component structure, in terms of processing memories, ability to switch and the ability to transduce. It gets bits from the physical world into the computing environment. The second great invention, which occurred at the same time, was the transistor and, subsequently, the IC.

CYBERIZATION

Everything that is cyberizable will go into cyber space. That's really driving the use of computing. It's the cyberization of the world. We can think of this as a series of universal networks into the worldnet, in the continents and down on into cars and homes. Everything will be connected with an IP address down to body-area nets. Cyberization is the encoding of all of the information into this one universal network. It's a coupling to all information. To information processors, it's pure bits like printed matter. It's bit tokens like money. It's the state of everything - places, things, just knowing where everything is. It includes the state of physical networks, knowing everything about highways and where cars are and everything like that.

LAWS AT WORK

Let's look at the collection of laws that makes all of this possible. We've got the two inventions. We've got this forest to fiberize the world to encode everything, put it in a network form and give it an IP address. What we have to look for are transducers to make this possible. Moore's first law governs what's happening. You should remember that with this kind of exponential growth, the past doesn't matter. So, 10X here, 10X there, pretty soon you've got some real thing. PC platforms have declined faster due to volume, learning curves and the demand curves.

Computer components must all evolve at the same rate. Amdahl's law says that for every instruction per second you need a byte of memory and a bit a second of I/O. Processor speed has evolved at 60%. If you ignore the fact that you can't really get that 60%. People haven't learned about processor speed at the application level as opposed to at the benchmark level. Storage is evolving at 60% and it is substantially more important than the semiconductor part of all of this.

A law that I have is a corollary from Moore's Law, but it's a law that governs computer classes. The computer platforms emerge based on chip density. They require three factors, namely, the platform itself, the form factor and cost. They require a unique network and they require some form of cyberization, the connection of the computer to something whether it's a car, a human or anything else. Then applications follow that particular class and then each class becomes a vertically disintegrated industry that we saw before based on the hardware and software standards.

Intel's strategy is to put more on a chip and try to keep the price constant for those chips. The other thing is to take the system, integrate all of that on a single chip and take the resulting cost-reduced chip form. With the 256Mbit chip coming out, the 32 Mbyte chip, why not just put a processor and some I/O on there and have a single chip system. You know what the cost is so you decide what the price is based on what operations are done.

What we see today is an industry that is a completely disintegrated level of integration and all of that has happened by standards. All of these levels can be done by itself.

ECONOMICS

The economics is the other factor that drives everything. As the price declines we get a doubling of demand. We have the learning curve that says every time the cumulative volume decline increases by 2X, you get a unit cost. Bill's Law defines the economics of software in that it says you have to do things at a large scale. Nathan's Law defines the virtuous circle. Metcalfe's Law defines the value of a network that allows us to sustain the market. These all work.

Bill's Law says the price of things is really related to the unit cost. Bill Joy initially talked about this and he said you can't afford to write software for under 100,000 platforms. That is if you want a $1,000 price. That all comes out of marking up the engineering expense and looking at the various lines. Bill Gate's Law is that you can't afford to write software for under a million platforms. I'd say that's probably 5 to 10 million platforms right now. You've got a fixed engineering cost or you've got a fixed cost and you divide by this number of units.

DISTRIBUTION

The marginal cost of distribution has gone from a cost of $1 or so for a CD and all the cardboard and air that's shipped with these boxes to zero, or near zero, as you distribute software on the Internet. We see what happens when you look at UNIX versus NT. A massive economy of scale just because of the units. If you look at how the price of Oracle software has come down, this gives you some clue as to why Larry Ellison wants to get rid of PC's because for some reason now the SQL-Server that used to cost $100,000 on a Sun box suddenly costs $6,000 on a PC. That will eventually will be shipped with NT and there will be another lawsuit. It's unfair to ship a SQL-Server to a non-database person. That's just a file system. Do you want a file system included in your computer or not? That ought to be the standard. It also explains why there are no spreadsheets or presentation packages on UNIX, VMS, MDS or those other machines. You just can't afford to design for that.

Nathan, the Chief Technical Officer of Microsoft, describes his laws of software that makes it all happen. He's right, software is a gas and it expands to fill whatever container that it's in. Software then grows until it is limited by Moore's Law, the amount of memory. That's why we can put so many features in. Software growth than makes Moore's Law possible because if you didn't have it, you wouldn't have to have these memories and there wouldn't be other application to fill that. Software is really only limited by human ambition and expectation and our ability to cyberize, to encode things into the computer.

That whole thing is a combined virtuous circle that starts with innovation which presumably has some utility and value otherwise people don't buy it. That leads to volume and volume leads to competitors which goes back to innovation. All of this is not possible without standards. Standards are the thing that makes that whole circle work. If it weren't for that, everyone would be able to operate in their own stratosphere as the UNIX vendors have stabilized by having their own little standard cycles by themselves and, in fact, are being driven back to a single UNIX standard. That will eventually happen as the PC standards drive them out of business or drive them to a single standard. We are seeing consolidation in the UNIX space and we are seeing people leaving the microprocessor. We see HP leaving it's thing. DEC will be next followed by MIPS and probably SPARC and IBM as the last to go in terms of the standard for high volume manufacturers.

The final thing is Metcalfe's Law which says that utility of a network grows as the number of users squared. The utility to an individual is the number of other people it connects to but it's really the sum of all users and that's how you get that square term in there. This says why you've got to have everyone on the Internet or why you want everyone on a single telephone system is because it grows exponentially.

HARDWARE TECHNOLOGY

Hardware technology makes it all happen. We've got a bunch of factors. Number one is we get more. There is a 256Mbit chip that will happen in the next few years. That's a 32Mbyte compute. These will all shrink down to single chip systems. LSI Logic has proclaimed itself the system-on-a-chip company with a very large number of gates. Megabit bandwidth by then will be as easy to get as ISDN is today. That's really a fairly small number of users. There are about a million ISDN subscribers in the US today. What we've got then is a collection of networks that are all forming. These new networks like the pager networks, the phone, fax and all of these have to be connected together to form this cyberspace. We've been on a Moore's Law curve for processing memory backbone and storage. Those lines going out there like 60% a year lines, and then we contrast that with the telephony there which has been about 14% a year.

The question is will that get up into the megabit level which is what's critical by this time. Semiconductors densities are increasing. The communications disaster. Processor performance is going up. If you run these out 50 years, you've got these curves. Things get cheaper. These were curves I did in 1975. I didn't believe those curves. We just couldn't comprehend those kinds of change. New and cheaper always wins but if it weren't for the law of inertia. The law of inertia says that data and programs sustain the platforms. That's why the mainframe will fill in operation in another 100 years. It will be probably be working at the 2100 year problem at that point.

You have to understand the goals of the hardware suppliers is the uniqueness and to differentiate and lock in. The goal of the software vendors is to differentiate and lock in but to operate on as many platforms as possible.

PREDICTIONS

I have a bunch of bets that there will be a non-predictable computer that none of us see today. The nice thing about this are surprises. Larry Ellison believes NC's will outsell PC's by 2000. This is really the whole business, scaleable network platforms. Some of my friends believe that you can make arbitrarily large computers from PC's. That's based on commodity hardware and commodity networks whether it's ethernets or other kinds of things. With that, you can put them together in arbitrary ways. For us, scaleability means reliability. Always upscaling in terms of number of nodes. Scaling in locations, putting it anywhere you want. Scaling it with machine generations. That means not having these multi-processors like the vendors are trying to push. How many microprocessors can I put in a big box? That's not scaling. You can put 10 to 20 of those in a box. To us, scaling is if you can put more than a thousand in a box.

The model is that various kinds of networks will be put together and integrated to perform as a single system. I don't think the vendors see this change coming because they've got these great businesses that wouldn't have been there without the Web. The Web has changed the whole computing server market but that's going to decline.

Chapter 12: Special Report: Inside the Mind of the End-User: What Semiconductor Users Really Care About

Jonathan Yarmis
Vice President and Research Director
Gartner Group

Agenda: The end-user,Technology, Year 2000, The Internet, Conclusion.

THE END-USER

When we look at what's inside the mind of the end-user, let's also understand that in many instances the things that they're most excited about are not necessarily your best selling opportunities. We see the hype cycle of any new emerging technology. What happens is that when it's first introduced, it's the most amazing thing ever to strike the face of the earth. It's going to cure cancer cells, global warming, render El Nino a non-issue and reinvigorate economy. Then three weeks after these massive pronouncements are made, people say that it has failed to meet expectations. We go into this trough of disillusionment that says that this product is never going to work. How could I have been so stupid to ever think this was going to work. This is a dead technology. This is usually about the time that people are going to start buying this stuff. We go into the period of reality where we tell ourselves we have this inflated sense of expectation about what this thing could do and it couldn't do any of that but when I went to dismiss it, there actually was something to that technology that's interesting to me and maybe now it's time, with realistic expectations, to go into that space. We see it time and time again.

TECHNOLOGY

At the peak, we've got Push Technology. The amazing thing with the Internet is things go from the hype cycle to disillusionment to reality in about a four week period. Now, you see the stories how Microsoft is backing off from Push. Point tests should have sold out to Rupert Murdock because they're never going to be able to have an IPL because nobody ever wants Push. This is a long term, essential, important technology. The good news is that now that we're going into that disillusionment period, we can begin to get real about the capabilities of the technology.

What do users want? Every once in awhile, we try a novel approach and ask them what they want. We've got a program called ITEP for the IT Executive Program made up of a couple of hundred CIOs. Every year we ask them what they want. There are two interesting and important things here that has some relevance to the transformation that's going on in this industry. The North American commercial market turned out to be 1 and 2. That is aligning IT and business goals and IT for competitive breakthrough, competitive advantage.

There are some profound business transformation issues that are interesting in the user community. This enables IT as a competitive differentiator. This competitive advantage thing is something we've said forever. In most instances in this industry, competitive advantage was a fleeting thing. You'd gain competitive advantage for about 72 hours and then everyone would see what you were doing and copy it. All of a sudden competitive advantage was lost. What we redefined was the level playing field. We believe that the next 10 years will be the greatest opportunity for competitive transformation that any of us are going to see in our lifetime. The way people use information technology to reach new customers and new markets. The way they transform their businesses and more tightly integrate themselves with their customers and their suppliers. There are massive opportunities to transform industries.

TECHNOLOGY ISSUES

There's a lot of technology issues in the user community. What do you look for in real estate? Location, location, location. Technology is Internet, Internet, Internet. You cannot talk enough about how companies are looking to transform the way they do business given Internet technology. It's replacing a lot of what they've already done. It's also creating incredible new opportunities. There's a side of me that says let's be honest about the Internet. This is the biggest step backward this industry has ever taken. What kind of breakthrough is this? We can display text and graphics on a computer screen, and this is what we are hailing as the greatest breakthrough? Somewhere along the way in obtaining this ability to display text and graphics, we decided that subsecond response time wasn't really that important. Maybe users will tolerate subhour response time. We've taken major steps backward in an opportunity to achieve a global connectivity. That is such a radically transforming idea that it is worth all of the pain and suffering that we'll be going through. From a technology side, Internet, Internet, Internet is what's on the users mind. Year 2K is also. They wish it wasn't on their mind. It's certainly a big impact. When they're looking at strategic directional technology issues, we can't talk enough about the Internet.

YEAR 2000

There's this looming Year 2000 problem. The bottom line is whether you are going to be Year 2000 compliant when it comes time to be Year 2000 compliant. About half of the companies out there will not be fully year 2K compliant. We define full compliance as not only your company being compliant but all your business trading partners and so on.

The very fear of Year 2000 non-compliance becomes a self fulfilling prophecy causing catastrophe in the global financial market. This is a truly terrifying issue. In a corporate marketplace, this is a big and growing issue. Anyone that is trying to figure out what's on the minds of corporate users, figure out a Year 2000 angle. It's a fascinating one. This industry hasn't done enough talking about Year 2000 issues. There are a whole bunch of things in the semiconductor marketplace where there are nonintuitive date issues. The whole issue of what the lurking liability is and what the impact on the industry is, is certainly something that's very much on the mind of end-users and growing more everyday.

THE INTERNET

This is the Paradigm shift. This is the single biggest business transformation opportunity any of us is going to see in our lifetime. The technology has enormous impact on everything we do. I'm a walking Internet access device. This one isn't Internet enabled yet but probably some of you in the room have Internet addresses on your cell phone. This is 2128595@skytel.com. We're going to start to see the Internet pop up in so many different things. Think about how that's going to change the way you buy goods and services in the next couple of years but, even more profoundly, going out 5 or 10 years. You cannot say enough about how the Internet is changing business. IT organizations are spending a lot of time looking at this. We are capturing the imagination of business users around the world in a way that's going to mean great things for those of us in this industry.

Where are people investing? We've given a whole bunch of things but there are only three that really matter. Number one, obviously, the Internet. Number two, Year 2000 stuff. If you have something that solves the Year 2000 problem, you're getting close to the ability to write your own ticket. This is one of those things where if your companies haven't jumped on the Year 2000 bandwagon already, how you're going to solve the problems, I'd love to know who you are because I can write you a really big invoice to help you get out of an enormous problem. There's a lot of money going into this space. The amount of resources that are available are going down. The demand is going up. It is a great solution place to be in. If you could just label semiconductors Year 2K compliant, you could probably price them 10% higher. That's probably being conservative.

I can shift the burden to someone else and sue them if it doesn't work. That's another terrifying statistic for Year 2000. I don't know how many of you saw Lloyds of London. If there are any publicly traded law firms, go invest in them. Lloyds of London estimates that the legal claims arising out of Year 2000 noncompliance in the United States alone will be a trillion dollars. This is a profound business issue that is very much on the minds of users. It is not just an IT issue anymore when you start talking about a trillion zeros. I'm not as conversant with all the number of zeros as the semiconductor people are, but it's the one with a whole lot of zeros at the end. Enough to put some of us out of business. This is very much on the mind of end-users.

There is a market in the world for NetPC's but it is probably not at the expense of personal computers. There is room in the world for both devices. A lot of people are going to want more functionality and more power at the desktop. There's a category of users for whom the PC is overkill and there's room for a large and vibrant NetPC market. After 10 years of harping about it, users have finally gotten on this total cost-of-ownership bandwagon. This says that in deploying this technology, the acquisition price is only about 15% of your 5 year life cycle cost of the technology. Contributing to reduce total cost-of-ownership is a big selling proposition and the user market is galvanizing very quickly. I give Sun, Netscape and Oracle a lot of credit for trying to use that, for obvious marketing reasons, but it's caused the users to wonder what it really costs them to deploy technology on the desktop.

On the software side, what users care about is organizational productivity. We've spent 15 years in this industry solving the personal productivity problem only to discover we don't have a personal productivity problem. We've got an organizational productivity problem and the flow of information goods and services is very much what people are focused on. Don't give me yet another solution to a personal productivity problem, give me Groupware, give me Enterpriseware, give me InterEnterpriseware. That's where all the focus is in the software market.

CONCLUSION

Users want to change their lives on the Internet while surviving the Year 2000 problem. They want to learn how to cope with the increasing complexity that we're delivering. How do we, at the same time, deliver an incredibly rich technology solution yet make it easier for them to manage? There's a side of me that's a representative of a technology company. It embarasses me that we give you these fat binders. The only reason why I'm not more embarassed is most of you, even if we gave you a CD-rom, aren't there yet with the technology. Think what it would be like in a couple of years when you've got your small handheld, handwritten, annotated device that has high resolution display graphics capabilities. We would beam the presentation to you and all throughout it, in real-time, you've got the information available to you. That, to me, is the fascinating transformation process. Apple referred to it as anytime, anywhere computing and we're really on the cusp of that. That the prices, the form factors, the supporting infrastructure and all of those things are taking place that enable incredible business change. It makes this remain the vibrant industry in our economy and a great place to work. .

Chapter 13: Special Speaker Panel: Bringing It All Together

Moderator:

Jonathan L. Yarmis
Vice President and Manager
Special Projects
Gartner Group

Panelists

Larry Bowman
Bowman Capital Management

Mark Melliar-Smith
President and COO
SEMATECH

Wally Rhines
President and CEO
Mentor Graphics Corporation

Gordon Bell
Senior Researcher
Telepresenec Research Group
Microsoft

Agenda: Cost vs. Demand; Information Technology; Money Impact; Interface Problems; Communication

COST VS. DEMAND

There's a lot of cost demand tradeoff. On the GPS side, it might be nice to have the GPS chips and have them tell you where to go. Any of you who've used GPS systems in cars, they're phenomenal. You drive along and it says you should turn in 200 feet and so on. The problem is, it's $2000 or $3000 added to the price of the car and the average person can't afford it. We have to do an awful lot of cost reduction to adjust not only the chips, but the displays, the interface, the speakers and everything.

The economies of scale that can be achieved are going to happen naturally. With some resistance because they require cooperation. We've seen it on the manufacturing side. Years ago, everyone did their own manufacturing equipment and everyone did their own process flows. Now, you buy a piece of manufacturing equipment and process recipes can be included. Everyone uses one of a few alternatives for each process step. The same thing has to happen on the design side. As long as 50 different companies design 1394 interfaces, everyone's got to debug them and everyone's got to put effort into it. The big leverage point is going to come when one design for 1394 is the one that everybody uses. When it's been in a hundred designs, it's debugged and it's high quality. When you just pull it up on your computer, use it and spend your time doing the things that are unique. Big economies of scale-end design yet to be realized. We already have the mechanism and it works to cause more specialization and more economy.

For the system on a chip, there has to be an integration at the same level as the PC. That is the original bus structure so the 1394 interfaces to something. You need the layering and architecture for system on a chip. I predict there will be three or four of those particular things. There will maybe be four 1394 interfaces for the various varieties. There will be an Intel chip. There's certainly a chip that has all of the integration out there that works with the software. There are handlers there so people can write software to those various system on a chip platforms.

Capital markets are really neat. If you remember the movie, greed is good. Where there's money to be made, there's capital available. Period. Where there's money to be lost, there's money to be made. The money's there. Period. It will be there. You talk about $10 billion FABs having to be built. The neat thing about high technology is that the real money makers are the people early on. The small design houses that develop innovative technologies, patent them, get licensing breakthroughs, get royalty streams off the intellectual capital. The revenue stream, the royalty stream doesn't even really kick in until '99, 2000, 2001, 2002. It's the start-up companies of today that will give you the homeruns and grandslams, the return on your invested capital, that will allow the designs and implementations to achieve whatever it is you guys were talking about.

INFORMATION TECHNOLOGY

Do you think there's an appreciation in the market of the transformation and role that information technology is going to take so that maybe those old rules of valuation, at least the ones that apply to this sector, don't apply anymore?

The rules have already changed. The highest quality growth stocks in the world, typically ones we'd look at, whether it's Home Depot or Coke, they trade at their growth rates. They're going to grow earnings at 30% a year for the next 10 years, they trade at 30 multiples. It used to be that technology companies, because they were technology companies, had a multiple because of that. You never got over 15 multiple for a PC company. You never got over an 8 multiple for a disk drive company. Intel never got over 15 or 16 multiple up until one or two years ago. Now, the markets are starting to realize that maybe they can count on PC's continuing to be around in 3 to 6 years from now. Whereas, before, we thought they were just a fad. Three or four years after PC's started up, in '85, they hit a wall because we ran out of all the single user app's. We had a downturn in units in the PC industry. Now, there's more and more long term confidence that technology is here to stay and again it is a growth industry. Confidence and multiples are certainly going up and people are much more willing not to just puke out because they had a bad quarter. It's a $26 billion company, growing at 30% a year and people are starting to recognize more and more that there are ups and downs within the cycle. Intel is up, almost 100%, in the last year.

Where are the discontinuities? That's something that as market forecasters, we're as guilty as anyone. We can draw these great extrapolations and then along comes something like the Internet. I'd love to sit here and tell you we predicted that. There were people out there who predicted it. You generally thought of them as lunatics before they predicted it and even though they were proved to be right, you still believe them to be lunatics. Where does the discontinuity come in? The things that say that Moore's Law is conservative and we're going to kick that curve up or that Moore's Law is over and we're going to ratchet it back 30% a year.

Moore's Law is not the phenomenon. It turns out that it's inconvenient to shrink feature sizes and grow wafer diameters and that has caused Moore's Law to be true. The real underlying principle is the learning curve- the cost per bid or the cost per function. That's what drives it. Looking out in the future, it may be that going to all these exotic technologies that Mark talks about isn't the lowest cost way to get there. In fact, molded chip solutions and going vertically instead of horizontally, shrinking and that kind of a thing becomes a more viable mechanism. If I looked to the design side and say where is there a discontinuity as we move forward, the one that's on the horizon the soonest seems to be at about the quarter micron point, maybe a little sooner for some, a little later for others. You just can't use the same toolset to verify a full chip solution. You get to the end of the design or somewhere in the middle of the design and you put it on the computer. You network every workstation and server that you've got and it runs for days and days and it never gets there. That's the breaking point in which you need a new tool, a new methodology. You need to reuse blocks so that you don't have to verify the individual things they're built from. You need to get a capability that does, for example, hierarchical verification instead of flat verification, which means you throw out all your old verification tools and buy a whole new set and a new methodology. That's going to occur in the next two years.

Chip designs of the past have basically assumed that the embedded software I put in ROM will work with the hardware and that if it doesn't, I can go back and make a quick change to the ROM and fix it. The simulators of today don't even verify enough hardware instructions to boot the operating system much less look for subtleties of function. The tool developers have to get tools that, in a reasonable amount of time, do hardware/software co-verification for that embedded processor. It's one of the faster growing areas of the EDA industry. A whole set of things like that is required that wasn't required before. That's what the discontinuity will drive.

Are people a barrier to implementation in technology? Absolutely not. I'm a simple person in a simple world. We need to get to the point where a person can do something easily. They want to talk to someone in Hong Kong so they do a teleconferencing. It needs to be that you walk in and turn on the television and there's all the data you need on a worldwide web. You say open the door, let out the dog, start the car, turn on the hot water, what's going on at work? What technology has done greatly today is provide horsepower, CPU cycles. What it has not done is provide band width, which I see as the great theory to implementation of all these technologies. Software certainly hasn't kept up with the hardware. When we come full circle, when Larry Bowman can pick up a big red crayon and accomplish everything he has to do in a day by drawing a couple little pictures on a piece of paper, in a hundred, two hundred, three hundred years, then we're starting to come full circle again. Technology right now is the biggest barrier to technology implementation. As far as capital markets, were you asking are the people that invest smart enough to invest in technology companies?

MONEY IMPACT

The neat thing about the industry is money attracts brains. Every one of these conferences I go to, I get a dozen resumes from bright doctor level engineers that want to work in this business. They see there's a reward and people are excited about the stock market. It's the world's biggest poker game. This is a lot of fun. There's nothing more chaotic. There's nothing more competitive. There's nothing farther out there on the risk/reward curve. Most entrepreneurial people, the brightest people, do two things. They go into venture capital, the stock market, or they start the next leading edge companies. We want them to stay inside companies. We don't want them going off and raising money. That's a waste of engineering talent. If you can engineer, stay in there. The challenges are far beyond any of that. There's nothing new about market. People have been buying and selling for a few thousand years. They did it with stones, now they just do it..and then we'll just do it with bits in a few years. Don't send your resumes. Stay in engineering. The challenge here is so much greater.

Here's a company that were it not for that ability to grow capital...Microsoft pays below market wages and the reason why they can do that is one, Seattle's cost of living is less than the Valley's, and two, you go to Microsoft knowing full well that when your stock appreciates at the rate it's been appreciating, your salary is an irrelevant to your total compensation package.

INTERFACE PROBLEMS

I personally believe that it is significantly a paper interface. I'm trying to trigger a lot of work that was aimed at products so that we can get better use in there. There's a tremendous amount of Faxing going on within an American corporation and there's about sixty E-mails. People are either doing Hiragana or Katakana depending on what kind of a keyboard they have. For me, that's a bear to do. They'd rather scribble. Computers can't understand the scribbling. Use the FAX, use interfaces there. Paper is a wonderful interface to buy, you should just never have to store it or transmit it. Once you get that idea that it's really a screen and it's a temporary computing media, then good things can happen. There will be more paper in the world.

Let's talk about the interface problem. We're in a world where most users are stumped by managing the information environment on their stand alone PC. We are now making available to them the collected wisdom and ignorance of the world. Statistics say that with the current rates of growth, every man, woman and child in the world will have a homepage within three and half years.

The bottom line is that there are people who can't navigate their desktops and we're giving them the world. How are we going to cope?

Frankly, it's going to go via TV. WEBTV is the kind of answer that if all of you who communicate or want to communicate with your parents or PC challenged friends, we'd just go buy them WEBTV. You'll find that it's significantly easier than those VCR's.

One of the ways you can ease this transition is to build in intelligence in the actual appliances as they are used. So, a telephone that can look up phone numbers on the Internet becomes more useful than one that can't, or even one that can dial a number and ask someone. Those kinds of things are a driving set of entrepreneurial ventures going on doing the Internet interfaces that are sort of transparent to the user. And use one form or another, either Direct Internet Addressing or Java apps or whatever. Build it into the electronics.

COMMUNICATION

You have to have a need for them to communicate. Need has to drive this thing. You have to be able to buy this thing. You have to be able to install it. I got two cameras which one has already been given away because I've got a non-monogamous PC. Plugging that one in happened to blow the machine so I said it had to go away.

I can define the need very easily. My advice is get married. When I was single, my clothing was probably 30% more expensive because I would mix colors and they would bleed and I would shrink things. I'd love to be able to throw it all in the washing machine and have it figure out what color clothes I have and what temperature it needs to be. It's to a point now that some of the jeans manufacturers believe, merely on step prevention alone, they can justify putting a microchip in every pair of jeans they ship. We're getting very close to this kind of stuff.

Usually this stuff comes as things transition from wants to needs. The things that are very want-oriented tend to take a long time because people won't pay for them. The ones that are needs-oriented. Really, when the PC became a needed part of business, it really took off. As long as it was a fun thing to have, it wasn't as big. The things you're talking about are in the same category. You'll probably live without the washing machine that can sort out your clothes for you for awhile, but you may get to the point where it's almost impossible to get by without being able to get phone numbers looked up by your cell phone.

Chapter 14: Panel Discussion--Internet Appliances: Are You Ready for the Cyber-Consumer?

Moderator:

Dale L. Ford
Senior Industry Analyst
Semiconductor Application Markets Group
Dataquest

Panelists

Kevin Fielding
StrongARM Product Line Manager
Digital Semiconductor

David Limp
Vice President of Consumer Marketing
Network Computer Inc.

Raj Parekh
Vice President & General Manager
Volume Products Group
Sun Microelectrics

Jay Freeland
Vice President of Strategic Marketing
Spyglass Inc.

Van Baker
Principal Analyst
Dataquest

DALE FORD

Those who are seeking to develop the market for Internet appliances have clearly identified many of the key markets they are pursuing, ranging from the home, to businesses, to government and education, and penetrating into the international scene. There are significant barriers between where we stand today and where we want to go tomorrow, ranging from important issues related to the display and the input and output of information.

The entire user interface issue remains a very open question, and developing solutions that will be consumer friendly, to the adequate communications technology for delivering the content and providing sufficient interactivity.

Standards continue to emerge and play a critical role in shaping all of the new emerging markets that we see in the electronic marketplace. Certainly the Internet is no exception, as well as the underlying semiconductor technologies.

The issue of lack of interactivity with many devices today with the installed base of PCs or lack of interactivity with other Internet-oriented devices, or lack of interactivity with more than one Internet service provider stands as a challenge for this space.

KEVIN FIELDING

The Internet itself is a phenomenon. It's real, it's very useful, it's providing real utility. The second point is that companies like Dataquest and several of the other analysts are pointing to the huge opportunities that the emergence of Internet access devices will provide to semiconductor suppliers. The third point is Moore's Law still applies within this market, be it consumer or computer-oriented.

In terms of considering is it going to be the PC or the Internet appliance, we within the StrongARM Group have been engaged over the past, almost two years with a number of different companies, be they compute-centric, PC-oriented companies or consumer electronics companies. We're firmly convinced at this stage that the consumer electronics, the TV type approach to this, the appliance approach, will certainly dominate.

The PC industry is very frugal in terms of how much they'd like to pay you for the various components that go to put together a PC. Their margins are pretty thin. If you consider the PC vendors are frugal, then in comparison, the consumer electronics companies don't bear description.

Intel recently launched its most aggressive Pentium type product for the mobile notebook market, the Tillermook. The Tillermook comes in at about a 200 megahertz processor and fairly extraordinary in performance, and it does that with a reduced 1.8-volt which gets it to about 3.9 watts, under four watts of power for a really beefy microprocessor.

There are three things that I find are fundamentally important, making something accessible, usable and enjoyable to the consumer. It must be instant on and always available, it must be intuitive and easy to use, and it must be predictable. Every one of them must be fairly similar. That model doesn't apply to a PC today. It does apply to a phone.

Who's going to win from the semiconductor supplier point of view? Is it going to be the traditional suppliers of consumer electronics products? We believe that it's going to require a mix of best practices from both of these companies in order to be successful in the consumer space. It's not going to be either the computer semiconductor companies or the consumer-oriented people but rather best practices in both of those firms that will eventually be successful here.

DAVID LIMP

Technology only matters to the OEM. In the end, NCI is an OEM company. I want to have a brand like Intel Inside. I want to have a brand like Dolby. We sell to the list consumer electronics manufacturers that were listed up on the screen just a second ago. The Sony's, the Matshusitas, those types of people care about technology and cost. They care about what goes in the box. They care about the chips that you're worrying about producing, the microcode, the embedded OSs. The consumer could care less. The consumer cares about simplicity and content, and that's it. Every time a device has been successful in the computer realm, it has had content driving it, plain and simple.

The technology is still important. We're still producing browsers that fit into 500 kilobytes. We still want to make sure they're portable to every kind of processor that's out there because the industry that we are all embarking in, has quickly become a subsidized industry. Subsidized industries are tough businesses. It's a very, difficult market, and we're beginning to have to drive that. They begin to show it through things like content, added-on applications and ways to get into some of the premium services as we move forward. Those are all content-based, so what's going to pay for this is content.

Our technologyies are all based on open standards, whether it scales up or down. This is going to make the industry a success because Internet as a content or an application environment makes a lot of sense on TVs and telephones and other things. People aren't going to be writing, word processing, Microsoft Excel or Power Point for your television. Those are passive-based applications. They don't work well on TVs, and people that try them will fail. People are going to be interested in chat along with television. They're going to be interested in cool interface for Pay-Per-View, and those can all be written in HTML, Java script, in Java, languages that are out there today, simple to access and easy to get talent that understands how to author it.

The other big problem about interactive television was that there wasn't necessarily a delivery mechanism. We have another dynamic going on in the industry right now. Content becomes, if not free, much less expensive because of the Internet, but competition has taken hold. All of these people in many cases are competitors of each other, and they all want to get closer to their customer.

Consumer electronics companies are very interested because Digital TV is one of the things they must and will do over the next 10 years, and there's 600 million devices to be replaced. Phone companies and everyone else will do this because they see competitors on their heels, and they're looking for that way to get that much closer to their customers. Content companies for the first time see a business model that can conceivably work and pay for this whole infrastructure, pay for the hardware that was just talked about, because they have fundamentally competition breathing around them in every space, and they want to make sure they can write once and deliver many so that their content keeps up with all the other content vendors.

RAJ PAREKH

Sun is not in the consumer business. We have not made a single chip, single system, single piece of anything for the consumer market yet. We believe that the whole dynamics of the market is changing. Internet came around and that provided connectivity everywhere. Unfortunately, pure connectivity is not as useful. You don't know which computer is sitting on the other end of the line, which equipment is sitting on the other end of the line, which operating system it is running, which application it is running, which version of software it is running, which version of hardware it is running, have they loaded the latest ECO or not. How can you communicate? That's where the Java came in. It came from Sun. It was purely accident. Someone had to do it, and Sun happened to be the company who did it.

Java allows you to communicate from anywhere to anywhere, from anything to anything, and you don't have to worry about which processor it is, which operating system it is, is it of yesterday or is it of tomorrow, you don't have to worry about it. It is all standard, completely open, available to all, including Microsoft. The important thing is everybody has to conform to the spec.

I go to the computer at my job and plug my card in and it will bring up my environment, but when I'm traveling to Japan and my local office or someplace else in the hotel, I plug card and in a few seconds it will bring my environment up. That's the power of the Internet. That's the power of Java. That's the world of tomorrow.

In the future the companies and organizations will be successful who on one hand look at the history and not create $500 processors for the phone. At the same time look at what value can be provided and what kind of alliances we must do with others so that we can leverage on their knowledge, their experience, their intellectual property and combine our strength with their strength, and come up with a uniquely different solution. With the world of the Internet there is no existing company today that is perfectly suited for it. We have to migrate towards that. How do we migrate? We can either learn ourselves, or we can do the partnerships with others. That is one of the key essence for the people in the future are going to be.

JAY FREELAND

The move to digital television and the whole concepts behind that have delayed a whole set of things in terms of the marketplace. If you look at the PC industry, there's a lot of people that are very proud of the fact that the PC has 40% penetration, but it looks like it's not going to get much more than that. Sixty percent of the population aren't going to use PCs.

In the last year more new content has been created than in the entire previous history of humanity. What's driving this? It's a combination of the Internet and more. You could argue that a lot of that content is not particularly valuable. We see there is this massive dynamic that's happening which is that there is indeed new content and that if you can convince consumers that there is some way that they will either be able to save time in their lives or that they will be able to be provided new information or very importantly, that their children will have some kind of advantage, that's what starts to make this seem compelling.

The world is becoming more embedded. The kinds of applications that are happening are chips are appearing everywhere. Today's average microwave oven has more computer power in it than computers of 10 or 15 years ago. That's a certainty. What's happening is devices are being created, but how do these devices all communicate with each other?

The other dynamic that is very important is no operating system, no chip--we don't have the same Wintel world being recreated that was created for the PC. Critical mass will come from a whole variety of applications. In order for this market to be successful, it's about getting consumers information and getting them compelling content for whatever they happen to be looking at that particular moment.

The other thing that we think is so important is a lot of this work is going to be done in the infrastructure where to the consumer it appears to be seamless.

We try to integrate core technologies and provide value to our partners, and we work with everyone. We announced a relationship with Nokia for cell phones on Tuesday. We work with people like RCA Thompson and Worldgate. Consumer electronics companies are struggling with how do we get everyone onto the Internet on their TV or the Internet in X device. Cnsumers want more information, better, faster, an easier way to live their lives, so they don't have to go to a telephone book.

VAN BAKER

The consumer is very overwhelmed by technology at this point in time. We do measure some standard things like PC penetration in the marketplace. When we first started this in mid '95 we had a little over 27% penetration. Last fall we got to 36% penetration. We are interviewing as on this fall's data, but preliminary indication looks like we're going to get to 41%, 42% penetration of PCs into households. Last fall we had about 17 million of the 35 or 36 million households accessing the Internet on an ongoing basis. That's about 17% of the households in the U.S.--that's going up dramatically. If you look at the people who agree or strongly agree with the statement, it's 75% of the people. Theoretically, penetration could be 75% of the market as of last fall. The benefits of technology are there. The penetration is increasing and the consumers are interested in this, but they're confused by it. They're a little overwhelmed by it. Most of them are nervous about it. They also tell us they're concerned about the impact technology has on their lives going forward, and that's even true in PC households in addition to the non-PC households.

What we're talking about with these devices is the convergence market the coming together of computing technology and consumer electronics technology. Some of the things that we think are important to pay attention to in these markets is what we've got is two different markets coming at each other, one of which is very data-centric, which is the PC and computing base, and still even in the home the majority of this far and away is a data-centric, a work-oriented set of tasks going on with this device versus an entertainment-centric set of devices, which is a much less interactive environment, a much more passive environment versus a highly interactive environment on the data side. That needs to be taken into consideration.

We've got two different industries, a consumer electronics industry and a PC industry, that have hugely different business models that are approaching this market. We've got the consumer electronics industry that introduces new technology at a relatively high price point, which in the CE industry approaching $1,000 is a real high price point, and they ride the price point down and have long life cycles and they don't really hit typically very high volumes until they get underneath the $500 price point.

The PC industry has had a history until just recently of keeping that price point constant and refreshing the technology on an ongoing basis to keep the price point high. Very different approaches to business. That has to be reconciled in some form or fashion in order for consumers to be able to realize the benefits of this convergence.

Relative to prospects in the future, in the near term we're relatively bullish on Internet telephones. The reason why we're bullish on that is we think it's going to take a cell phone model. We think the telcos will use the screen phones to get services revenues out of the consumer and in essence give them the phone for free. Most of the folks working on the phones are at least in negotiations with the telcos about that, and we think they'll wake up to that, look at it as a revenue opportunity, and implement those that way.

In the intermediate term as we get home networking possible, home servers and Internet television become much more viable products.

In the longer term as we get bandwidth from the curb to the home, then the network computers being scattered and proliferated throughout the house and all of them accessing things over a fat pipe coming straight into the home that lets you customize your services and everything we think is very viable.

Chapter 15: Panel Discussion—Bandwidth: Who, What, When and Where?

Moderator:

Gregory L. Sheppard
Chief Analyst
Semiconductor Application Markets Program
Semiconductors Group
Dataquest

Panelists

Brett Azuma
Director and Principal Analyst
Public Network Equipment and Services Program
Telecommunications Group
Dataquest

John Armstrong
Dataquest

Don Miller
Dataquest

JOHN ARMSTRONG

There are a lot of interesting and innovative things happening in the networking space. There are some standards activities that are taking place from a number of different bodies that are dealing with IP issues, MultiCast, RSVP which is a reservation protocol which is a quality of service protocol. ATM development has had a lot to do with that. Asynchronous transfer mode is a networking access method that's been around for awhile that was designed from the ground up to support multimedia. We have enhanced local area networks now. We have a Gigabit Ethernet, which is 100 times faster than the old Ethernet, 10-megabit Ethernet, 10 times faster than Fast Ethernet. We have wide area network capabilities now that are very complementary to the local network, xDSL and ATM. ATM's actually a LAN and WAN technology. All these things are happening at once.

We also have another phenomenon that's occurring in the networking space and that is the predominance of IP or TCP/IP as the protocol of choice in the network. Ten years ago when traditional routers were being developed, they had to be able to support multiple protocols. New products coming out the chute, such as Layer 3 switches, don't have to do that necessarily. Most of them support TCP IP or IPX.

What's happening in the Intranet? Web server growth is increasing both on the Internet and Intranet and the software growth is increasing almost in a hockey stick effect.

That puts pressure on the local area network backbone. We have demand for bandwidth coming from sources that didn't really exist in any significant fashion a few years ago. A lot of push technology, is also entering the LAN backbone.

We see an inversion of the old 80-20 rule. The 80-20 rule was that 80% of the traffic occurred out in the workgroup, out into the user areas, out of the wiring closets in the LAN space and only 20% of the traffic was concentrated in the backbone. That has reversed itself pretty quickly and now we have the opposite situation. A tremendous amount of pressure on the backbone in the network and that's the network center, the confluence, the convergence of all the traffic that comes into the network.

Gigabit Ethernet is an extremely high-speed networking protocol. The standards should be completed by mid-1998 and we have ATM which has been available now for about five years commercially, which has actually matured in technology and is being adopted fairly widely in the backbone now by customers who are looking for technology that has had a presence in the market for some time. Gigabit Ethernet is not quite ready for prime time.

For the Gigabit Ethernet the pros and cons are that Gigabit Ethernet is really fast. We're talking 1,000 megabits per second and it's viewed as an extension of the existing 10 and 100 Ethernet. Same frame size fits in there. The big guys are not out with a full set of products to support gigabit. Also in Gigabit Ethernet we see added value schemes such as quality of service capabilities still being implemented on a proprietary level.

There are a lot of advantages to Gigabit Ethernet — migration, superior economics. Gigabit Ethernet is also absorbing ATM-like features. Quality of service capabilities—things that were previously the domain of the ATM product area are now being adopted by Gigabit Ethernet vendors. The challenge is for these vendors to make these added value capabilities standard across the board so that there's interoperability between products. Nevertheless, these products can be used to augment routes that exist today which, as we all know, are fairly costly and complex.

DON MILLER

The remote access market is tough. It isn't easy anywhere in the technology sector but it's extremely competitive. In 1996 we saw some shakeups in terms of a couple of small players as well as large players. Shiva and Gandalf may have been dealt the terminal blows in terms of what happened to their business. Bay Networks, a large player, was well established as a leader in the enterprise space but has also been impacted by a shift taking place in the industry from access servers to access concentrators.With the explosion of the Internet, the small office home office channel, residential users dialing into the Internet is also accelerating the growth and interest in this market.

If you're not in the top five because of the consolidation in this industry, you're going to be relegated to the walking dead. There's a certain amount of companies in the industry that achieve typically under $100 million in revenues and they plateau. I call them the walking dead. They never get beyond that state. Consolidation continues to drive the very top level companies, so you'll see a handful of very large players at the top that are going to be the big dogs and that are going to drive what the future is going to look like.

Outsourcing is viewed as a very big opportunity. This remote access stuff is voodoo in many cases. It's not straightforward, it's very labor intensive, you have to deal with a large user community and it's an IS organization's nightmare. You have to be able to communicate with them over a long distance and provide support.

As an ISP, how do you build a network that is going to be resilient enough to deal with these traffic patterns that are changing with demographics and various geography? Customers want consistent service across their network. They don't want hot spots where they've got a lot of problems, especially if they've got a high concentration of users. There's other mechanical things. Service agreements need to be hammered out. How do you bill? Some of the technologies to enable the deployment of these outsourcing agreements are referred to as virtual private networks. Most of this hasn't been deployed on a large scale and it's coming. This is going to emerge as a fairly viable market but not until 1999 and beyond, so there are going to be some pioneers.

The predominant amount of connections for remote access are going to remain analog modem of some kind or another, mostly the higher speed one. We'll see low speed DSL deployed primarily in the small office locations, a lot of remote offices and folks that are subsidizing telecommuting connections. There will be options. We're already aware of some companies that are going to offer DSL types of connections to corporations. Like in the Bay area, they'll go to 3Com, Cisco or Oracle, any of the large purveyors there that have a large population of people who telecommute. Standard Oil will be another one. Instead of offering them ISDN, they'll offer them a higher performance line at a more competitive rate. You're going to see this type of thing emerge in the not-too-distant future.

What's ahead? Deal-making. We see more consolidation and a new form of carrier class—remote access company. This will be a new form of competitive local exchange carrier that will emerge to take advantage of some of this.

BRETT AZUMA

Things don't move nearly as quickly as everyone expects. It's a somewhat inert market. Not inert from the standpoint that nothing happens, it's just that the market continues along the same trajectory and at about the same speed over a long period of time. It requires a fair amount of effort to create momentum to move the market forward.

A couple of things about the services market and what happens with the carriers. There's a law of three that describes the primary driver for massive carrier deployment of any technology. There are only three things that move the market forward. Item number one is regulatory fiat, some type of new regulation that comes out. Item number two is some type of operational cost savings and item number three is market opportunity. The first two and arguably the first one, regulatory fiat, have been the biggest drivers of evolution within this market.

Unbundling is what people have been talking about doing within the public network to enable further competition, to enable new service provides to enter the market and utilize the same assets that the carriers have used for years. In theory it ought to accelerate deployment to new services and in theory it should move things along. We're finding just the opposite. When you can't necessarily predict what's going to go in the plant, it becomes much more challenging and becomes virtually impossible—to maintain service integrity to your existing subscriber base.

There are two pieces to capital exposure. The cost of being wrong in this market is extremely high. If you ship product, like a PC, to a store that has relatively low traffic, you can actually redeploy that PC to a store that has high traffic at some point and the cost in doing so is really only transportation. If you place fiber to the wrong neighborhood and your take rates are real low, there is no way to recover from that. You can't redeploy the fiber terribly well because most of your costs are labor costs.

Distance is always a problem. If you take a look at the technologies that are being offered, you have about 15 kilofeet that go from the central office. There is about 30% of the subscribers that you would try to target that won't actually be able to get your service. It's difficult marketing a service when you can't offer it on a wide scale basis. People complain about that and you end up coming up with relatively expensive engineering solutions to be able to provide that type of coverage. The challenge in deployment, as we look at these things, is to provide for wide geographic coverage with relatively low density.

ISDN—last year's dead product is this year's booming product. A 40% growth rate over the forecasted period. Why is ISDN moving? It's got momentum. It's got product in the channel. It's got a large footprint of deployment. It's the only access service that's available that can give you anything above 28.8 on a somewhat reliable basis. The ISDN people are doing some smart things to enable their service. They're packaging CPE together. They're developing technologies that allow you to get to take e-mail and push technologies without running a four-figure usage bill every month. We believe ISDN has a lot of legs.

ISDN is not dead. ISDN has got a lot of legs to it. It's got momentum. We see ISDN growing for a very long time.

xDSL is an exciting technology and over the long term it's the ideal data product. However, it's not there yet. You can't order it. You'll see a couple of announcements for deployment this year but you're not going to see anything that looks like serious deployment until about mid-to-late next year.

Remember this: "Momentum precedes significance." It takes time to develop momentum in this market but we believe that significance will follow but it follows momentum, not precedes it.

Chapter 16: Panel Discussion - Major Device Trends and Forecasts: Looking Ahead

Moderator:

Joseph Grenier
Vice President & Director
Dataquest

Panelists

Nathan Brookwood
Principal Analyst
Personal Computer
Semiconductors & Applications Program
Semiconductors Group
Dataquest

Bryan Lewis
Director & Principal Analyst
ASICS Worldwide Program
Semiconductors Group
Dataquest

Jordan Selburn
Principal Analyst
ASICS Worldwide Program
Semiconductors Group
Dataquest

Jim Handy
Director & Principal Analyst
Semiconductors Memories
Worldwide Program
Semiconductors Group
Dataquest

NATHAN BROOKWOOD

If you look at it from a marketshare standpoint, Intel walked away with 96 cents out of every dollar that was spent on microprocessors that go into PCs, workstations and Intel-based servers—96 cents and everyone else argued over the four pennies that remained.

Revenues over the next six years are going to increase by a factor of three from a little bit under $15 billion in 1996 to almost $44 billion in the year 2002. It means that the average compute microprocessor is going to go up in price. That's slightly a counter intuitive thing since we all know that the prices of these things comes down on a regular basis. Anyway, there is a document that we have coming out at the end of the month that goes into the mechanism that allows Intel to keep lowering its prices while people pay more for what they're buying.

The non-Intel portion of that market—this is a split based on our projections of what Intel will be selling, which is the solid area and what the clones—AMD, Cyrix, IDT, others to be named at a later date—will be selling during the same time period. Clones will be taking more and more of the revenue. They will go from about $770 million a year ago to $7.7 billion in the year 2002. That's not a bad business. It's not quite what Intel will be doing at that point but still $7.7 billion is enough for people to chase that market.

Last year, Cyrix tried a pretty disastrous experiment early in the year. They decided that since they now had Pentium level performance, they could charge Pentium level prices. They have now learned their lesson. They are charging this very noticeable discount from what Intel charges. We would expect that Intel Inside will continue to command a fairly reasonable premium over the cloners.

We think of Intel as a semiconductor company but really, they're a printing company. Intel has addressed that issue to a large extent by accelerating the rate at which it brings MMX from the high-end and mid-range which is where it started out $400 or $500 processors down to the low-end $100, $150 processors. To blunt the initiative of the network computer—that $500 thing that really cost $700—they introduced what they call the network personal computer or NetPC which is really a PC that they took the floppy off, put a lock on the box so you couldn't open it up and then put in your own add-in cards and said, "Here." Organizations can use this and heaven help the poor guy at his desk or girl at her desk who tries to install some software that MIS hasn't authorized because that's where a lot of the support costs come from in PCs is people get stuff in the mail and they install it on their machine at the office.

AMD had been suggesting that everything was under control, the design was complete, performance looked good. They launched the product in April. It really did look like a very competitive product. They die is smaller than the Pentium II, the performance was arguably in the range of the Pentium II. They were going to do 1.5 million in the third quarter. They ended up doing about a million, throwing away half a million units literally. They didn't work at all. That translates to $100 million in scrap as opposed to bottom line. Really disappointing. They're trying to do $2 million this year.

Intel is going to drive performance. They want to drive performance as high as they can to put some competitive distance between them and AMD and Cyrix, to encourage all the people who bought a machine a year or two ago to say, "This machine is too slow, I need another one."

AMD's goal is to scale the K6 performance to match what Intel is doing as best they can. They also better address their yield problem so they can ship something or else nobody will care about their performance. They're going to increase their performance by adding a level to cache onto their chips.

The new ingredient in the microprocessor scene for the desktop—compute microprocessors is competition. It's there, it's struggling but it's there. More than we could say last year. AMD and Cyrix have a window of opportunity. The degree to which they succeed will be very much a function of how well they can execute over the next year or so.

BRYAN LEWIS

System level integration. You've heard a lot of talk about it over the last couple of days. What's the real key ingredient as to what's making this market start to take off?

The concepts of putting an entire system on a chip have been there for years but the market really hasn't taken off. One of the key ingredients in the manufacturing technology. We now have .5 micron, .35 micron and we're really moving into the .25 micron strong and we'll shortly have the .18's. With the .25's, we can really get in to having over 3 million gates on a chip. That gives us a lot of real estate so now it is practical to put the microprocessor on the chip, put the memory on the chip and the logic and really have the entire system on a chip. The manufacturing technology is the key enabler that just came around just recently.

There are plenty of applications. When it comes down to time to market, if you don't have design reuse, you really don't play. The emerging markets that are really starting to take off right now is a set top box, a lot of the wireless communication stuff, the DVD, — the digital video camcorders—they are starting to adopt technology rapidly because they must shrink the system, get the performance up. System on a chip is here.

When you look in the year 2000, you see that system level integration is over half the marketplace. This is in terms of revenues. If you're to look out in terms of designs, we'll see the crossover closer to 1998.

The core standardization. They're all working on standards. If you have different intellectual property from different vendors, in order to mix and match, you have to have some standards so that they can communicate with each other.

There is a battle going on. It's the system suppliers, the OEMs, the semiconductor manufacturers and the third party suppliers. Some of these can even be EDA suppliers as well. The bottom line is he who offers the values will get the profits.

JORDAN SELBURN

The concept of value in system level integration is really driving a fundamental change in the structure of the ASIC industry. If you go back a number of years, the value an ASIC company would bring would be in their service, the applications engineers people, people working on layout because the product itself was actually fairly generic, even though the name is application-specific. All the parts were five volts. There was no such thing as an embedded function. The value came from the people at the company helping you do your design. Things are changing quite a bit as we go into the area of system level integrations.

Value is a great thing if you've got it. It's not so good if you have to buy it from someone else. If you were to ask Compaq or Dell what they thought of Intel's value, privately, they're not so pleased as Intel is.

We've recently completed our 1997 user survey. One of the questions we asked is in your system level designs, what makes up the core area of the dye? For the current year, we see it's a split—about half random logic, about one-third memory's and one-sixth cores or system level macros or slams. Going into 1998, the memory stays pretty well constant but we're seeing a shift from random logic to cores. Cores are going up substantially—actually about 40% relative to where they were in 1997. This is a very significant room because the entity that adds the value in each of these areas can be different.

There is a battle going on for the control of intellectual property and the value that it brings. The OEMs want to use the intellectual property to sell more boxes. The semiconductor companies want to retain ownership of their IP and use it to get more of the dollars that the system designer was going to get. They don't care which semiconductor manufacturer or which system designer. They want to distribute it as widely as possible so they can capture the most of the value. VSI is moving in this direction where it's becoming a marketplace where people can buy and sell system level macros, particularly, mainly it's going to be the semiconductor vendors and system house doing the buying the third parties doing the sewing.

System knowledge is king, really understanding the entire system. We're going to look at each one of the vendors in this marketplace and look at a strategy. The semiconductor manufacturer—this is a heavyweight battle. This is no room for people who don't have deep pockets. It's going to be key to have very large economies to scale and you're going to have to be very efficient. This is a very cost-competitive market and if you don't have the cheapest silicon, then they'll go look elsewhere. They're going back to this having this application focus. This even applies to the semiconductor manufacturer, even if he's just a pure manufacturer and doesn't do design. Each one of these application markets has different elements requirements. If you have DRAM on there, that'll have a certain set of mass and you have to be able to do that. If you have flash on there, that's another set of mass. If you have analog on there, that has incremental mass sets. You really have to be flexible to move with this marketplace and really understand where your customer's needs are.

The system vendor or OEM vendor has one of the bigger challenges in the marketplace. Everyone is attacked—the intellectual property, the value. They must focus on only intellectual property or value that differentiates their product. The focus has to be for all their internal people is to really differentiate their product. Innovation is going to be key because they're going to have more of a commodity type of product. They're going to have to find innovative ways to market that as well try to get more features in the marketplace. Distribution channels, can be a good way to try to find new avenues of revenue out there. Bundling with other complimentary products target at that application can be helpful too. That doesn't work for everyone. I think Microsoft has proven that bundling can be a little risky. Innovation and being flexible and watching out for garden hoses is going to be critical for some of these system suppliers.

JIM HANDY

Does it makes sense to look at the DRAM market as an embedded market that's going to be disappear? I have the misfortune of watching over the DRAM business during a time of probably the biggest slide in history. A lot of people are staring at what their business is and all this fab investment that they've made and they've said a) are we going to be able to hang onto this market or is it going to become completely embedded—which a lot of trends indicate that certain parts are becoming embedded and I'll go over those in a minute. Then b) is there any way that I can capitalize on embedded DRAM to take advantage of the fact that I've got all this excess capacity?

The embedded SRAM is a good alternative to DRAM. There are a number of good alternatives to DRAM. Embedded SRAM is extremely well understood. Anyone who can make a logic process can make an SRAM. It can be manufactured on an ASIC process using a standard ASIC six transistor cell and it's much faster than DRAM. If you use an external SRAM as a competitor in competition internal embedded DRAM, it still might be cheaper than using the internal DRAM.

There are a lot of good reasons to embed your memory. You can cut the system costs possibly because DRAM vendors don't want support the old low density DRAMs. If you don't need four or 16 megabits which are the current generations for DRAM, then it might make sense to use a one megabit internal DRAM inside your chip. It's like motherhood and apple pie. You could also reduce the number of I/O pins you have devoted on your ASIC. You can increase the bandwidth that you get out of the DRAM significantly in that kind of an application.

The average density of DRAM chips increases 62% per year. Moore's Law allows you a four times density increase every 18 months or so. DRAMs have been on that forever. It looks like there is a possibility that things could accelerate because usually every three years at a Premier Design Conference, I SSCC, R&D labs show the next generation DRAM. NEC waited only two years after the one gigabit DRAM was introduced before they did their four gigabit introduction. There is a possibility that we could be accelerating rather than just going with a straight 62% per year density increase. DRAM is being ported over into the graphics accelerator in a 64-bit wide signal path. The signal path could be increased in width however that would swell the size of the ASIC to the point where it would probably be no longer be cost effective.

For these other markets that we're looking at—DVD, DBS but also other set top boxes for wired networks—hard disk drives are being looked at very hard by Seimens, and some other companies. There is a possibility that because of the complexity of the process that as the average density of SRAM moves into the domains-- these static two megabyte sizes or whatever is required by these things—that the technology could move over to SRAM. That's something to watch out for. If SRAM does displace embedded DRAM in these applications, you can be guaranteed that there will be other larger applications that will want to use embedded DRAM. There are a lot of difficult issues though that have to be confronted.

Chapter 17: Panel Discussion - Will Big or Better Win the Flash Market?

Moderator:

Bruce Bonner
Principal Analyst
Semiconductor Memories Worldwide Program, Semiconductors Group

Panelists

Bill Howe
Vice President
Intel Corporation

Miin Wu
Asia Pacific
Founder and President
Macronix., International Co., Ltd.

Dan Auclair
Senior Vice President
Operations and Technology
SanDisk

Bruno Beverina
Vice President, FLASH Memory Operation
SGS-Thomson

Wally Maghribi
Group Vice President
AMD

Dr. Hyung-Kyu Lim
Senior Vice President and General Manager
Memory Division
Samsung Electronics Co., Ltd.

BILL HOWE

Flash memory has been an EPROM replacement used for many code stored applications. To trigger the growth of the flash memory business, we really have to break into the data market. There were some early predictions that flash memory was going to replace hard disk drives, even eventually DRAM. We've toned down our expectations from that from those original days and now are looking for various applications that will use and have an interesting usage model with flash memory but nothing at the size of the disk drive market or the DRAM market itself.

Handheld devices are not replacements for your personal computer but are adjuncts or companions or complimenters for your system. We see a combination of using your flash memory for both code storage and data storage. That is not a separate device to store your digital data while you already have a flash memory for code storage in your handheld device but one device that can do both. That's what we call this new emerging code plus data market.

Intel has introduced a new flash memory called Stratoflash, using multilevel cell technology, which gives you the advantage of a high density, low cost flash memory.

In this emerging code plus data market, we see Stratoflash providing the read speeds that customers are accustomed to for the code storage market with the low cost per bit, with the high density that they want for the data storage in one device.

WALLY MAGHRIBI

Today's market, driving flash, is code storage, about 95% of the market. This will continue to drive the market at least through 2001.

The flash market is totally different than a DRAM market. It requires multiple architecture. The flash market is a summation of several markets. One of them is mass storage.

Over 50% of the product being shipped today is shipped with single power supply. For the application Intel is talking about, mass storage, we believe NOR MLC will not win because it does not fit the market and every market requires different attributes. The data storage requires high reliability. Portable market requires low voltage and low power. Even video and data storage requires small sectors. Very few applications require 64 megabit.

AMD and Fujitsu are partners. We have grown our market share from zero in 1991 to 33% in 1996. We ask the customer what he needs and we build the product to fit the customer requirements.

MIIN WU

As a new company in the flash market, we create new applications and the new fields. By doing that, you have to create new market segment and then create the technology for that. Not only can we serve the code and the data storage but also we can make a chip very small and we can also make it fast.

BRUNO BEVERINA

Since their first appearance, we all knew that flash would be an enabling technology. I see a clear path to a widespread integration to enhance the responsiveness and the interactivity of the overall electronic appliance. The question is: "What is the segment driving the technology?" Global communication. The flash designer should learn how to better communicate the low energy with the speed, with the density in an environment of increased value-added at lower and lower cost. Cellular phones are the biggest eaters of flash. Later on, the expert system is really what probably will drive the future.

Flash memory is fundamental to giving more life to system. It allows them to remember what they are and what they should do. It gives to this system the life because the system code of the behavior can be changed by the system itself. Like life, it allows the system the propagation and the production.

DAN AUCLAIR

CompactFlash has won the first battle for the small form factor cards. Another new technology that we agree with Intel on is that double density is going to change the vision that we have of flash memory in the future. It will drive down cost for data storage.

HYUNG-KYU LIM

For the mass storage market, cost is the most important factor. This application typically requires a millisecond program/erase endurance. Cellular handsetClower VCC capability is very important. For automotive, it's the reliability issue. The communication infrastructure requires high density and high speed flash memory.

We can utilize all the events of the DRAM process technology and the manufacturing facilities for those mass storage flash. We expect early next year to achieve $2 per megabyte with our 64 megabit chip. In late 1999, we can break $1 per megabyte. We don't need new process development for flash. To win the market, you should assess yourself by your customers. Big is not enough. You should be better all the time.

FULL TRANSCRIPT

Table of Contents

Chapter 18: Panel Discussion - Major Device Trends and Forecasts: Looking Ahead

Nathan Brookwood
Principal Analyst
Personal Computer Semiconductors & Applications Program
Semiconductors Group
Dataquest

Bryan Lewis
Director & Principal Analyst
ASICs Worldwide Program
Semiconductor Group
Dataquest

Jordan Selburn
Principal Analyst
ASICs Worldwide Program
Semiconductors Group
Dataquest

Jim Handy
Director and Principal Analyst
Semiconductor Memories Worldwide Program
Semiconductors Group
Dataquest

Chapter 19: Panel Discussion - Will Big or Better Win the Flash Market?

Hyung-Kyu Lim
General Manager
Memory Division
Samsung Electronics Co.

Chapter 1: Conference Opens

Gene Norrett
Corporate Vice President and Director
Semiconductors Group
Dataquest

Gene Norrett: On behalf of our 70 semiconductor analysts spread around the globe, it is my pleasure to welcome you to our 23rd Annual Semiconductor Conference. We are truly very glad to have you here at this fabulous location in San Diego at the Los Coronado Bay Hotel.

This is the closest that we will ever come to having a Dataquest semiconductor conference in Hawaii. In our questionnaire that we will hand out or in the evaluation form, if you'd like to write in that you'd like to go to Hawaii, write it in.

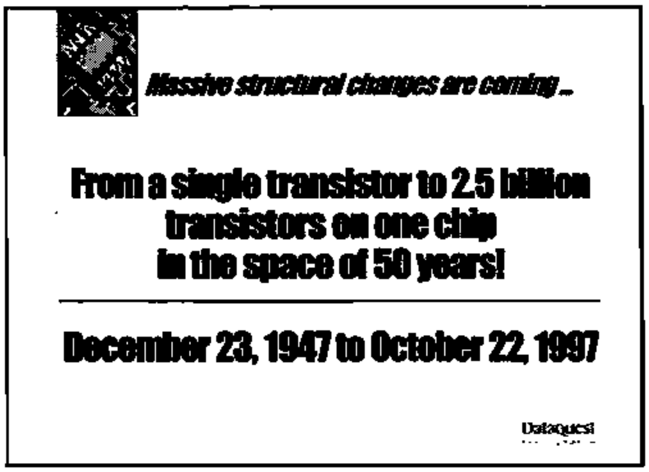

This year we have beat all of our historical attendance records over the last 22 years and over the next two days we expect to have over 500 people hearing about 50 experts in their field giving us their very informed opinion on the outlook for the economy as well as their slice of the semiconductor food chain.

This year's conference will examine the theme, "Massive Structural Changes Are Coming, Are You Ready?" from a wide variety of perspective and presentation styles. You will hear 14 industry executives give their presentations about the dramatic changes in their segment of the information technology industry. We also have eight fireside panel discussions featuring 40 of the industry's brightest minds who will be tackling the key issues that you folks told us that you wanted us to tackle at this conference here and we hope the sessions are very useful and informative to you.

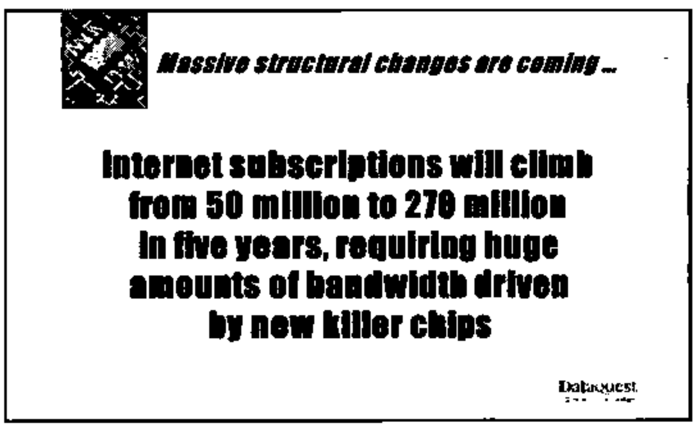

We again hope that this conference will exceed your expectation in terms of the content, the venue, quality of the delivery, networking, and of course, deal making. Please send us an e-mail message if you're able to do a deal. I'd like to keep track of these as we move forward. It seems like there's an awful lot of deal making that gets done and we don't hear about it.

At this time I'd like to thank our sponsors for their support of this conference. Their names are shown very clearly in the binder as well as on the tables outside this conference room. You all know these sponsors very well and we are very proud to be associated with them. For all of our formal sessions, both this morning and tomorrow morning, we're going to be meeting in the Commodore Ball Room. For our networking breaks it's going to be outside in the foyer as well as on the Bay Terrace.

We're going to be breaking down this room as usual into two sections. These afternoon sections will be, track one and track two, so it's going to be very important for you to take your binders with you at lunchtime.

The whole purpose, is to get to know you a

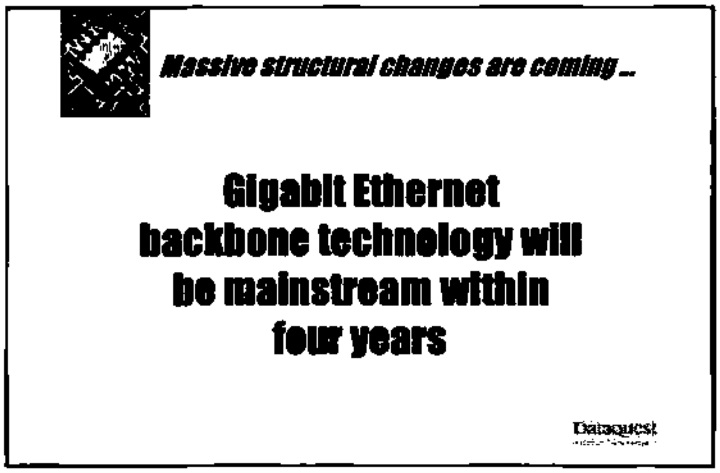

little bit better, to help you solve some of your critical issues that you don't get answered by calling up our analysts on the phone and just talking to them about their opinions on this and opinions on that. We have a very large team of five people at the conference room, our consulting group, that's very focused on helping you out.

Your feedback is extremely important to us, it helps us pick locations for conferences, pick the themes of the conference, and the issues associated with the conference so please fill out your questionnaire and give us a rating on our performances and then we will do our best to improve and help you have a better conference.

Chapter 2: President's Welcome

Manny Fernandez
Chairman of the Board
President
Chief Executive Officer and Director
Dataquest

Gene Norrett: I would like to now introduce my friend and big, big boss for the last 12 years, Manny Fernandez. This guy didn't follow Horace Greeley's directive, "Go west, young man, go west." On the contrary, he went from the silicon valley to Connecticut to reach fame and fortune as the Chairman and CEO of the Gartner Group. He served as Chairman since April of 1995, Chief Executive Officer since April of 1991, President and Director since January, 1991, and husband to Joanne for about 30 years now.

Before joining Gartner Group, he was President and Chief Executive Officer of Dataquest. Before Dataquest, Manny was President and Chief Executive Officer of Gavilan Computer Corporation and Zilog. He certainly is no stranger to the people who are also having the same problem I do, a lot of gray hair on the temples. He's been in this industry a long time. Manny is going to give us his unique and valued welcome. Please welcome Manny Fernandez.

Manny Fernandez: Good morning. It is terrific being here again. Last night was like good old home week for me. I started to hug people in the audience there, it's been like a long time since I've seen many of you and it's wonderful being here today.

It is clear that, as Gene said, this is the 23rd Annual Semiconductor Conference for Dataquest. We were comparing notes last night and some of us have attended some 20, 19, 17. I've seen a lot of people here. It happened here last night with Dr. P. June Min. He tells me he's been to a record 19 in a row. Maybe there are others of you that beat that record. It's been a long time and it's been great and over these 23 years we really have seen an incredible change in this industry.

Unquestionably we have gone from 3 inch to 300 millimeters. We have gone from 3 micron to .18 microns and we have gone from about $2 million fabs to about $2 billion fabs. It was kind of interesting as I was reminiscing last night when I had come to the first conference I had just finished building the first 3 inch fab for Fairchild in San Rafael in California. I remember Corrigan being all ticked off at me because we had overspent. We went over budget and we had spent $2.3 million to build that fab so there's been a lot of change over this period of time but a lot of things have remained the same. It's incredible that the fast pace, the change that has taken place, continues to be the same. Obviously the title for this conference is "Major Changes Are Coming, Are You Ready?"

I'd like to begin my quick remarks by trying to give you, at least from my vantage point, that this industry has really come an incredible way. After some 50 years, this industry, for the first time about a couple of months ago, the lead story on *USA Today*, the front page, was a semiconductor processing story. It almost tells you that this industry now is really truly in the forefront of American business and the world economy. There are going to be a lot of people here in the next couple of days talking about that same thing. I will leave it to some of the speakers that will follow me.

Before I turn this over to Gene and all of the wonderful speakers of the day, I would like to be able to talk to you about three or four different things. First and foremost, I would like to say thank you. We at Gartner Group and Dataquest, we work hard. We think we deliver the very best product in this industry, but it would all be for naught if it would not be because of you and your loyal and incredible continuing support that you have given us over these years. We really truly appreciate that. So from the 2800 people at Gartner Group, thank you very much for your continuing support of this corporation.

Gartner and Dataquest have kind of a very interesting vantage point that most people do not have. We today do business with some 8,500 different companies around the world representing somewhere in the neighborhood of about 250,000 different information technology and semiconductor professionals around the world that receive and communicate and interact with our 600 analysts on a worldwide continuing basis. It is that window, into how products are being implemented and utilized. That is rather a unique position to have that hopefully we can make it available to you and let you see through our eyes what is happening out there. Sometimes we, in the semiconductor business, get ingrained into what's happening in fab seven and do not look into what is happening in the user's user world on what is the pull through of that implementation that is going to eventually take effect at the semiconductor world.

What I would like to do today is I would like to talk to you about six items. Six things that we think are trend setting, things that are going to have a definite impact in the user business and therefore will have a significant impact in the semiconductor business. It is clear that information technology, meaning computing and telecommunications by themselves, are not the only areas that consume semiconductors. Of course, automobiles and consumer electronics, but when is consumer electronics will not be really truly part of IT and no longer just consumer. If you look at it, the computing and telecommunications sector of the marketplace, it is really the overwhelming piece that is going to effect your growth and definitely the bottom line of your companies. Let me give you, my version at least from where we stand, and look at some of the things that we see.

First and foremost, let's begin with the expenditures in information technology. Expenditures in information technology will reach over $650 billion dollars in 1997. That will represent somewhere in the neighborhood of 7% of revenue of all of the companies. 7%. Just to put this in some perspective, looking back the average IT expenditure used to be about 3.5% only four years ago. The forecast for the next four years, by the end of the decade, it puts IT expenditures over $1.3 trillion by the year 2001, $1.3 trillion, and the percent of revenues that are going to be spent in information technology will exceed 11%. These are humongous numbers. It obviously creates an incredible amount of opportunities but it also puts an incredible amount of pressure upon the companies who are now trying to invest that kind of money into information technology.

Let me put in the negative side of this expenditure. From 1997 through the year 2001, 20% of all of those expenditures will not be spent into any new technologies. It will be totally redirected strictly to fix the year 2000 problem. 20% of all of the expenditures will be redirected. That means that there is going to an incredible amount of trade-offs. Every chief information officer, every chief technology officer in the world is going to be asked to all of a sudden do more with significantly less because of the redirection of the budgets to fix the year 2000 problem.

By the way, it will be interesting for you to ask yourselves how you're doing. Most of the time we find that more of a technologist you are the less compliance we see. The problem is incredibly severe. I'll just give you another number. I wasn't going to talk about it, but now that I have brought it up let me just tell you that by December of 1999, we expect about 80% of all of the companies in the U.S. to be compliant. We expect about 65% of all the European companies to be complaint. If you take Japan, Southeast Asia, and Latin America combined we only expect 27% of those companies to be fully compliant by the year 2000. This is a huge problem and you should realize that that may not be on the top of your semiconductor list, but it is going to effect your place of employment so keep it in mind.

The trade-offs that companies are making is going to be kind of an interesting situation because what is going to happen is that many products, many projects, and many companies are going to be left behind in the cutting room floor. They will be trade-off victim by the Year 2000. So be aware of that that forces into who you do business with, who you form partnerships with, who those partners are that you are going to have, who your clients are going to be, who are you going to rely upon. A very critical element is going to take place so the first thing is the whole issue of mushrooming budgets but be aware of the black cloud, that 20% of it will be utilized as only to fix old problems.

The second major trend is that 60% of all of the existing software today will be obsolete by the year 2000. That means that every platform, 60% of all of the platforms that are either custom developed, or that are bought through a shrink-wrap process software, will all be obsolete by the year 2001. For most of us who are in the hardware business, this doesn't mean a heck of a lot because we have seen the shrinking life cycles of 10 years going down to 18 months in a year. In the software marketplace, this is a huge revolution. The average platform in software life cycle was 13 years, only a couple up to three years ago, and we are now forecasting the platform of software to shrink to about four to five years. This is a revolutionary new thing that is taking place and this again is going to take a significant amount of the budget that exists that will be reassessed or reallocated to be able to try to take care of the obsolescence problem that will exist. It's a problem that has not been handled, yet the second major place where trade-off of this $1.3 trillion dollars will go to.

The third trend, mobility. In the mobile environment today, if you take a look at every single dollar that is being spent in information technology, about 3% of all of the IT budget is going into mobile environment, to work outside the main place of work. We are forecasting that by the year 2000, that will be 25% of all of the IT budgets. 25%. That means that over $300 billion dollars will be spent in outside the office environment to be able to work with the mobile users. At the same time, we expect that over 50% of every single one of the telecommuters of today will never come back to the office. We expect over 50% of the telecommuters to be firmly relocated outside of the office space. These are significant trends and there are significant implications to the semiconductor industry as we enter the whole issue of mobility.

The fourth trend is a no-brainer. It wasn't a no-brainer three years ago. I remember when I came back after my four years of absence at this conference, I asked the audience, "How many of you have an Internet address?" and very few people at the time raised their hands. Today if I were to ask that question, everyone in this audience would raise their hand. The Internet, I told you at the time, will be a major inflation point of the industry. It will be the thing that will drive the next huge demand. It is clear that that has taken place. Last year alone there was 45 million Internet users. 45 million in a single year of Internet users. The growth that has been projected is 20% of growth per month. 20% a month for this next year. Estimations of having 270 million Internet users, which is a number that has been thrown around as a good place to begin, I believe is totally underestimated. We believe that the numbers into the 500 million to 700 million users by the end of the decade are very real.

The Internet has a lot of problems. We understand the problems. We still have a huge bandwidth problem. As a matter of fact, by the Year 2000, we are still expecting only 50% of all of the remote connections to be modem, so think of the incredible room here for growth. Unquestionably, security is a big issue in the Internet, the ability to create commerce on the Internet so there's still some limitations and unquestionably we have not even began the whole process of search. The whole issue of search is well behind the eight ball and it has to be if we are really going to be able to make the Internet the useful tool that it will be as we move into the next decade.

The fifth trend that I would like to talk to you about is financing. There are some issues here. In 1997 more money has gone into venture capital funds more than ever. 60 to 65% of all of the venture capital funds that have been invested have gone into information technology related companies. Which is a lot, 65%. Think about all of the medical and all of the other stuff that had a significant influx of capital, IT is taking 60%-60% of all of the financing.

In 1997, only 34 IT companies have gone public raising somewhere in the neighborhood of $1.6 billion dollars of IPO capital. In 1996, just to give you a reference, over 91 companies had gone public in the same period of time, raising over $5.5 billion dollars. The exit strategy is not matching the entry strategy so keep that in mind. Life always come back into some kind of a balance so it is an important piece, another trend for you to keep in part of your mind as you're continuing to do your planning process.

The sixth trend is the whole issue of industry consolidation. It is clear that we have seen an incredible massive consolidation that is taking place. British Silicon, or is it World Com, no, it's GTE, no, maybe it's AT&T will buy MCI but there's one thing for sure, MCI is for sale. We have seen Compaq by Tandem, we have seen 3Com buy U.S. Robotics, Ascend buy Cascade, SBC buy Pacific Telesis, and Cisco and Microsoft buy the rest of us.

It's clear that there is a race or a massive industry consolidation that has taken place. I believe that there are no fences around any of these things that just fences semiconductors out of this equation. I believe that you will see this drift itself to the semiconductor business equally as well, it has to. It is clear that the days that people could go at it alone are over. We are starting to see enemies cooperating, people that never even talked to each other. I remember being in the semiconductor business and some of these people would not even be in the same room together. They are building fabs together now. This is a hell of a business but it is just the beginning.

Dataquest
A Gartner Group Company

It is clear that the consolidation of the industry will move itself into the semiconductor world, it has to.

Let me close by summarizing what I believe are the three biggest opportunities. Clearly the whole issue of the Internet and its pull through capabilities will have a profound impact in what you do and in your lives of the future. The whole issue of bandwidth, bandwidth extensions, and the whole telecommunication opportunities is probably the biggest niche that we have seen that is going to effect all of us. Mobility and the distant worker will bring all kinds of opportunities to be able to expand and extend many of the product lines that you deal with today.

You as an industry, I believe, we as an economy together, we have some issues to deal with. Capacity and the industry balance, and I'm sure Gene is going to speak about this in his remarks, the whole issue of technological breakthroughs, race. Two years ago if we would have asked, "How many of you had a new program going," and all of a sudden everyone has one, right, so there's a race.

The whole issue of financing, cash flow, and consolidation. How do we bring all of these things into balance? Is there going to be enough money to be able to continue to move forward into this incredible industry that you have created? With that I would like to introduce my friend and partner, Gene Norrett, the MC for the rest of the day today and again, from all of us at Gartner, thank you very much.

Chapter 3: Technology and the Global Economics

Donald H. Straszheim
President
Milken Institute

Be wary of the trends and the issues that are involved. Now, I know what you're all saying. You're saying, "Oh no, an economist. How did he get in here?' Well, it is true that we economists have as bad a reputation as politicians, probably as little personality as accountants. Note I didn't say as computer jocks. Now that's called playing to the audience but what I want to do today is talk about how technology is changing the way we live and work and some associated issues.

My agenda is to first talk a little bit about the computer industry in 1967. A lot of people in this room weren't even in the computer industry, weren't even in business at that time but a historical perspective is interesting. Then I want to talk about a host of what kind of impacts of technology that you folks drive and what it's meaning to globalization, to retailing, to how the rich are getting richer, how it's effecting the regional location of economic activity, how it's a part of cost cutting and restructuring, and on and on. At the end I'll try to leave time for a few questions if we've can do that.

Ten years ago I was sitting in my Merrill Lynch Chief Economist chair when the market crashed. I wasn't responsible, but I was sitting there sort of watching that. Last week I had an occasion to give a couple of presentations about a reflection on the market crash of 1987. In the process of doing that research, I was looking at some old newspapers from that time and so forth and I came across a document that I thought you'd find interesting. This is *Time* magazine's cover on January 3, 1983, in which they cited the computer as the "Man of the Year" if you will. I looked at that document and I want to read just a couple of things.

By the way, much of this document is in the handout that was on everyone's chair when you came in this morning. We didn't get it in time to put it inside the notebook but let me just read three sentences from *Time* in 1983. They say, "Indeed by October, 1983, the entire *Time* editorial operation will be using the latest generation of word processors." Now they thought this was notable enough to put on the very first letter from the publisher in *Time* magazine in January of 1983. One other line. I quote, "When the final figures are in for 1982, according to Dataquest, a California research firm, more than 100 companies will probably have sold 2.8 million units, PCs, for $4.9 billion dollars." That's called playing to your host. I thought if I could find a quote like that in there I ought to read it. The point is how far we have come and how primitive some of the stuff that you'll find in this 1983 *Time* magazine seems now.

Now let me go back to 1967. We got hold of a copy of a Merrill Lynch computer report in October of 1967 and these were the companies that Merrill Lynch had picked in October of 1967 as leading the computer business for the future. If you scan through that list of companies, at the very top of the list, IBM, is the only company that remains a real major factor 30 years later. General Electric, a great success story, hasn't made a computer for a decade. RCA, long since out of the business and so forth. Sparior, Rand, & Burrows combined and I could go on and on. You notice a whole host of companies in there that have evaporated in one way or another. University Computing, Farrington Manufacturing, Lease Co, Levin Townsend and on and on.

This phenomenon of investing in the computer industry in the late 60s. One of those companies, Scientific Data Systems, the stock skyrocketed, probably looks familiar. You can take that name off and put one of your popular names right now. Here's another, Mohawk Data Sciences. Straight up. They dwarf some of the movements you have seen in these businesses of recent. An awful lot of those companies now don't exist, some others have made a lot of money one way or another but the interesting thing to speculate about is this. You say, "Look, I know from that list of 25 companies that it's hard to pick the winners from the losers. That's probably a tough game so instead of doing that, I'll buy them all. I'll buy the whole industry." If that was your strategy in 1967, it still wouldn't have worked. The reason is the growth in the industry has largely gone to new entrants for the last 30 years so

Merrill Lynch in 1967 had it right. They picked out the industry of the future but the growth went to people who weren't even in the industry for the next 30 years.

Here is a list for today. A lot of people in this room are employed by companies who's name you're now reading and the question is — and I'm not going to answer it, the question is which ones of these are going to be the winners over the next 10, or 20, or 30 years? If we were to do this again in 2027 at a new hotel, how many of those companies would we look back at and sort of laugh like we have today. Now let me go on. A lot of people in this room have made a lot of money in the markets in the technology businesses in recent years so here is a list of the Dow Jones industrial average for the last 25 years, from 1971 to 1987. I've got a little note there, the August to October 1987 period when the market went from 2700 to 1700 surrounding the crash in October of 1987, looks like just a little wiggle in a sort of a broader context and technology of course has lead the way in this advance. Now let me give you one other just quick piece of historical perspective. This shows the market from 1928 to 1934. It hit 381 on September 3, 1929, then you remember the crash, it went to 198, half. That's like going from 8,000 to 4,200 in a period of a few months. Then by 1932, it was down to 41. That's like taking the market over a three-year period from 8,00 to 1150. My guess is if we go to 1150 on the Dow, I won't have a job but you won't either so we'll see what happens here.

To just give you one other perspective on how perspectives change, there was a period in 1973 when the market was 1050—it went down to 578 from 1973 to 1974. At that time, the mentality, if you went back and read *The Wall Street Journal* or whatever, people's mentality was, "I'll never own a stock again." The mirror image of attitudes in the last few months of, "I want more stocks." In the run up to that earlier period, there was a time in the late 60s, early 70s in which people talked about the "nifty 50." These are so-called one decision stocks. That is, you buy them and you keep it. You never sell it. In that, "I never want to own a stock again," period, this is what happened to the prices of a bunch of those companies. You can pick out your favorites. Disney went from seven to one. That's on an adjusted basis with probably the current stock price. Coca-Cola, from three to one. Automatic Data Processing, three to one. HP, three to two, and so forth.

What I think is striking about this list is how attitudes change among investors and these were the leaders of the next 10 years or whatever. Just like in the markets now people are picking out those companies that we all know and love that are going to be the leaders for the foreseeable future so here's some more. These were more than nifty-50 companies. A lot of these you don't hear much about even more. Here's even more of them, and again, some of these have fallen on hard time in one way or another. K-Mart's in that list and I'll talk about that in a little bit relative to technology. Here's the last part of that group, and Walmart's the last one on that list.

There are times in which attitudes change a great deal and it's often hard to know really what's going to be right around the corner.Silicon Valley beats Detroit. Let's talk about technology in particular, not just technology in the markets. This chart shows that in 1973 — cars is the white jagged line that goes flat across that chart. Cars were three times as important as computers in terms of industrial output. Now computers, and this is a very narrow definition, computers are 50% more important than cars. It is in this context that your industry is just utterly changing the way we live and work in the business landscape.

Here is another measure of the same thing. It shows industrial output indexed to 1977 and you see information processing straight up, you see nondurables manufacturing up a little bit, and primary metals production no higher now than it was two decades ago. That's what's going on and it is in this context that you should keep in mind what's going on in your industry. When I used to get together the analysts at Merrill Lynch and talk about the companies, I used to say, "The single most important thing that distinguishes a winning company from a losing company is technology." The winning companies in every single industry are those who understand the power of these new technologies. Those that know how to adapt those technologies to their own particular corporate purposes and those who proliferate the technologies the fastest, cut their costs, and beat up on the competition. Nothing is nearly as important as the development of these new technologies. To understand, to adapt, and to proliferate them.

Let me just give you a couple of simple examples of how this technology is intruding on our lives. You go to a PGA tour golf event. Five years ago the sign says, "No cameras." Now it says, "No phones." You go to a store and you make a purchase. You give them your credit card, they swipe it through the little machine, and make three phone calls. The first phone call is to the credit card company. Does Straszhiem really have any money or not? Should we sell him this sweater? The second phone call if that one's okay is back to the warehouse saying, "Send us another, we're down one in inventory." The third phone call is to the supplier, "Make another one and send it to the warehouse." All I did was buy a sweater. Three phone calls and we can see this everywhere.

If the rental car company at the airport doesn't have one of those little remote deals that checks you out right there when you're getting your bags out of the trunk, get yourself another rental car company. Your bags on most airlines now they're all bar code scanned. It shows who checked you, it shows what time and what your name is and all that. It's everywhere. You can see it as much as I can. It's impossible to get away from the office and quite frankly I think this technology is going to ultimately drive us all crazy because the office is no longer a place, it's a capability whenever you've got your laptop and a phone line you're at the office. Hotel on a trip, vacation, where ever you are, an airplane, and man's not made to work seven days a week, 24 hours a day, but that's what's going on around us.

Let's talk about another aspect of this technology — how it's changing, the companies, the factors in our economy. Here is the top 100, by market cap, firms in the country by 1955, 1975, and 1997. Oil simply represents heavy industry. I could've put all those up there or chemicals or coal or railroads in the 1930s or whatever. You see pharmaceuticals driven by technology, of course, finance, the whole computer industry and so forth.

Let's talk about one of these industries in particular, the pharmaceutical business. What you see is the number of companies in the top 50 market cap has gone from none to four to seven over the last 40 years. We show the total market cap of these companies and how large they are in all of industry. It's our view that the 20th century was the century of physics and you all know that in your business, and the 21st century is going to be the century of biology driven again by the technologies that you people generate. I want to talk about regulation a little bit in our economy. There are important regulatory issues that are in the news right now in the technology business so I thought I'd go back to an episode that happened earlier in the 90s.

What we've done here is used July, 1992 as a basis 100, the left-hand access of that chart, and I plotted the S&P 500 stock price in Merck, the leading pharmaceutical company in the world. From the summer of 1992 to the summer of 1994, the stock market overall went up 9%., it's a dark red line, and Merck went down 43%. In the presence of a material market Merck got crushed. What happened? Summer of 1992, remember your recent history, the political campaign looked like President Clinton was going to win, healthcare was a big issue in the public mind. After the President got elected and took office, we had this big hubbub about healthcare and it dominated the news for a year and a half and then finally those plans largely fell from their own weight for whatever reason about the summer of 1994 and people thought, "Well, that's not going to happen." Since the summer of 1994, Merck has made up most of the ground lost but Merck is a different company. All of the pharmaceutical companies are different now. There wasn't a single law passed, there wasn't a single bit of new regulation but that company is a very different company. They dramatically cut back on their research and development so if you look at this from 1992 to 1997 you say, "Nothing happened." It's simply not the case because Merck is a very different company now than it was in that earlier period.

I believe there's a potential for changes like this in the technology business again as regulation is considered threatened, or whatever terminology you want. Here's a simple chart that shows how high tech dominates capital spending. In 1970, high tech was 7% of overall capital spending, now it's 46% by this chart and this is a great understatement because there's a lot of high tech content in that low tech capital spending. The other 54% is a lot of those machine tools considered low tech that are digital controlled and all the rest. By the way, there's no software in here anywhere so this just shows you how your industries are dominating and revolutionizing our economy.

Let's talk about a different dimension. I call it "consolidation in banking." This chart goes back to 1950. There were 14,500 banks as of 1985, now there's about 10,000. What's going on here? The staff and computers of one bank can handle the transactions of two so you make two into one, cut the cost, cut the people, and you go do it again. It's gone from 14,500 to 10,000. Where are we headed? 1,000? 200? Canada's got seven banks. What do we need all these banks for? The answer is we don't and you will see driven by technology, a further enormous consolidation in banking. If you put up the chart that shows the insurance business, it looks similar, they are just about five years further behind. If you put one up of the big four accounting firms, used to be the big eight. If you put up securities business or the advertising business, any of these others, enormous economies associated, it's all part of the flattening of organizations as data is made available throughout an organization. You don't need all these layers of management so you're getting these pyramids that are becoming much flatter eliminating all these people. These companies have a much broader reach and are able to accomplish things in their own business that they couldn't before, a lot cheaper than before, a lot more efficiently than before, with less errors than before, and it's going on everywhere.

Here's another one. It shows that back in 1985, total retail sales in the red and mail order. Now this mail order is the low tech analog of electronic shopping. There's not enough data yet on the electronic side but trust me, electronic shopping is going to become a big deal. 1800LLBean.WWW.LLBean, and on and on. Some people say, "Well, I think this is just utterly change the face of retailing in America, electronic shopping." Some people say shopping is a social experience. Well, I guess that's right. Quite frankly, I've thought about this, and I have a long list of social experiences that I enjoy more than going to the shopping center. That jar of peanut butter that you're going to buy when you push your little cart down, you don't need to look at that jar. You know exactly what's in it. All you want is the right price and have it show up at my kitchen door. That's all you really want. People say people won't buy produce because they have to look at the produce.

How many people here remember the guy who used to be in New York on the tv early in the morning called The Green Grocer? It was about 6:00a.m. Anyone remember the Green Grocer? Okay, there's a few people. The Green Grocer used to be on for about two or three minutes in a little plaid shirt he'd hold up these peaches. "Boy, the peaches are great. They're just coming in from Georgia this time of the year. They're sweet and juicy and they're terrific." Then he'd hold up this old pear, "Pears are getting kind of woody, we're past the season and so forth. You want to buy the peaches now and avoid the pears. We had our pear season. www.green grocer. com." I'm going to have a standing order 4 pounds of fruit every Thursday and if he starts sending me those woody old pears, no deal but if he keeps sending me those peaches or whatever, that's fine. I'm allergic to bananas so my order is going to say, "By the way, always send him a nice assortment of the best fruit, but never any bananas. Straszheim is allergic to bananas. This next week I'm going to be in Las Vegas, zero, don't send me anything. Next week I'm having a party, send me 10 pounds." I think this is going to overwhelm the retailing business and it calls in the question the little micro-economy that grows up around all these shopping centers. You just watch what this does.

Another dimension, the rich are getting richer. On this chart there's three different bars. They're for 1970, 1980, and 1994. The one I simply want to show you is on the right-hand side, the $75,000 incomes and over, a big increase in the number of people at that highest end. This is being driven again by technology. Those who've got the requisite job skills are finding their wage being bit up and those who don't are falling off the bottom. This is inherently dangerous. I believe our country's most important problem, not just more important economic problem, but most important problem is the increasing mismatch between the skill mix of people that are out there and the skill mix that companies want to hire and we fail to address this problem at our peril. It's going to take the public sector, Federal, State, local. It's going to take private companies, the non-profit sector, the academic institutions all together to solve this problem.

You go to a ghetto in any city and you run into this 13-year-old mother carrying her little baby around pregnant with her second and you ask yourself, "What chance does this mother, this girl, have of ever being integrated into the mainstream of our society and really contributing to our economy?" The answer is, "What chance? Low." Then you ask, "How about her offspring?" The answer is, "Less." The technology I think provides the opportunity to solve some of these problems just as it is in the short run, aggravating some of these problems but too many people in my view are being left behind so we've got to figure out how to resolve this. Along those lines, I just put this chart in because I thought it was interesting. Related to education and what you see is the school year for 13-year-old kids. Taiwan, Korea at the top and the U.S. at the bottom. A U.S. college graduate has spent less hours in a classroom than a Korean or Taiwan high school graduate. Maybe that explains in part, contributes to some of these issues.

Let's talk about a different aspect of how technology is changing things. Here we see employment going back to 1970, index to 100, for the U.S. in total and for Sioux Falls. Sioux Falls is not your retirement haven. It's not your resort destination. What's going on is 1970 Citicorp, Citibank at the time, decided to put a credit card processing facility there. That was possible by the technologies, put it anywhere they want, and you see the rest. They had an available labor force and Sioux Falls has been for the last quarter century one of the fastest growing regions in the country. Now you've got Gateway and some other things as well, Sioux Falls.

You can find similar kinds of examples of this regional impact all over the country. I mentioned before, Silicon valley beats Detroit. Since 1980, the employment in Santa Clara Country California is up 300,000 people. Employment in Wayne County Pennsylvania is down 300,000 people. That shows what's going on. People can work where they want to live and Manny talked about mobility before. It is important in this process and nowhere near over. All these little what I call "freestanding cities" rapid growth. People go to Sun Valley, to Aspen, to Pinehurst, to Scottsdale. Hey what a great place. San Diego, what a great place. I think I'll put my business here. That's no longer really a pipe dream, that's increasingly possible and increasingly is a reality.

Another one. Here's inventory sales ratios in the manufacturing business. From 1958 through 1983, this thing has chopped up and down and since then straight down. That's because of inventory control. Every manufacturing firm has just-in-time inventory control features, every retailer has point-of-sale data capture and so forth. Also driven by technology. I think this says something about the size of future economic downturns of our economy driven by technology.

Globalization. This chart simply shows the share of trade in our overall economy back to 1929. For 40 years or so, from 1929 to 1973, it averaged 9.5%. Now it's up to 24%. Increasingly it doesn't matter where you headquarters is domicile. Companies' source, raw materials, intermediate products, labor all over the world. They sell all over the world. That's increasingly going to be the case in the future. One anecdote a year ago when I was still at Merrill Lynch, I went out to see one of our clients who thinks they're going to double their business in the next three years, got to put up a new plant, where should they put it. They'd narrowed it down to three places. The first place was right next to their current facility here in the States, the second place was India, and the third was Ireland. That's what's going on in our economy now.

Here's another one. There's a lot of information on this chart so let me just dwell on it for just a minute. The red bars show in immigration to the state of California from outside the U.S. from 1990 to 1996. Average over 200,000 people a year. The green bars show what we call domestic migration in and out of California versus the other 49 states. It shows this massive out migration from California during the period of the 90s beginning to slow and technologies growth was not able to keep California on a growth path earlier in the decade. There were other issues that drove that, but California beginning to do much better now. Let's talk about these green bars. What I show is for each of the different states from 1993 to 1996, I left a bunch of them blank because they were small numbers. This shows minus 363 in California shows an average of 360,000 people a year migrating out of California to the other 49 states. Then those pluses, 87 in Arizona, Nevada 54, and so forth, that's immigration to those states from the other 49. You see of course this massive growth in the Rockies. Then you see the continued out migration from the Northeast and you see the big immigration in the Southeast. So the regional location of economic activity is being changed dramatically by these technologies.

Let me end there, but first just give you a quick commercial about the Milken Institute and what we're doing and then try to answer some questions if we have the time. We're a not-for-profit economic think tank that studies issues we think are important to the future of the economy. Education, labor markets and jobs, financial institutions, capital markets, corporate finance, how technology is effecting the economy and society, globalization and trade, demographics, all these sorts of issues. If we can simply get our views out on these topics we think we'll have a more informed public, better public policies, and hopefully improved economic outcomes. That's what we're up to. If any of you are interested in the material that we've presented this morning, feel free to get in touch with us and we would be happy to work with you because technology is one of the areas that we think is more important than any other in the economy in the future.

Q: Do you agree with the current popular economic theory that U.S. worldwide productivity from technology is being grossly unmeasured?

A: The stock market seems to validate that proposition. I think the answer to that is yes and I'll give you a couple of simple examples. You take a money management firm. What are your mutual funds? Probably everyone here has some money in a mutual fund. Mutual funds are managing, in many cases, 6 and 8 and 10 times as much money now as they did 6 or 8 or 10 years ago with no more people. That's productivity gain and is largely unmeasured. The reason is in many industries where our government gathers statistics, they don't have good output measures. In the car business all you've got to do is add up the number of cars, and they change a little bit, the nature of the car, but it's a fairly simple proposition but in a lot of industries there's not a good output measure so they use inputs as a proxy. In those cases, the growth in productivity be definition is zero. So the answer is yes, I think there's a great understatement of our productivity growth and I think the stock market has suggested is a reflection of that.

Q: Technology is a two-edge sword. Will government and social institutions evolve fast enough to accept, adapt, and proliferate the technologies needed?

A: I think that's sort of a relative question. It seems to me that some areas in our economy we're adapting very rapidly and in other areas very slowly. I was going to make a point earlier when I mentioned Walmart and K-Mart, those two companies are good examples of disparate appreciation of technologies importance. Those two companies were about the same size 20 years ago, now they're not. I think one of them understood, adapted, and proliferated the technologies and the other didn't. The last point about this question, if you think about where in our economy they really have made a lot of progress and where not, I would cite in the where not category, government that in education still, the blackboard and piece of chalk, a technology of 150 years ago, still remains one of the key devices that's used. I think that's going to change rapidly in the coming years but it hasn't really as of yet.

Q: What's the implication of flattening corporate structures on income distributions?

A: I'm not sure that the flattening of corporate structures—I think implicit in the question is that's going to maybe steepen the income distribution even more. I'm not sure that that's the case at all. We are eliminating a lot of jobs from those corporations, layers of middle management, and since 1980 or so, with all of the rationalization of our economy and the adaption of the technologies, we have destroyed a huge number of jobs, but we have created far more. I would argue that one of the reasons that we're at an 8,000 market and one of the reasons that people see our economy as leading the world is because we are more competitive than anywhere else. So those people who got forced out of those organizations in the flattening by the technology are finding jobs in all different walks of life and many of those are very good jobs at current pay levels with great long run opportunities for income growth in the future. To me, the flattening is not so much suggestive of changing the income distribution as it is suggestive of how much more competitive our economy is with a lot of new opportunities for many people.

I appreciate the chance to come and talk to you this morning. I've got to leave but I hope and trust you'll have a successful conference for the next two days. Thanks very much.

Graphics were unavailable at the time of publication.

Chapter 4: Organizational Opportunism – Utilizing Change as a Competitive Advantage

John McCartney
President and Chief Operating Officer
3Com (U.S. Robotics)

Moderator: Our speaker today prior to the merger he was President and Chief Operating officer for U.S. Robotics. His recent accomplishments include growing the U.S. Robotics international operations from $6.5 million in 1990 to $176 million in 1995 and then stepped up as its President and Chief Operating Officer. John holds an MBA from the Wharton School and a BA from Davidson College. Please welcome John McCartney.

John McCartney: I'd like during the course of my presentation for most of you to try and keep in mind one of those first slides that Don put up about the leaders in the computer industry in 1967 because most of my remarks will be related to that slide.

First of all, I want to thank Dataquest for the invitation, I'm delighted to be here. It's always nice to come to San Diego, especially on the cusp of winter in the Midwest but I have to say there's always some trepidation in addressing an audience that significantly consists of our suppliers to 3Com. I wondered of course when I received the invitation what I did to deserve it. I can tell you that after extensive analysis, it does not appear to be a reward for our price negotiations with many of you in this room. However, my purchasing staff is hopefully wishing that the next invitation will be from a customer group and not a supplier group. It's also a pleasure to speak to an audience that has a lot of financial analysts in it that don't cover 3Com.

I've been a customer of the semiconductor industry for 14 years so I can shorten a lot of what you'll hear today on a semiconductor outlook into a few short remarks. Today we buy a lot more devices that are much higher quality and much more powerful at much lower prices than we did 14 years ago. Over the next decade we'll buy a lot more than we buy now at much higher quality and much lower prices than we do now. I'm not going to tell you which companies we'll be buying them from or how we're going to survive and what the devices will be, but in terms of growth, there's no question that the engine that you are fueling with technological innovation is going to continue. It's that growth and what it means to our organizations that I'm going to talk about today.

I want to talk to you not just about change in society and technology, but in the impact that change should be having — I emphasize should be having, on the organizations in which we work. I'm going to talk about three elements: organizational structure, leadership, and attributes and behaviors of successful companies. I'll try and throw in a number of examples, but many of my examples will come from my own history at 3Com and U.S. Robotics.

There's no question that the pace of technological change is accelerating with great impact on the society around us. Even high school students, while they may not be in school for very long each day, know Moore's Law and they understand in some fashion the impact that it has on them. Micro chips are every place from computers to cars to microwaves to smart cards. The real story is not just the advances in the underlined technology but how they're being used, particularly in communications.

Personal computers in the last five years up 100%. Cellular telephone subscriptions in the U.S., 227%. Internet hosts — 1900% just in the last five years. To follow on a theme that Don mentioned, do you remember when a day away from the office was in fact a day away from the office? How many of you, when I conclude and we start our break, will check voice mail and/or e-mail to find out what's going on? Communications and how it effects people's lives I think is the most dramatic form of a change that technology is bringing. This is a dramatic one for me. 2.7 trillion. That's the number of e-mail messages that will be sent this year. That's 15 times the number of messages handled by the U.S. Postal service. 2.7 trillion messages.

I can remember that if you were active in international business, on your business card you had your telex number and on your company stationery you had your cablegram address. Even five years ago it was extremely unusual to find e-mail, Web site hosts, pagers, cellular telephone numbers on your business cards. Now I'd wager that many of you have a business card that's so crowded with contact numbers, people have trouble finding your name.

AT. The original title for this slide by my marketing guys had this trademark. I don't really think it can be a trademark so we changed that a little but think of what this has come to mean in the last few years and how people in our businesses are using it. Companies like IBM, GTE, and Compaq are investing tens of millions of dollars in corporate image campaigns related to this icon or its cousin the "e" for electronic commerce. Tens of millions of dollars, not directed toward the sales of a single product, but toward image and solution.

I think it's easy to envision a time when the Internet access symbol will be more ubiquitous than McDonald's golden arches or even Coke's contoured bottle. The changes in PCs are even more amazing. The dramatic increase in power and performance and communications capability is driving underlying fabric of change. I just happened to read last night that the price performance ratio for computers over the last 30 years has changed by a 100,000%. I don't know whether the source is reliable, but think about what that means and if the next 30 years changes much how instantaneous communications will be around the world. This technological change, what's changing in the way we use things, shouldn't blind us to another important issue and that is how fast are our organizations changing to adapt to do this market environment. We individually seem able to change our behaviors to take advantage of these advances and products but our companies changing as quickly.

It's always dangerous to use equations when there are a lot of engineers in the audience, but I'm going to try anyway, venture into difficult territory.

Think about corporate energy with the focus on the most famous equation of law, energy equals mass times the square of the speed of light. Now rework this a bit to this one. Force, and here I mean competitive force, equals resources times velocity at which they're applied to market opportunities. Competitive force. The idea of competitive force is ages old. I'm sure that had he thought about it, Caesar's legions would have had some formula like this in mind. My favorite quote from Napoleon certainly does. On beginning a new campaign, one of his marshals came in and said, "Victory is assured, God is on our side," and Napoleon responded, "That's right. God is always on the side of the big battalions."

We have focused for years since the age of the robber barons on increasing competitive force by using this formula, but most of the focus has been on the "r," on the resources. Building up huge piles of controlled resources. Mergers, product extensions, new factories, new fabs, globalization, all spreading tentacles out building force with more and more resources even in high tech industries.

Think about what's happening in defense. Lockheed Martin, Raytheon Hughes, Boeing, McDonnell Douglas, focused on more "r," and not just manufacturing industries. Just in the last few weeks we've gone from Morgan Stanley, Dean Whitter to Coopers & Price Waterhouse and now Ernst & Young. All focused on "r", how we get more and more resources under control.

Closer to home, the same assessment has been applied to 3Com and U.S. Robotics. But increasing competitive force by adding towards resources is extremely risky. For one, the spread makes it more difficult to manage, and two, particularly when that increase in resources is accomplished through mergers an incredible amount of energy is focused on internal issues. Integrating companies. Doing things that have absolutely no impact on the marketplace or on competitors. Gaining more "r' has been the traditional approach, but it's really counter-e formula, where is the leverage? The leverage is in speed. The same resources applied with the same velocity give much, much more leverage, it's the power of the square. My point is that today's organizations, if they're going to be successful, if their name is going to be on the list of survivors 20 years from now, need to figure out ways to move with speed. Not necessarily bigger, although bigger can be good, but move faster.

Certainly, we can see the impact of speed if we look at sources of innovation in our industry. Why is it that David can beat Goliath? It's because David got in the first shot. He was there while Goliath was still winding up. It's why so much innovation comes from start ups. There have been loads of proof points in Don's presentation, let me give you one. I have a friend who is also a customer who's the senior executive of one of the value added networks in the world. Before he moved into the Internet business, he spent 25 years managing R&D organizations at IBM.

Let's go back to the early and mid '80s and think about the wide area network in the mainframe and the backbone. There was an emerging new technology, technology we now call "routing." IBM had everything that was needed to succeed. The products were already designed. They were working in the lab. There was a sales force to sell them, there was a brand name, and there was an understanding within this organization on how to explain this to customers.

Cisco was a start up with some really smart engineers and a lot of venture money. No other positive qualities to indicate they'd have success. Today Cisco has revenues of $6 billion dollars, a market cap of $50 billion, nearly half of IBM's, and unfortunately, is a networking infrastructure powerhouse that I have to compete with every day in the marketplace. Why? Because they moved quickly. They moved with speed. They had the advantage that so many start ups have and that they were focused on a single target and they could move quickly.

How many internal meetings do you think they had about whether to bring a router to market? How many internal meetings do you think they had about whether it would cannibalize some existing business? How many internal meetings do you think they had about whether or not this was a good idea? It was their only idea. It's what they came together to do, and they took very, very small "r" but they applied it with incredible "v."

Now I used Cisco as an example because it's been tremendously successful and it's close to home for me abut you guys could list 35 more. Structure can work against you and why? Because a lot of our companies are still operating on this model. The old network model. Hierarchical structures.

In the semiconductor industry in particular, a lot of the foundation in early growth came from supporting defense and government customers. If your customers are organized in a hierarchical fashion there's a tremendous incentive for you to organize in the same way. What does this represent? What does this structure, whether it's a network or organization, represent? It represents mainframes and dumb terminals. Data collected at a huge number of points, tremendous reach for data collection funneled up to the top where the data is collated, analyzed and decisions are made and those decisions or pieces of them are funneled back down. There's no empowerment at the bottom. There's no decision making at the bottom. There are vast resources that are not applied to the marketplace but only applied in data collection. As their structures get bigger and bigger and bigger, they get slower and slower and slower and it creates a culture. I'll give you an example of this culture.

I spent the last five years living in France and working in most parts of the world outside the United States. Here's what I think about France. It has a generally strong economy, a tremendous history of technical innovation, tremendous educational resources focused on technology, and God knows, a reasonably large and well protected home market. In France we had the first broad scaled deployment of interactive technology to a large portion of the consumer face, Minitel. We had Minitel in France long before the Internet exploded, long before PCs were used in the United States for communications tools — 15 million French families had an interactive device in their home, and yet today, Minitel is light years behind the worldwide Web. There's been no leverage from France Telcom to improve it. It's a whole process and technology is totally obsolete. Why? Because of this model. Because Minitel could only be used up and down, up and down, no peer-to-peer, no crossing. My suggestion is that we have to adapt our organizations to a much, much simpler structural model.

What have PCs done to us? What PCs have done is make communication across any point in the network possible and it's created an environment where if allowed decisions can be made down at the nodes, information can be shared. Tremendous amounts of power, tremendous amounts of velocity can be generated throughout organizations and not just at the top if organizational structures are adapted to use it. I'd suggest that the greatest change that PCs have brought about is just beginning now. If you think about it, what were the PCs of the '80s and even the early 90s. What drove the revolution? Certainly at the beginning, anyway.

Spreadsheets. Think about the first use of an Apple computer. The first use of a PC, multi-planner, Lotus 1-2-3. Not a communications tools. It's only in the last two to three years that PCs as communications tools have changed the way that we work and have started to generate those 2.7 trillion e-mail messages. They've dramatically changed the possibilities for organizational structure and yet how many of our organizations have actually changed to take advantage of this? You can communicate with your office in Beijing in a matter of seconds yet it takes you six weeks to confirm my purchase order. You can go out and work with your engineers, no matter where they are distributed around the world, on a project and yet it takes three weeks for the HR department to approve a raise or a promotion for a key contributor. Why? It's all culture and it's all organizational structure. There's no reason why it has to do that.

With today's tools in front of us, and think about how these tools are gong to change, we could change the "v" in our organizations dramatically. There are some companies that are doing this. Some companies have dramatically changed their focus in a very short period of time. There's one in this industry, Texas Instruments. Texas Instruments is a completely different company today than it was 18 months ago. I have no idea what the internal organizations and hierarchies have done and if speed has picked up, but certainly the speed at which a corporate transformation took place was dramatic. I can't believe that there's one of you in this room that competes against Texas Instruments in their various parts of business that isn't more worried about them today than you were 18 months or two years ago. That opportunity is there in front of all of us.

What we need is much, much more speed. Let me give you some assessment points, indicators of whether I think that there will be speed in organizations. What's the tenure of the senior executive team? All 20 years in the same company or as an infiltration of new blood? Are employees rewarded for failure if they took a risk that seemed reasonable? When you use cross-functional teams, are they quickly evaluated and is the evaluation fed back to the next group on a cross-functional project for improvement or is there just check mark so that they can go back to McKinzie who helped with the last reorganization and says that we use cross-functional teams. Do you use 360 degree feedback at every level of your organization including senior executive staff? Do you spend more time focused on internal organizational and development issues or more time focused on the marketplace and your competitors? Those are indicators I would say of structures that indicate speed is there or can be there within your organization.

Besides changing structure, what needs to happen? I would suggest that the next big challenge is leadership. I think that we have moved because of the power of technology from a phase where management is the most important thing to leadership becoming the most important thing. It is impossible for a manager, a director, an executive vice president, or the president of a large business unit to be involved in every decision and it is not difficult to find qualified managerial talent or to develop it. What is difficult is leadership because of all of these resources even focused at high velocity don't do anything but generate energy, not force, but energy unless they're well directed. I would suggest that a major responsibility for leaders today is directing that force.

I'm not going to go through a whole list of what managers are versus leaders are but I would suggest a few key attributes of leadership of business organizations and they relate to that vision thing. The senior leadership creates the vision. What is it that we're trying to do? If we don't know, we're in bad shape. It articulates the vision in terms that people within the organization can understand and it energizes employees to achieve the vision. It's just one aspect, but it's an incredibly important aspect and you know you can do a lot of this with e-mail. You do it through communication. You do it through constant interaction with employees at all levels, across, up and down, and sliced because that's what generates change.

George Bernard once said that progress is only possible with change and people who cannot change their minds cannot change anything. I would suggest that the ability to articulate this vision and move an organization forward is one of the key aspects of leadership today.

Let me talk to you about some of the changes that I see happening. If I'd been here a year ago I would have been talking to you as the President of U.S. Robotics. U.S. Robotics by almost any measure is an incredible success story in technology and the communications industry. We went from $7 million dollars in revenues in 1984 to more than $2.5 billion in our last 12 months as an independent company. We went from 4% marketshare in North America and zero outside North America in 1986 to more than 50% in the U.S. and more than 40% in Europe in 10 years. We were public for 23 quarters and reported 23 consecutive quarters of record profits. In today's environment, the most important thing, we met or exceeded expectations regardless of what the results were. Out of 5,000 companies over the last five years ending at the end of last year, U.S. Robotics ranked ninth in total shareholder return.

Just as a commercial I'll point out that 3Com ranked first. We were suppliers at both ends of Internet access, a business that Manny says and I agree, is going to explode and change all of us. Why did we in January we pick up the phone and call 3Com? This chart explains why. What we have is the last 12 months of revenue growth and we've only put in large companies.

The chart wouldn't be that different if we put in smaller network companies. On a vertical axis, how much in dollars revenue grew by in the last 12 months and the growth rate in the last 12 months. We could plot this a lot of different ways with market cap, with size, with whatever you want, but what this chart says is that a number of companies were making great progress in a battle for number two while number one, was getting further and further ahead at an accelerating rate and the entire networking industry was threatened by overwhelming leadership of a single company. An overwhelming leadership by a single company can be great for that company and very, very bad for an industry. That's why we believe that consolidation was coming and the best way to deal with consolidation was to grab a hold of it ourselves.

3Com/U.S. Robotics is definitely about resources but I suggest it was a lot more about speed and vision. The speed by which we get to this vision and that is the network. I'll point out that basically 14 years of hard work has been reduced to one big red hemisphere. The network is out there at the edge, the big balls, then the next layer, the access layer, whether it's from the land or the wind and we see them coming together, and then the core. What we believe is that to be effective in networking you have to play at all layers. You have to have major presence, you have to be able to take your customer, whether it's a consumer or a huge enterprise, from one end of the network and all the way through and out the other end regardless of the technology they're using or the route they want to go. We believe that companies that position themselves to do this will be successful in the networking environment of the future.

We believe that the networking environment of the future will be much, much different than the networking environment today. We believe that the term "pervasive networking" will apply to it. We are not a representative sample of the population in this room. We are basically, I bet, overwhelmingly 90% plus network enabled. Less than half of the people in the world have ever made a telephone call. The world will become network enabled.

Communication as we know it today will be extended everywhere over the next few decades and band-with will increase dramatically. We believe that if you're everywhere around the edge, that if you drive a network vision that's simple to use, incredibly complex in the backbone, but simple for people to use, whether that's an individual consumer in Malaysia or whether it's a network administrator at the largest healthcare organization in the United States, if it's simple for them to use and manage it will grow. We believe that multi-media applications are right around the corner. We know that this business will be global and we hope that it will be deregulated. That's the vision that we see of the network of the future.

Let me just throw out a few things on behavior, characteristics that I would suggest fit extremely successful companies. There could be lots. There are loads of attributes. I'm going to talk about three because they are three that I identify with and which we use internally. The first is ingenious. What we do is all about innovation and invention, the creation of intellectual property. If we don't create, if we don't innovate, then we will never be successful. Companies that are not satisfied with incremental improvement but expect and organize themselves for breakthrough are the ones that are ingenious and I think that ingenious is a key element to success. It's not just in technology. It's not just Netscape creating a browser. Think of Home Depot — tremendous ingenuity. Taking a formula applied to other parts of retailing, delivering it to the home improvement market and then doing it fast, store after store after store crushing little mom and pop hardware stores.

I would suggest that successful organizations are agile. You need a strategic vision but you need to be able to adapt it and change it because the market will change for sure. A couple of big examples of what I would consider agile companies? Probably what's recognized as the best managed company in the United States today, General Electric. I'd suggest that there have been more big screw ups while Jack Welsh has been in charge of General Electric than in almost any other company. Kidder Peabody, NBC, the jet engine turbine problems that they've been having. Big, big screw ups, but this is a company that is incredibly agile, that reinvents itself time after time after time and moved very, very quickly to do it. Another from a country not generally known for organizational innovation would be Burtlesman in Germany. From book distribution to record clubs to CD-ROM to a leader in Internet access and content and now to interactive television. A dramatic transformation of a strategic vision over a short period of time

Finally, relentless. There's a great quote in *The Wall Street Journal*. It's talking about Doug Ivester taking over Coca-Cola and his attitudes but the quote was from Ray Crock. "When you have a competitor that's drowning the thing to do is turn on the garden hose and shove it down his throat." That is the ultimate definition of relentless. Are your energies focused on your competitors in the marketplace? Is everything that everyone in your company does focused on winning in the market and when you hit a stumble, do you give up or do you move ahead? An easy metric for relentlessness. If you ever sit through an internal presentation that doesn't have a slide that shows how what you're talking about is going to help you defeat a competitor, warn the presenter once and the second time fire him because there's nothing that matters except defeating the people you're competing against and relentlessness is a key attribute in doing that.

If you don't like these three or don't think they're accurate, and certainly there are lots of others that can be used, let me throw out a proof point. Let me take a real easy one. The biggest business success story in technology, you might argue about technology, but the biggest success story in business over the last 15 years, Microsoft. Is Microsoft ingenious? Gaining control of an operating system they didn't invent? Convincing IBM to use that operating system as a platform, immediately propagating it across an entire industry and out and surrounding Apple, and then leveraging that in every conceivable fashion into marketing power. That's ingenious. You can argue about whether the products are any good or not but it's ingenious.

Agile. From operating systems to application software back to operating systems with graphical interfaces then operating systems for servers and now, I lose track, are we on the fourth or fifth Internet strategy. No remorse that the first one was wrong. No remorse at all, we're just on to the next one and we're going. Relentless, I don't think I need to provide any Microsoft examples about relentless. I would suggest that these are three attributes of successful companies and the companies with hose attributes are able to focus on velocity and use their resources to generate competitive force in the marketplace.

I'll close with a little story to generate this. I took my family a few years ago to a nature preserve down in Georgia. In the brochure describing the preserve that you got before you went out on the trip it described, in order to prepare people, particularly children, some of the things that you would see happen. In the preserve there were alligators and there were egrets. The brochure described essentially that you might see during the course of your visit an alligator devouring an egret and the explanation was that this was nature's way of calling out the species, of getting rid of the sick, the old, and the less aware.

That's the one I like, the less aware. If you worked for a company that has the capacity for a terrific change for moving for applying your resources whatever they are with great velocity, if you have leadership that has strategic vision, and if you are ingenious, agile, and relentless, I congratulate you and I hope that I'm a shareholder. If that doesn't sound like the company that you work for, I would suggest that you run, not walk, from this room at the break and get in touch with a head hunter because the alligators are out there and they will in fact devour your organization. Thanks very much and I'd be happy to take questions if we have time.

Q: About the size of the number of e-mails spent, and by the way, I thought what he said that was the number of e-mails sent across the nation I thought it was to me. Regarding the 2.7 trillion e-mails messages sent, do you have a sense of how much of the e-mail today is junk or what he calls low value messages?

A: I have a sense it's totally anecdotal and it derives from logging on with my son at home and logging on at the office and I would say an incredibly high percentage is junk with little value added. In fact, I think that this is a great area in which, and this one isn't necessarily leadership, but the theme of leadership could be applied. Just think of the increase in productivity and speed at which you could move if you got rid of all that crap coming through your machine and you generated a way within your organization to focus on your important value ones but I suggest that lots of it is crap just like what comes through the postal services.

Q: What is the networking industry doing to improve the Internet bandwidth, hence improve speed?

A: Lots and lots and lots of things. I was delighted, by the way, when Manny mentioned that the Internet was going to be a big driver and then we heard later that Gigabit Ethernet was going to be mainstream, that's terrific. There are a number of things. You have to look at bandwidth in a lot of different areas. First of all, right now there are some big congestion points in the backbone. Whole new generation of backbone technology is in the process of being designed and installed that will free up tremendous amounts of capacity.

In addition, we have the huge issue of data redirect in the central office. That's an issue now when we're using standard voice lines, congestion in the central office. In fact, probably the most public cognizance of bandwidth issues comes in Northern California where at certain times of the day you have a huge delay or inability to get a dialtone. Not where there's peak voice traffic but where there's peak data traffic. That has to be solved before there's broad scale deployment of things like interactive cable or DSL technologies. It has to be solved by putting Dslams into the central offices and redirecting those. In terms of commercial bandwidth, I believe that we will see substantial progress back in 1998 and early 1999 in things like DSL and cable modems but I would suggest that well into the next century standard analog at 56K downstream, probably 56K both ways by then, will be the man access method.

Q: One last question, why did the call go to 3Com and not Cisco?

A: Because we're alligators and at Cisco we'd be egrets. I won't go back on the slides because of time. The real reason is we felt there's going to be tremendous change in the configuration of the network and the way that it works. We believe, importantly, that power at the edge, and you have to think about the power that the new 3Com has at the edge, overwhelming marketshares in edge connections. We're the only company in the networking industry that touches customers. The only company. Any of you have Cisco routers on your desk? You have 3Com or U.S. Robotics modems. We touch customers and we believe that the architecture of the network will be such that the ability to touch customers will be very important in designing and selling next generation network architectures and that our fit in the promise in the future of 3Com is much, much better than Cisco. I want that $50 billion dollars in market cap. Thanks very much.

Graphics not available at time of publication.

Chapter 5: Tracking the "Food Chain:" Dataquest's Worldwide Outlook for Key Technologies

Moderator:

Gene Norrett
Corporate Vice President and Director
Semiconductors Group
Dataquest

Panelists

Gregory S. Sheppard
Chief Analyst
Semiconductor Application Markets Program
Semiconductors Group
Dataquest

Joseph Grenier
Vice President and Director
Semiconductor Device Group
Dataquest

Clark J. Fuhs
Director and Principal Analyst
Semiconductor Equipment, Manufacturing, and Materials Program
Semiconductor Contract Manufacturing Services Program
Semiconductors Group, Dataquest

Gregory Sheppard: Good morning. I haven't figured out exactly how to deal with this "food chain" analogy yet. We can take kind of a terrestrial view, looking at grass and cows, I guess. There's also a marine way of looking at it where you start out with protozoa and go up to bigger fish. Ultimately, you have these big fish that are predators. I guess sharks and...does anybody have any idea what those might be? It's an offbeat lawyer joke there.

It's my pleasure to talk to you about my favorite topic, which is trying to find out what's going to drive the chip market from an in-use or demand perspective, and trying to characterize how the semiconductor industry can position itself to best take advantage of that profitably. That's the hardest part, to make sure the profits are flowing.

My structure will be to start out at the 10,000 foot view and drill down into what we think are the hot applications that will be driving much of the above average growth for the chip industry, then wrap up with a few conclusions.

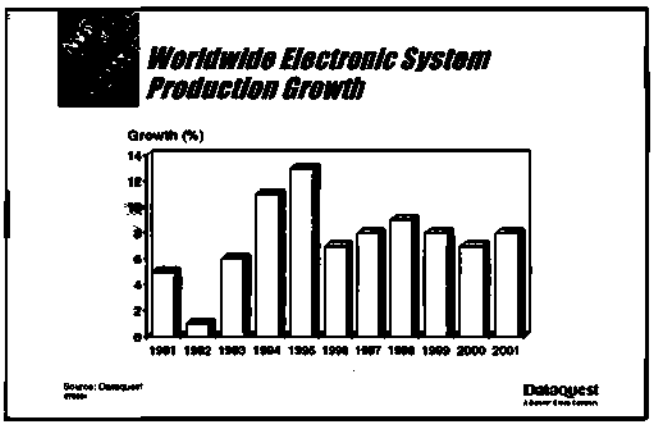

Perhaps I can summarize what is basically our core assumption about what is going to drive the ongoing demand for chips into the future. More importantly, how do we get this industry from the $150 billion level it's at today to a $300 billion level by the year 2001?

I am just a small proxy, if you will, for a hundred-plus analysts at Dataquest who are looking at the systems world, and we do have quite a few experts

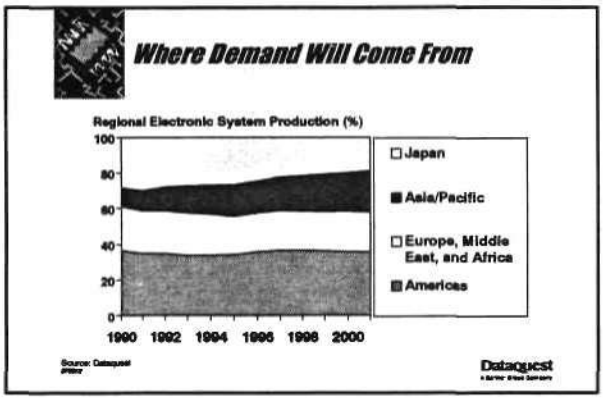

running around if you're interested in further drill-down information in these areas.

The most prominent one that comes to mind is the information or IT technology upgrade cycle. This is the idea that, back in the 60s you had an IBM 360 and that got upgraded by a 370 and so forth. Now we've moved into the client server model of computing and, in some ways, a lot hasn't changed although certainly the processing power has proliferated more and more through the desktop.

We are assuming that we'll continue to see cycles of IT upgrade. Just as a warning, the "Year 2000" thing is something to consider as a bit of a hiccup, where you're directing investment money toward software issues at getting the "year 2000" problem fixed. On the other hand, there are always new markets that seem to emerge to counteract this type of negative effect.

The Internet is going to five-times itself in terms of usage over the next five years. There are all kinds of numbers floating around. If you look at it one way, almost all of you probably have two Internet addresses, one at home and one at work. Something

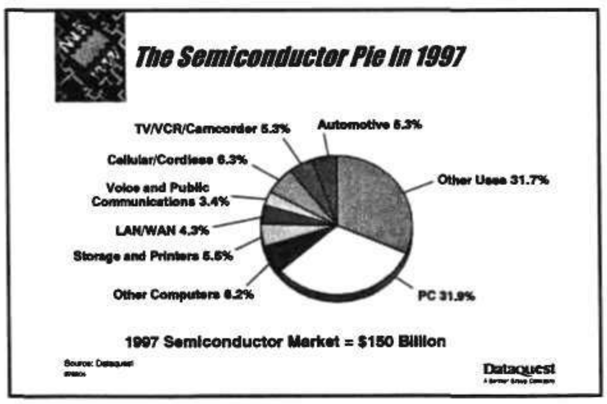

along a quarter billion users is not unfathomable in 2001, particularly as other regions of the world, moving outside the developed countries, kick in.

It's still very much a U.S. centric phenomenon to date. It still has a way to go and there are a lot of chips and equipment to be sold around the world to implement that.

Wireless pervasion: really what more can we say here? It's exploded. We're projecting that we'll see over 300 million users of wireless services of all kinds, cellular telephony, paging, the wireless data satellite systems, etc. This is certainly going to have an impact on the companies that are positioned in supplying DSP processors as well as RF IF circuitry in this area.

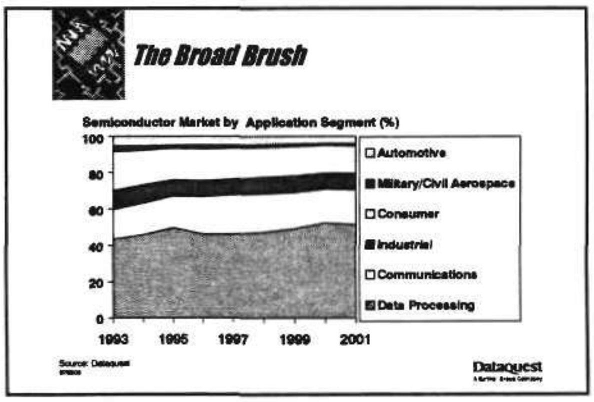

Wireless will grow to be about 1%0-12% of the market by 2001, basically up from 1%-2% in the early '90s so it's been a dramatic addition to the market.

Digital Consumer: This is really the concept of all the purchasing that goes on in the home. If you include the greater consumer spending and throw in PCs and communications technology along with the audio visual and interactive entertainment equipment, we're seeing spending that's starting to rival what's spent on food and clothes, the family car — something quite unthought of a few years ago. Likewise, the huge developing middle classes of Asia/Pacific, Latin America, and Africa remain greatly untapped, so there's a lot of head room, at least another 10 years of head room, if not more, in the consumer area.

One of the modeling methodologies we use for projecting the demand for chips is to start with the production of electronic equipment. On a global basis, this shows the growth rate we are projecting into the future and, as you can see, there was a great deal of build-up in 94/95 in the terms of production of systems. Certainly, there was inventory building and we also had currency translation propping up the numbers. We've come back down to a more realistic state in 96 and moving forward; but if you take the currency translation effects out of the numbers, we've been in a steady state of growth of electronic equipment production of around 7%-8%, starting in 96 and going forward to next year.

We don't see any serious hiccups in next year's plan for production of electronic equipment nor out through the year 2001. The only possible area that there could be some question marks around is the

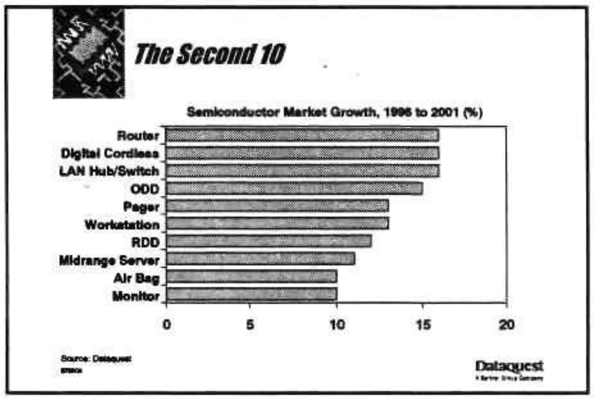

impact of the currency crisis we've been seeing in Asia Pacific with the newly emerging economies like Thailand and Indonesia; but, for the most part, those countries are production centers for export to the rest of the world, not necessarily having a huge impact on consumption, although they do factor in. It's the fact that they are producing for the rest of the world and are viewed as a shipper of chips. Determining where you need to route your chips to, this area will still remain a vibrant production area.

As this chart, which is basically an over-time pie chart, shows, the Americas region as a producing area of electronics remaining roughly flat. I should qualify what we mean by "Americas region." This includes Mexico and Latin America, and we are anticipating more and more production of electronic systems in those area. In general, we feel that what will be consumed locally will be produced locally. There will still be some areas in the commodity

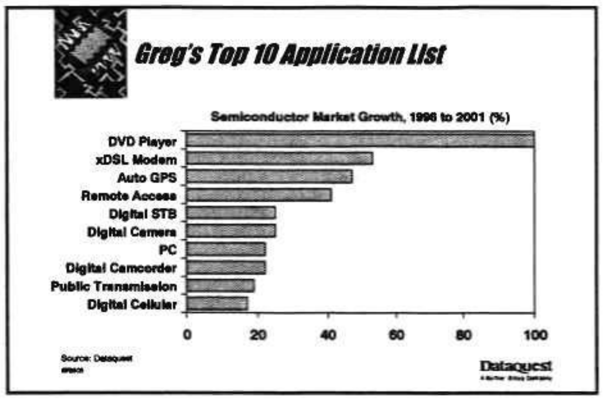

space, where the labor content is higher or perhaps there is an infrastructure in place to do so, that will keep production in a particular area versus another, but, in general, the rule of thumb is what's consumed is produced in the same region.

I've also set up this chart so you can see the trade-off between Japan and Asia/Pacific. As you can see, we are anticipating them to essentially switch places in the worldwide scheme of things as a production ground for systems.

Now if we model the production of this equipment back into the semiconductor forecast, and look at it from an application standpoint — in this case, it's the very broad brush categories — we note that data processing remains the largest sector and we're anticipating it will continue to grow. Of course, this is where the PC is, and the semiconductor intensity of the PC has a great deal to do with that. You can see a DRAM cycle going through the green area for data processing, recognizing the fact that prices did decline a bit there, and you can also see we are expecting another slow time for the data processing segment come the year 2001.

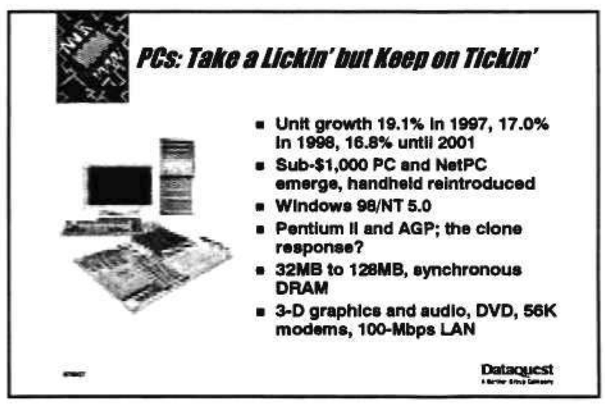

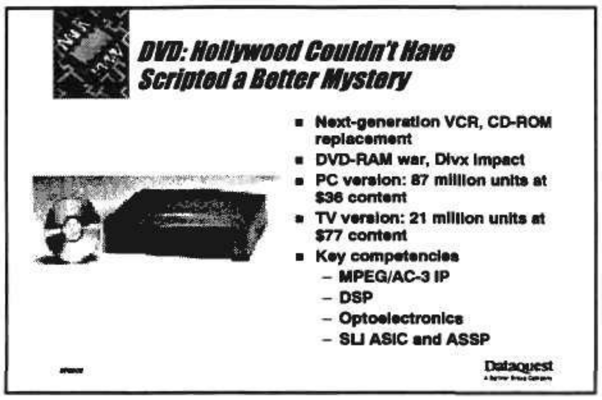

The second largest area is communications so this is LAN/WAN voice as well as wireless communications and public Telecom.

The third largest area is industrial, which is basically the collection of thousands of applications that measure and control. The medical electronics area is also included here.

This is followed by consumer, which includes audio, video, interactive electronics, things you find in and around the home but which are not a computer or communication device.

Then we have military/civil aerospace, which is wedging down from its glory days, particularly the Reagan build-up era which was quite an active market. In the early days, some of the initial IC projects were for military use. I suppose that's a good sign, that it's gone from 50% of the market to where it is now, in the low single digits.

The next area is automotive, which is quite active, but in the large scheme of things it still remains a small percentage of the pie.

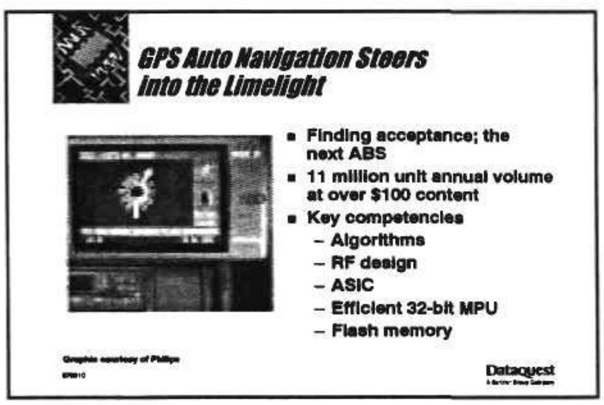

If we drill down a little further and map it into specific applications, we see that the pie chart indicates that the PC is definitely *numero uno* as far as being the largest application out there. This is essentially all the chips that would go into the mother board, after market memory, graphics, audio, those sorts of things that would generally apply to a PC. This is down somewhat from where it was in 95, for obvious reasons. You take nearly half the content and cut its price by 70%, then the BIP growth has been there but it hasn't been quite enough to overcome the overall dollar value being allocated to PCs. As you can see, the other applications are pretty

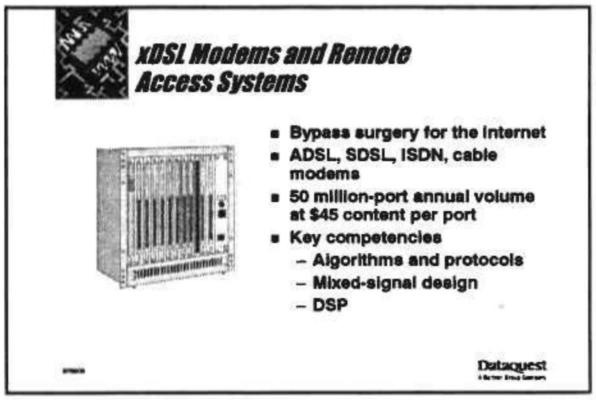

evenly spread. Most notable is cellular and cordless and other computers. This is the mid-range mainframe super computers, pretty much down in the 6% range, having flip-flopped with PCs over the last 10 years — still quite a sizable market and a lot of investment in new server technologies, multi-processing capability. Within each of these, there is quite a bit of complexity with aspects which are rising and declining. The prudent company is tracking all these.

If we drill down a little further and look out into the future at what are going to be the main hot applications, our semiconductor forecast is in the mid-teens. What are the applications that will be greater than the mid-teens in terms of growth rate? I've captured those in the next two charts, and I'll go through them quickly if you're not familiar with the terms.

DVD player: which is being positioned as something to replace the home VCR for playing back movies. This is now being introduced into the market place. I'll discuss this in my presentation later.

xDSL modems: DSL stands for Digital Subscriber Loop — just think of it as digital modem technology. It's higher speed and it's there to help us uncork the Internet.

Automotive Global Positioning Systems: used for navigation. I'll show an example of that.

Remote access: this is back to the telecommuter theme and is the type of equipment that would go into the business or central office to enable high speed web surfing and telecommuting, and digital STB, which is Set-Top Box digital camera, which is the still camera that doesn't use film.

Of course, the PC is a little further down on the list but due to its weighting, it's obviously quite important.

Digital Camcorder: another example of a piece of consumer equipment that's going digital, and with it, semiconductor content is rising quite rapidly as well.

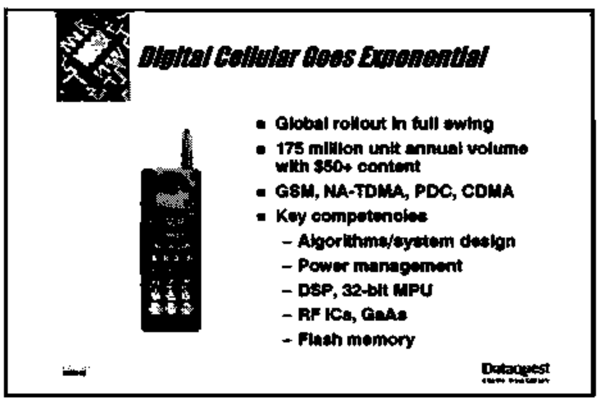

Public transmission systems: these are the infrastructure put in place in the public network. That includes a variety of systems that connect the local loop into the backbone of the telephone company as well as both the local and long distance telephone companies. In this area, obviously the Internet is impacting this as well. We're seeing a surge in the deployment of fiber optic systems. STH technology taking place.

We'll have a panel tomorrow that will address the topic of bandwidth in greater depth and we can explore these issues. On consumer items, we'll have a panel looking at the cyber consumer to explore those issues as well, and Digital Cellular rounds out the list.

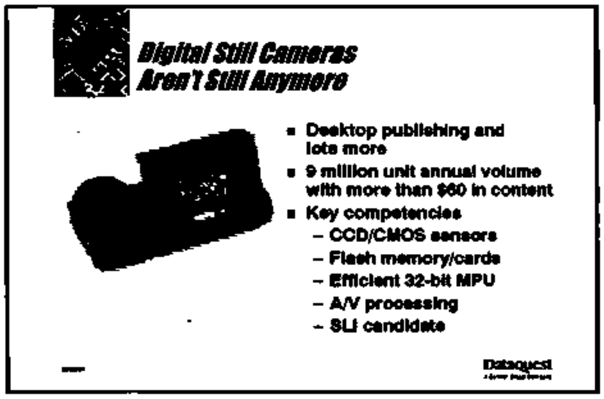

Routers, which we just heard the story of Cisco, it's still going to be a very nice market — 15+% range in terms of chip market growth.

Digital Cordless: these are the phones that are replacing the analog cordless that you find in your home. There are also versions like the personal handiphone system in Japan that can be taken around into public forums, such as this hotel for example.

Local area network hubs and switches. ODD is Optical Disc Drive. RDD is Rigid Disc Drive.

Airbag refers to airbag systems for cars. Even though we're starting to enter a more mature phase of deploying airbags, the fact that you are now having them on the side, on the roof — you'll be completely smothered in airbags at some point — still continues to expand the market. Lord knows, if Lady Di had sufficient airbags in her car, maybe she'd be here today.

Let's drill into a few specific applications. The PC holds about a third of the market. This year we have it projected to grow at 19% on a unit basis, with some slowing from the previous year. Next year, we're looking at a 17% growth factor, still very

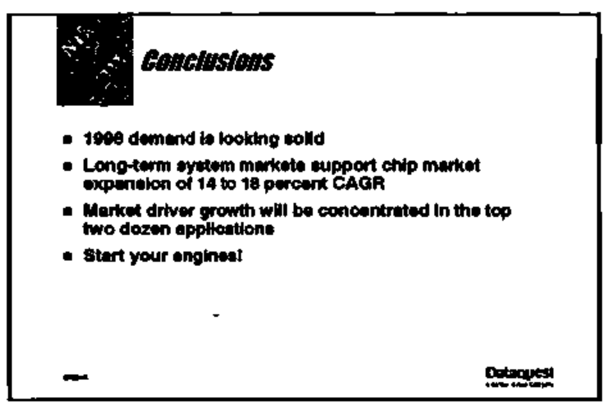

healthy unit growth. Our PC group measures this on a quarterly basis. They feel quite confident that we're on track. Then all through the out years, through 2001, we feel that something just a little bit south of 17% will be the average.

What's really hot out there today? The sub-thousand dollar PC is catching the world by storm with a lot of recent articles about success in this area. Likewise, we're seeing the emergence of the NetPC. This is targeted at companies. Essentially, it's a stripped down PC. In most cases, it has the disc drive taken out and is basically oriented to make system administrators happy in terms of acquisition costs and manageability issues.

We're on the cusp of Windows 98 being introduced — I guess we all hope it will be introduced in 98 rather than 99, but so far it looks like 98 is in the works — as well as the server version of Windows called MT in its latest generation, 5.0. These have significant impact on the chip industry in terms of memory content, processing power required, network infrastructure that can be built into a system in terms of network management, and drivers that are incorporated into the software.

Pentium II systems are rolling out with abandon now. There will be a panel later on in the conference that will look at the whole microprocessor area in much more detail.

AGP is Accelerated Graphics Port. That's now being incorporated into systems. It's a new, high speed way of dealing with graphics and, of course, the whole clone response in terms of Cyrex and DMD and IDT and what they will be doing to do battle with Intel.

Memory will continue to go up — no big news there. 32 megabytes today, heading toward 128 megabytes by the year 2001, and we do a lot of balancing of where we see demand versus supply in our group. Actually the quintessential thing that this "food chain" chews on is something like DRAM, where we're looking at how it's manufactured, supplied, and demanded.

Audio and 3D graphics are being rolled out with abandon over the coming years. DVD ROM: this is the read-only memory version for PCs. It's being introduced as we speak, and 56 kilobyte modems and 100 megabyte fast ethernet technology.

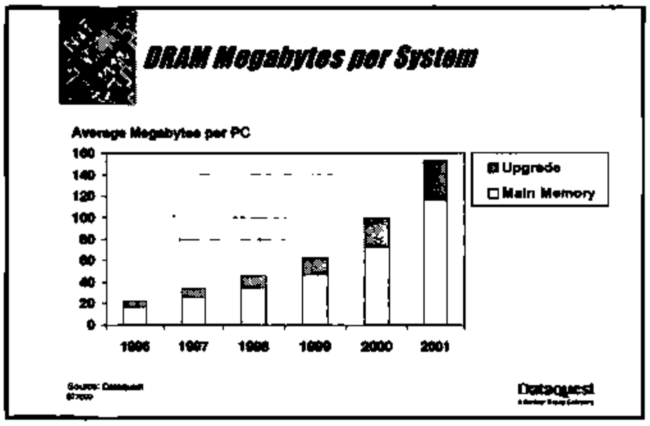

These are all the things that over the 1998 time period will be the big opportunities.

Just to take a look at some of the other hot applications, DVD, which as an acronym doesn't really stand for anything — you can use Digital Versatile Disk or Digital Video Disk, take your pick

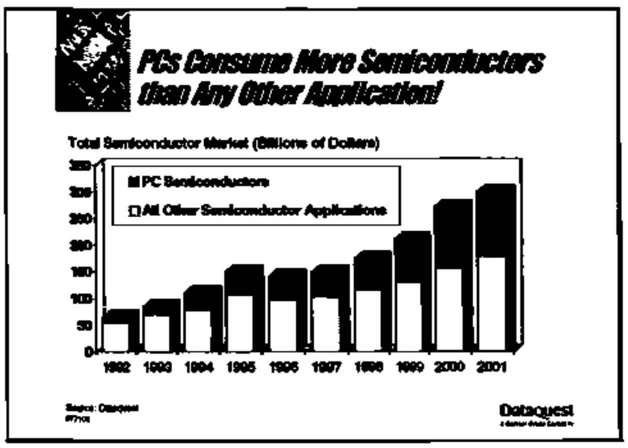

— it's really being positioned to ultimately replace the VCR and the CD-ROM. Its main advantage is that it has much higher capacity, also the capability to play back and compress movies.

There is a division in the ranks, however, where we've seen the DVD-RAM, or the version that's able to record as well, now take on various fighting factions, and it's going to waylay some of the progress.

There's a new technology called Divx, which is essentially a low cost DVD that can be purchased with a movie on it. It's encrypted and you're given an encryption key: you watch it once and throw it

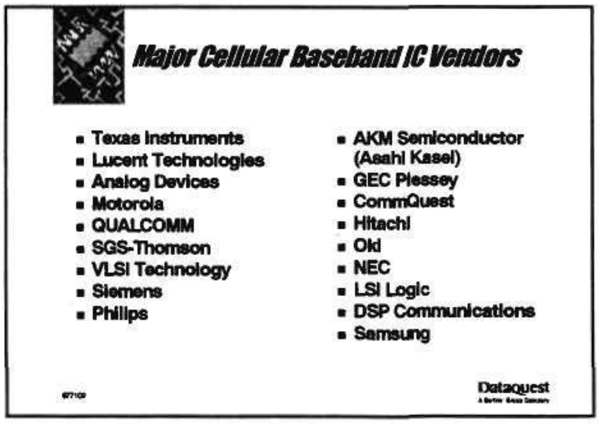

away, or you can actually buy the disc. That's a competitor to DVD.

I've noted year 2001 projections here: the PC version, 87 million units at $36 a content; the TV version, we're looking at 21 million units at $77 a content, so we're very rich in content. Key competencies I've noted here. If you want to be a player, you've got to have these capabilities.

IP refers to Intellectual Property, in this case, and I think it's important to point out that this is an application that's going to be best served by system level integration. The company that can put together all of the above elements, including the processors, so it can be a programmable solution, is going to be the company that's going to do the best.

Moving into the Internet area: we have the xDSL area where we have the modems, and we have the remote access systems that would in some cases incorporate these modems inside. This is essentially what's going to accelerate the 56 kilobyte on-and-off ramps to the Internet up to megabyte and-be-gone. It

really is here to rotorooter out the arteries for getting on and off Internet.

Basically, there is an alphabet soup of different approaches to this area. ADSL, which is an asymmetrical approach, takes a good old twisted pair of lines and makes them run ten megabytes per second or greater, and this is technology that's now in evaluation and early deployment in some areas of the world. Symmetrical DSL is SDSL. That's one to keep an eye on — one we're bullish on in that they've gone and addressed cost issues, and it provides upward of a megabyte per second capability still using the same twisted pair line.

Of course, the ISDN and cable modem are the other hot areas in this space. We could spend the whole conference talking about opportunities in this area. We are projecting by 2001 this will be a 50 million-port business. That's the way you look at it. You

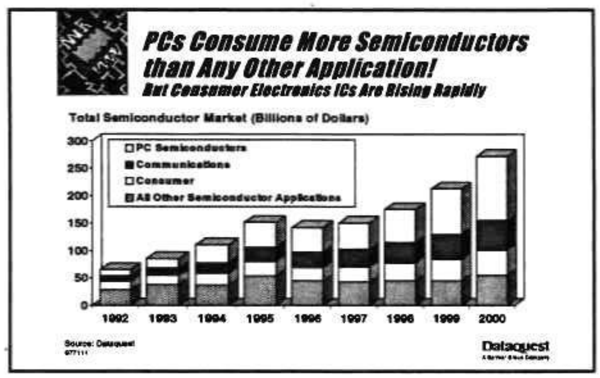

write things in port terms when you're looking at infrastructure, with each port having $45 worth of chip content associated with it. If you want to play, you've got to be good at the algorithms and protocols, mixed signal design, DSP — those are all going to be fundamental. Ultimately, all these combined in one chip will be what is required to play here.

To dive into another hot area, we have Global Positioning Systems used to help drivers navigate. I was pleased to see Philips actually running national TV ads with this system on it, so it's quite timely for this conference if you happen to catch it. It's now starting to take off, becoming either a standard or a standard option in many luxury cars as well as rental cars in the U.S. The big market has been in Japan and Europe to date. This will be an over 10 million unit market annually by Year 2001, with some hefty chip content.

Key competencies are noted here. This is an area that has even attracted some start-up company activity with companies putting together complete integrated solutions to implement the GPS engine, as they call it.

Digital Still Cameras are taking the world by storm. It's hard not to open up a *Time* magazine or watch a *MS/NBC* and not come across an instance of this technology. Currently, it's in the early adopter phase of purchase. Certainly, the desktop publishing crowd has latched onto it, as have certain vertical markets. We feel that the technology has to improve in resolution and the costs have to go down before it becomes a true replacement for the chemical, single lens, reflex film based market. Nonetheless, given a middle of the road scenario, we feel we could be

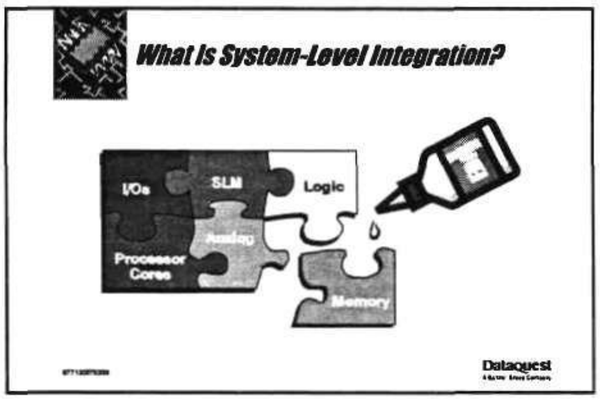

approaching 10 million units in the year 2001.

Here's the list of key competencies.

Low cost sensors: this is an area where CMOS sensors are being kicked around as an alternative to the traditional CCD. Of course, this is a big user of flash memory 32-bit microprocessors. This is also a system level integration candidate.

Digital cellular: what else can you say? It's huge. In the out years of our forecast, it's rivaling PCs in terms of unit volume, but with a lower semiconductor content, it doesn't quite have the impact of PCs.

Probably the new kid on the block is CDMA. We're here in San Diego, which is the world headquarters of CDMA because of being the key licenser there.

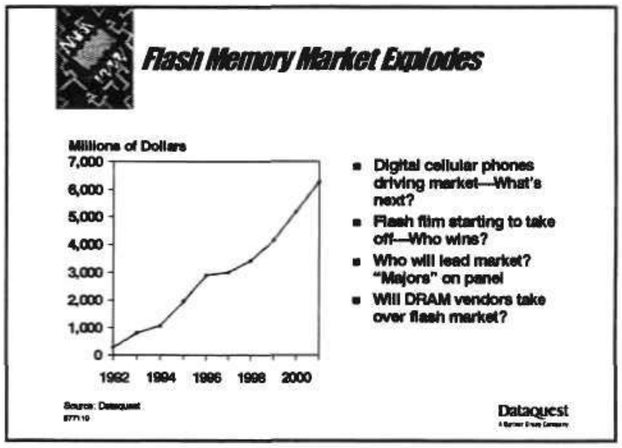

CDMA is going to be adopted. It already has been in Korea and it's moving to be adopted in Japan, as well. It is going to be positioned to replace PDC.

In key competencies, once again, it's important to note that power management in a battery-operated application — companies that can master that both at the hardware and software level — will be the ones best positioned. It's also important to note the trend toward the use of RF ICs.

To wrap up, basically 1998 is looking pretty good on the radar screen, with no major bumps in the road that we can foresee. There's a slight slowing of unit shipments in the PC. Of course, the main issue is if there is a worldwide recession, it will obviously impact this. Or who knows what *El Nino* or what weather related...not to blame *El Nino* if the forecast is wrong, but certainly there are other issues that loom.

Semiconductor Rankings

1996 Rank	1995 Rank		1995 Revenue ($M)	1996 Revenue ($M)	Percentage Change	1996 Market Share (%)
1	1	Intel	13,172	17,781	35.0	12.5
2	2	NEC	11,314	10,428	-7.8	7.4
3	5	Motorola	8,722	8,076	- 7.4	5.7
4	4	Hitachi	9,136	8,071	-11.6	5.7
5	3	Toshiba	10,076	8,065	-20.0	5.7
6	7	Texas Instruments	7,831	7,064	- 9.8	5.0
7	6	Samsung	8,332	6,464	-22.4	4.6
8	8	Fujitsu	5,536	4,427	-20.0	3.1
9	9	Philips	3,900	4,219	8.2	3.0
10	14	SGS-Thomson	3,398	4,112	21.0	2.9

Source: Dataquest
977125

Dataquest

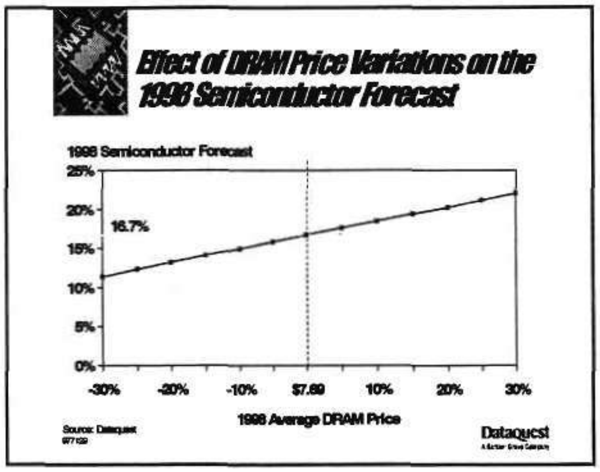

In terms of the bottom-up look at applications and how we feel they're defusing into the marketplace, we look like we're on track. The long term production numbers I showed you earlier, 6%-8% range, support a 14%-18% long term semiconductor market, barring supply-side pricing factors, of course, so manage yourselves well. Don't build too much capacity out there.

I've showed you what we think are the top applications that will drive a lot of growth. Certainly, there are some others, depending on how you categorize and if you look beneath the surface, but I think we've captured the bulk of them in this list.

Final comment; if you're not playing now, you better get going, and there's still room to get in. Within each of these applications there are new waves of innovation that create a window for players to get in there.

Joseph Grenier I'm going to take up right where Greg left off and continue down the food chain.

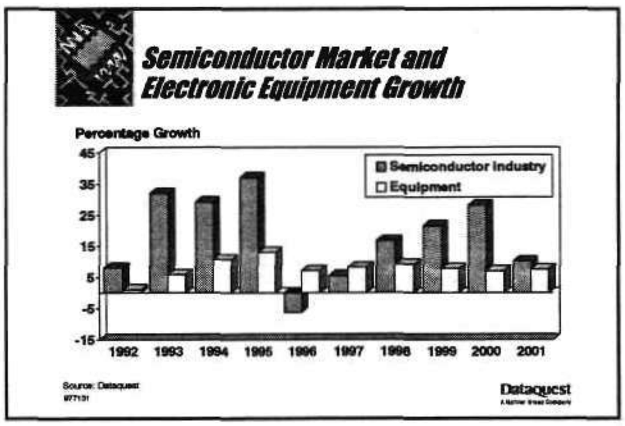

Worldwide electronic production equipment is on top. Greg just talked about that. I'm going to talk about the semiconductor market in the middle, and Clark will follow and talk about the equipment and materials industry. The reason I'm showing this is to indicate where the food chain is going.

Currently, U.S. elecronic equipment production accounts for about 4% of our gross GDP. Vladi Catto, a Chief Economist at TI, estimated that could rise to as much as 15% to 20% over the next 15 years, so I think if any of these predictions are reasonably close, the food chain's going to yield a really big crop of semiconductors.

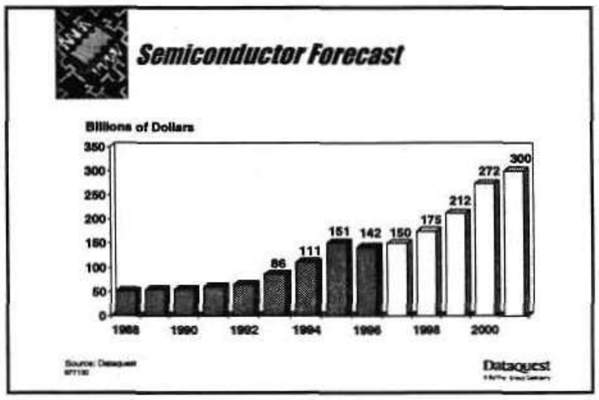

The task right now is to look at what is going to happen in the next five years. Here's what I'm going to address. The equipment Greg talked about, the "big three" drivers are PCs, communications equipment, and consumer electronics, and I'm going to translate what Greg talked about into impact in the semiconductor industry. We'll also look at some of the hot areas in semiconductors, and we'll wrap up with a five-year forecast.

The main driver of semiconductors is PCs, of course, and here is our current forecast. Just in the last day or so, Dataquest has published its latest PC forecast and has reduced 1997, which previously was 84 million units. We reduced it to 83 million units, so that brings that 19% down to a little over 17%. There's been no change in the outlying years and that 83 million units of today will double to 152 million units by the year 2001.

PCs drive the DRAM market and here's the impact of PCs on megabytes per system. 1997 is going to grow from 32 megabytes to 152 megabytes per system. Greg mentioned 128 and what he meant was

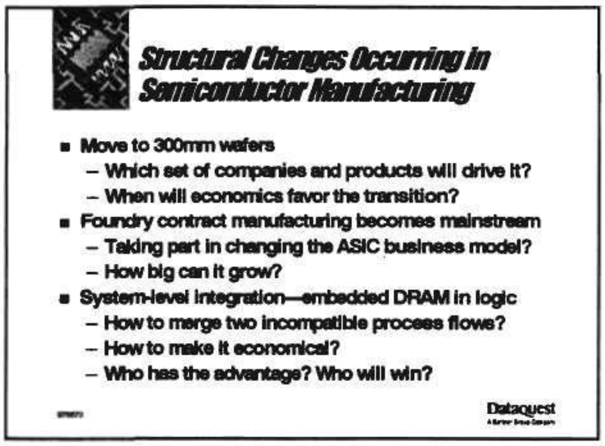

the high-end, not including the upgrades. Driving this will be all the software enhancements that are going to come, and particularly graphics and multi-media enhancements.

There is an interesting arena happening in the green area of the upgrade. That's about $6 billion today and that's going to grow to about $16 billion by the year 2000. Up until now, that's been a market that's not been much addressed by the PC OEMs or by the DRAM manufacturers. It's been mostly the province and the arena of third party DRAM manufacturers like Kingston, probably the best known. You see their ads in the airplane magazines. The PC OEMs and DRAM manufacturers are going to get very much involved in this as this market continues to grow.

PCs drive the microprocessor market. Here's the X86 forecast. It's $15 billion in 1996 and that's going to rise to $30 billion by the year 2000, and then $45 billion. That's an awful lot of microprocessors, and, by the way, CAGR is 19%, just a few percentage points higher than the growth rate in computers. At Dataquest, we really keep these two tracking very closely. This also shows the clone,

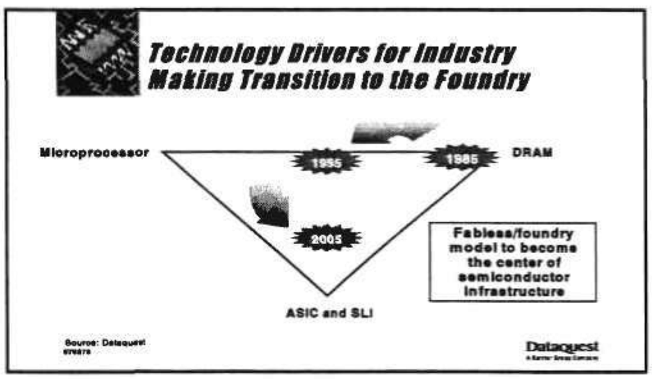

which are buried into this forecast. They're not broken out here, but if you come to Nathan Brookwood's panel session tomorrow, he'll break this identical forecast down into the clones and you'll see that the clones, for the first time, have a really good chance of taking some market share away from Intel, but we'll leave that till tomorrow.

PCs drive the multi-media market. PCs have been evolving from a productivity tool for spreadsheets, word processing, and other data processing tasks, into an information, communications and entertainment tool. PCs now have MPEG2, AC3, audio, telephony, fax, modem, and other features. To do all this requires a lot of processing power, and that increase in processing power, and complexity of chips, is going to drive the semiconductor industry.

I'd be interested in knowing by a show of hands how many people use their PC primarily as a communications tool rather than what I'd call typical spreadsheet analysis. Anyone here? It's quite a few. Interesting. It's to the point where multi-media will

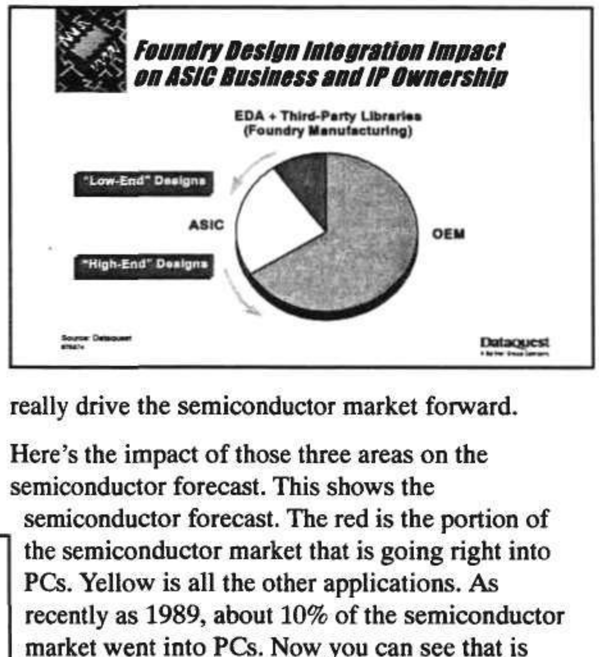

really drive the semiconductor market forward.

Here's the impact of those three areas on the semiconductor forecast. This shows the semiconductor forecast. The red is the portion of the semiconductor market that is going right into PCs. Yellow is all the other applications. As recently as 1989, about 10% of the semiconductor market went into PCs. Now you can see that is substantially more. In fact, back before the DRAM crash, the dollar amount was in the 30s, and even out by the year 2001, it was approaching 40%. So it's pretty significant.

In fact, that red bar out in the year 2001 is $120 billion just going into PCs. It wasn't very long ago when the whole semiconductor market was only $110 billion. We don't break out multi-media in here, but roughly, according to Geoff Ballew, our analyst here, roughly about 15% of that red bar, or of that $120 billion will be multi-media chips. So that's 15% of $120 billion.

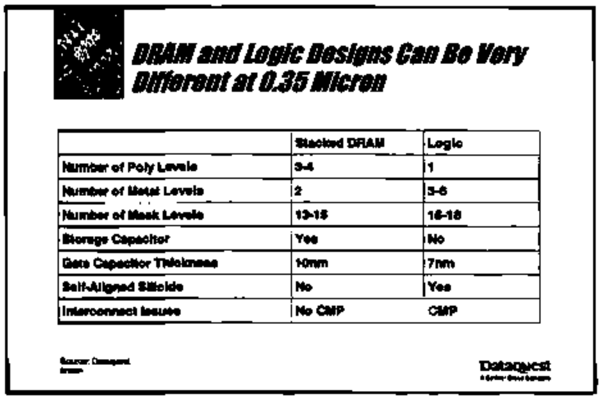

	Stacked DRAM	Logic
Number of Poly Levels	3-4	1
Number of Metal Levels	2	3-6
Number of Mask Levels	13-15	16-18
Storage Capacitor	Yes	No
Gate Capacitor Thickness	10nm	7nm
Self-Aligned Silicide	No	Yes
Interconnect Issues	No CMP	CMP

We know who the movers and shakers in DRAM and microprocessors are, but who are they in multi-media? First of all, almost all the companies here are U.S. companies, with the exception of Creative and Yamaha. All of their research is pretty much done in the U.S. It's an area that is absolutely dominated by U.S. technology.

This is a real tough market to compete in. There are probably about thirty companies, all with 3D chips. The interesting thing is they all have great technology. There has to be some consolidation. In fact, there's probably a violation of Marketing 101 going on here. Usually, the market share leader is the most profitable company and is leading the price, but S3, who probably has about a 46% market share in this area and just reported their earnings, is doing dismally. Something has to shake out. There are 30 companies in graphics and another 10 in audio. Almost all of these companies are fabless using foundries in Asia.

Greg has coined the acronym LAW: "L" for Local Area Network, "A" for Access and infrastructure, "W" for Wireless. I think this very nicely categorizes the entire communications industry. Not only is communication equipment exploding but the ICs that support that industry are also exploding.

It's really the special ICs that are allowing advances to happen in the communications industry. In other words, it's really the semiconductor industry which is driving the communications industry, and because there are so many special ICs, there are a lot of new companies emerging to fulfill these requirements.

What we'll look at are some of the movers and shakers in each one of these three areas. First, let's look at the total impact of the communication market on the semiconductor market. Here I've added to the forecast the chips going into the communication equipment. As recently as 1990, communication chips accounted for about $8 billion and that's going to grow to nearly $60 billion by the year 2001.

Here are the movers and shakers in the LAN area. The Gestalt of all this is there is only one non-American company here, Datacom, which is a Taiwanese company. Literally, the Local Area Network, the special ICs, the ASSPs as we call them, are absolutely dominated by U.S. companies.

What about the movers and shakers and the remote access and WAN chip area, again, this is almost all U.S. companies, except for two European and the Canadian company, Mitel. So it's another area where U.S. companies absolutely dominate. Many of the emerging companies on the right are fabless companies. The other point I would make on here is that Motorola is on the emerging list, not an established player. Motorola has been very strong, particularly in the analog IC area for wireless, and they are really beginning to get into this chip area for digital.

In the baseband, which is part of the digital cellular phone which does all the data converting, there are

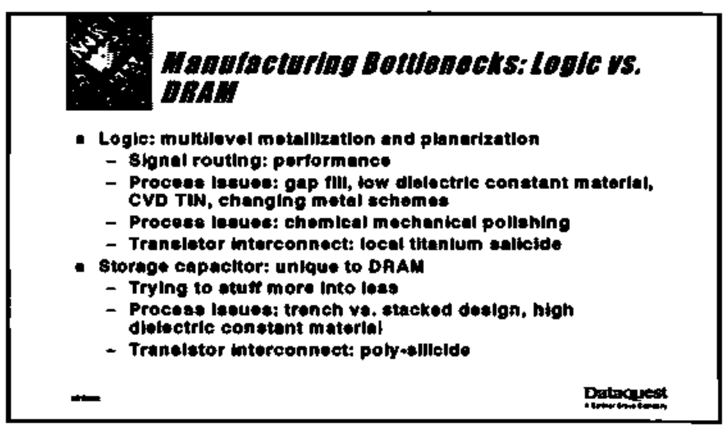

two parts of digital phone: the baseband and the RF or radio section. Again, the significant thing on this list is that the first four companies on this list are also the top four DSP manufacturers in the world.

It's really cellular phones that have spurred the DSP manufacturers as cellular phones are the major DSP application right now. In the past, Motorola and NEC dominated this market when it was analog but the move to digital has begun to reshape the competitive landscape and you're seeing a lot more companies in here.

By the way, the move to digital is also reshaping the consumer chip landscape as we'll see in a few moments. TI has been highly successful in this area because they, like other semiconductor companies, have reorganized and are focusing their business on applications rather than on the device itself. We are seeing more and more companies realigning their product marketing and manufacturing toward applications rather than devices.

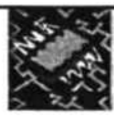

Trench Capacitor DRAM Design More Compatible with Logic Flow

	Trench DRAM	Logic
Number of Poly Levels	2	1
Number of Metal Levels	2	3-6
Number of Mask Levels	13-15	15-18
Storage Capacitor	Yes	No
Gate Capacitor Thickness	10nm	7nm
Self-Aligned Silicide	Yes	Yes
Interconnect Issues	Some	CMP

Source: Dataquest
976693

Dataquest
A Gartner Group Company

Digital consumer electronics is going to be coming our way and these are some of the markets: advanced video games, set-top box, digital cameras, digital camcorders. Collectively, they will make a big impact. In addition to these, there are going to be a lot of applications that we don't know about today making an impact at the turn of the century. One possibility is voice recognition. All of a sudden, every time you dial one of these 1-800 numbers, you can speak yor numbers or whatever it is. I think they are a little premature, because most of the time, they don't work.

Here's a familiar chart. I've added the consumer chip portion to this. By the year 2000, that will approach $46 billion so if you collectively look at these three big IC drivers, they roughly account for about 75% of the total chip market.

Finally, we have a slide on the movers and shakers in the next generation digital consumer area. This market has been dominated by Japanese chip manufacturers and effectively still is. You can see they still dominate digital camcorders, digital cameras, video games, etc., but because there is a move to digital, other companies are now getting into this area because they are leveraging the digital technology they've been developing for so many years into digital space.

For instance, MPEG and AC3, are all U.S. standards. SGS-Thomson and Philips are European companies. Another thing, look at where LSI Logic is in all these spaces. The most successful companies so far, like LSI Logic and SGS-Thomson, are really companies that are leveraging their digital expertise across many different applications.

Let's look at some of the hot areas in the chip arena.

First, a not so hot area: DRAM price in 1995 reached it's maximum at $3.47 per megabit. It has skidded downward to $.49 in August,which represents an 86% price decline. A little error on this slide, I think that skier should be tumbling down rather than skiing down.

But bit growth has remained strong. Prices have fallen, but bit growth in '95 was 83%, 96 it was 78%, and so far this year for the first eight months, bit growth has been 100% over the first eight months of last year. However, DRAM revenue has been flat.

Back in '95, the prices were constant. Bit growth was rising, so we had rising DRAM revenues. In '96, bit growth was rising, but prices fell at a much faster rate, so no amount of bit growth could have kept the prices from falling. This year, that 100% bit growth is just keeping revenues flat with the price declines.

This is the inexorable law of DRAM pricing and shows a long term price curve for DRAMs which historically is about 28%. Down in the lower right, you can see where it went horizontal for a few years, kind of strayed away from the inexorable downward path. Then it fell precipitously to where it is pretty much back on that line. So if the price of DRAMs were allowed to fall as normal a few years ago, we'd probably be right where we're at today, but in the meantime, it caused a lot of havoc.

Now for some of the hot areas: Flash is going to go to $6.3 billion. The key thing here is this single curve masks a lot of dynamics that are happening. The flash market is really trifurcating into three very distinct segments,which are all very different — different by players, different by application, different by technologies.

System integration is another hot area. Everyone is talking about embedded DRAM. In our recent May semiconductor conference in Japan, every Japanese manufacturer got up and made a presentation and talked about embedded DRAM. SLI is the glue that holds all this together. A little note of caution, semiconductor manufacturers very blithely talk about putting 100 million transistors on a chip and.18 micron technology, but the people who design those can't do that. There are some real problems with design software which need to make some strides ahead to be able to hook all those 100 million transistors together.

Here's our forecast for SLI. You can see that it will be the dominant ASIC solution by the year 2000. Actually, designs in 1988, 1998 and 1999 that are SLI based will actually represent more revenue than ASICs, but since they don't reach market for a couple years, that revenue doesn't show up. There also starts in generating SLI revenue because SLI has a higher revenue per design. However, this whole area of SLI is such a hot area, I think there are four different presentations going forward from now to address different aspects of this.

DSP is taking off as well — another hot area. You can see the growth phase of MPU is pretty much the same as the growth phase for DSP. The question is, what's going to happen after the year 2000? There will also be a talk by TI on DSP. DSP area is dominated by three U.S. companies: TI, Lucent and ADI have 75% of the world market.

End-Chip Revenue per Square Inch Varies with Product and Technology

Cost important for SLI: Revenue per square inch can increase only marginally

Revenue/Square Inch ($)	Leading Edge	Mainstream	Lagging Edge
MPU	300-600	150-250	80-150
ASIC, Logic, MCU	100-140	80-90	50-60
DRAM	80-90 (Micron: 130)	60-75	45-50
Power/Discrete and Analog	30-35	25	<15

Source: Dataquest

Dataquest

What's the impact of all this on the semiconductor market? The key thing is Intel now is 12.5% of the world semiconductor market.

What I've done is taken our forecast and break it out in the non-DRAM and DRAM components. You can see the non-DRAM forecast is nicely behaved. The growth rate from '92 to '97 is historically about 18%. From '97 to 2001, it's 17%. This is very much in line with the way we do our forecasts. This I think is a very reasonable forecast.

However, DRAM forecast looks a little different. We're not going to see DRAM market reach the 1995 levels until 1999. There's going to be a slight over capacity situation in the year 2000, such that the market will drop a little bit to 2001. But still we see the DRAM market reaching nearly $70 billion from today's $26 billion. There are ups and downs, but still there is a lot of growth.

You put both of these together, this is our semiconductor forecast. Two numbers are very easy to remember: 1998 is going to be $150 billion. The year 2001 is going to be $300 billion.

Next year will be about 17% growth and we have some good growth years beyond that: 20% and 28%; and then they there's a down of 10% in 2001.

DRAM ranges from 15% to 25% of the total semiconductor market and it's really the swing factor in all of our forecasts.

This give you an idea of how sensitive the forecast is. For the 1998 DRAM forecast, we assume $7.69. That gives us a 16.7% growth rate for the semiconductor market.

Basically the sensitivity is for every 5% change in our forecast DRAM price, that affects the semiconductor forecast by 1%. It's clearly hard to accurately forecast the total semiconductor market in times of rapidly fluctuating DRAM prices.

This puts in perspective the semiconductor growth rates with the equipment growth rates. The thing to mention here is that electronic equipment growth rates are pretty steady throughout the period. We had a little adjustment period for '96, but the electronic equipment industry was still up 8%.

In summary, the message is that the U.S. semiconductor industry is poised very nicely for the emerging areas. We expect good growth in all areas of electronic equipment, PC units, and semiconductor market, and Moore's Law is still intact. There were a lot of questions about what's happened to that. The industry looks pretty good for the next four years.

That concludes my presentation.

Clark J. Fuhs: Our group studies most aspects of how chips are manufactured. We study the fundamentals of capacity silicon consumption, process infrastructure and how that relationship ties to the chip industry.

Supply side fundamentals have been a major part of the semiconductor industry of the last couple of years and we expect that to continue going forward.

I'll basically go through a forecast overview and report card. We'll take a look at some of the short to intermediate term details in terms of dynamics, covering 97-98. Then we'll go into some key technology issues. Our group's forecast cycle is January to July and the forecast I'm presenting was actually created in July. However, I'll add a couple of comments that take into consideration a few things that have happened in the last three months, and I will make at least one outrageous conclusion.

Here's our capital spending forecast for the year 2002 and the next slide is our Wafer Fab Equipment forecast. Both of these show a profile calling for a two year pause in 1997 and 1998, with 1998 at a single digit growth rate before resumption of accelerated growth in 1999.

Notice that we are calling for a cyclical downturn in 2002 and this is consistent with a chip downturn in 2001.

In 1996, we had a transition year from growth to decline. Backlogs were taken down during the year and that set up the two pause years of 97-98. We expected acceleration of growth and the capacity buys in 1999 and technology buying modes to take us through most of this and next year.

Our general outlook has remained unchanged over the last two years. We've always expected a slowdown in 96 leading to a decline in 97 with single digit growth in 98.

We've been able to maintain this record by focussing on the fundamental issues associated with the supply side of our industry.

Here is our range for the possibilities for the Wafer Fab Equipment forecast over the next few years. Let me emphasize that 1997 has more upside than downside potential, which is associated with the extension of the technology buying phase of the cycle and Taiwan manufacturers. Taiwan is actually the wild card there.

In 1998, there is actually more downside risk than upside potential. Two things contribute to that. First, the over capacity fundamentals are still weighing on the market and 1997 is coming in a little better in capital spending so there is a somewhat more capacity added than we had originally thought.

The second thing is that we suspect the issues going on in Southeast Asia will put a throttle and a governor on the capital availability in Asia, which very well could put 1998 into a decline. We're not calling that yet but it is a risk we'd watch.

Our quarterly forecast continues to have a W profile, which we originally presented in January. I think we're the only market research firm to actually be calling a W profile in capital spending and Wafer Fab Equipment.

In the middle of 1996, when everyone was adding capacity and the blood was starting to flow, the valve was shut off capital spending quite severely. Toward the end of last year, we saw a normal event called, "Let's cut our operating cost by investing in technology and shrink these devices." That led to an incremental increase in spending associated with technology upgrades.

That part of the cycle normally lasts nine to 12 months. What usually happens is that you have an equilibrium state where over capacity remains in place. We are near the end of that stage right now. Then over capacity takes over, the technology buys are basically in place, and everyone sits and waits for demand to catch up to the supply base on a silicon level.

Some updated information: the second quarter of 1997 is actually shown as the peak there. The second quarter came in a little lower than is shown and the third quarter is coming in a little higher. The fourth quarter is actually a little higher than that, and we expect the first quarter to be down, so it's actually been delayed a month or two relative to our July outlook, only because Taiwan has continued to invest.

This profile is also supported by the most recent equipment purchase surveys which show a 4% and a 16% sequential decline in orders for the fourth and first quarters. We don't think it's going to be that severe in the first quarter but we do expect some sequential declines to come into play on the order side.

One other item of note, the first quarter of 98 is only 4% above the first quarter of 97 in that same survey, so that also supports the single digit growth going forward into 1998.

When can a recovery be sustained? The first point we'd like to make is in taking a look at the raw silicon market. This can provide some good hints on how the capacity is actually consumed in this business.

The silicon wafer forecast appears to be more stable, and after a couple of years of slight growth we see resumption of high single/low double digit growth rates. However, that masks what actually happened, and there was a severe correction between the second quarter and fourth quarter of last year that saw silicon consumption into fabs go down by 20% on a run rate basis.

We are now just getting back to the peak levels of the second quarter of 1996, and in that six quarter period a lot of capacity has been added. This is actually shown as, and can be tied to, the DRAM cycle. The green bars show the silicon consumption into the DRAM segment. This was caused by the migration of the 4 to16 meg DRAM, which caused fantastic silicon efficiency in the market.

The combined bars represent the total capacity in the DRAM segment, so the yellow bars represent the over capacity that is going on. The mode we're in today is that silicon demand is increasing into the DRAM area, but so is capacity so that the pricing pressures will remain intact.

Another metric which has us concerned is capital spending as a percent of revenue. We are now in the third year above 25% which hasn't happened since the 84-85 time period. With capital spending coming a little better than we anticipated in 97, and the semiconductor numbers being revised down, we actually believe that this percentage will be between 26% and 27% for 1997. We believe the equilibrium level is closer to 22%, so there has to be some fundamental correction in the market. We expect that to occur in 1998.

Our next concern, when we look at the number of new fabs being constructed, is a question as to where the fabs are. There are not that many new fabs coming on line. That means that 1998 will be an upgrade year rather than one driven by large new fab orders, and, therefore, a somewhat frustrating year.

Another issue that has us concerned is the foundry market supply and demand. The concern is not a decrease but that it is not going to be fuel for a boom. Prices for foundry wafers have dropped across the board, 0.35 micron pricing has dropped 17% to 20% in the last six months according to our survey completed in September. While this industry is growing and is good for stable growth, it is not the fuel for an equipment boom.

Narrowing down to 0.35 micron supply and demand, we see that supply is leading demand, but it is a fairly tight market. The way I like to describe it is there is about one fab available today for every 0.35 micron. However, that demand is growing so steeply that a fab is filled up in five months. Thus, allocation of capacity is still the norm but it also represents a pricing environment which remains soft because the supply is there. This is probably the most ideal balanced market we could have.

The fundamental story is intact with several events that have occurred in the past several months. If anything, there have been fabs added and fabs taken away. I would suggest that lead time is now closer to four months than five, but the fundamental story is still there.

That is a rundown of the issues that lead us to believe that 1998 will at best be a single digit growth year. Now I want to steer into some of the technology issues.

I'll skip over the technology mix quickly. During a slowdown you normally get a redirection to logic oriented products and technology enhancement, so we see memory coming down and logic ramping up some of the hot equipment markets that will drive such as Deep UV lithography. I think everyone is well aware of the move to quarter micron here. CMP is associated with logic capacity and multilevel metallization. The one area not on here is copper deposition. The sputtering seed layers and electroplating systems that will be used to feed in copper pilot lines over the next few years is going to be a hot market. Let's take a look at a couple of these in more detail.

This is our lithography technology forecast. Deep-UV is ramping to 300 units this year and about 400 next year. This is partially driven by supply side issues and lens availability, which we expect will throttle the market through about 1998.

Even if you look in the 2000-2001 time frame, there is going to be a three-way mix and match between DUV, I-line and G-line. One thing in the Epi area I'll mention is a new defect which is cropping up. They're called COPs, i.e., crystal originated particles. Basically, it's not a particle at all but a hole in the silicon. It's starting to create a yield impact as you get into 0.18 and 0.15 micron technologies. The big issue today is yield impact into the DRAM area and whether or not you need Epi as a way to fix this problem.

There is a lot of uncertainty in the market when we take a look at the low side and high side scenarios for Epi consumption into the DRAM market. I will say that the silicon wafer manufacturers are ramping up Epi toward the higher end of that scale so I can see a situation, perhaps in 1999, where they will cut back capital spending in this area to evaluate whether demand is really going to be there.

In terms of 300mm wafers, even by the year 2002, we only expect 300mm wafers to be about 9% of the market. We see 200mm wafers ramping pretty much with being the key technology and key motive for production in this industry, but there will be a lot of activity in the next couple of years to implement this.

We always get the question, "Who will be first?" It's an irrelevant question because this time really is different. This is going to be an industry move with many companies participating, not just one leading the charge.

Because of this we expect a traditional "double hump" equipment sales pattern to emerge, where the first set of equipment will come into play in the 98-99 time frame. We'll get to a point in the year 2000 where there will be a lot of pilot lines going which may not be economically viable when you compare them against 200mm fabs.

The question will be raised by many chip companies, "What plant do I build to support production in the year 2001?" Our guess is that it will continue to be 200mm wafer fabs, so the race to build the last 200mm wafer fab will have to wait until the year 2002 or 2003. This is our "double hump" equipment sales pattern as a percent of the overall equipment mix.

Which companies are likely to lead in 300mm? We've split it into three waves of companies: first wave-second wave, and others.

There isn't really much difference between first and second wave, perhaps six to 12 months. The first wave companies are typically viewed as the leaders, those that will actually buy equipment in 1998 for early 1999 start-up, whereas the second wave are more late-1999, early-2000. There are some that may fall off the second wave list and there are some that may switch lists, but, for the most part, we've gotten pretty good indication that this is a good list.

Bottom line, expect about eight to 12 pilot line facilities to be on line by first and second wave companies combined by the end of the year 2000.

One comment about the fabless/foundry models: they work. The foundry industry will be a driving force going forward because of concentration of capital and concentration of capacity bringing economies in scale to bear.

I'm going to skip over the slide representing the true nature of increase in the cost of a fab. They're getting bigger as well.

One issue that is going to be a subject of the afternoon session is the technology drivers for our industry. In 1985, the primary process flow driver was the DRAM. By 1995, we had a mix of both microprocessor and DRAM, and we see that moving to an ASIC and SLI imbedded DRAM and logic process flow really starting to take the lead now, going forward.

The manufacturing companies that are positioned to provide this are going to be the technology leaders going forward, and we expect the fabless/foundry model to be at the center of these developments.

One additional issue: how does IBM and the copper announcement play into this? From a strategic positioning perspective, we would not really classify IBM as a DRAM company, even though they're in the DRAM market. We would not classify them as a microprocessor company even though they're in the microprocessor market. We would classify them as a custom logic ASIC business.

One other fact is that IBM is the second largest foundry in the world, and they have mentioned that they will put the copper technology available in their foundry business. Thus the comparison of IBM being a year or so ahead of Intel and AMD and so on, is really the wrong comparison to make when you want to study the impact on the market. When they introduce copper into their foundry later in 1999, it is very possible that they will be three or four-plus years ahead of their competition in the foundry market. None of the other companies that have copper are in the foundry market.

This will have a tremendous ramification with regard to availability of the types of chips to compete in the ASIC and SLI categories.

In summary, we see the end use demand picture for semiconductors remaining strong. However, the large over capacity today means a significant lag for the equipment market into 1999, and we'll be covering the ramifications of imbedded DRAM, SLI and the foundry market in our afternoon session.

Chapter 6: Emerging Technologies in Equipment Processes

Moderator:

Clark J. Fuhs
Director and Principal Analyst
Semiconductor Equipment, Manufacturing, and Materials Program
Semiconductor Contract Manufacturing Services Program
Semiconductors Group
Dataquest

Panelists

Dr. Peter Chang
President
United Semiconductor Corporation

Dr. Michael R. Polcari
Director of Silicon Technology and Advanced Semiconductor Lab
IBM Microelectronics

Dr. Dale Harbison
Vice President, Semiconductor Group
Manager, Manufacturing Science and Technology Center
Texas Instruments Incorporated

Inseok S. Hwang
Senior Vice President
Semiconductor R&D Division
Hyundai Electronics Industries Company Limited

Dr. Eiji Takeda
Department Manager, System LSI Development Office
Semiconductor and IC Division
Central Research Laboratory
Hitachi Limited

Clark J. Fuhs: I'd like to welcome everyone to the panel discussion on "Emerging Technologies in Equipment and Process Technology." The theme for this conference is the structural changes that are occurring in the industry. What we'll specifically talk about is semiconductor manufacturing on the front-end. The next panel will talk about some of the issues going on in the packaging and assembly part of the world. There are 3 main issues that are happening and structural changes that are occurring.

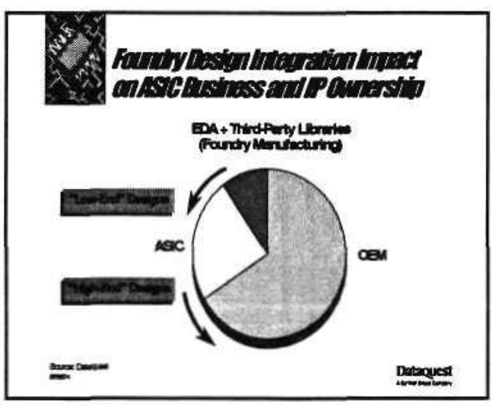

One is the move to 300mm wafers. There are issues there which we covered at a presentation quickly this morning. I won't talk about it here as an introductory phase but the panel is prepared to address the move to 300mm's here. The second issue is the emergence of foundry contract manufacturing becoming mainstream and what that means in going forward. The third is system-level integration, specifically, from a manufacturing perspective, we want to concentrate on embedded DRAM in the logic. There are specific process flow issues that need to be addressed.

What I would like to do is give a Dataquest opening statement about the foundry and structural changes. This is a chart, compiled by our ASIC group, of the intellectual property split between the system OEM, the ASIC provider and Third-Party EDA Tool Libraries. What's happening is that the foundry manufacturing offering third-party EDA tools is starting to eat, on the low-end designs, into some of the markets that the ASIC manufacturers have typically played in. The ASIC suppliers are responding by trying to take a larger part in what has typically been the system OEM design. How foundry has come into play as a manufacturing power is basically the cause behind some of these migrations.

If you looked at the world in 1985 in a different way, doing an XY chart on manufacturing power versus design power, the OEM system providers had all the design power and the ASIC suppliers were uniquely positioned in providing the manufacturing power. The foundry model entered in the late 1980's and by 1991 had established themselves, but they had manufacturing technology that lagged behind many of the ASIC players. They were moving towards the ASIC space though. What has happened

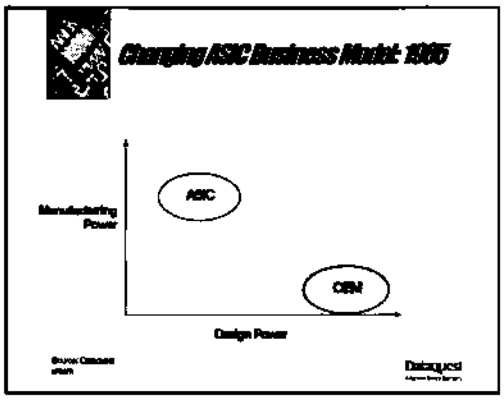

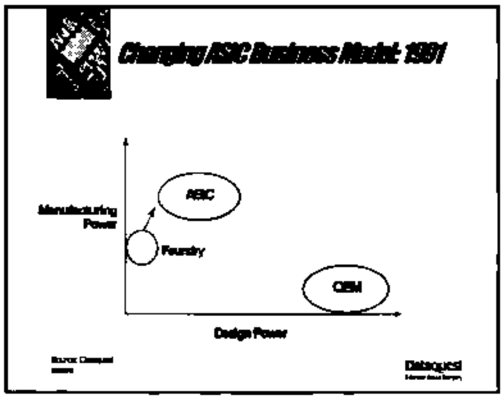

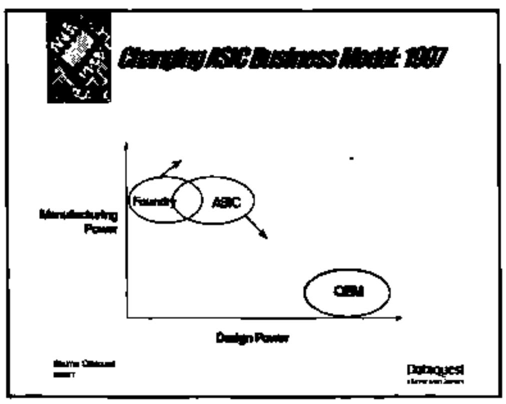

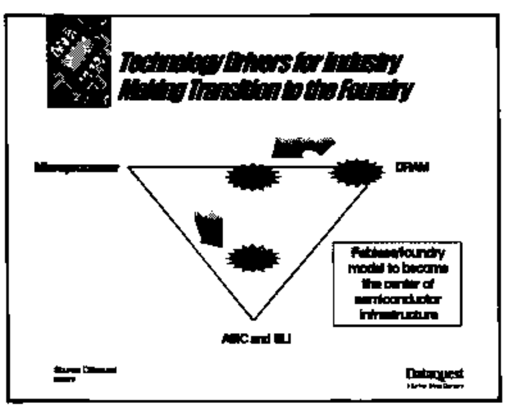

is that it has now translated into an overlap situation today.

The traditional ASIC supplier has a choice to either move to a business model which is more closely tied to software and chip design that the OEM does, or they can choose to become a foundry. The changing ASIC business model is currently underway as we speak. The fab announcement last week by VLSI Technology could be interpreted as a first step of a major fab company becoming fabless.

System-level integration and embedded DRAM in logic. The fabless company, as was presented this morning by Joe Grenier and some of the major players, are part and parcel some of the key demand issues. The foundry has become the natural supplier to some of those designs.

The manufacturing perspective on SLI is that you have DRAM and logic in the same fab. We talked about how that positions foundry to be a process flow-driver going forward. From the straight manufacturing perspective, what you're really talking about is putting logic and DRAM on the same chip. Since they have different process flows and different requirements from a device-design standpoint, you're going to optimize one of those at the expense of the other. You cannot optimize both at the same time. If you have a little DRAM on your part, you're going to optimize the logic. If you have a lot of DRAM on your part, you're going to tend to optimize the DRAM.

The process flow requirements are not trivial to look at either.

When you look at process flows for a typical logic chip today and a typical DRAM process flow using a

DRAM and Logic Designs Can Be Very Different at 0.35 Micron

	Stacked DRAM	Logic
Number of Poly Levels	3-4	1
Number of Metal Levels	2	3-6
Number of Mask Levels	13-15	15-18
Storage Capacitor	Yes	No
Gate Capacitor Thickness	10nm	7nm
Self-Aligned Silicide	No	Yes
Interconnect Issues	No CMP	CMP

Source: Dataquest

Dataquest

stacked capacitor design, you can see the number of poly levels and the number of metal levels are not compatible. There are many other issues such as self-aligned silicide. If both of these have mask designs of 16 or so mask players, the combined set could have as many as 20 to 24 mask levels, which increases the cost per wafer of the embedded DRAM within the logic device. There are ways you can take care of that. These are the particular process flow issues which I'm not going to go into in great detail.

There is one design of the DRAM which is better suited to be included as an embedded DRAM design. This is the trench capacitor. There is still a net cost-

Trench Capacitor DRAM Design More Compatible with Logic Flow

	Trench DRAM	Logic
Number of Poly Levels	2	1
Number of Metal Levels	2	3-6
Number of Mask Levels	13-15	15-18
Storage Capacitor	Yes	No
Gate Capacitor Thickness	10nm	7nm
Self-Aligned Silicide	Yes	Yes
Interconnect Issues	Some	CMP

Source: Dataquest

Dataquest

adder, however, even with the trench capacitor

End-Chip Revenue per Square Inch Varies with Product and Technology

Cost Important for SLI: Revenue per square inch can increase only marginally

Revenue/Square Inch ($)	Leading Edge	Mainstream	Lagging Edge
MPU	300-600	150-250	90-150
ASIC, Logic, MCU	100-140	80-90	50-60
DRAM	60-90 (Micron: 130)	60-75	45-60
Power/Discrete and Analog	30-35	25	<15

Source: Dataquest

Dataquest

design. Your basic problem is that, as a manufacturer, you can get a certain amount of revenue per square inch from a logic device. Advanced logic devices are generally about $100 to $120 per square inch that you're gaining in revenue by making a standard logic part. You are putting DRAM on it, which has generally been a lower revenue per square inch and you are increasing the cost. In all of these, you are decreasing the return per square inch and you're increasing the cost per square inch.

The net is that there has to be a very good premium and a very good reason for putting these together on the same chip. These are the issues that we will want to address through the panelists.

We've collected a very good and very diverse set of people. This is the order in which they will be presenting. We've asked each one of them to make a five-eight minute opening statement on their view and their positioning on some of these challenges in manufacturing going forward, and then we will open it up to questions. As each panelist is presenting, we ask that you write down your questions immediately and give them to the people that are walking around with the question mark signs. This way we can have a set of questions as we get started with the Q & A session.

This will be the order of speakers. The first to speak is Mr. Peter Chang. He is the President of United Semiconductor Corporation, the first joint venture fab that is managed by UMC and is a dedicated foundry.

Dr. Michael Polcari is the Director of Silicon Technology and Advanced Semiconductor Lab for IBM Microelectronics. He'll be talking to us about issues from the IBM perspective.

Dr. Inseok Hwang from Hyundai Electronics is Senior Vice President. He is in charge of technology issues transferring to manufacturing.

Mr. Dale Harbison, Vice President and Manager of Manufacturing Science and Technology Center at Texas Instruments, provides that same function, where they're looking at the economic design of process flows and how it goes into the chips that are designed.

Dr. Eiji Takeda from Hitachi performs that basic same function as the Deputy General Manager of the Semiconductor Technology Development Center.

We welcome to the stage at this point, Mr. Peter Chang.

Peter Chang: Ladies and gentlemen, my name is Peter Chang. I'm representing the UMC group. My first topic is a bit busy. Everyone should have a copy. We will address the first issue related to the embedded DRAM. The embedded DRAM and UMC group as a foundry source. As this particular slide will show, on the far left-hand side, is a logic chip and on the far right-hand side is a DRAM chip. Both technologies are being manufactured in the UMC group.

In between them, we have two approaches to the embedded DRAM situation. One is we call the 0.35 embedded DRAM high performance. The other one we call 0.35 micron embedded DRAM low cost. If you look at the cells, the low cost cell is closer to the DRAM process cell. For example, the metal layer is only 3 layers on a low cost cell. Basically, it's only one gate outside for all the process cell. Then, most of the technology on the peripheral is very close to the DRAM. Even the peripheral is on the poly side. You look around this thing, that's why they reduce the mask set to close to 20 layers. That will reduce the overall cost.

In some customer's cells, we found they are very cost sensitive in this particular application. Their application may be concentrated not on the particular performance but on the power savings. That is the application we are using for our foundry customers. The other sets for high performance are mainly applications in the graphics area. They want the logic to be super fast. There is also a very large bandwidth along with everything else. They want the performance. That is the application for overall. In this application, we have dual layer of the gate outside thickness from 70 Angstroms to 100 Angstroms. Also, there are four metal layers in the cell.

However, you have to pay the price. The price is 25 layers. You have to put on all the layers. The cost is relatively high. That's our strategy at this moment. This technology right now is in pilot production. This is the 0.35 embedded DRAM. This is a slide showing recently that the 8" in UMC FAB. The complication in the fab itself is that we have several different kinds of technology. We have logic technology, DRAM technology, embedded DRAM technology as well as SRAM and volatile memory. Those are the requirements for a foundry business if they wish to survive for the next 10 years.

We'll get into 0.35 micron later; 0.25 is in the pilot stage now and 0.18 is in the development stage. It's very obvious that we are using logic as the technology driver to continue to drive this technology down. This is the full planarization of 0.25 micron of 60 transistor, six layers of metal of processed cell. This is six layer of metal processed protection and SRAM, and DRAM.

This answers the question, in general, about the 300mm schedule. This is our internal schedule. We have made a plan to invest heavily in the southern part of Taiwan, which is Tainin Science Park. We plan to break ground on our first 12" fab in 1999 and plan to pilot as close as possible to the beginning of 2001. That's our current plan and it looks like it's conservative. We believe that there's still a lot of technical difficulty and cost problems related to 12" wafers. This is the total UMC group wafer output forecast. By the end of fourth quarter, 1999, we'll put out close to 492.

In summary, in the embedded DRAM situation, we have a deal with our customers. Right now, it's 0.35 microns in pilot production. In our point of view, in the future, the trend is going to be to go in that direction, but the question is, "How big the market?" That will answer of whether the cost could go down. Our customers are telling us that they are really under pressure as far as the price goes. The big question is whether or not it will become a major force. Second, for 300mm, in our minds, it's not mature yet. We'd like to see the wafer cost down, the raw material cost down and also the equipment should be expediting their fee to get up to the progress as in other areas. That is my presentation. Thank you very much.

Dr. Michael R. Polcari: I'd like to thank Clark for this morning. After seeing the slides on the microprocessor business and how much is controlled by Intel and how bad the DRAM business was, Clark put my mind at ease by telling me that I'm not in either one of those businesses. I don't need to worry about anything.

I'd like to talk about our view of system-level integration and what we think are some of the issues and some of the important things that need to be focused on. One thing I'd like to say, is that cost is a primary concern in all this. There are a number of packaging solutions that may, in the interim, serve to relieve some of those constraints.

In the long run, we feel, as most do, that system-level integration is something that's coming and will be here soon. If you look at what's been driving it, our ability to go to smaller dimensions and increasing the performance, and power reduction which is allowing us to integrate everything into a single chip. Process technology itself is not going to allow us to get to where we want to be. We need to combine the design skills necessary for system-level integration, as well as some of the design tools and methodology to be able to get to where we want to go. This is a difficult task, putting all of this into a single chip, and we need to martial the right forces and do a number of things like reuse the hardware and software objects.

IP is a major issue, in terms of having access to the cores that one would like to invent. One would like to have the ability to have cores from a number of different users, which leads one to be concerned about being able to license cores. Also, the interoperability of these cores so that one can mix and match cores and come up with a correct type of system-level project that one is working on. We feel that alliances, partnerships and licensing are going to be important in terms of doing this.

The next line illustrates something that we've done. I'm going to touch on a couple points here. What you see is a chip that we've built which is the heart of our chess- playing grand master. It's a Power 2 super chip. What you see is from 1995 to 1997, we've taken the chips that were on a rather large multi-chip module and combined them into a single chip processor, with the performance and power improvements that one gets by putting that all together. The first thing that anyone in manufacturing will notice when they look at the chip is the size of it. It is an 18mm by 18mm chip which is rather large. Larger chips are something that we have to be concerned about in the future. As you put more and more functions on a chip, the tendency of the systems people is to make the chip larger and larger. We are continually talking to our systems-level people about trying to do things that make sense.

The other thing that I'd like to mention, relating to this charge, is the fact that putting all these cores together raises the issue of tests, and how one is going to be able to test the kinds of system-level chips. Will we be able to have a test strategy that will be able to test cores from different people? Will we test at the core-level or the system-level? Those are things that will come up as we go further down this road.

Mixing all of these technologies together to bring a system-level design out is a challenge both in the development and in manufacturing. In manufacturing, introducing different technologies into a fab brings in variability, which is always a cost issue. We've been manufacturing memory and logic in the same fab together and we need to be concerned about the variability, to make sure the costs are consistent with either technology.

Mixing these two together in an integrated process flow brings together another level of cost issues that you need to address. You need to understand what is the real driver and what is the product that you're aiming at. Is it something with a high performance logic and a lower level DRAM or a DRAM base with some logic that comes along with the DRAM? Those are issues that will be driven by the customer as well as the solution and what you are trying to accomplish with that.

The last slide is to address what we see as some of the advantages of having an integrated model, where one can optimize across all of the different types of issues, that one should think about, in terms of bringing a system-level chip out. By having a vertically integrated type of system, you're able to optimize across the boundaries much better than if you have the individual horizontal type model. You still need to worry about things like licensing but that's probably the best solution to being able to bring a solution to the customer. With that, I'll have the next speaker come up.

Moderator Fuhs: Next up is Inseok Hwang, Ph.D from Hyundai. I should point out that the number of Ph.D's on this panel is very rich.

Inseok S. Hwang: My name is Inseok Hwang. I'm talking about the embedded DRAM process flow challenges and process issues in embedded DRAM and finally 300mm issues.

It looks that the embedded DRAM approach improves performance, power, die size and cost. However, in order to really meet the cost reduction, we have some issues to be resolved. We have experience that an application requiring a small size

Process Step in Embedded DRAM

Item		DRAM (4P2M)	LOGIC (1P5M)	EML	Process
Mask Step		18	21	22-23	P2/P2C/P3/P3C/ P4
Total Process step		120	135	160	
Process Step	TR.	50	60	60	Salicide & Dual Gate
	Bit Line & Capacitor	40	0	40	CMP & Ta_2O_5
	Metal Interconnection	30	75	60	Metal Bit Line & Storage Node

Semiconductor Research Division — HYUNDAI ELECTRONICS INDUSTRIES CO., LTD

DRAM is not cost effective with the embedded DRAM approach. It may have to use a very expensive strict combination of DRAM and logic process.

One of the challenges of process flow for embedded DRAM is how to deduce the increasing mask and process steps. For the combination of DRAM process into logic process increase mask steps and process steps. If DRAM process with 18 mask steps and logic process with 21 mask steps strict combination resulting about 25 or 26, however, if you use metal layer for bit line as mode and effective number of mask steps will be 22 and 23. In terms of mask steps with a stacked capacitor could be comparable with a trenched capacitor.

However, the process steps increase proper process steps and right now if normal process step for DRAM is 1 than logic will have about 13% more

Process Step Comparison

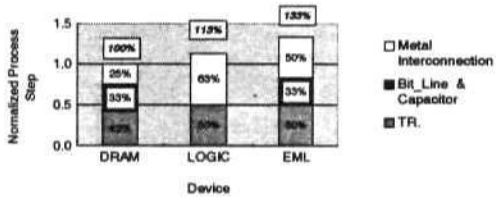

Semiconductor Research Division — HYUNDAI ELECTRONICS INDUSTRIES CO., LTD

Process Comparison for DRAM vs. EML

Process	DRAM	LOGIC	EML
TR.	No Salicide	Salicide	Salicide
Bit_Line & Capacitor	Use	Not Use	Use
Metal Interconnection	2 Level	5 Level	5 Level

Semiconductor Research Division — HYUNDAI ELECTRONICS INDUSTRIES CO., LTD

process step and resulting MLM will have about 30% more steps. The other challenges we have are to further reduce proper process steps for MLM.

The next question is should the process be compatible with DRAM or logic. My answer is the MLM process should be compatible with DRAM, especially DRAM with a stacked capacitor. This is because of the management of thermal bodies in the process.

The next area I'm addressing is the process issues in embedded DRAM. There are several issues which we have to solve in an permanent manner, like resolution in technology. Especially in embedded DRAM, there is a challenge to implement in the regions between cell and random logic patterns. Planarization is not easy because of the high topology difference between cell and peripheral area.

Metallization in the embedded DRAM process has a very small lithography margin; this is especially true with embedded DRAM with a stacked capacitor. However, we can resolve these issues. Finally, when integrating the MLM process, one of the top issues is the capacitor process because the capacitor process usually requires a formal process. If we use a capacitor, we can handle these issues.

For a year now, we have had to use the trenched capacitor and so the precondition of capacitor process should provide no degradation due to

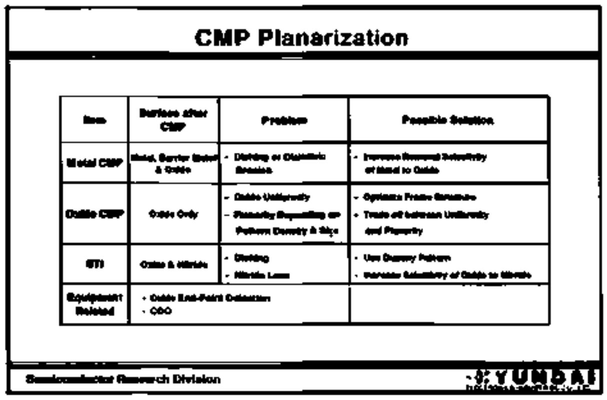

capacitor process. Right now, we have candidates for capacitor materials. This table shows the compatibility of capacitor process with MLM. If we omit it and it turns out that we can omit it and provide for ILD depredation of capacitor is kept below 650 degrees C. The conclusion is stacked capacitor process is compatible with MLM manufacturing process.

Next are the 300mm issues. First of all, the timing for industry conversion to 300mm is not certain. This

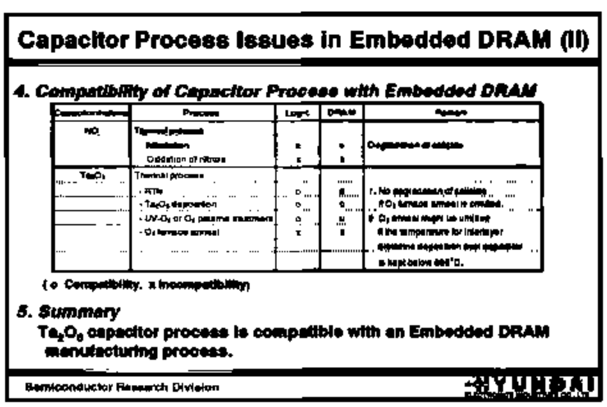

is because the initial pilot line start-up cost is very expensive due to poor utilization. All 300mm equipment will not reach the same delivery of maturity. Also, the economics of product revenue and investment may not match. Initially, the wafer cost is very high and secondly, there is a question of whether global standardization can be really successful. If so, how fast will global standardization be completed? This depends on industry cooperation and I300I and J300I construction cooperation. Thank you.

Moderator: Next we have Dr. Dale Harbison from Texas Instruments. He reminded me that he has a Ph.D. as well.

Dr. Dale Harbison: My presentation is going to be a little bit different than the previous three speakers talking about embedded DRAM. I decided that all the speakers would probably show some comparison of the DRAM flow and logic flow and DRAM steps and logic steps so I didn't need to repeat that. What I believe the real issues to be are the cost and the economics. The economics in this situation is what gets in the way. The barrier is certainly not the ability to integrate the two technologies. We can do that; we have seen that in the previous speakers. What you have to do is find a way to justify the additional cost to do that. You need some application with the performance need that will justify that, and also that will have the volume to be able to justify the cost.

As an industry, we're facing a major cost challenge. I

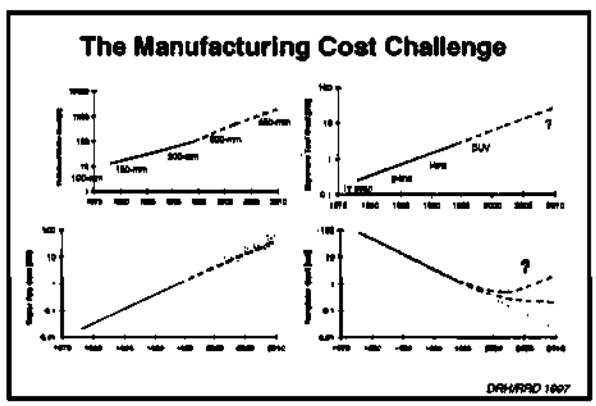

tried to illustrate that in this slide with four different graphs. The one at the top left shows 200mm wafers costing about $100 per wafer and 300mm wafers costing about $1000 per wafer. Maybe, when they get into volume production in the 2000 to 2002 time frame, they may get down to $600 or $700 per wafer. That is still going to be a major increase in the cost per square centimeter, and then you have the cost to build FAB's for that technology.

If you look at the top right corner, just the lithography tools, for example, as we go from i-line to DUV, will jump in cost from $2 or $3 million to $4 or $5 million, depending on whether it's a DUV scanner or stepper, and maybe a couple of million dollars more than that by the time they get to be 300mm tools. At the present time, people are complaining about factory costs being $1 billion or $1-1/2 billion for a 20K to 25K wafer factory per month. With 300mm tools, that cost will continue to rise. I'm not sure if it will stay on the same slope or whether it will even get worse.

The bottom right hand graph shows the transistor cost. For our industry to grow, we need to stay on that productivity learning curve. That curve has been shown many times by Semitech people and other industry speakers. If we can't stay on that green line and have that green line continue to move down at that rate with the transistor cost, then the question is whether or not the industry can continue to grow and how fast it can grow. We have a serious trouble if we begin to get off of that line -- like some of the red-dashed lines show there with the question marks.

How much can we save with 300mm? These are some comparisons that are fairly typical industry estimates of the cost of 300mm versus 200mm at maturity for a 20K wafer per month factory. We expect the cost per square centimeter to improve by 25% to 30%. Because you get a significant improvement in the usable area, especially with large chips, depending on the chip size, the cost per chip may be in the neighborhood of 30% to 40% improvement. We've given the suppliers a goal of no more than 1.3 times the cost going from 200mm to 300mm, with the same footprint and the same throughput. Some suppliers are claiming they can't do that. If that cost increases significantly, when we will convert to 300mm, and whether or not we will be able to convert to 300mm, becomes a serious issue.

There are other things that are going to drive up the factory costs. The weight of a 300mm carrier or front-opening pod with 25 wafers is, I believe, about 17 pounds. This is heavier than typical and is heavier than ergonomically feasible for someone to carry on. Instead of having bay-to-bay automation like we currently have, we will need a lot of intrabay automation which will require a very sophisticated CIM system and mini-environment.

Some general observations are that it seems like the transition to 300mm certainly can be a win-win situation if the IC makers and their suppliers can get the costs in line, to be able to sustain the continued growth of the IC market, and if we can meet these productivity goals that we've set for the equipment suppliers and for ourselves. The use of larger wafers will continue to drive the progress of semiconductor manufacturing just as it has in the past in going from 4" to 6" to 8". The 300mm will cause significant requirements for improved CIM systems, single wafer processing, process control, automated material handling systems and other things that, at the 200mm level, we can sometimes get along without. The interface standards supporting both the ultimate solution as well as having to have another interface that supports some backup or more conservative solutions.

If we can't agree on a single set of standards then that will cause additional cost increase and cause

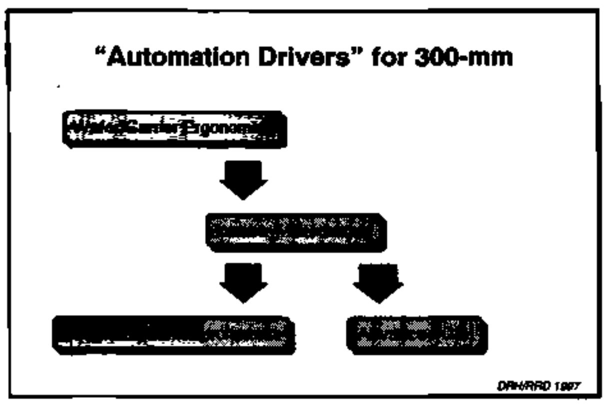

barriers to further productivity improvements. New levels of collaboration are being fostered in the global 300mm effort, as Dr. Hwang said. We've got I300I activity and the J300 activity. These are all critical things that are needed and we need a very good global cooperation force to be able to stay on this productivity learning curve. The next speaker is Dr. Takeda.

Dr. Eiji Takeda: Good afternoon, ladies and gentlemen. My name is Eiji Takeda from Hitachi, Semiconductor and IC Division. Today I would like to talk to you about the challenges for process and device technology in embedded DRAMs.

As you know, the new paradigm shift is occurring in system LSIs as follows. For example, megatrends of

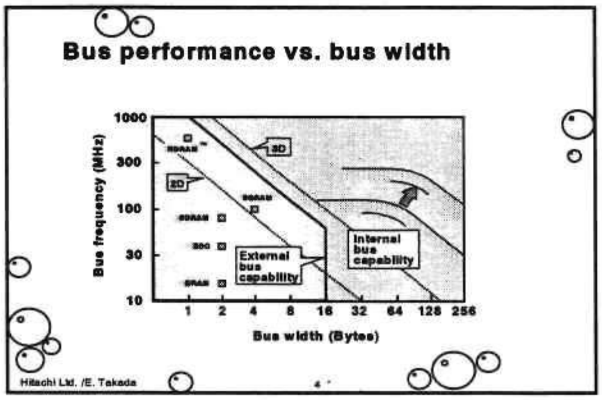

nomadic computing ending up PDA, and multimedia computing focusing on the three-dimensional computer graphics and so forth. These trends of new technology are strongly required; for example, embedded memories, 300mm wafers, production lines and so forth.

These are market fields. There are three market fields in every DRAM. One is the embedded system

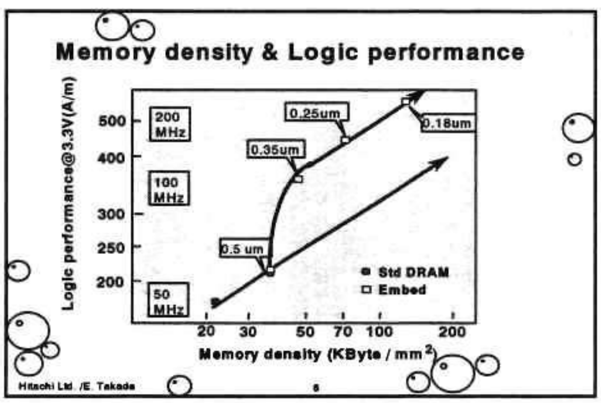

such as a DVD and inkjet printer. Second is the HPC and PDA aiming at the low power. The third application is 3D graphics and games in which we are aiming at the higher performance level.

Here we have the bus performance versus the bus width. These very high frequency region achieved by conventional DRAM approach. This region can be achieved by DRAM with high bandwidth.

This slide shows a road map of the DRAM. We are now developing 0.35 to 0.25 micron technology. In

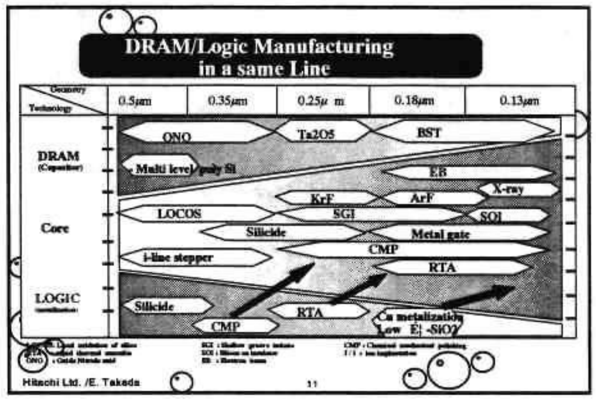

the near future, FLASH is also embedded in the single chip. In every DRAM, logic performance must be at the same level as the conventional logic performance. In the standard DRAM, transistor performance is a little bit behind the conventional logic.

Here is an example of RAM and logic densities. In the case of 0.25 micron, 1mm gate logic and 40megabit DRAM can be fabricated in a single chip and there's an assumption of die size equals 10mm square.

This is regarding time-to-market. In higher performance, the RAM-on-chip approach can

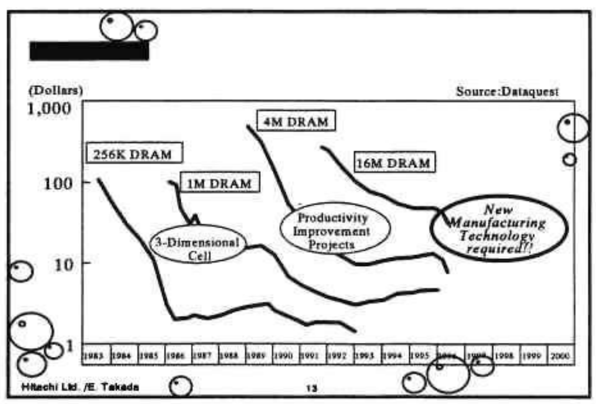

provide the shorter time-to-market because the memory architecture, bus design and card design cannot be neglected.

Concerning the modular layout in embedded DRAM, we have to provide the architectural DRAM capacity

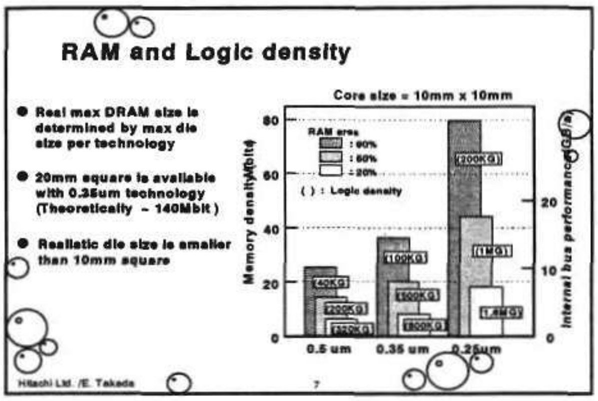

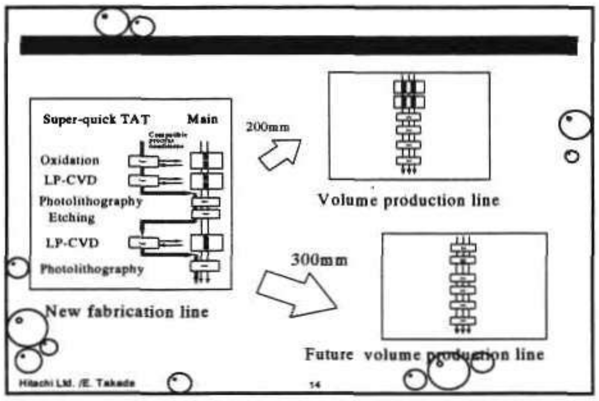

like a basic method. We are using the multi bank architecture. One bank consists of 256Kbit and we use the CRAM, DRAM technologies.

In the near future, we have to fabricate DRAM, logic and embedded DRAM in the same line. DRAM-oriented and logic-oriented processes must be added to the core process.

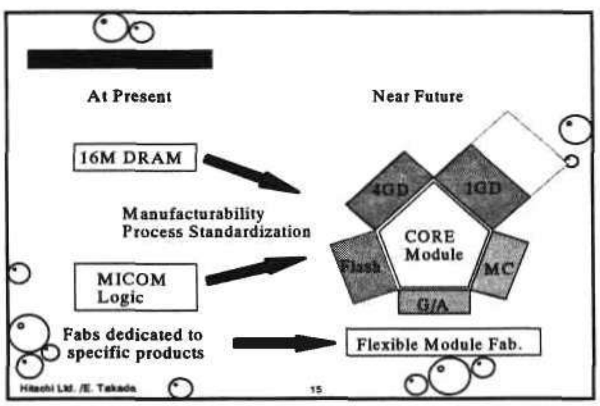

Here we have a comparison of the memory cells, trench and capacitor, in embedded DRAMs. Both memory cells are in the same level. In terms of flatness and gate performance, the trench capacitor is a little bit better than the stacked capacitor but by using the CMP and low temperature process, the improvement can provide better performance in stacked capacitors.

[Remainder of Dr. Takeda's presentation not recorded]

Q: inaudible

A: If the 300mm will have process, I think the cost problem is not so bad.

Q: inaudible

A: Well, yes, cost in general is proportional to the number of levels obviously, because you have price steppers. And, of course, DRAM is usually poly level driven and logic is metal level driven. By the time you add those two things, you have a significant increase. I can't remember exactly what the question was but it was in regards to those two.

Q: If there is a cut-off point between optimizing toward a logic versus optimizing process for a DRAM in terms of amount of memory.

A: I'm not sure. There's not a cut-off that I know of.

Q: Do you see one, Dr. Hwang, or is there a specific set of applications that is driving embedded DRAM?

Dr. Hwang: First of all, the embedded DRAM approach requiring more than 50% of area for DRAM -- maybe the process may be optimized with respect to the DRAM and far less than the DRAM requirement, far less than 50% in embedded DRAM. The approach may use logic process with pseudo DRAM cells and the direct number really depends on the available density of available technology and applications.

We have experienced this in the area of format. Format was 16mg area. Combination of this controller and memory is not a good candidate because the portion of DRAM in an embedded DRAM approach is quite small. We decided to use pseudo DRAM approach. Regarding the optimization and controllability issues, when you use a DRAM with stacked capacitor, certainly, for technical reasons, unless you don't use it, you don't gain the degree of performance in the EML approach. In order not to degrade you have to somehow manage a thermal process above it.

In EML, a DRAM with a stacked capacitor, is the reducing of thermal bodies in the capacitor process. I mentioned that the use of oxide looks like it can resolve the issue. Regarding stacked capacitors and trench capacitors, there are more apparent challenges in using stacked capacitors, but in using a trench maybe a lot easier in near-term or middle-term applications. There are other issues with trench capacitors in the long run. How to increase density. It may be 15gigabit or 4gigabit beyond.

As far as the size of the memory determining whether or not you optimize around the memory cell or the logic, it is more driven by the application than the actual size of the memory. Lower-end application may be okay, with a low performance transistor, but certainly there are real computer system applications with a real limited bandwidth between the processor and the memory. The more memory you can get on a chip close to the processor, the greater bandwidth you're going to have. In that case, some of the designs require very large amounts of memory. The issue becomes how big does that chip become and can you effectively yield that kind of a chip at an appropriate cost. If there's significant performance, then there's a premium that will be paid.

Moderator: Anything to add to that, Dr. Chang?

Dr. Chang: I certainly agree with Michael. It's not driven by the density. Actually, the density is driven by what kind of embedded memory you're using. We know that a lot of people actually are embedded with SRAM. Those are very low density - 250. With a shrinking scale going down, 16 transistor, you can't even get 1mg. Very easy to get improvements. That process is much simpler. It's exactly the same as logic. No problem at all. Even right now driven to 2mg. We have some customers that put 2mg SRAM, 16 transistor inside the logic chip. The performance is very good. The density, when you go to 4mg and above, you want to think about DRAM. This is basically driven by the application itself. Your focus is on the power saving or the performance of the bandwidth.

Q: What applications are you seeing embedded DRAM most likely to be set into. I'll only address this to the two on the end. Dr. Polcari going first.

Dr. Polcari: We see the graphics application as one that seems pervasive. There are system level applications for getting the memory close to the processor.

Q: Any other applications that you see, Dr. Chang?

Dr. Chang: Those questions still depend on the companies' goals. For example, there's a company in this field that's very famous and very successful. Their focus on the graphics itself is on a portable. Total purpose is saving the power. It's a graphic chip. It's probably the number one right now put out on the DRAM. Embedded DRAM situations, applications and graphics. I think everyone knows the company. That's purely the performance. It's not as good as any other independent one but their total dedication is for the power. The battery life could be 3 times longer. That's the selling point for any portable. They don't care that much about the speed of the graphics itself. Some of our customers are really in tune in the future for the speed of the graphics, 3D and those kinds of things. They want very wide bandwidth and high speed.

Q: One last set of questions on the embedded DRAM, and then we'll shift to a couple of questions on 300mm. The last set of questions has to do with the testing challenges associated with system-level integration and embedded DRAM. Dr. Hwang, can you please comment. What sort of testing challenges are there? Is test equipment available?

Dr. Hwang: In the embedded DRAM approach, there is one important problem which can affect the cost of the EML approach. Ideally, the EML may require 15 tests and scan tests. Otherwise, the testing costs could be very high. Right now, the initial EML application is mostly around graphic applications. A bit later, maybe impact plus memory with DRAM, that kind of application. Later, then, we may have application like single chip PC, as when you have DRAM. Then, certainly you can build a microprocessor with today's main memory requirement. In that case, testing is becoming really important. In the long run, definitely EML consider very much about testing issues and probably requiring testing facilities.

Q: Dr. Takeda, can you comment on the testing challenges?

Dr. Takeda: We are now talking about the test style. At the first stage, we have to test the memory part and the logic part separately. In the near future, we are thinking about the test, a system to reduce the test style. It will take a little bit of time from now.

A: The tester issue is very significant, because typically testers are optimized for memory or optimized for logic so now perhaps you have to have a memory tester and a logic tester, and then some integrated tester to see if it's working together as a system. Self tests and built-in tests are crucial for this application.

Q: Are those sorts of equipment available today or is that something that is still on the drawing board of many companies?

A: I don't think they're readily available to do everything that you want them to do.

Q: I assume that's on some tech roadmap somewhere.

A: Probably not, I'm not sure.

Q: Dr. Polcari, anything to add there?

A: I would agree. The test issue is significant. Particularly not only solely from memory and logic, but for more different types of functions you add to the chip. The different cores and the ability to test them is a very complicated issue. The question is, "Do you do a system-level test or a core-level test?" Some kind of built-in self-test is essential.

Q: Anything to add, Dr. Chang, on the testing issues?

Dr. Chang: I have very similar comments. The only other thing I see is that for many graphic applications, in general, the memory requirement is not as stringent as independent memory itself. The test procedure therefore is not as severe. You have to be done with the standard memory. Internally, there is a certain forgiveness factor in there. You factor this in, to help you reduce the test time.

Q: We didn't get too many 300mm questions brought in, but I do have one that came up at a dinner meeting I had last night, so I'd like to pose this one to the panel. It has to do with the relative capital costs associated with 300mm versus 200mm. There's been a guideline of 1.3 to 1.4; that's been earmarked as the guideline for economic return. I'd like to propose a hypothetical situation to the panel, and then ask the question as to what the most likely decision would be to the question.

The hypothetical situation says that we're at the end of the year 2000, entering 2001. We have a set of pilot lines that have proven feasibility for 300mm wafer processes but the economic, the yield issues and the learning curves, are not yet there, in order to economically go into 300mm production in the year 2001. You have to build a fab to meet current design or current capacity needs. I'm going to further put in this hypothetical situation that the decision is to run 200mm wafers in this fab that you're going to build to start up in the year 2001. The question is, do you also put 200mm equipment into this fab, built in 2001, that is 300mm capable for a planned upgrade in the future? That is the question I'd like to propose to the panel.

A: Can we change the scenario?

Q: Can we change the scenario? I will put one other constraint in. The constraint is that there are equipment sets available that are 200/300 equipment that do meet the 30% to 40% requirement for capital cost. The question is which kind of equipment do you put in this 2001 fab? Do you put dedicated 200mm equipment in, at a slightly lower capital cost, or do you put 200/300mm equipment in this fab? I'd like to start off with Dr. Polcari.

Dr. Polcari: First of all, you have to take into account what business you're in. If you're in the memory business, I'd go into the year 2000 with 300mm equipment or you won't be able to make money. That's an overriding constraint here. Those people in the memory business will make sure your scenario doesn't happen. If you have to follow that scenario exactly, and you really believe that you can't make 300mm work at the beginning, you'll start at 200mm and try to change.

Q: And try and change in the middle. So, 200/300 is the scenario. Dr. Chang?

Dr. Chang: I think there are two related questions. One of the questions is, when you were saying "not meeting the criteria," is how far away is it? If you say it is 1.5, it's not too far away. It may be extrapolate the curve if you said I predicted years down the road it's going to be happen. In that situation, the choice is very easy. If it is far away, you talk about the four times higher or those kinds of things, you say it's hopeless within a couple years or three years for this to reduce the cost. When your choice, it is natural to go to whatever is available right now.

Q: So, if the economics were to be based on the learning curve and the 2002 or 2003 timeframe, the equipment you buy in 2001 is 200/300 equipment.

A: Another factor that is very important is whether you'll see the equipment as mature. If it is 200/300, whether its performance is as good as 200. This means if the 200/300, the equipment, is not mature, then the uniformity is a disaster. It's not going to perform. Then you take a bigger risk instead of your doing just only the cost itself. It's not only expensive, but also the risk is very high that your process is not going to be there. That's a problem. It's not a manufacturable process. Maybe you have another choice that you have to think about. The third one is very important: whether later on they're going to support this 200, from an equipment vendor point of view. You have to look into all of these three factors. These are very important in my mind.

Moderator: Dale?

Dr. Harbison: First of all, I don't think your scenario is going to occur. I feel like we've got enough data to show that. I don't think there are any technical challenges in going with 300mm equipment. There may be a few, but we've seen enough data already to feel confident that we can do it. The guideline is not 1.3 to 1.4, it's 1.3 period. It's got to be 1.3 with the same footprint, with the same throughput in order to be economically viable. Our suppliers have gotten our message, and I believe they're progressing down that curve to get there.

It would be very difficult to have a 200mm and there's also going to be very little equipment that's 200mm upgradable to 300mm. You can probably pay the extra, get the 300mm and run 200mm wafers but I don't think you could justify doing that. You can't spend 30% more and still run 200mm wafers. You couldn't survive. As Peter said, if you're close and you feel as if you can do that, then you would probably go ahead and do the 300mm or you would continue to delay. What we at TI would do, we'd just delay, as long as we possibly could, having to do that. We'd try to get additional productivity out of existing factories, or somehow cover the extra six months (or whatever) it took to meet the metrics that you have to have to do the conversion. I don't think anyone wants to have the last 200mm factory.

Moderator: Dr. Takeda?

Dr. Takeda: There are a couple of 300mm wafer consortiums in Japan, in which research is being done on the many equipment configurations for the 300mm wafers. After looking at the results from these consortiums, we have to decide when and which equipment is to be used. In my opinion, it takes a little bit longer time to decide. It is difficult for us to change the 200mm wafer equipment to the 300mm equipment. It is very difficult to say at the moment.

Q: So, just to make sure that I understand the answer and we all understand the answer, you're suggesting that 300mm readiness is going to be delayed, so you believe that you would probably buy a dedicated 200mm wafer FAB equipment set in the year 2001. Dr. Hwang?

Dr. Hwang: Assuming the scenario on the memory side, maybe, right now, I think 64mg does not desperately require 300mm wafers. Right now, the technology volitional variable for 64mg DRAM can produce enough of a number of 64mg dies on 8" wafers. However, if we go to the 256mg DRAM area, then 300mm wafers amy be required. I make an analogy to the tradition from 6" to 8" and at that time format to 16mg DRAM. Actually, this is a complex problem to be optimized, because there are many variables involved, and variables circled around. Unless fixed on one thing, it is a difficult task.

Q: To summarize your answer so we can get a consensus, do you believe that the 2001 scenario will probably be a 64/256mg DRAM fab?

A: Actually, I would like to use the product rather than the year for the timing for conversion to 300mm.

Okay, so it's possible that for a 64/256mg fab, you would buy a 200/300mm equipment set. A couple of years ago there would have been a consensus that there would be a dedicated 200mm fab and a dedicated 300mm fab. Today, we have a split decision as to whether or not that would be the scenario. Actually, it seems to be edging towards the 200/300mm equipment set.

Moderator: That wraps it up. I hope it was enjoyable to everyone. I thank the panelists. I appreciate their time dealing with a series of difficult issues and questions.

Chapter 7: Semiconductor Interconnect: The Market Differentiator

Moderator:

Mohan Warrior
Director
Strategic Final Manufacturing,
Sector Manufacturing, SPS
Motorola

Panelists

Ed Fulcher
Director of Development
LSI Logic

Steve Anderson
Senior Vice President
Amkor Corporate Product Marketing

Tom DiStefano
Founder
Tessera

Dr. Rama Shukla
Manager for Components Technology Development and Advanced Interconnects.
Intel

John Novitsky
Vice President of marketing
MicroModule Systems

Moderator: The last session has been on demand from our clients for a long time. It is an in-depth discussion about the role of semiconductor interconnect. I am very pleased to have with us one the key figures in this industry as our moderator, Mohan Warrior from Motorola. He has a very distinguished panel that he'll bring up and then conduct the panel session in the way that they have decided among themselves to conduct it. Some of the panel sessions have been a little bit different throughout the conference and we like that. We like some creativity here depending upon the market and the people involved.

Let me introduce Mohan. Mohan is the director of the Strategic Final Manufacturing sector at Motorola. His principal assignments include managing an advanced bipolar wafer lab, directing the advanced interconnect effort in the product sector to drive the bump technologies and develop new multilevel structures. Mr. Warrior has authored several publications in these areas. He has a master of science degree in chemical engineering with 13 years of diversified experience in the semiconductor industry. Please welcome Mohan Warrior

Mohan Warrior: I am privileged to introduce my co - panelists this afternoon. First I'd like to introduce Ed Fulcher. He is the director of development at LSI Logic and he has authored papers on multichip modules. He has received several patents; he has an electrical engineering background from the University of Florida and a business background from Stanford.

Next I'd like to introduce Steve Anderson. Steve is a senior vice president at Amkor Corporate Product Marketing. He has been about 20 years in the industry; he also has a bachelor's degree in EE and an MBA.

Next we have Tom DiStefano who is the founder of Tessera. Tessera is getting to be a home word in packaging. He has several patents in the area of packaging. He has a Ph.D. in applied physics from Stanford.

Dr. Rama Shukla is from Intel and he is the manager for components technology development and advanced interconnects. He has a Ph.D. in materials science from the University of California. I am privileged to have him here.

We have John Novitsky representing MicroModule Systems. John is the vice president of marketing there. He also has a background in computer science from Michigan State University. So that's our panel.

It is always difficult to do the last session on day one of a conference because the weather being very good outside this session is what is preventing us from getting outside. However we hope to make it very exciting as we are going to talk about some of the topics related to how we can differentiate in the advanced interconnect arena. What we as packaging technologists and people who bring the interconnect world and be able to connect to the silicon technology we talked about in the earlier session this afternoon and make a difference for making the industry move forward. That's the topic for this afternoon.

I'll start off by just talking for about 8 to 10 minutes, putting the framework of the discussion. Firstly the domain we are going to talk about is chip-to-package and package-to-mold. In some areas we will overlap into the first interconnect let's say the bumping and those areas which enable us for what we want to achieve at the system level. Let's rapidly go through what we think are the advantages and try to delineate as I go through my presentation on the performance issues, costs and some manufacturing concerns.

Some of these areas will overlap as I show some of the analyses for the industry. Finally I pose some questions to our panelists and as I illustrate some of the barriers on cost performance and manufacturing in their presentations. I'll be addressing some of them and highlighting what's possible.

Firstly let's just set the base line very clearly because we do understand that certain functionalities are what a customer buys a majority of the semiconductors for, so the differential we can offer in terms of the basic functions that package serves is listed there in terms of production and convenience and how for example if we have, as you heard earlier, copper and locate electric etc. we don't know that interconnect advantage at the package level. Also a set of definitions which we will be going through in the next hour and a half. We'll be using these words: peripheral and array packaging; most of you know about it but I just wanted to clarify that that is what we mean when we talk about those issues.

What we are seeing is obviously a trend towards array and fine pitch to increase density. I'd like to highlight the terminology. It is getting pretty complex because we have BGAs, microBGAs, TBGAs, CSP. Obviously those are array technologies which are being driven as a function of what you see on the two axis, they are the Lead Pitch and the Pin Count Density. You can see a trend all the way down to a Pin Count Density of 300 to 400 pins per mm square. Those are the three things that I'll try to bring out and the other panelists will talk a little more about that.

In terms of size the issue is that this product focus that's driving us for smaller and smaller outline packages like we started from the SORCs and BGAs and now the chip scale packages and then finally going to direct chip attach. Just to illustrate that you can think of a CSP site as compared to a standard either a pure QFP site and you can see about a 75% reduction in board size by using chip scale packages which at the system level brings a fairly high degree of cost advantage in the end application. If you think of DCA it's been a controversial issue for if the industry will go directly to DCA will go through a CSP DCA type of scenario, I think I'd like to highlight it's an application specific domain in the sense that there are some spaces where DCA makes a lot of sense and it's worth going over the hurdles of DCA and making it happen.

On the other hand where there is frequent redesign driven by shrinks in geometry then you have one of the star interfaces which will have to retool on the board level all the time. So the interface approach of a microBGA or a CSP may be favorable. You can see some of the applications. Also one enabler for the DCA is the sort of Bump technology for flipping from the computer generated graphics to a foil but this is again one of our systems standardization problems. The Intel onslaught got us because the "xx" logo and the "xx" won't transfer very well.

I want to talk a little bit about the solder bump technology because if you think about it bump fabrication is more akin to wafer fab technologies. The cost issues are similar. In a sense if you think of populating a wafer fab, a large amount of capital equipment goes into litho equipment and vacuum processing equipment. Whether is PVD, CVD or implantation equipment it's vacuum processing equipment.

As we can eliminate both litho and vacuum related processes you will save cost in terms of the manufacture. If you think of the old C4 type process from IBM way back into the '60s to a mixture of evaporated or spotted and electroplated, you are going from a lot of vacuum processing single mask technology, metal mask in the old evaporated case, to full lithographically generated resist and still spotted, so you've still got litho and vacuum. Finally on the bottom line you've got a Bump technology which is maskless and does not use any vacuum technology. So that is where the cost differential will come through.

Just to conclude on that, the DCA should relate to bumping which is not for every one but for specific needs it will have very good win at the system level. In terms of array packages if you think of the BGAs, as you know there will be widespread applications whether it is conventional BGAs or emerging micro BGAs that I'm sure Tom will talk about. There I think the issue will be cost driven for what I'm calling wire bond technology based on finalization efficiency. If you think of ordinary packages the final efficiency will dictate a lot of that because substrate costs are getting to be pretty enormous. Then as we get into the high density interconnect work that will become even more dominating.

We talked a little about the size in terms of the performance and the cost issue. The electrical performance; this is comparing from a QSP to a CSP in terms of taking any one of those parameters. You can take self inductance or mutual inductance, you can get a fairly substantial performance improvement and that will drive where you have very high speed logics where we're spending the money in the wafer fab to gain a low inductance circuit both with copper conductors as well as low K directors you don't want to blow it into the package. We see the other reason why we are looking at it so heavily. The mold array package, you'll see versions of this probably in Tom's presentation, but this is generic definition and again I want to highlight what we call the MAP package which is the molded area final package is driven primarily because of the efficiency you get at the final per cost savings.

The issue is standards. One of the barriers that you may have in this scenario is that you can see here almost three times as many variations of chip scale packages of microBGAs. So I think there will be a need over time for having some kind of standardization which will help us work the same way that we did in the wafer size or in the final size. The technology that we're favoring is something called JACKS-Pak. Again this is an acronym driven word that stands for Just About Chip Size, it is not quite chip size. We are looking at that because it is giving us a little more reliability improvement especially in fairly high chip size in the sense where you are thinking of a die size which is fairly large.

Lastly I want to end on the reliability before turning it over to the panelists. I will talk a little bit about design also before I finish. On the reliability slide I did not chose to put the various packages, I just listed them as a ABCD etc. Various published data are shown in terms of the number of I\O and the temp cycles that it can stand and I've taken the extreme conditions let's say -40 to 125. As you know some of the automotive guys will want us to go from -40 to 140 but it looks like we have to match about 1200 cycles. That will be a requirement for robust applications. That will be a challenge that we will have at board level reliability.

What is driving us of course is the miniaturization. You've seen the hands-free type radio and you've seen the BGA and the wisar radio; and the last one doesn't exist by the way. That's the kind of thing that this package technology enables us to do.

Lastly on the design I think we have unfortunately not as much standardization in design tools on the package side as we have on the silicon side. The earlier we can standardize that given a particular chip layout configure the BGA or the CSP fast enough and be able to give both electrical specs and hopefully over the time the veracity of our mechanical bond still leaves a lot to be desired but where we can ascertain that, that will drive the design issues to a simpler level.

I want to conclude on two notes. I think that reduced size and improved customer solutions are basically key enablers for us at least on the product side in our company but in other companies there will be other drivers. The choice between direct chip attach and array packaging will be application dependent and will have to match benefits you have on current QFPs and things like that. There will be a timing of insertion that will take place.

Moderator: Now I'll call upon Ed Fulcher from LSI Logic.

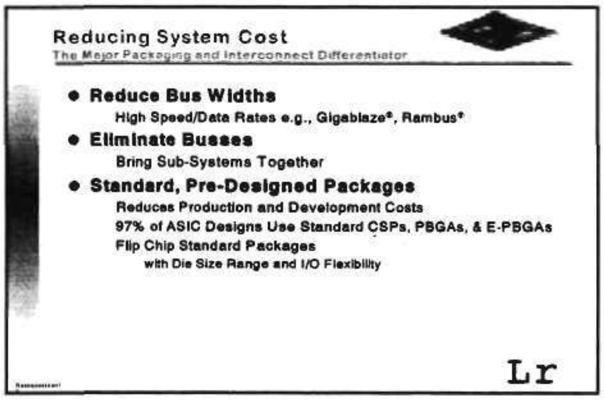

Ed Fulcher: No matter what your perspective is the system is really a massive array of interconnects. Occasionally you find an active element in that massive array; it looks this way when you look at your rack when you look at your back plane when you look at your board when you look at your package. It actually looks that way when you look at your chip. So this panel is going to address at least some of those levels of interconnects.

I was asked by the moderator to look at what is a major packaging and interconnect differentiator. In my way of thinking it is reducing the system cost.

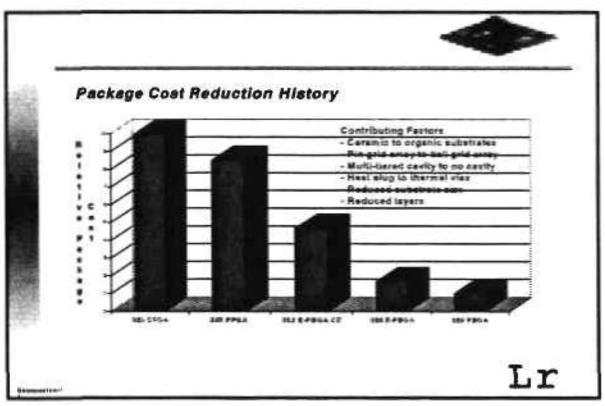

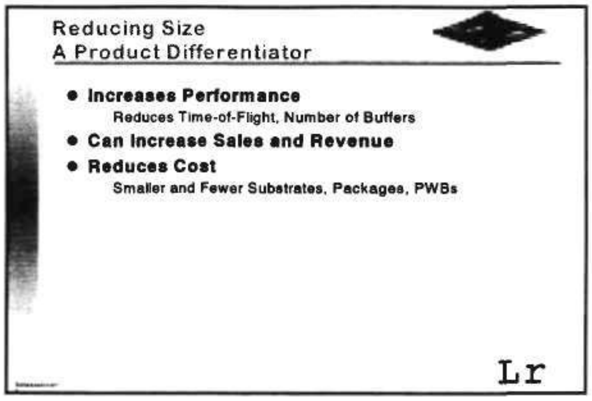

That is what drives the businesses, that is what drives the customers. It is to get more in a system for less money. One way packaging can do this is by reducing the bus widths. When you reduce bus widths you can do it with higher speed, so that you can transfer more data on fewer lines. Most of you have heard of Rambus some of you have heard of Gigablaze that's a gigabit on a pair of differential lines.

These technologies are already in place, already being shipped. One customer now has up to 12 Gigablaze circuits on one IC. We 're designing with up to 32 and then 64 and going single ended. So this is going to reduce buswidths while dramatically increasing the data rate, significantly reduce the cost of high data rate transfer. Meanwhile the speed will go up; it's a gigabit now, next generation two, next generation four. It will continue to go just as silicon density goes.

Another way is to eliminate the buses all together. We've heard talk about that system on a chip, multiple systems, multichip modules; anything you can do to eliminate buses either eliminate them between racks, between boards, on the board, on daughter cards or in the package or bring them all the way onto the chip. Packaging plays a role in virtually all those aspects. A third way that the interconnect in packaging can reduce system costs is especially for the ASIC business that we are in, we can't afford a custom package, for every design, otherwise we'd be in the application specific package business, the ASP business, so you must have standard predesigned packages. We find that over 97% of all our designs go into a standard package. This helps to reduce our cycle time also and get you to market faster. If you don't have to tool a package which typically can take longer than tooling an integrated circuit.

We once thought that when we moved to flip chip we wouldn't be able to do this, that the marriage

"System-On-A-Chip"
Reduces Both System Cost and Size

- **Capable of Multiple:**
  - Processors
  - Memories
  - Core functions
  - I/O Types and Mixed Signal Functions
  - Plus Custom Logic
- **Systems Need Fewer**
  - Silicon Chips
  - Packages
  - PWBs
  - Internal Interconnects

Lr

between the die and the package would be so great

Leading Challenge
Greater Interconnect Density

- **Smaller Systems with More Capability Need Greater Interconnect Density**
- **Denser Silicon Die with More Functions and I/Os Need Greater Interconnect Density for Die and Packages**
- **Smaller Packages with More Leads Need Greater Interconnect Density for Packages and PWBs**

Lr

that we would have to have custom packages. We find that we have been able to use design techniques to have standard packages and recently announced the full family of standard flip chip packages up to 1700 leads with a 1000 I\Os. These packages even give the same kind of die size range flexibility that you have with wire bond standard packages. You can have an infinite variety of die sizes in each package. The other thing is you there has been a proliferation of I\Os all different kinds, voltages, speed standards to satisfy the broadening applications of the communication and consumer and computer markets. These packages have the flexibility to provide low inductance sources for a variety of I\Os and voltages.If we didn't do that we would have serious trouble in the future.

Greater Interconnect Density

- **Surfacing Needs for New Packaging Materials**
  - Packages Moved to Relatively Known Set of PWB Materials and Processes
  - New Class of Materials and Processes Needed for Both Packages and PWBs
    - No fibers
    - Moisture Resistance
    - Thermal Expansion Closer to Silicon

Lr

Package cost itself has been reduced significantly over the past five generations. If you look seven to ten years ago we were in ceramic packages that the first bar, it's a relative scale of cost, cost of the assembled package. Then we went to pin grade arrays, then we went to a plastic pin grade array which was a printed circuit board laminate type and reduced the cost there about six years ago. Then we moved to a ball grid array from a pin grade array allowing us to be higher density, smaller package and reduced the cost again.

Then we even took the ball grade array and simplified it, eliminating the multi-tier wire bond shelves, changing from a heat slug to thermal plug and reducing the number of layers; so we now get more performance and lower cost in what is called an enhanced PBGA, and then to a two layer PBGA that you are also familiar with. If I were to plot the next three bars on here I think you would see low cost flip chip as one of them. We are all working on that and it has not emerged yet but it will soon. Chip scale packages would fit in, especially if they are also low cost flip chip. And finally chip on board are likely to be the next three bars on this curve.

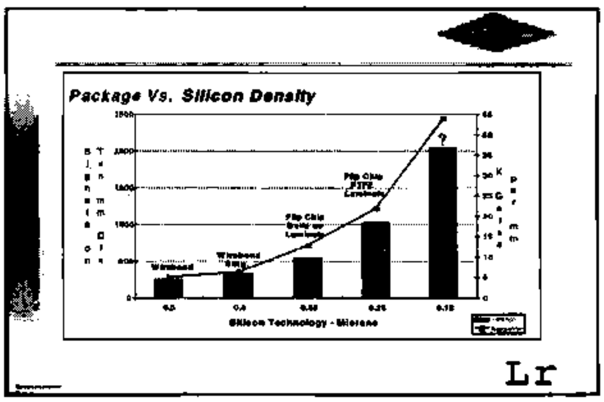

We were also asked about differentiators. What is the product differentiator for packages? You saw in the earlier talk and in many talks that reducing size is the product differentiator. You saw a number just before me on how it increases the performance. Reducing size anywhere increases performance not just in the package. It can also increase sales and revenues. In many products people want it to be smaller. The first people out there with the smallest application will get the highest number of dollars and the most customers. So you can get more money for a product that cost you less. Isn't that interesting! That is really the way to go in size and it doesn't apply to just portable products. It applies to all products. All our customers from the highest end to the very lowest want less size in their next generation.

One way to reduce both system costs and size is system on a chip. You've heard that by many speakers so far today. We are no exception.Possibly the first semiconductor company to call themselves a system on a chip company. We are shipping systems on a chip now that have over seven processors, over 40 different custom designed memories and numerous core functions which includes the intellectual property that was covered earlier today, a variety of I\O types, mixed signal if you're going to be a system on a chip you must have mixed signal on a chip also, and then custom logic. It is interesting that in an earlier session today it was mentioned that by the year 2005 the driver of the industry could be system on a chip products. I had not heard that before so that was very interesting to us.

The other thing that I was asked by our moderator was what the leading challenge for interconnect. The leading challenge we face is greater density. I should have added the words "greater density at lower cost" because you can always have greater density if you are willing to pay for it. Greater density is the challenge. As we go to smaller systems, denser silicon, more I\Os even though we may be reducing the number of busses and shrinking those busses a form of Rents rule still applies and the customer still want more I\Os than they had before in the previous system. So more I\Os, smaller packages, denser die, they all ask for the same thing: greater interconnect density for less money.

When we initially moved from ceramic to laminate type packaging, organic packaging, the industry moved along the easiest path it could to a known set of printed wiring board materials. It already existed worldwide. The infrastructure was already in place. Just make it a little denser, cut it up and call it a package. You may have to pass a little better reliability test because you are closer to the silicon die and you have to protect it. So you evolve into those things and that's what we all did. That's the barriers that we have difficulty crossing now on density.

What is needed and what we are seeing now in the companies we are partnering with for the next generation of packages are for new materials, new processes and this will be true both in packages and printing wired boards. You hear of pieces of this like microvias and board materials that have no fibers because the fibers get in the way of the vias, materials whose moisture does not get absorbed and does not cause problems, very important when you are going to be next to silicon die.

One of the most important of all is that we go to great pains in extraordinary complex solutions because a printed wiring board does not match the coefficient of thermal expansion of silicon and you've seen in the literature and you'll hear of some today of solutions that require pin springs, underfilled materials and all sorts of things we do to try and match this when the answer could be well lets make all our packaging materials match silicon, rather than go to elaborate pains to take up the slack in between.

The companies who are working in this area are the ones that are going to come up with the revolutionary solutions. What we are doing now is evolutionary solutions for the printed wiring board industry rather than the next revolution. You can see touches of this happening; a company called Gore has announced a material that meets many of these criteria in packaging. TSP recently announced a high temperature coefficient ceramic which by the way is also a low dielectric constant, which has been one of the nemeses that plagued the ceramic industry and why they haven't been in quite so many of the new designs. By covering both of those at once you may see ceramics rise again to the forefront, or things like ceramic.

This curve shows a comparison of density of silicon versus density of packaging. The red curve on the top is the density of packaging in thousands of gates per square mm across five silicon technologies from.5 micron to.18. The left axis and the blue bars are the density of number of input/output signals not counting power and ground, they are extra, that we can have on a 10 millimeter die and can effectively escape in our package to the printed circuit board. By effective I mean with a package that is cost competitive. We are not talking about a forty layer ceramic package, we are talking about very competitive packages and how we can do it. You can see that we used wire bond and then we were forced to go to staggered wire bond to keep up with the silicon density of the.4 micron generation and now flip chip using the term "built up laminate" which many companies are developing around the world as the first major flip chip technology, doesn't quite keep up.

When we go to the PTFE type materials such as Gore has announced with the next generation we certainly make a huge step but we still don't quite keep up. Look what is about to happen to us in .18. What technology is going to get us there. We've already used all the known ones in the first four bars. We've done fine pitch, we've done staggered wire bonding, we've done flip chip area with interconnect, we've done micropitch via lines and spaces. These are all products coming to market now. To keep pace perhaps we need to take on and track silicon interconnect technology. But just a few years ago and you've seen this in conferences and articles where the package may take on some of the aspects of the top couple of layers of silicon. It needs to take this on just to escape the signals. So you'll see technologies like silicon blended with these new materials that I discussed for the next revolution in packaging. Thank you.

Moderator: Next is Steve Anderson.

Steve Anderson: I am from Amkor working with Anam semiconductors. We are involved in the segment of interconnect called packaging traditionally a not heavily tracked and typically not an afterthought for most companies. Everyone focused in the past years heavily on the FABs and what you'll find now in the future is what Ed talked about the packaging or getting that silicon to a circuit board to get to a product is going to be a major limiter. What I show here is two charts. If you look at the one that angles up we play in the red circle

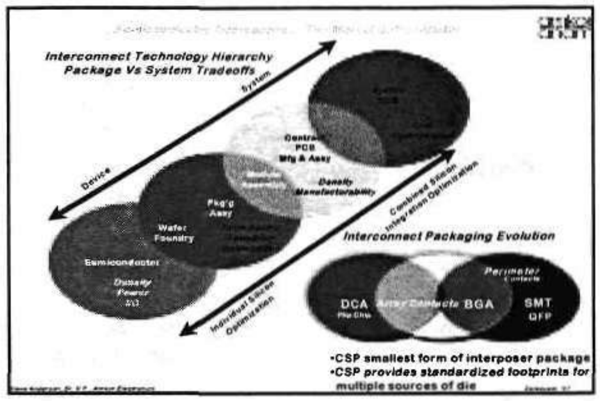

which is titled the packaging and assembly. We do the assembly and test part. Our job is really to work with the semiconductor companies who may have their own factories or may be foundryless and come up with a package that will get the right price performance for the OEM.

What we've seen happening as was discussed this morning very heavily is the compression in the food chain, happening very rapidly. Our company and others in the packaging industry are being pulled or pushed or thrusted in these two new areas where you see the overlap between the purple and the yellow circles. Where we are talking about purple or semiconductor we are starting to see alliances forming with wafer foundries with packaging people.

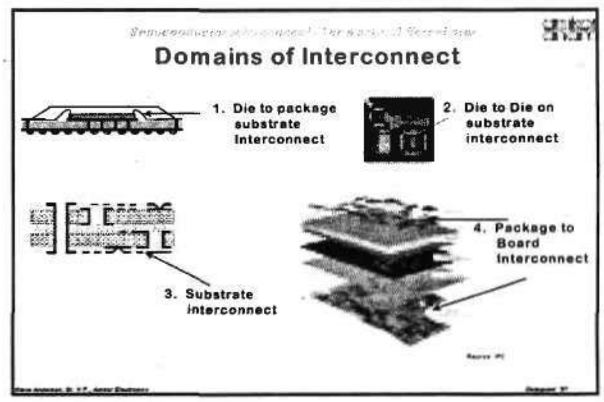

Our company is involved in that and you'll see many others in the future which means that we are going to understand much better hopefully the issues that Ed brought up. What are those interconnect technologies trends going on in the wafer, how do we help semiconductor companies address those with packaging. In the yellow circle we are going to see that the typical board industry is going to change

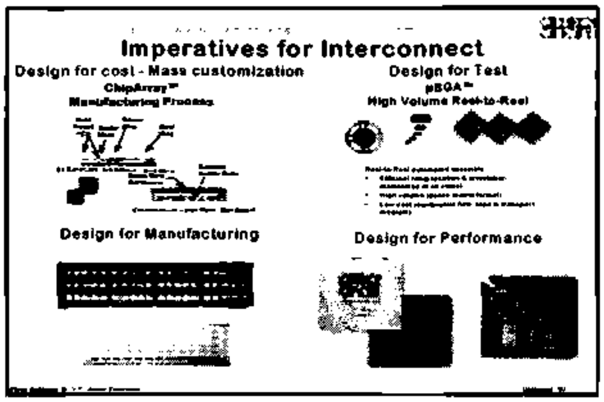

radically too. These big mother boards that you used to see are all shrinking. All of the end appliances you are buying are smaller, the phones, the PCs, the laptops. What that's forcing us to look at is some different packaging approaches that may not be your traditional one chip package. One of the speakers today is going to address the module assembly.

Down in the bottom we going to see a proliferation of packages, what we call in our company the third wave of packaging. Early in the start of semiconductors you had dual in live ceramic packages. The second wave came about mid '70s, late '70s to early '80s was a surface mount. Now we're really in the part that we talked about earlier the array packaging. Underlying trends of all of these and what will make this packaging very challenging for all of us in the next few years is the infrastructure.

The infrastructure to us is the support of the laminates be they tape or PCB type materials which doesn't matter whether you are doing direct chip attach or flip chip or whether you are doing chips scale package or doing a BGA. That industry right now has to grow and develop to be able to supply millions of packages. To us that is an interconnect that we will get into. We as a company have only dealt with number one interconnect, that was the die to package substrate. Most of these were done to a ceramic or metal lead frame, now we're forced to deal with four levels of interconnect. The typical die to package we are doing multichip packages and things as low tech as 8 leads SOICs, on up to 800 leads. I think our densest is a 1600 or 1700interconnect 4 chip solution that goes to a thousand external leads. The substrate technology is really been underdeveloped and we are pushing hard, as is

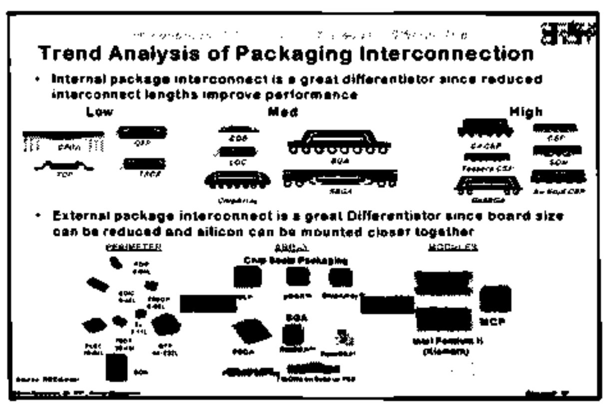

Intel, Motorola and many of the other people in the BGA area. You've got to do a lot more interconnect in that board.

That is an industry that typically was doing very large panel mother boards, not semiconductor grade technology. We learned the hard way with BGA that we were not able to get the buys we needed to support a lot of our early customers. and ramped the way the BGA should have ramped as the semiconductor people want them to ramp. I think that that's one that we have to pay attention to, it's one that's not as glamorous as a lot of these other technologies. The fore point is the board, the package and where that's involved is understanding the reliability. Mohan talked about meeting 1000 to 1200 joint reliability. Our company never looked at that. Most of that was done by the OEM prior to this, it was well understood. On the TQ of Ps, Q of Ps, PLCCs or any of those technologies we now are asked by all of our customers that as we supply those new packages, no matter what the version, that we give that level of reliability. Now we've had to add surface mount capability to all our factories and this was foreign to us in the past.

Our imperatives for interconnect deal with four areas and they are not necessarily exclusive. Number one would be design for cost, and Mohan touched on that with what I think he called the MAP package, we call it chip array. That's nothing more than taking a panel, it could be a two inch square or four inch square, overmold it and cram as many die into that as possible, putting them as close as you can and then wire bonding or ultimately flip chipping. Using standard wafer saw technology and sawing so you get the smallest

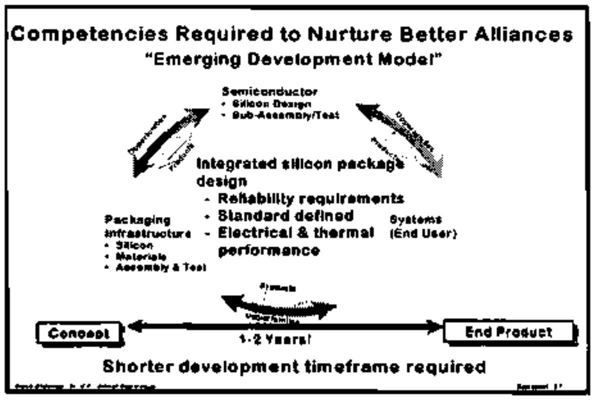

possible package or waste of laminate. The next thing that would come up is design for test. We see that falling in two areas for packaging. One where you are dealing with gravity fed testers and the other where you deal with pick in place. In the gravity fed where you are typically dealing with small lead count say eight through 60 leads we see the industry moving smaller or moving slower possibly to chip scale unless we have a good cost effective test solution.

A lot of the feedback we get from customers is don't talk to me about this chip scale until you come with a good test solution that does not raise my test cost from what I am now employing which is typically gravity fed parallel test. We are going to struggle a little bit as we get in that low end and as Tom will talk about with the Tessera package we are trying to look at approaches like reel to reel that will meet that objective. The design for manufacturability. We've got to try to move these new packages we are bringing up into existing infrastructure until new infrastructures can be developed. You are going to find us doing panel or strips that match most of the equipment built today in the packaging industry. Then the area we struggle the most with what MMS will address is design for performance. It's probably about 10% to maybe 30% of the market, not real high numbers which means that a lot of the substrate suppliers don't get very excited until you can really work with them but it is going to be hard to build materials sets from our viewpoint.

The interconnects get very complicated. We see packaging broken down into three areas and this is driven by cost. We get a lot of debate with our customers as to give me best interconnect you can to get the best speed but at the lowest cost. Many times they will take a step backward at technology in order to keep the cost low.

So in the low area we brake into two areas internal to the package and external to the package. The internal to the package in the low end we use standard wire bond. In the medium area we would use gain wire bond for the most part but we use a board technology to do a little more interconnects underneath the part. In the high end we are working on a number of approaches such as flip chip technology but still use substrates where we think the flip chip on substrate will be key. This is so that you can deal with multiple sources of die and yet provide a common footprint to the board assemblers. On the bottom side we see that we are moving from the perimeter to arrays clearly a direction where it is going to be where you can give that array package at a cost effective level.

What does that mean? Our customer would like to see it at the same cost or cheaper than they are now buying the lead frame or the standard technology. That is not always possible because the infrastructure and volumes are not there today. So what we typically do is ask a lot of questions. If you have a 15 square mm for example BGA versus a 24 square mm lead frame or perimeter product can you shrink your board. If you can the right solution is the BGA. If you can't get any cost reduction out of your board you may want to stick with standard lead frames till the BGAs come up to speed. We see a lot of customers playing with those types of options.

We see a new competency need to be developed. This is a much tighter working model with semiconductor manufactures OEMs and the people involved with packaging. That would be companies like ourselves as well as the material suppliers. We've got to move much faster

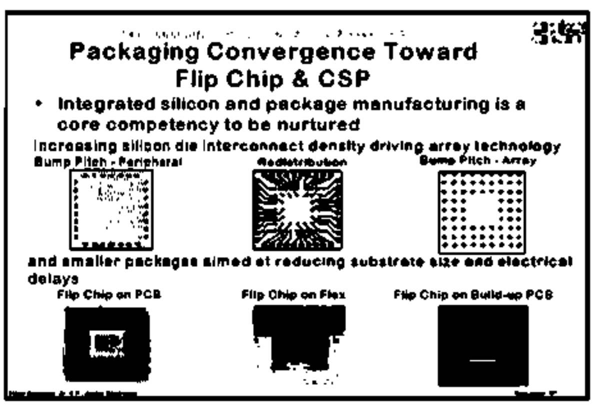

to understand what the reliability requirements are. Mohan talked about standards. Those are absolutely key. We can't be doing 15 package versions and still be making money. Then we have to understand the electrical and thermal performance because there are trade-offs there. We see flip chip coming but we see it coming in probably some phases. A lot of it has to do with the technology. If you deal with some of the stuff that Ed was talking about where they are moving very fast as system on a chip obviously the flip chip will be moving much faster whereas when you are dealing with hand held appliances where cost is key then you are going to be moving slower.

So we see three areas, first is just bumping a perimeter pattern where you have 100 microns or 200 micron type die, typically works well with memories or a lot of the larger dies. Next step would be redistribution again with a standard die and this approach is popular where a customer may want one die design but he wants ten to twenty percent flip chip. He'll want two die designs. The third area, the bump pitch array is usually where to go and redesign your die and it is for flip chip. Obviously then matching all these have to be the substrates, which is a big problem today in the market. It is finding enough volume suppliers who can build the technology at semiconductor quality.

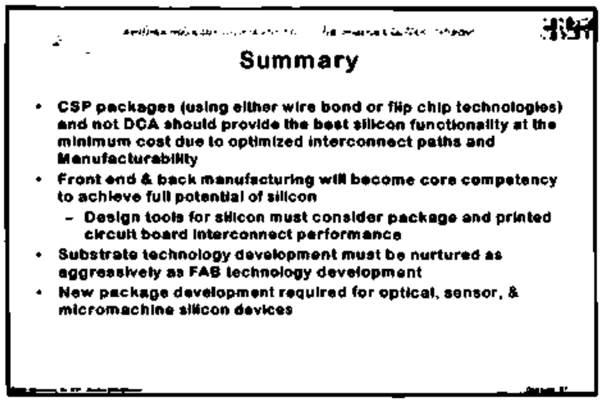

In summary we see these major trends:

CSP packages from our vantage point, we are looking at about 200 semiconductor companies, will certainly fill the void for the low lead products until the direct chip attach solution are more cost effective.

Front-end and back-end manufacturing whether it is one company or not must be integrated and there must be better design tools. We deal with three design tool sets right now: at the wafer level, at the package level and at the board level. If they could all be working off one cell library you would see much better products and they would be much cheaper. Today that's not true.

We see the substrate technology development must be nurtured.What is happening in BGAs, we could have built many more and shipped them but we could not get laminates that met the quality requirements. That is a major problem if companies want to move faster to this BGA technology. It doesn't even matter either tape or laminate, both are going to be limited in the short term.

The last thing that we see moving up in packaging is we are involved now in micro machine optical and sensor packaging. A lot of it off the BGA trend but it is a thing that needs to be understood better and nurtured. We'd like to see even Dataquest do a lot more work because we think that's going to be very exciting as optical packaging grows.

Thank you.

Moderator: Next Dr. Di Stefano.

Dr. Di Stefano: It's been three years since the CSP (Chip Scale Packaging) first surfaced. That's a very short time in this industry and judging by how far CSP has come in those three short years it's difficult to look ahead and see where this field is going but I'd like to look ahead a little bit from where we are today.

Recycling to provide a low stress link between the solder ball and the die pad. That bond can be anywhere, it can be in the perimeter, center area, right, anywhere. It is the same basic mechanics. That bond ribbon allows you to attach the package to any substrate of any expansion coefficient. The next key part is the relief of the mechanical stress on the solder ball. This shows a stress map of a solder ball. You can see by the red that the stress is really pushed up into an elastomeric layer on which the solder balls float.That decouples expansion of the substrate from the chip. If you look at the structure of this package you look at the guts of it here there are no laminates to trap waste or to delaminate, so the package is very moisture resistant. The first CSP packages you are going to see are personal electronics where the driver is size or form ones used for.

In 1997 you see small lead count packages being used for flash memory and some of these personal electronics products. The next step in the propagation of CSPs is for standard DRAM whether it is a high speed bus like a RAMBUS company name or synchronous DRAM where standards are being set for chip size packages of high pin counts. Here the drivers in the memory area are performance for RAMBUS and the high speed data busses and pin count for the high pin synchronous DRAM products. From there we expect the CSPs to migrate to higher pin counts really being gated by the availability of high density substrates to mount these high density packages to. Then for small dies you see some fan out configurations or fan in\out configurations where there are more I\O than will fit under the shadow of the die. This rounds out. the evolution of CSPs into the higher pin counts and the higher performance in the standard products.

Standardization is really an important part, not just for the package but for the infrastructure. Then concerning the infrastructure of microBGA we're moving to establish a broad infrastructure supply that's partitioned into the first piece which is the flex or tab tape on which the package is manufactured. Next piece is converting that tab tape into a package by applying elastomeric die attach film and putting the flex or the tab tape into fixtures that will run on a standard automated assembly line. The last part of the infrastructure is the assembly itself. A very important part of this is that the package must fit the existing infrastructure in order to ramp up manufacturing very quickly. That's of course supported by standardization of the equipment from wire bonders, dispensers, die placers to support this growing infrastructure. Looking ahead a little further into the future you see that the high pin count end, area array pads which Ed mentioned, the drive toward higher I\O especially for processors and ASICs is moving us toward area array contacts the chip itself.

At this point there are a lot of benefits to assembling this chip size package on the wafer itself. Since the package is chip size you can build the package up like a high rise building on top of the chip in the wafer form and simply extend the move toward batch processing instead of 30 parts in a strip or 50 or 100 the natural evolution is lets put these on a wafer and process them on the wafer. This is where the sheep are separated from goats and chip scale which is nearly chip size or 20% larger than the chip and they are truly chip size because the chip size packages are able to be fabricated directly on the wafer to give you manufacturing efficiencies, cost efficiencies, faster cycles of learning.

This shows a schematic of wafer level packaging.Hitting the highlights, the package is fabricated on a wafer size dimensionally stabilized piece of polyimide. That is married to the wafer one complete wafer at a time. One hundred thousand to a million contacts joined at once, gang bonded and then after joining the package it is molded onto the face of the wafer. Again the whole wafer at one time Next just like the CSPs the solder balls are attached on the wafer, one hundred thousand solder balls at a time. Chip test at the wafer level again paralleling the batch processing that was introduced by Steve's talk but now the batch is the wafer. It is only at the end of the process that you actually dice this up. When you dice that wafer you have the end product which has been tested ready for use. You have actually tested the chip when you test that wafer. What this shows is that the evolution of chip size packaging will have a relatively profound effect on semiconductor manufacturing.

The labor content has been basically taken out of the process. We are not making wire bonds one at a time. We are making one hundred thousand at a time; it's more like wafer fab than it is mechanical assembly off shore.

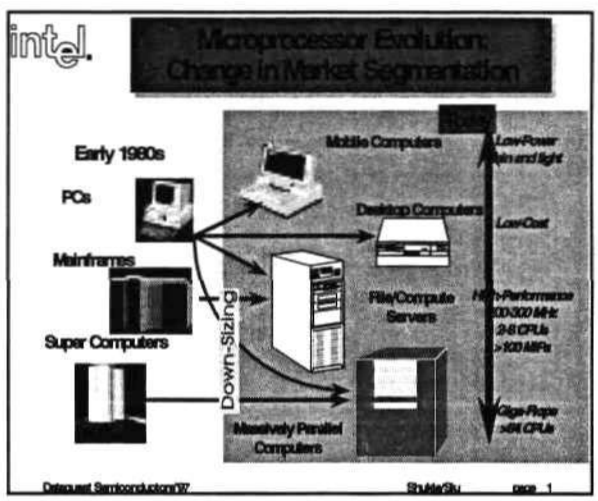

The implication of closing this kind of packaging as it goes to chip size, as it goes to wafer level packaging processing will come back closer to the fab. Basically it will be the back end of the silicon fab. That gives you better cycles of learning, better inventory control and testing efficiencies along with just the progression of more pins at lower cost. Looking inside this wafer level package this slide shows the cross section stylized or schematic cross

section of an area array of CSPs fabricated on the wafer. You see the same flexible bond ribbons that connect the chip up to the copper/polyimid layers and decouple expansion the copper base system from that of the silicon based system. This allows you to use standard copper interconnect with expansion coefficient of 17 ppm and decouple that from the silicon by that flexible layer. The part is completely encapsulated with the molded encapsulation between the polyimid and the face of the chip. There is something very interesting about this structure and that is that you see the opportunity to start putting intrachip wiring on the wafer.

The first and the obvious benefit is it lets do power/ground right on the package. You have a very low impedance power and ground connection through that low inductance bond ribbon. I should mention just stepping back a bit the bond ribbon self inductance is about three tenths of a nanoHenry down to about .22 depending on how it is formed. That is an order of magnitude less than an equivalent wire bond and it is one of the reasons we went to this bond ribbon rather than the longer wire bond. These very high performance bond ribbons or links connect the copper intrachip wiring in the package down to the chip. So that you can start thinking of doing things like clock distribution in good low resistance shielded clock trees on the package rather than doing it on the wafer in resistance lies that are poorly shielded.

The next thing you can think about doing is to put some of the wiring on the package. Some of the critical long nets to avoid the RC countdowns that you have with the long aluminum wires on the wafer. This is an exciting extension and it is a natural extension of the initial CSP products now surfacing in memory applications. That same technology, the same flexible bond ribbon, the same compliant package will extend all the way up to thousands of I\O chips fabricated directly on the wafer. Looking at how this integrates with flip chip we can see here that starting with flip chip C4 you are now able to put layers of intrachip interconnect up in the package

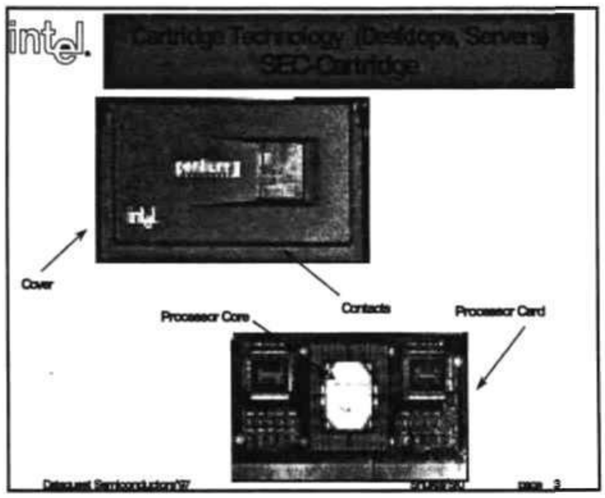

to do some of the global wiring. The wire is much more sensitive to noise resistance RC time constant delays. That gives you added functionality. What we've done is take you in this short talk from a package that gives you size performance cost advantages, now looking ahead to where that package will provide functionality.

You don't have an integrated circuit until the additional layers are added to that chip and that functionality will give you performance, it will give you a reduction of the pin count coming off the package because you do some of the power/ground distribution right in the interconnect layers themselves. Looking ahead where all this leads you make a few observations and they are really logical observations starting with the premise of the chip size package. First of all, CSPs will replace permanently the parts. The talks you've heard earlier show that the drive is to higher I\O, area array small packages. The next observation you make is that these CSPs are based on the existing infrastructure. There is no new equipment, there are some modifications that the materials sets require but fundamentally it operates on the standard assembly infrastructure in place today. Upgrading the equipment, modifying the materials somewhat are necessary to make CSPs. The result of that is that we expect the adoption of CSPs to be extremely rapid.

The one thing that is pacing the adoption of CSPs is the availability of high density substrates, the standard grid pitch being half a millimeter to match the existing processors and ASICs today. The next step is progression to wafer level packaging, for improving manufacturing efficiencies and lowering cost and to get packaging into a productivity industry rather than being a penny a pin static technology. I have mentioned that the requirements for more pins at lower cost is going to have learning curves more akin to wafer processing than it does mechanical assembly of individual piece parts. Looking into the future the drive will be to put in more wiring because into the packaging layers. That will move packaging closer and closer to the integrated circuit, in fact we foresee that wafer level packaging and semiconductor circuits will merge. Looking 10 years, out the integrated circuit will be the semiconductor part plus the intrachip wiring added by the package.

What we've done is to take a look at the immediate and then looking ahead 10 years to where this industry is going. I hope it is helpful but I think it is relatively straight forward. Thank you.

Moderator: Next Dr. Shukla

Dr. Shukla: I think the panelists so far have done an excellent job of covering all different types of interconnection technologies and the outlook for the future. What I'd like to do in the next ten minutes or so is build up on that with a very specific focus area and that is how does the semiconductor interconnect with the package and what are the market segments. How does that play in the world of microprocessors. In other words, I am giving you a rather limited but very important segment which is microprocessor technology and how that market is impacted by the differentiation and the interconnections.

Some of the slides I am showing you are from a colleague of mine. First I am going to hit upon the microprocessors evolution, how that has created some level of market segmentation and what does that mean for us currently and looking into the future. I'll touch upon the hierarchy of interconnections, what are some of the current approaches and some challenges again focusing on the microprocessors.

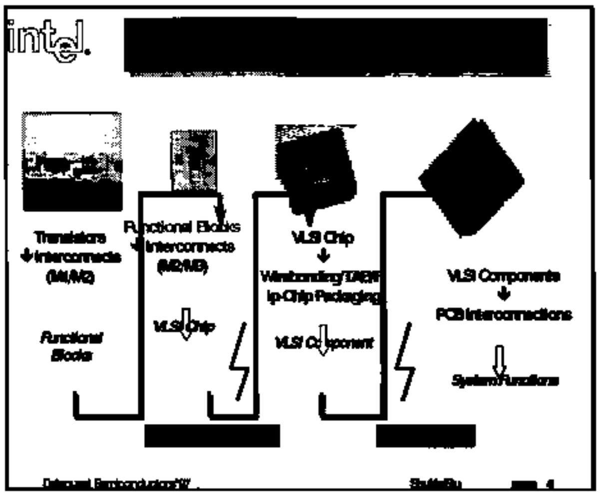

If you take a look at the computer industry from the eighties to today there has been a phenomenal change in terms of market segmentation. In the early eighties we had three universes, the PCs, the main frames and the super computers. They are three different domains, very different technologies. PCs use CMOS or NMOS, main frames use ECLs and super computers use gallium arsenide They are really three different universes and they did not talk to each other very much. They were fragmentations as opposed to segmentations. The evolution of microprocessors changed dramatically. We have pretty much similar technologies in terms of silicon interconnections covering a broad range of products all the way from mobil to desk top computers to servers, work stations and the massively parallel processors. While the technology has evolved to this new type of segmentation, the key here is that it is really not technology which has created the segmentation but the specific needs of the market itself. I mentioned mobil, multimedia PCs, servers etc. that's the current segmentation of the market and a key point is that all these markets if you believe some of the talks earlier today are growing very rapidly which is good for us and these are segmented markets meaning if they continue to grow rapidly we have both technical and business opportunities in those areas. There is also a challenge and that is in all these segments there is a drive to continuously improve the performance and the cost. It's similar to" you can never be too rich or too beautiful"; you can never be fast enough and cost effective. There is continuous drive to do that. The fact that these are segmented markets, there is a challenge to keep standardization in terms of technology, in terms of the infrastructure because if you don't do that we'll go right back to what I call the fragmented market. So you will have on one end personal computers which will use completely different technologies versus servers or work stations and that will take the cost right back up. That is a very key challenge in keeping those markets growing without losing the standardization.

Just looking back over the last several generations of PCs the differentiation has been based on packages. I'd like to spend a minute on it. I think there is a point here which is critical to CPUs which is the difference between socketed solutions like ceramic or plastic arrays versus the PQFPs or BGs which are surface mounted. If you take a look at the current trend in an industry like the much talked about Dell computers built just in time or built to order, the socketed solution is absolutely demanded here because you want to take a CPU which is fairly pricey part of the system and put that into your system at the last moment when you are building it, as opposed to surface mounting it. The surface mount although it's great for manufacturing it really has been for CPUs in the area of the lower priced end of all computers. For the mobil PC, the drive has been the size, the miniaturization. I remember the time of the first Compaq PC where they basically took a PC box, put a handle on it and called it a mobil computer. We've come a long ways from there with all these notebooks and clearly that required a new technology.

The last point which is actually a big theme of my talk is having a package which provides some level of flexibility and head room for growth in terms of performance and features. That is exemplified by Pentium II processors which I'll talk about in a moment. Let me show you to contrast two different market segments here is what is known as the TCP or Tape Carrier Pack or TAB based on Pentium CPU for mobil computers which is barely a millimeter in thickness and maybe an inch square in size. It is really a great product for the notebook market. I want to contrast this with one of the latest products which happens to be the connection cartridge or SECC (Single Edge Connection Cartridge) which is what Pentium 2 is based upon. This is actually the real pictures. The lower one is what's inside if you're an engineer you want to open it at your own risk and look at it; the top one is the cover which is what marketing would like to show. Basically what I want to show is that thing in contrast with TCP packages is really big. It will never fit into a notebook. Look inside. What we have here is the CPU which happens to be on a BGA package and it has a bunch of other components. So why would we go in this direction? That is something I'd like to spend some time on because this shows the evolution of the interconnections.

Before I go into that segment let me spend some time on the interconnections from a different perspective which is after all what are we trying to do is take these transistors billions of them as Carl Sagan would have said, and connect them on a chip using the odd chip interconnect then that has to be put into a package which is your wire bond or TAB, then you have the board level. If you take a look at this continuum as I mention earlier the technology for these at the high level, we want to try and achieve a level of standardization where we can move back and forth

Today that is not the case. What I mean by that is if you take a look at the interconnection densities for the chip level which is illustrated as metal one through metal four to bond pads package board and connectors. If you look at the distance the electrical signal has to travel and the pitch, the kind of densities,.it is a huge range. It is no wonder that the technologies used on silicon are completely different than what is used on the package versus what is used on the board. Again as has been said before the board technology is very different from semiconductor fab and they are driven by these different geometries and requirements. So they are in the technological sense different universes. What is happening is that it is evolving very rapidly.

In the past assuming that they are different universes there was local integration and local optimization. Silicon people optimized interconnects based upon multilevel aluminum, silicon dioxide dielectric structures but what is happening is now partitioning is moving more toward market segments. We can only segment so much using silicon interconnects. So the mobil form factors, sockets, slot one, slot two, these are the real segmented segmentation enablers for the CPUs, not he aluminum and silicon dioxide.

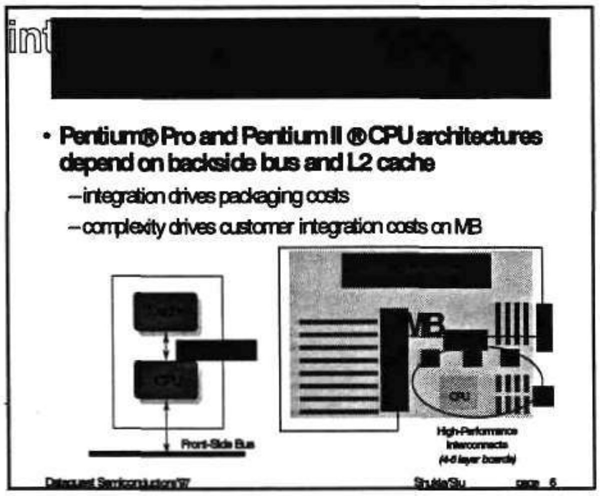

The future is going to be even more different where the partitioning rather than applications specific will be more based upon economic partitioning and distribution of the logic meaning globally optimizing all the interconnects. For that to happen there are many brake through paradigms which need to happen. I want to present here three views of how you would partition for economics. The point I want to make on this slide is that depending on which universe one comes from the viewpoint can be different. If I am a chip designer and a VLSI wafer fab person my goal is to build denser and denser chips and put more functions on it. What you get is a very high performance CPU ultimately a system on a chip. That would be the ultimate. If I am a package person the ultimate is you give me all the chips you can and I'll put them in one package called (Multichip Module) MCM package and give you the best of both worlds. PCB (Printed Circuit Board) center viewpoint is we really don't need these packages, these component packages between silicon and the board are a nuisance, necessary nuisance today because add cost and performance. Get rid of them; the best case is direct chip attach; you take the CPU put it directly on the board or as you heard there are many barriers doing that so you go to chip scale package which is the poor man's direct chip attach.

Now which of these views is correct? I would like to show you that today based on where the technologies are headed none of them actually survive in the long term as the solution for mass market. In the interest of time I'd like to leave you with a thought that system on a chip has some very key barriers such as power, handling.

If you put the entire system of a CPU on a chip you may have a hundred watts to dissipate and when a chip becomes very large you get into electrical problems such as global interconnect; aluminum lines which work great today if you take it on a very large chip and a global interconnect from one corner to the other of the chip. The RC delays or the parasitic delays on those lines will lower the performance of the chip. You don't gain the performance as one might have expected. With the MCMs the problem is slightly different In multiple chips how do you manufacture test burn in, it is a manufacturing and cost issue. With PCB integration the direct chip attach as we've heard there are some enabling issues with substrate capabilities. The bottom line is that any of these technologies don't have in themselves the capability in the next several years to go to high volume as a CPU interconnect primary option

Let me skip this and go to the problem statement. What is the best way to continue to improve the performance. Do we do it on silicon, package or system? I think what we are seeing to end the discussion is the real challenge is to achieve global optimization rather than local optimization. In doing so we have to blur the boundaries between silicon and the package on the board. By blurring the boundaries there are possibilities that you take power bussing and put it on the package if you can get the maximum advantages out of that. I'd like to illustrate this by an example which is how we have approached it on the CPU evolution today. The current approach is to continue to focus on high performance, performance has to improve with every generation but provide compatibility in terms of backward compatibility and very important standardization. So that we can go into high volume and still be able to contain the cost.

For example one CPU design which takes years and hundreds of designers and a lot of validation so they won't have bugs in it. It is a very finely tuned machine to produce a high performance CPU. So you produce that but you put it into different interconnect schemes such as Slot 1 and Slot 2 to serve different markets in an economic way. What we are doing is partitioning the architecture based upon economical considerations. Let me explain what I mean by that. If you take a look at Pentium Pro and Pentium 2 CPUs, the primary architecture is that you have the CPU engine here and you have the level two cache. As the CPU becomes more and more demanding you have to put more memory around it. In the Pentium Pro Case we put the level two memory in the same package as the CPU and that was done at the expense of package cost. The ceramic package used to produce this is actually very expensive, as you can guess.

If you tool the other approach of taking the CPU and the cache and put that on the mother board we are trying to manage a very high speed bus. In the worst case this connection between the CPU and the second level memory could be at the clock speed of the CPU itself. If you have a 500 megahertz CPU you could potentially have a 500 mega transfers per second exchange between the CPU and the memory. If you try to do that on a mother board, you'll burden the entire mother board with the complexity and the cost of managing those electricals because everything else here the DRAMs and a bunch of other components don't require that kind of very high integrity electrical connection or the density of interconnections.

The approach has been to separate the cache memory from silicon; that creates flexibility because you can have a cache like one meg or two megs depending upon the application, desk top may require less service and work station require more and more memory. It avoids large die syndrome so you don't have to put the whole thing on one single silicon and it's cost ineffective. You also now allow memory to be purchased from the companies that are experts in that area, high volume, high performance memory and you lower the technical cost on the mother board industry. If you take a look at the Pentium Pro processor in the ceramic package it was a two- chip MCM. Technically it was done, it is a product today but in going to higher volumes looking in the future as we speed up the backside bus we want to go use this format where we have taken a portion of the mother board the local portion which is between the processor and the memory and made that as high a performance as possible so you don't have to do it on the whole mother board.

The message is that rather than going completely to one side versus the other in terms of system on chip versus doing everything on the mother board we try to blur the boundaries and do it for the CPU and its local environment where the high performance is necessary. Looking into the future, how would we continue to progress in that area. First of all there are some breakthrough paradigms which are necessary, meaning we need some enabling things here to make us go further. The system optimization. Once you tell the designers that they have this cartridge type of approach now they can start thinking about marrying the memory and maybe other potential chips which could work in concert with the CPU. This will change the whole approach on design.

So you start doing a system cost and complexity optimization. I call it vertical integration in an industry which is very much horizontal. It requires very significant partnerships among companies who are in this food chain. There is a demand for high performance substrates although it is not localized to CPU cache memory subset of the mother board to provide that

It is an enabler which is needed. I did not spend time on the technical aspects of area array interconnection namely flip chip. They are a must for this kind of optimization to take place. This cannot be done with wire bond.

Lastly portability will require a different type of approach as the notebooks continue to get smaller and lighter.

In summary the interconnects are at a very interesting threshold right now. If you look at the CPU business how it has evolved from a single component to two chips in a package to a cartridge it tells you something about we are at the threshold of rethinking how a CPU should be designed and enabled for mass production at very competitive cost and still have some head room for performance growth. Interconnects play a very significant role and they are the foundation for this market segmentation which is created by the CPU revolution. In the near term focus will be to take the existing technologies and evolve them to maximize the local integration between CPU and cache subsystems. In the long term as the area array technology become more pervasive and as the high performance substrates become available then other stuff such as partitioning functions from the chip to the package or the cartridge becomes possible. That will be the way to go.

Moderator: Next John Novitsky from MicroModule Systems.

John Novitsky: We have a fab in Cupertino, California in Silicon Valley and we make very fine pitch substrates for single chip packages or future packages or multichip packages. We also make some special test contactors that are used as well to test very fine pitch devices. Some of the opportunities in directions that I'll address quickly are that I\O density as is happening with the flip chip are deriving an entirely new class of materials. I'll show you some examples of that. Afterwards if anybody wants to come up I brought a collection of different physical samples of things like a wire bond chip of 59 micron, a flip chip of 250 micron and some CSPs at .5 millimeters and .75 millimeters. So you can actually see and touch these things and get a sense of physically what's happening.

I'll talk about future packages and point out what is happening with the IC metallization. The SIA trends in the packaging workshop are reflected here. If you look across the row that says package I\O count you'll see that within the next few years the SIA road map is creeping up near 2000 diodes. With the wire bond pitch that's dropping down to 50 microns, a flip chip pitch dropping down to 150 microns and a CSP pitch dropping to .4 millimeters To give a quick idea of the implications, these are all drawn to scale. On a 10 by 10 millimeter die you'd also explain some of the differences that the number of I\Os that are physically possible with those types of geometries over the next few years and likewise if you go to a 15 by 15 millimeter die again using some of the future projected items on the SIA road map. The key with making these fine pitch substrates is the lines for most of the manufacturers like us are not the problem.; it's the vias. They kind of get in the way.

There are a couple of things going on here. I have taken for the next few foils from left to right commodity PCB boards using five mil line and five mil space and a 25 mil capture pad with a 13 mil drill. The middle one represents more or less today's state of the art SLC type of process, continental circuits, microvia type process and the one on the far right hand side represents a class of materials that I happen to be most familiar with my own company's design rules so I used that as the example but there are several other companies that have comparable types of things as well.

As you see if you look at this in the same area there is roughly 15 times improvement going from the far left standard PCB out to the fine pitch of what we can do. In between again if you take it from the perspective of what do you need to get in 10 lines again it is roughly 13% of the area. If you look at the vias the fine pitch boards today still are at a point where they are barely able to capture some of the half millimeter CSP type geometries and some of the fine pitch materials do much better than that as well. When you put it together with four nets and four vias it looks like this.

If you put this together now combining both the line pitch and the via size if you take the same die and this actually is a live writing example that we had done using an Intel 120 megahertz processor die as again the difference between a chip on board using commodity PCB, using chip on board, using a very fine line FR4 and again the same chip on board using a very fine pitch, a thin film like we have. Using the half millimeter CSP pitch again with traditional PWB you cannot even capture all things because the CSP pitch is beyond what two balls are. If you go to some of the fine pitch like again the SLC and continental circuits you are able to capture the two balls and rout one escape trace.

What that means is you could have two roads before you would have to spray out multiple layers going down into additional layers below and additional layers increase the cost and increase some of the test concerns as well. Now when you go to flip chip on the vertical axis is represented the flip chip pitch the circles represent the vias that are in production or about to be in production at these three classes of suppliers.The horizontal axis represents the number of I\Os that are possible given the combination of the routing density the via geometries and line spacing and the line width.

What this is showing is that using state of the art just about the best you can get in production today of the fine pitch things you are at a point where for a 15 millimeter die you are able to do on the order of using two layers 230 I\O escapes compared to what's happening with this new class of materials which is about five times better.

It's basically a horse race. The way the industry is shaped right now there is a class of new companies like MicroModule Systems, Gore is another example of this that are pursuing this very fine pitch set of materials and we are trying to build an infrastructure fast enough to service an industry with companies that are represented here as well. On the other hand the PCB guys are trying to figure.out how to, in the past they've always taken drills and plate them and used them for their vias, they are having to develop some very different techniques and it is a horse race between those two. For those who have historically made ceramics the horse race there is try to get a different set of materials, different than ceramics for some of the dielectric and routing pitch and other historical concerns as well. I think it is too soon to call but again as an industry landscape this is kind of the rough shape of the industry that is trying to service the needs that have been pointed out.

Rama briefly touched on this and in our business again we don't care whether the people use the dense substrate for single chip or multichip but I'd like to point out an opportunity here in that in the next few years there will emerge such an infrastructure of a very high volume probably pretty low cost dense capacity to do these types of substrates. If and when it exists, it allows us all sorts of opportunities and particularly to pursue few chip modules. The previous session today talked about the integration of the DRAM technologies with random logic and talked about what some of the trade-offs were.

Going forward it is sometimes appropriate to decouple. For example the analog technology from the rate of change in the memory technology or to decouple the DSP from the memory but putting them on a common substrate like what Intel has done with several of its most recent introductions of its processor family allows you to each take those separate semiconductor technologies and they can evolve at their own separate rate of change while on the outside you just put down the common footprint and away you go.

If it makes sense over time there is a path to forward integration where you do combine those technologies and bring them back together a few years later. If on the other hand it does not make sense either technically or economically you can still allow yourself to keep that footprint down and allow them to still change independently.

For some of the interconnect future directions up to maybe .15 micron or so we will probably see several things happening. There has been some discussion recently with IBM's announcement, Motorola's announcement and then last week at the microprocessor forum Jerry Sanders of AMD also unveiled some of their directions as well going to copper for the on chip interconnect and also going to a lower dielectric insulator beside silicon dioxide which has historically been used and that was touched on briefly here.

Dr. DiStefano also mentioned that as you think of this it's also possible to partition at least some of that dense routing into the package particularly the power and the ground and the clock distribution. It's a different way of getting from here to there. Below .15 micron or so it becomes inevitable that it is very difficult to separate the chip design from the package design from the board design and from some of the system interconnect designs. It will probably inevitably happen that there will have to be vertical alignment between possibly broad horizontal companies and broad horizontal industries that are servicing each other in a tightly cooperative way, in a manner that has not happened in the past. So that is one of the other trends that I point out as well.

Q: There was a general theme of questions around the wafer level packaging proposal that was put together by Dr. DiStefano. I'll address the question to him. In terms of the process complexity and where do you think things like dicing the wafer, damaging the package part and then what happens when you shrink the die. Those kinds of aspects. Maybe you want to comment on when you think that will become more of a reality than a conceptual part.

A: The wafer level packaging paradigm is very similar to the one that Steve showed of processing packages, fabricating packages in arrays. Those packages are processed beginning to end in that array including test. At the end of the line there diced apart just like a wafer today. The next logical extension is to say lets do that right on the wafer, and skip the stuff about cutting apart. The problems you run into with wafer level packaging that you don't have with the individual package parts processed in these batches is that you have to be careful of dicing damages. It is very similar to area array. I would say it is a natural extension of the trends

Chapter 8: Surviving and Thriving in the Chipless Business

Gary Smith
Director and Principal Analyst
Electronic Design Automation Program
Online, Multimedia, and Software Group
Dataquest

Good Afternoon. I want to start out with a correction to Joe's foil. He had an upside down triangle and he forgot the tip of it is $2 billion and it's called the EDA industry and it supports all this other stuff now.

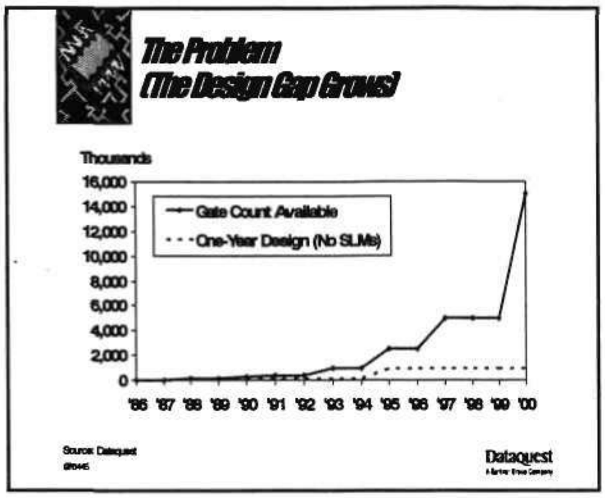

Most people forget that little tip of the industry but I'll try to bring you up to speed. I'm an old semiconductor guy, never actually worked in the EDA industry. I ended up as a design methodologist, spending a lot of time in the ASIC industry, so I come from your side of the fence. That sometimes bothers the EDA people, but we are making progress. In fact, a lot of semiconductor guys are moving into the EDA industry, while the EDA people are moving into the embedded systems industry. Go figure.

A question I was asked back in 1981 by a customer of mine changed my career path and has continued to be the most important question in the industry today. That was "What in the heck are we going to do with all these damn transistors?"

This is a fairly familiar chart, and that is the design gap. What's going on and has been going on since the late 80s is that the semiconductor industry has continued to exceed our capability to design silicon,

By the way, if you hear of the ever-shrinking design cycle, that's all marketing BS. Don't believe it. It sort of flattened out to one year, and it's been that way for most industries now for the last three years. Automotive is the last one, they're still working about a two year cycle.

Basically, it looks like we're going to see all the designs, except for the very high end processors and custom designed DRAMs are going to level out to a one year design cycle.

We have been trying to address the silicon problem by changing methodologies. Every time we change methodologies, we get a 10 times increase in productivity and that just isn't working anymore.

As you can see, unless you use SLMs, we're losing the race big time right now. The answer is System Level Integration, which is the obvious way we're going to address this problem. SLMs are the blocks we feel will be filling the void of silicon, increasing our productivity enough so we can actually use 5 million, 10 million, 15 million gates — whatever the semiconductor industry tries to throw at us.

The definition of system level macro, it's changing, getting larger — initially, this is the definition we use: 2500 gates for those who haven't been in the design community is a UART. That's the simplest slam. A lot of people are calling this "IP." Don't use "IP." That's Internet Protocol. It's really not Intellectual Property if you really look at it. If you look at all 50 members of the newly emerging "IP" industry, none of them have legally definable Intellectual Property, which is another issue we could talk about.

Now think of a world where 40% of all microprocessors and DRAMs are SLMs, no longer pinned out. Virtual components — 60% of all DSP microcontrollers, SRAMs are SLMs, and 80% of all ASPICs are now SLMs.

Most of us who've been in the industry for awhile know this cycle. It started out in the early 80s with disc drive controllers. They were all ASICs. The ASSP guys came along and made them standard product and got their price down. Those standard products became macros, basically today they would be SLMs, which would fit into another ASIC design, which was multi-SLM design and became an ASIC again.

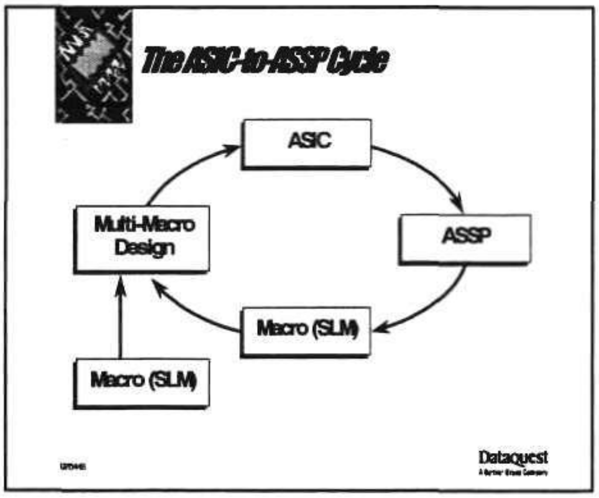

Every jump in technology is made by ASIC implementation and is then consolidated into an ASSP. The volume of shipments is done as ASSPs. The initial guys in this technology Chips & Technology Of course, Cirrus is a big player in this. There are a lot of ASSP companies out there right now. Most of them are fabless, but that sort of business detail I'm not going to address.

What did I say in the slide before? I said we're shortcutting the pinned out silicon on 80% of most of the ASSPs and a lot of the other designs, so suddenly the silicon business does not become an IC business. They start separating, and there will be an entire industry based around producing designs that never get to silicon as a stand-alone pinned out IC.

That's the major shift we're seeing going on in the industry today. By the way, that's a shift most companies are unprepared to address. Therefore, we're going to lose a lot of companies that are household names today.

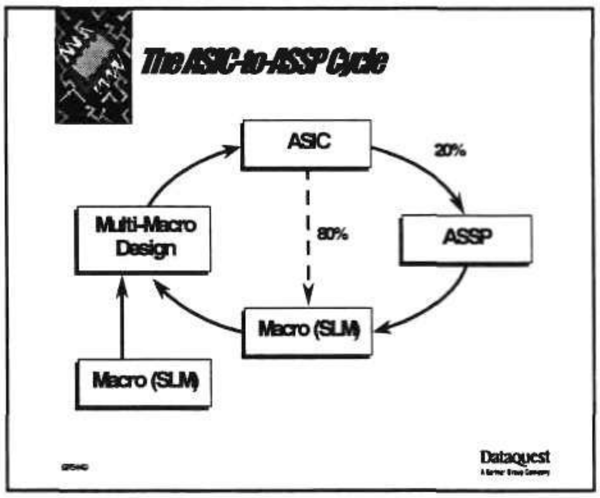

Last year, the top 10 lists were very popular. They've died off now which means I can do one and not feel trendy.

I won't go through them one at a time. I had fun making them up but you can read them in your files. The ones I liked though include number seven, probably my favorite, Libya just got into the DRAM business. The really interesting one is number five, customers are saying anything under a quarter million gates is SSI.

I'm working on the EDA road map, part of the semiconductor road map, and one of the proposals on the table is that in the year 2004, the average SLM will be 150,000 to 250,000 gates. Whenever I get calls from these ASSP guys, and they say things like, "I've got this great design—is it going to be hot!" I ask how big is it and they come in under 100,000 gates, I say, yes, it will be a great SLM but it'll never make the silicon. That blows a lot of minds but it's the sort of thing we're starting to look at right now.

If your design is under 100,000 gates, in three years it will never get pinned out. It will go into a larger SLI.

Of course, the one we know and love, number three: "Intel just integrated my last hot chip into their processor." Intel is doing to silicon what Microsoft is doing to Windows. It just gets more and more of the uses integrated into one system. Only I think Intel is doing it right, Microsoft isn't but that's my bias.

I usually take about one consulting job a year and I've been doing some methodology studies for some of my clients who are end users rather than companies. It's interesting how many people fail this test. For some reason we've lost track of the fact that we're an engineering driven profession.

Very few people can tell me what methodology they are using within the engineering department of their company. Standard methodologies for SLMs? Very few companies know how to do this. Of course, those that do are doing very well.

What are your engineering problems? Late design, yield issues, keeping your engineering talent: these are all indications of obsolete design methodology.

I got this from my real estate lady, but it works anyway. Methodology is the driving factor in design today. It was yesterday, and it will be tomorrow. If you're in upper management in any semiconductor company and you don't understand design methodology, you should start reading up on it because that's going to be the difference between whether your new fabs are going to be full or not.

What we're seeing today is a newly emerging theme in systems design automations. Actually, the first system design was done by Boeing for the 777. They cobbled together a bunch of tools, tied them together with glue and bailing wire, and damn that thing flew the first time out the door, which had never happened before in a major plane design operation.

There are now six system design automation study programs or projects going on in the United States and around the world: one in Japan, one in Hong Kong, two in Europe and two in the United States. Four are automotive programs and two are in the aircraft industry. That's something to look for and the direction we're going.

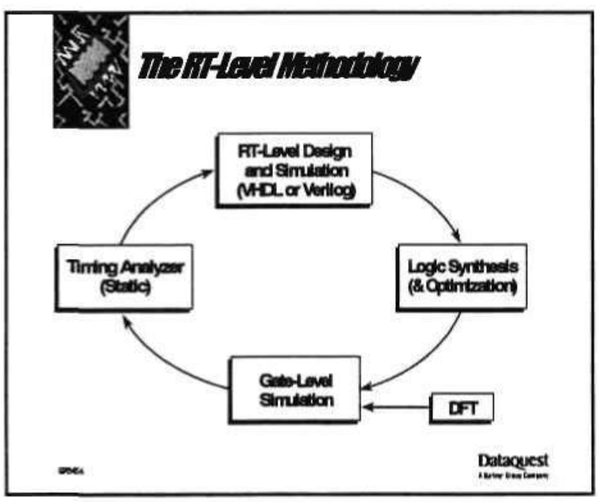

In 1994, we put out the first ES level designs. The Iridiam program for Motorola was probably the best example. You've probably heard of ESDA? That was trashed by some marketing people in the EDA industry, which is not uncommon.

We are now at the electronic system level. At that level of complexity, we can knock out about a million gate design a year. That's not using SLMs. RT level came out in 1987-88. Synopsis was a big driver in that. Actually, the user was a big driver in RTL and that's where most of the designers are living today.

Gate level is pretty much going away as a usable methodology and then you go down to CAD.

A little about the RTL methodology: it was once called VSMQ for Verilog Synopsis — the "M" was for Motorola. Those were the only two ASIC companies that knew how to do RTL design except for IBM. IBM is an exception to most rules by the way.

Quad design was the timing analyzer that was used. That comprises a fairly simple four tool toolset. If you throw in a DFT tool, you have five.

This is the new ES level that came out in 1994, designing at the behavioral and architectural level. You're doing hardware, software, co-development, co-design, partitioning. This is the hot area today and has a possibility of solving many software problems because the main problem with software today is we don't know how to partition our designs into hardware and software. Hardware gets everything you know, and software gets everything that's left over. The software is always late. It's sort of a "go figure" type of thing.

Unfortunately, this whole thing started coming apart a couple of years ago as we got down to 0.5 micron. We found that the RTL design was becoming extremely silicon sensitive. After you're over 80 Mhrtz, there is no such thing as a digital transistor. We've gone all this way with the concept of a transistor being Off or On. Other than that, we don't care; but at 80 Mhrtz, you start to care.

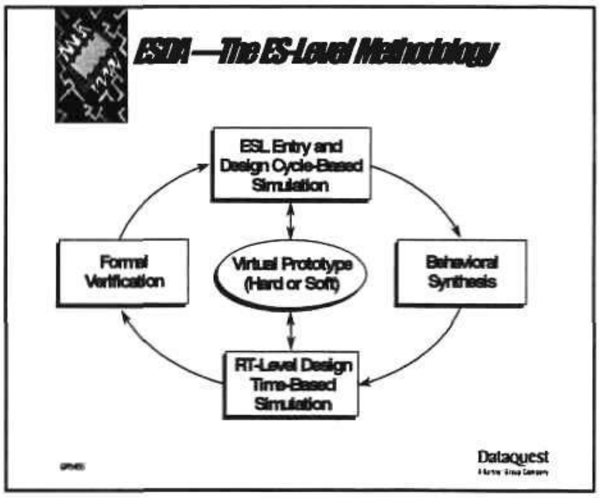

The other issue is that the systems level was really separated from the RT level, so passing designs from one to another was an inefficient mapping operation. You saw that the ES level got stymied insofar as growth potential.

The current idea is using virtual prototypes, which started at the ES level, to start to tie these designs together from the system to the RT level then down to the physical. People really involved in the semiconductor end of it need a completely new physical verification suite because we're starting to have very serious silicon problems.

Virtual prototype at the RT level is an RT level floor planner used as a cockpit for the design, with block place, SLM place, macro place. You're creating a virtual route, which is why Cooper and Chan was such a hot property awhile ago. Their router worked very well for this problem.

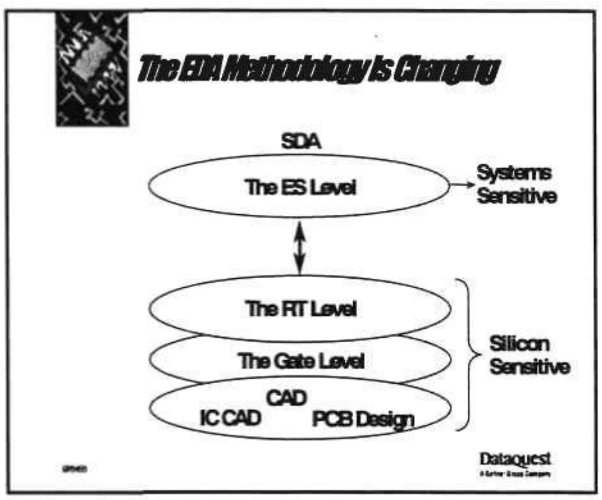

You also need to look at clock trees, power nets, scan chains and by metal layer. You also have to have an accurate delay calculator and simultaneous analysis estimates. There's not a tool on the market today that can do that.

The new RTL methodology and its entire tool set is extremely complex. Right now, the verification side has four tools. At one time, it had one. The design side is similar. It's all feeding into a virtual prototype which is getting simultaneous estimates of timing, power, Sig-Int, EMI, metal migration and thermal. You're feeding the whole thing into the hardware/software virtual prototype so software can be run on your design prior to ever making silicon. This is getting really complicated, and we can't do it yet.

One of the big problems we're seeing in the semiconductor industry is a dramatic change in the challenge set everytime you give us a new process. It used to be area

How fast can it go? At 0.8 micron, the power user is saying "Speed first, area second." At 0.7-0.6 micron, we have to look at the power side of it which is getting a little scary. At 0.5 micron, we start burning holes in our designs. There is a very famous one that came out with a hole in the right hand corner. They had to use a power tool to figure out how that happened.

At 0.35 micron, we saw the first metal migration. I had three reported instances of metal migration. Metal migration is the older term for electro-migration. I use it because I'm old but also because people were getting electro-migration confused with EMI which is a completely different phenomenon.

We run across it in RF back in the early 70s when you started seeing a lot of migration problems in 3 Ghrtz transistors. What happens is your metal line moves, and it tends to short out your design. For those in marketing, it's really embarrassing when 100% of your designs come back within a six week period, which is what was happening back in the 70s. When all the customers call up and tell you your design just died, you have to explain things a little.

You catch most of this in burn in. In fact, the three we saw last year, were caught at life test, nothing ever went out the door — everything was great; but I'm not sure how many people do life test anymore. That's scary because this year we've had multiple reports of metal migration. If there are people doing 0.25 micron designs and they don't do life tests, we're going to start seeing some of this stuff in the field and people are going to have to explain why all the ICs are dieing. This is an extremely sensitive reliability issue.

We believe the low yields we're seeing in 0.25 micron, which are much lower than expected when we move to a new process, is signal integrity. We are just getting the tools on the market now which can extract the information to give to an analysis tool so we can see if it is a signal integrity problem. We're a long way from solving this and we may have to live with low yield for another six to nine months before we start solving some of these issues.

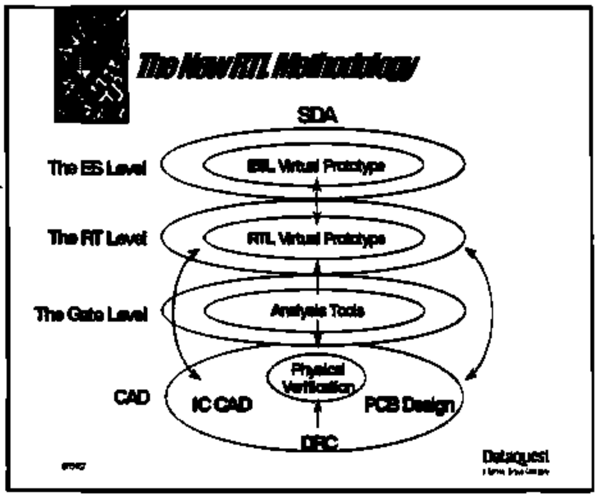

At 0.18 micron, we have to look at inductance. Right now we extract resistance and capacitance, and we are going to have to figure out how to extract inductence figures. There is absolutely no one in the industry that can do that. Unless we put out some tools in the next nine to 18 months, we are really in trouble at 0.18. We could slow down the world on this one.

For the new physical verification technology — basically you have a set of tools where now that the thing is really laid out, you don't have the problems you were trying to design around at the RT level. Everything is complex today, and people have been throwing out tool sets with some regularity and starting over again.

There are five problems we need to solve before we can do system level integration. There are people doing SLI today. There are about seven ASIC houses capable of SLI, which will be the survivors at the turn of the century. There are three more either there right now or close. That makes about 10 ASIC guys I know are going to make it, with two that are probably going to be bought. Maybe we'll get two more, maybe five, but I think five is maximum. All th cent-per-gate guys are going to go out of business.

If you don't bring something to the party in design, there's no place for you. You can use an EDA tool set, you can go to TSMC, you can do everything the old cent-per-gate guys did for you without any problem.

What needs to be created is the new physical verification tool set, or we won't get the designs into silicon. Semitec has two problems. (There are actually three CHDs programs.) One is to solve the verification issue. Lucent was given that contract. Lucent, some companies like Simplex, are working that problem very hard and I'm feeling pretty good about that one.

RTL virtual prototype was given to Synopsis and IBM. I'm nervous over that one.

The SLM socket standard is being addressed by VSIA and the ASIC Council. Now that the politics are over, they can address the real problems. I think we'll get there on that one, but you're not going to be able to mix and match macros or SLMs from different sources until these problems are solved. This is why the seven companies that can do SLIs today can do it only with there own stuff. If you go shopping in these companies and tell them you want an XY23 and that ASIC supplier doesn't have it, you have to move on until you find one that has it.

Formally provable ES design language or no RTL hand-off creates certain issues. One is that in the mainstream there will be two types of designs. Power users are going to use these major ASIC providers. In the mainstream, a lot of the companies are going to bring tool sets inside and use foundries which will solve their problems.

At the lower end, those companies aren't going to be able to afford, or don't do enough designs to justify, bringing the talent inside. A SLI engineer is going for about $120,000 in the valley. The tool set is about $155,000 a seat. There is an investment that has to be made and right now, the mainstream is used to paying $15,000-20,000 a seat for their tools. There are decisions that have to be made in the user community on what model they are going to use.

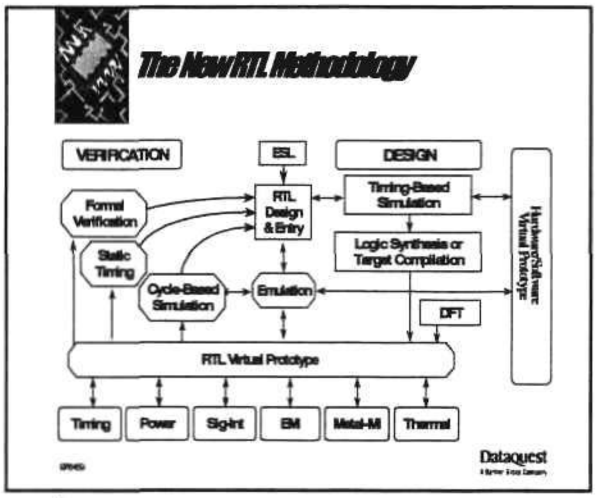

On the low end, if they are going to hand off to, say, Cadence — which is pushing to get into that market — you have to have a design description that will work. Andy Rappaport, in the mid 80s, did that famous report that stated that 50% of all ASICs designs fail in the system.

The primary reason for that was the inability to express a specification in an adequate manner. The designs and silicon weren't bad. You couldn't say "This is what I want built," hand that over a wall and have it built. They were using gate level designs at 20,000 gate complexity.

Think of doing this in Verilog, e.g., at five million gates. The problem of handing that type of design off without a formal approval spec, and you'll get lucky with a 5% success rate. Until we solve the language problem, this hand off thing is nearly impossible.

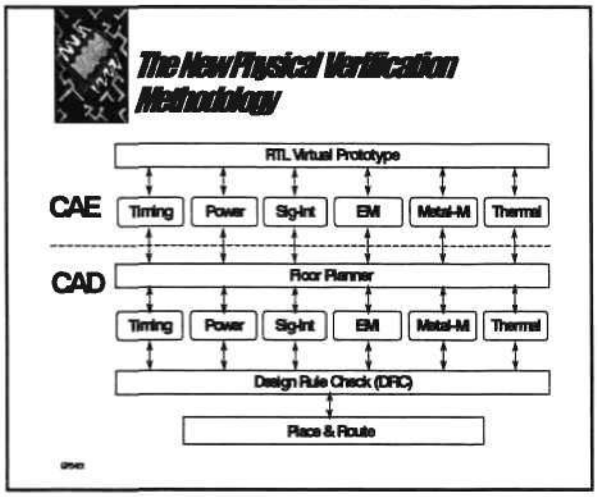

We also need a new test methodology. This is going pretty well. It's being driven by the industry, LogicVision, Mentor, and the new company working on analog, OpMaxx is really getting into BIST which is going to be the primary way we'll solve our test problems.

Obviously, without solving the test problems, we're not going to make the silicon. "Here's my design. Trust me, it will work" doesn't hack it in the ASIC business.

Most companies are missing the point. That is, SLI is not an engineering issue. It's a company culture issue. If your company does not change their culture, they won't make it in SLI.

The winners will use the latest methodology. They'll continue to invest in the latest tools and develop internal tools to fill the voids. Power users spend about 40% of their design dollars developing their own tools to fill the voids in the commercially made tools available.

Lucent has done the best job in making this switch. A lot of the big companies spend a lot of money in maintaining the latest and greatest schematic capture. Almost nobody was using schematic capture anymore and Lucent went in and decided that every tool not 10 times better than something available outside was dead. They then shifted all those resources into developing the tools they needed to complete the tool set. Today Lucent is a power in SLI. I'd put them at number two; they are really doing a great job.

Have a solid SLM development methodology driven by the president or COO. The problem is always with your best designers; they always know how to do it a little bit better. You have to explain you don't care about performance or size optimization, don't touch that transistor. I need a reusable core. That's what you have to drive from the top.

The losers will not understand how you design products, won't understand why the engineers want all those tools. The minimum cost of one design week lost in the design cycle is $155,000 and that does not count late-to-market costs. There are curves out there that people love to show but this is cost: salaries and equipment. Minimum we've found at Dataquest is $155,000 per week, so if your engineers are reshipping designs or sitting idle for a week, you just threw that amount out the door.

Losers will not insist on SLM methodology to collect and reuse internal designs. That way their engineers get to do the same designs over and over again. You want engineers working on new designs, not reinventing the wheel.

The interesting issue, and I've checked this out, is that there are engineering departments, an amazingly large group, that haven't seen the president or COO for nine months or more or the companies who is president says "I saw them at the Christmas party. Does that count?" are in trouble.

That's what's going on in the wonderful world of design. It's sometimes scary but always thrilling. The world is changing and a lot of you guys won't survive.

Thank you.

With that, we'll take questions.

Q: What's the difference between virtual prototype and a simulator?

A: A virtual prototype is a mockup of what you're trying to do. Simulators simulate the design. Basically, you're saying I want this logic function to happen. You throw a vector at it and simulate the logic function. We're talking about actually developing a software prototype of the silicon so you can check the layers, the metal, the actual layout of the components, how the routing is done. You can actually look at it and say I have a timing problem in this area and I'm going to solve it without creating any other problems.

One of the issues in SLM and probably one of the problems with the DRAM or any embedded memory is where that memory is happens to be extremely important. If you're handed a block of memory that's sitting in the right hand corner, your design might not work.

Q: Give me your top three applications, specifically, that are likely to most quickly move to SLI.

A: Graphics will be number one based on the bandwidth problem. "Gazintas and Gazoutas" kill graphics chips. They want to sit right next to the memory. We're currently seeing lots of SLI in communication programs. The irridium program was one of the first SLI designs, so telecommunications has already driven it. Work stations has driven it having always been the leading edge guys. Now we'll see graphics get in there.

Q: Can you differentiate between the design of an ASIC and the design of an SLM?

A: Should be none — the same methodology.

Q: How about the test site?

A: On the test, we should be getting SLMs that are already "BISTed." We aren't yet, but that will be a demand. If a semiconductor guy wants to use another macro or the client wants a SLM from some other source, the first question is how to test it. Within the near future all SLM compaines will provide not only the SLM but the BIST solution will already be embedded.

As the designer of the SLI you'll hook up a scan chain to all the BIST.

Thank you.

Chapter 9: DSP: Be All That You Can Be

John Scarisbrick
Sr. Vice President, Semiconductor Group
Worldwide Manager, Application Specific Products
Texas Instruments

Good Afternoon. In the handouts that go with the presentation, which you either collected on the way in or can collect on the way out, you'll see a copy of the paper I presented at the 1988 Dataquest Semiconductor Industry conference, and you can see what a difficult business forecasting the future is. I don't know who said "Forecasting is a particularly dangerous thing to do, especially if you are forecasting the future." If you don't have a laugh reading through some of the things that were forecast in 1988 for DSPs, then nothing will amuse you this week.

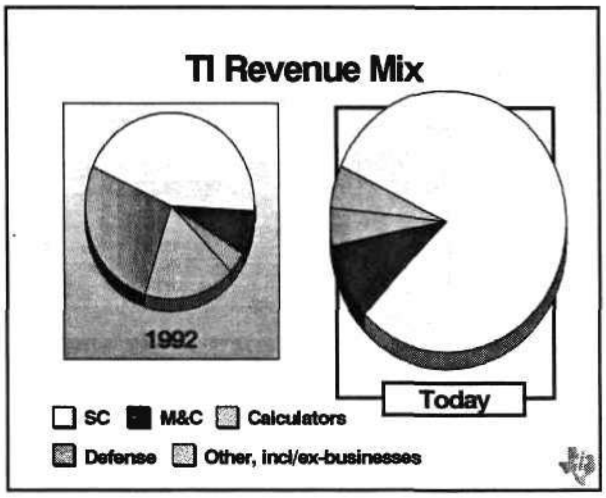

Thank you very much to Dataquest, our hosts. They're well known for their forecast expertise and I'm happy to report that in their September 17th *Semiconductor Alert*, they forecast the Digital Signal Processing market to be the fastest growing market in semiconductors. I'd like to personally tell Tom Starns and Joe Grenier that I couldn't agree more with them.

Anyone who knows Texas Instruments these days knows we're pretty bullish on Digital Signal Processes. I'd like to talk today about the DSP marketplace, a very hot semiconductor market, and given the nature of the semiconductor industry, which is pretty hot by overall industry standards, that's saying quite a lot.

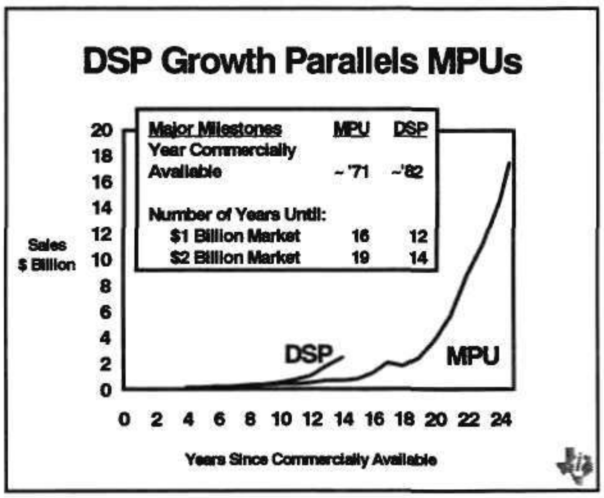

Second, I'm going to talk about TI's leadership position in the DSP marketplace; and, third, if I do my job well, I'll convince you that TI's future DSP leadership is as close to a sure thing as you can have in this uncertain industry.

Before I get into the DSP market itself, let me take a moment to illustrate how we've changed our focus at TI to concentrate on Digital Signal Processing solutions.

Looking back five years to 1992, you can see that TI had a portfolio of businesses, the largest of which was semiconductors, but we were in notebook computers, computers, defense systems, software, etc. However, even though Semiconductor was the largest business at that time, it was less than half of TI's revenues.

Today, we have focussed our efforts and resource in one area, Digital Signal Processing solutions. Semiconductor revenues now represent nearly 85% of our total revenues at TI, and within Semiconductor, Digital Signal Processing solutions represents about 40% of that and growing. It's a dramatic change for TI.

As we look at the DSP market, I think you'll understand why we've chosen this strategic position. As I mentioned, the DSP segment is the fastest growing semiconductor market that Dataquest forecast in their recent *Alert*, and I think this was echoed in a previous Dataquest slide today.

We've already seen tremendous growth in DSPs, which follows the same pattern that microprocessors experienced early in their life cycle. The first single chip microprocessor was commercially available in the early 70s. In 16 years, the microprocessor market reached the $1 billion milestone, and three years after that, the market reached $2 billion.

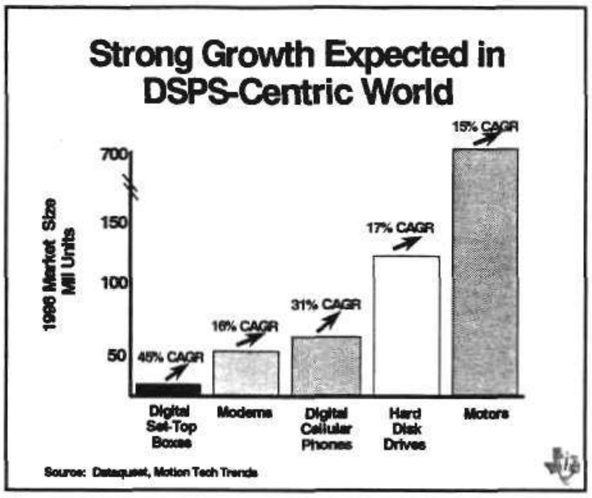

The first programmable Digital Signal Processors became available in the early 80s, and in just 12 years, in 1994, the market for DSPs reached $1 billion and just two years after that, in 1996, the DSP market reached the $2 billion milestone. The DSP market has actually built up more momentum in its first 15 years than the microprocessor market did in the same time period in its life cycle.

While the growth trend is very similar to microprocessors, the DSP market is not dependent on a single end equipment like a personal computer. DSP growth is driven by very different trends, and DSP applications are widespread and diverse. They include wireless communications, server control in hard disc drives, modems and other data communications devices, set-top boxes, imaging and multi-media, industrial and automotive control, and literally hundreds of other emerging applications. It really has become a DSP centric world.

Looking across the critical market shown here, we see double digit growth in every area. DSPs are not dependent solely on any single end equipment market. No one segment today represents more than 25% of DSPs total usage, and this provides multiple opportunities for growth and a relatively stable foundation for the DSP marketplace in the future.

With such broad base support, we believe the market demand for DSP solutions will explode. This includes both the DSP core and the mixed signal and analog components that make up Digital Signal Processing solutions. Since 1988, the market for Digital Signal Processes has grown by more than 40% per year, up to around $3 billion this year. We expect continued growth, well above that of the semiconductor market, for the next 10 years. We believe that the Digital Signal Processing solutions market, together with related mixed signal and analog devices, will reach $50 billion over the next 10 years.

The growth in Digital Signal Processing solutions means a corresponding growth in the mixed signal and analog components that go along with the DSPs. These provide the interface between the analog world and the digital world, and are critical to delivering a DSP total solution. You need to deliver both the DSP and the analog technology, and you need to do it well, and offer customers a total solution.

We are the only company with a strong leadership position in both Digital Signal Processing and mixed signal and analog markets. We are number one in Digital Signal Processing, with around 45% market share, nearly double our nearest competitor, and Dataquest shows TI as a close second, and closing, in the mixed signal and analog market.

We're the only company with the level of expertise both in Digital Signal Processing and mixed signal and analog to deliver DSP solutions that are both innovative and affordable. That brings me to the second thing I wanted to talk about today, which is TI's position in DSP solutions.

In 1997, TI's DSP solutions can be found in a number of areas. One out of every two digital cellular phones in 1996 used TI DSPs. We outsold all TI's DSP competitors in the digital cellular marketplace, working with creative engineering and systems companies like Ericsson, and we're now shipping around a million DSPs a week to the digital cellular market in what is a very high growth market for us.

Our DSP solutions can also be found in nine out of 10 high performance disc drives doing the server mechanism control and 1 out of 3 high speed modems. In fact, the high speed modem market we are very confident we can now declare victory for the U.S. Robotics and X2 technology standard. Last month, our data shows that X2 56 capable modems captured more than half the U.S. resale retail market in the 56k modem arena.

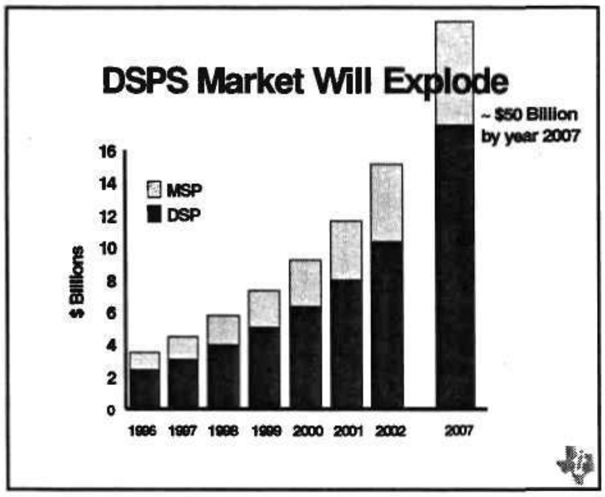

It is TI's DSP based platform and high speed modems, and our collaboration with U.S. Robotics, that has been, and continues to be, a very winning partnership and proposition.

At TI, we want to develop more winning relationships with innovative DSP designers like U.S. Robotics, and to that end, last month we announced a $100 million DSP venture fund, financed by TI and administered by Hambick & Quist. Its sole purpose is to fund the next wave of DSP applications.

As the world leader in DSP solutions, it's vital to us that the best minds in the industry are pushing the envelope on software technology on DSP applications, and we want to support the development of those applications with our DSP technology. We also want to put our investment money where our future is, in DSPS. We want to grow the entire spectrum of our applications using DSPs.

Since we announced the venture fund just one month ago, Hambick & Quist has received hundreds of inquiries from aspiring DSP entrepreneurs, and if that doesn't demonstrate the momentum behind the DSP market, and also TI's commitment to it, I don't know what will.

As this venture fund illustrates, building and retaining leadership in the DSP market requires much more than just developing great technologies and architectures. Perhaps even more important than the technology is the web of support you need to build around your products. By that I mean software programmers, universities, hardware and software companies, etc., all adding value to the architectural proposition we make to our customers.

One measure of a truly mainstream technology is the high level of support available from third parties. By this measure, DSP is certainly mainstream. We count more than 30,000 programmers writing billions of lines of DSP code, with more than 300 third parties adding value through software and hardware development.

For example, today, thanks to Dr. Aaphas of Tech Online TI is the only DSP company where you can get the design tools for our mainstream DSPs free on the net in a light format for a free test drive if you like — a great example of third party support bringing value to our customers.

Beyond third parties, there are more than 900 universities teaching engineering students DSP design methodologies on TI architectures. Today, we're fortunate to have a disproportionate share of these DSP industry resources focussed on TI architectures.

As I said, more than 900 universities worldwide teach DSP specific or computationally intensive engineering subjects at the undergraduate and graduate level. If you learn DSP engineering in a university, the chances are you'll learn it using TI languages. To accelerate these efforts, two weeks ago, TI announced a $25 million DSP university research fund. Its sole purpose is to support the software applications research, the high performance Digital Signal Processes at universities around the world into the next wave of DSP applications, things like wavelet technology and so forth. It complements our own investments like the Kilby R&D Center, in which we invested $150 million in Dallas, as we announced last month.

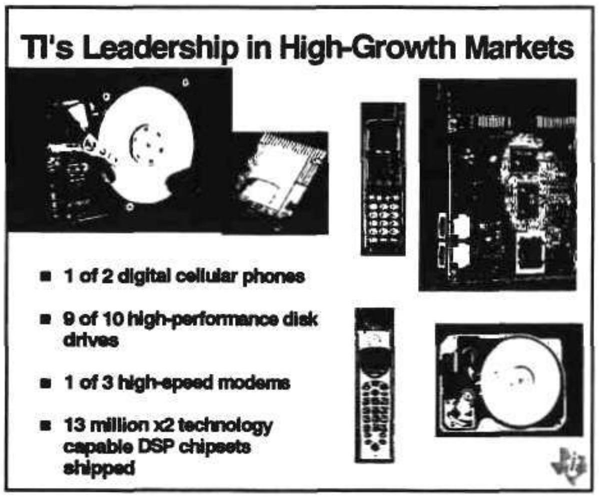

With these investments, we are building a strong future for the DSP market and for TI.

Our third parties are also building that future. I'd like to highlight just a few of the efforts of our third parties to give you a flavor of the innovative and valuable work they are doing. Spectrum Signal Processing and their CEO Barry Jinks have created innovative new boards, modules, and software to support TI's newest fixed point DSP, the C6X, which we announced at the beginning of the year, and the success of these efforts means they've already landed Department of Defense contracts for multiple million dollars worth of DSP products.

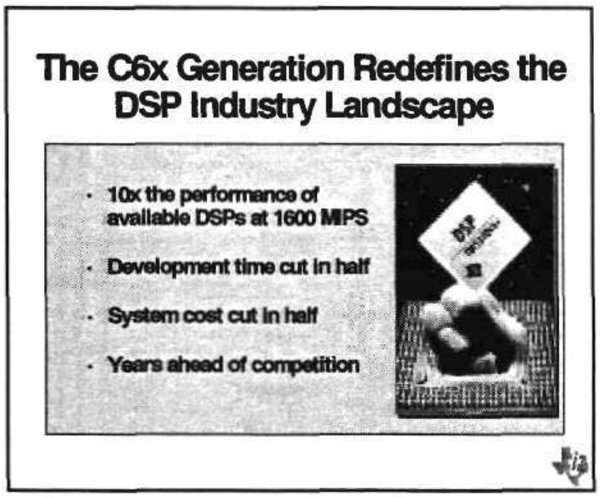

Spectron Microsystems and Bob Frankel, Spectron's Vice President and co-founder, developed SPARCs 8 years ago, the industry's first real time DSP operating system for the C30. I actually talk about it in that paper you have a copy of, which I presented in 1988. They now support SPARCs on many of TI's products. Currently, we're working with Spectron to develop an industry standard DSP software environment named DSP biles.

Hothouse Technologies and their President, Ross Mitchell, provide telecom software for TMS 320 customers. We're currently partnering with Hothouse to market and distribute basic software building blocks to the mass market for telecommunications applications. Although DSP is the industry leader for state of the art DSP de-bugging and emulation software tools, their code composer product sets a new standard for ease of use. It's two generations ahead of current DSP de-buggers and supports TI DSP solutions.

Finally, DSP Software Engineering in the northeast is one of the original TI third parties. They develop a variety of audio and telecommunications algorithms for all the S20 platforms.

This is a pretty impressive group of third parties, and it's really just the tip of a very large iceberg. If you're trying to solve a DSP algorithm problem, the chances are that a TI DSP third party had a major part of the solution already on the shelf in algorithms available.

With the efforts of people and companies such as these supporting the DSP market and TI's DSP architectures, the future looks more promising than ever. And, of course, in addition to this value web of third parties and university partners, TI also offers the highest performing DSPs on the market.

We have the broadest range of DSP cores of anyone in the industry. Rather than forcing designers to adopt a single architecture, we offer them a choice, not a compromise for their design.

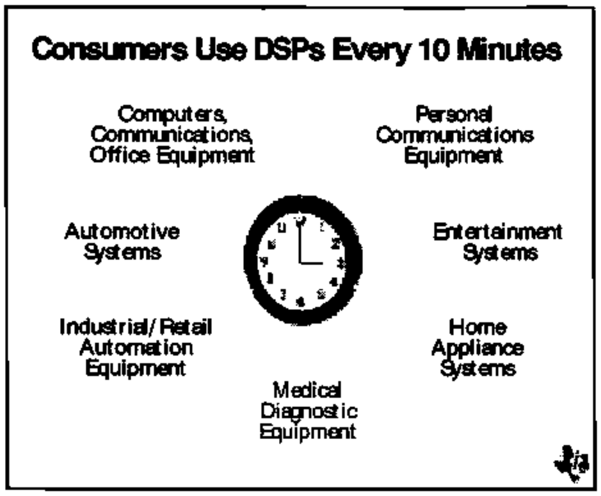

Recent evidence of the TI DSP advantage was the introduction earlier this year of the VLIWC6X DSP, which is going to change the way DSPs are used and the way electronic products are designed. At 1600 MIPs, the performance of this DSP is an order of magnitude that surpasses existing DSP chips. It will enable operations, which previously only happened in sequence, to happen concurrently.

For example, if you're an Internet service provider, the C6X can handle signals from 24 calls at once, whereas a single DSP could previously handle only one, two or possibly three.

This progress is reflected in the broad market acceptance of TI's newest digital processor. Leading Internet access providers, ADSL product developers, telecom switch vendors, and base station manufacturers have already chosen the C6X for their advance designs, and the total design ends for the C6X DSP are going faster than any of our previous generation of Digital Signal Processors.

Judging by the number of tool sets we've shipped to date, we have hundreds of developers working on the C6X already. Overall, we estimate around 100 new programmers a week are signing up to develop software for TI's Digital Signal Processors, and several hundred million dollars worth of business has already been identified in multiple market segments for that same C6X product, including communications and the mass market.

However, our customers told us they wanted more. They wanted a floating point equivalent to the C6X, so just last week we announced the latest addition to the ground breaking VLIW architecture group of DSPs, the C67X.

The C67X is the most powerful floating point digital signal processor core there is, with one billion floating point operations per second. That's again more than an order of magnitude faster than today's floating point DSPs. Designers can begin development on this core today, using the existing tool set for the C6X.

So where will all the power of the C6X and the C67X take us? To a world where two-dimensional imaging is a thing of the past. DSPs will handle the number crunching for advanced 3D imaging in digital photography, in high precision consumer video displays, and in medical and industrial imaging.

We'll see hand-held TV become a video conference forum, allowing parents to check in on their children in daycare, all due to DSP technology. We'll see voice recognition and identification technologies allowing us to control access to long distance calls from our cell phones by recognizing our voices. Things like retinal scan identification at an ATM machine will also make for better security to the electronic money network of the future — all, again, enabled by DSP technology.

Consumers already touch DSPs every day. Soon, it could easily be every minute.

In summary, the next wave in the semiconductor business is digital, and DSPs will be making most of the waves. The winner in the exploding DSP market will be the company that has the architecture, the process technologies, the installed base of knowledge and the product leadership. TI is that company.

Our strategy is to increase our world leadership in DSP solutions, and to continue to widen the gap between us and our competitors. We'll do this by sharpening our DSP focus, and by putting our resources where our strategy is.

Thank you very much.

Chapter 10: Big Opportunities in Wireless — Are RFICs Up to the Challenge?

David A. Norbury
President and CEO
RF Micro Devices, Inc.

I'm sure that none of you know very much about RF. I also expect you don't believe it has any impact on your business, but, quite frankly, there are some things I want to talk about today that you need to understand because some of you are in deep, deep trouble.

The world of wireless is exploding in anybody's terminology. The question is are RFICs up to the challenge? We ask that question because traditionally RFICs have been pretty much a mom and pop industry. Most RF activity over the last decade or two has been military. That's where cost means nothing and volume is low, which is not exactly today's world.

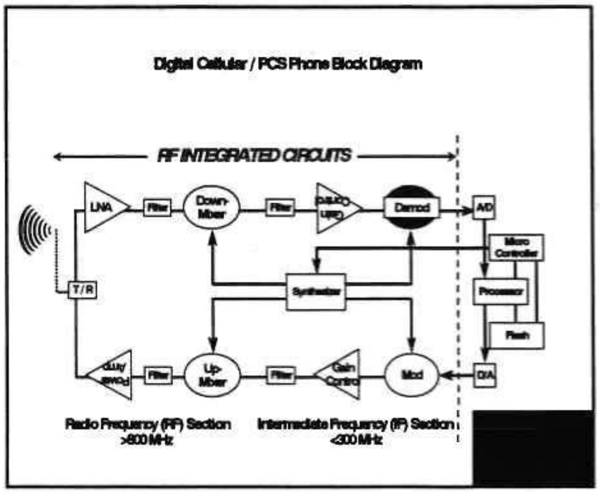

For instance, any self-respecting silicon billion dollar fab today starts more silicon area in 48 hours than a gallium arsenate fab does in a year. The scale we're talking about is tremendous. The bottom line is that the RFIC industry doesn't ship. Those of you that have a lot of silicon in these wireless applications don't ship either.

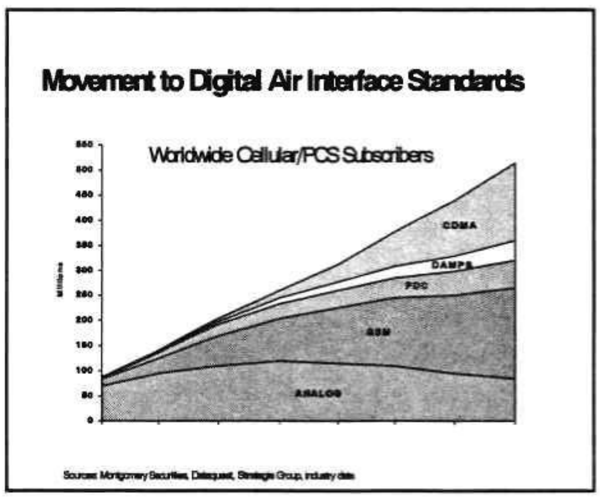

If you look on the far left between what looks like dotted lines and an antenna is what we typically call the RF part of a phone or any other wireless application. Off to the right of that in green is the world most of you are probably much more familiar with, the baseband side, where the processing goes on.

All the blocks in red are the kinds of products we're talking about here today.

If you look across the top, first part on the left, LNA is a low noise amplifier, the top part being the receive chain. What happens is the signal comes in, is amplified by the LNA, is sent to the down mixer or down converter, that takes the high frequency signal which is typically 800 Mrtz or so for cellular, 1.9 Grtz or so for PCS. The down mixer converts it down to a much lower frequency that's easier to process, which we cleverly call the intermediate frequency, which is typically 100 Mrtz or so.

The next red block is gain control. This thing has to have a wide range of gain depending on whether you're near the cell site or at the very edge. It then sends it to a demodulator. The demodulator strips off the 1s and 0s from the RF signal and sends it into the baseband part of the system.

On the bottom is the transmit change which goes in the reverse direction and pretty much does the same thing in reverse. The power ramp on the far left I might point out is probably the most difficult part of any RF system. That's where most of the distortion is generated.

The kinds of applications we're talking about today are essentially these. Today the big horses are essentially cellular and PCS. They account for something like three-quarters of the wireless world today.

A new one just coming out is called wireless local loop. It's essentially a cellular phone that doesn't go anywhere. That's where you can move a phone system into a third world country very quickly, typically six months as opposed to a wire line version that might take two to three years.

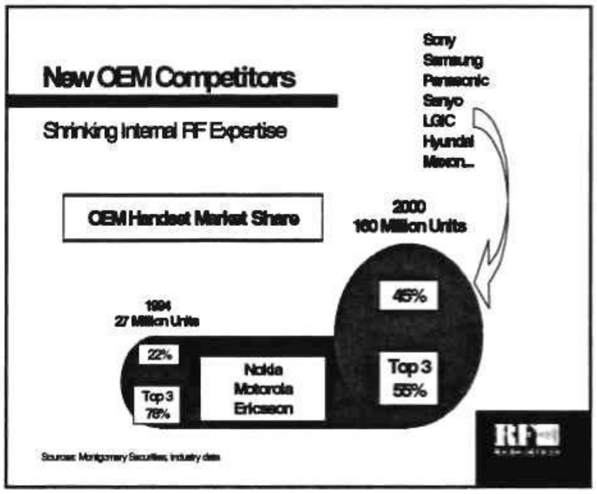

What you're doing basically is putting a cell site in the middle of a third world village. You hand everyone one of these local loop phones; they nail it to the wall, and they're up and running. They've found that it works better because in a lot of third world countries where they've tried to bury copper to run lines, the natives dig it up and sell it.

Some of the other areas that aren't quite as big but which we're seeing interesting growth in include things like cordless telephones. If you've bought one of the newer 900 Mrtz spread spectrum phones, and you're used to the old 49 Mrtz that have been around forever that are total junk, you'll hear a difference. The signal is better, there's no static and you don't lose your call, and your neighbor can't hear what you're saying.

The RFIC world is helping implement a lot of these new technologies which everyone is appreciating at this point.

Some of the market drivers pushing the RFIC world around are here and we'll talk a bit about each one.

First, the end-users, that's the people in this room. I expect most of you have a cellular phone or a pager or something like that. We're pretty demanding people. We want our next phone to be smaller, lighter, sound better, battery longer lasting, etc., and cost less and do more. We're doing a lot to drive the OEMs crazy at this point in time.

Another trend is the move to the new digital air interface standards. An air interface standard is a long name for a method of gluing your voice onto a radio signal. There are lots of different ways to do it. Some listed here, GSM, PDC, the alphabet soup type — in the bottom you see analog in purple type. All the other standards are what people call digital air interface standards. The reason the new digital air interface standards are so popular is that the carriers can take the same space the FCC gives them and shove two to 10 times the customers down that same pipe as compared to analog, so for them it's free. It's a wonderful upgrade to their system.

From our point of view, in the RFIC world, that makes our job a lot tougher. When you stick more signals in the same space, effectively every signal has to be cleaner to keep them from interfering with each other, especially in the power amp. I mentioned that's one of the trickier parts of an RF system.

The power amp is where most of the distortion is generated and is typically the bottleneck for digital systems.

Another trend is the move to higher frequencies. The reason is simple. As you go up in frequency, there's more space. You get a bigger pipe to shove your data down. At lower frequencies, 900 Mrtz, maybe you're talking about short message service, as you get up to a couple of Grtz and more, you get to the point where you can now talk about some pretty decent Internet access. The trend to higher frequencies and wider bandwidths is going to continue, and as you go up in frequency, things like RFICs make it a lot easier than if you tried to do it the old-fashioned way, putting individual transistors and resistors on a PC board.

At the OEM level there are a lot of problems. The life cycle of a phone is now roughly 12 months which is shorter than the development cycle for the phone, so what you see for OEMs is they'll have two and three generations of a particular phone in development at the same time. If they're a big OEM, they may have a dozen different phones. Compounding that is the fact that the new signal protocols, the CDMAs, the TDMAs, the more complicated protocols are making the hardware job tougher. This is driving a lot of the OEMs to use outside suppliers like us to supply RFICs.

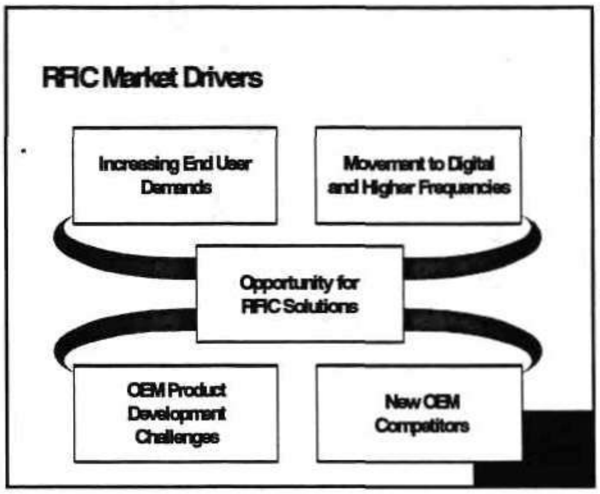

We live in a very "nichey" field where the practitioners are tough to find so it's very common to now outsource these kinds of parts.

The last market trend we'll talk about is very interesting. If you look a few years back, the big three, the Nokia's, Motorola's and Ericssons owned about three-quarters of the handset market. Each of them had tremendous internal RF expertise. They really understand system design.

If you look out a few years, it's forecasted that roughly half the handset business is going to be owned by a brand new class of competitors, pretty much what you'd call consumer electronic companies, the Sony's, Samsung's, Panasonic's, etc. These guys generally don't have anywhere near that internal RF expertise as the big three. What you find is the drive for outsourcing to RF suppliers is going up and up as these new players enter the market.

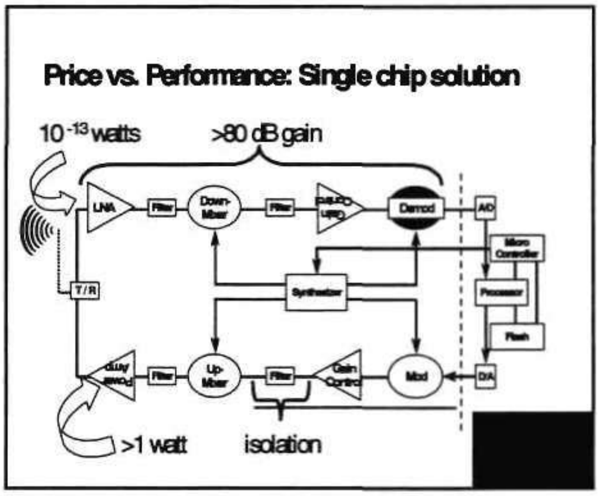

The point of all these drivers is that the market for RF integrated circuits is exploding. There are more opportunities than we know what to do with, but there are some problems.

I'll talk briefly about industry capacity and some packaging issues. I want to spend a little more time though on the tricky issue of price vs. performance. This is especially tough in the RF area.

In terms of capacity, there really aren't any big RFIC gap suppliers. There is a handful of relatively mom and pop shops compared to the mainstream semiconductor world. The other thing that makes it tough is that RF designers are hard to find. I've been in RF and microwaves for better than 20 years. There have never been many of them and they've never been easy to recruit no matter what the economy does.

Lastly, the production ramps, and for people coming out of the military world this is gut-wrenching. To be told you got a design in, you get all excited, then they tell you they need 50,000 a week starting in two weeks. Ordinarily, in the military world that would be a 12 year span, so you'd have a little more time to prepare.

Packaging. Most RF products are low pin count parts. The majority are in the 8 to 16 pin packages. We've been real lucky because we can use widely available, low cost, SLIC packages for most of our products. That's the good news. The bad news is we're moving into areas where the pin count is going up, the power levels are going up and more RF issues are becoming more important.

We talk to the packaging houses and they haven't the foggiest idea of what we're talking about. They're there mainly for you guys and you don't care about the same things we care about, so we're finding this to be a much bigger issue than we ever expected. My conclusion is that the fragmented base really needs to get together and create a more forceful impact on the packaging world.

Price vs. Performance. Every day, just as you do, we fight the same issues. It has to be lower price and you have to have better performance. In the area of lower price, the number of interface standards is a problem. You can lower price by going to the holy grail, which is a single chip solution which I'll talk about in depth, technically and financially, and put that to rest.

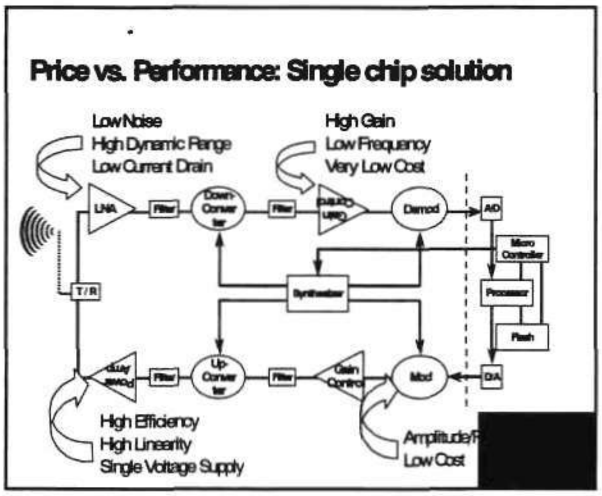

On the performance side, you have customization of RFICs for different applications and then multiple semiconductor processes for various functions.

There are too many standards in our opinion. Obviously, we don't call it but every one of these pieces of alphabet soup requires different RF components. In fact, I suspect that the only way you can get a standard approved now is to prove that no parts that work in any other standard will work in this new standard.

Even when you're talking about the same standard, Company A and Company B both insist on starting with a clean sheet of paper and they want to do enough different to where you can never sell two guys the same part. It's very frustrating.

In the attack on the single chip solution world, you see more and more people implying that's the answer: let's put a cellular phone on a single chip. Great idea. I'm not one to say it will never happen but there are some issues, the first one being some cross chip coupling problems.

Across the top on that receiver chain, for instance, there's about 80 dB of power gain, front to back. That's a ratio of 100 million to one. As you shrink that chip down all those interconnects start acting like antennas. With that kind of gain you turn this thing into an oscillator very quickly.

Similarly, across the transmit to receive supply, the power amps pumping out something like a watt and the poor receiver on the edge of the cell site, he's looking for signals at a level of 10 to the -13 watts. Thirteen watts of magnitude and you want to put those within a millimeter of each other. You're nuts! That's why it's not done today.

Another issue the thing I've called "isolation" down there. Our customers pay a lot of money for soft filters for this application, extremely sharp skirts, rejection 50-60 dB, they've got to have it. If this was a single chip implementation, the only way you'd be able to do that is to pin out the output of the gain control and the input of the up converting mixer. How much isolation do you think exists between two pins? You might as well throw that filter out. It just won't work. Thus, from a technical point of view, there are significant issues, at least in the short term, to keep us from coming up with a single chip solution.

From a financial point of view, the various different elements that make up the RF part of the system have very different requirements. It's got to be a low noise amplifier with a very high dynamic range. At the edge of the cell site it's looking at a 10 to the -13 watt input. At the near edge, near the base station, it may be seeing as much as a millowatt. That device has got to be able to handle 10 db of magnitude of range.

Further down the receive chain, it's got to have lots of gain and cost nothing. There are some very conflicting technical requirements with large implications on the process technology you use to implement these.

On the single chip solution, you've got no choice but to pick the one process technology that will do all of this. Unfortunately, that's always the most expensive one.

Quickly talking about another way to optimize performance is to come up with a custom IC for every application. You certainly get better performance. The downside is it stresses your development cycle and, you probably know this, in higher volumes you generally get lower prices.

A lot of our customers don't apply this to the RF world for some reason. In fact, if Intel followed the RF business model I suspect there would already be about 45 versions of the Pentium II. It's a tough business. Customers won't let us use standards for the most part, standard products.

From an industry supply alone, we need a lot more volume to get the supplier base to be a stronger, healthier business.

When I talked about the different process technologies for the different parts of the RF system, I referred to a term we coined at our company, "optimum technology matching." Power amplifiers are notorious in their requirements for efficiency, linearity, and they want a single voltage supply. For us that means gallium arsenate HBT, Hetro Junction Bipolar Transistor.

At the back end of the system where frequencies are low, silicon is a great way to go. Our approach is to match the process technology to each part of the system and we think we get the overall best performance and cost from that point of view.

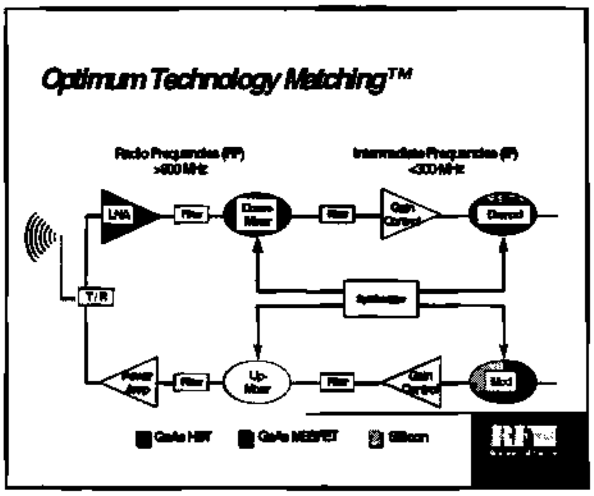

The last technology is one we have in R&D. We are willing to look at just about anything. We're technology agnostics from that perspective. Although we've been fabless in all these technologies up to now, we're just completing a gas HBT fab in Greensboro, so we'll have a mixed HBT model at this point.

This is a graphic description of the different kinds of technology and where they make sense. Again, the gas HBT is a wonderful technology. At low frequency, silicon makes lots of sense. Gallium arsenate MESFET has some nice applications for amplifiers and control devices. When you talk about a single chip solution, again you really have to pick the highest cost technology, which will drive the price up.

Some people when they talk of single chip solution really mean you take all the baseband processing and all the RF stuff and put it on a chip. That would be a wonderful thing but I don't think you'd like the cost. The area of all these chips is probably about 5 square mm. The process technology on a per unit area basis is probably on the order of 5 to 10 times more expensive than baseband process technology. Everything is tilted in the wrong direction.

Some day, I suspect we'll get there but I don't think it's going to be short term.

In summary, RFICs are absolutely required for the complexity of wireless products today. You couldn't jam enough discrete components into a cellular phone to do what a CDMA phone does today. You'd literally be carrying around something the size of a television.

The demand for RFICs is high. You need to know that because you're paired with us now in a lot of handset applications. Our supplier base is pretty fragmented. Our levels of integration are increasing. For instance, we have complete transceivers on a single chip today. That doesn't violate the single chip argument I just made. They're pretty poor transceivers, but for low cost applications they work great. For high performance, the single chip is a way off.

Last but not least, I think new RF-friendly packaging technologies absolutely have to come into being. We can't afford to keep going the way we're going now.

Are there any questions?

Q: What about silicon germanium as an alternative to gallium arsenate technology?

A: Great stuff, I think. We think silicon germanium has some potential but we don't know enough about it yet. We've talked to IBM on and off over the last few years. It typically has a fairly low breakdown voltage which limits its application for power amplifiers but for some of the small signal applications we think it has some pretty good potential.

Q: Wireless manufacturers demand more power, lower operating voltage devices to prolong battery life. What's the voltage migration trend for the RF portion of wireless handsets?

A: The trend in voltage is clearly down. We're driven by the same things the baseband chips are. The power amplifier is the toughest one; it needs the most voltage. What everyone is shooting for today is a three volt power amplifier. Some of the other components we think we can get down a little lower but three volts we think is going to be about the limit for power amplifiers for awhile.

Q: How many radio frequency chips are now in a cell phone if one chip isn't currently possible?

A: It depends on the kind of air interface standard you're addressing. Typically, all the red RFICs I had on the diagram would be in somewhere between 2-4 chips and the ASP would probably be $8-15 in total for all of those.

Q: Do RF devices have their own waver fab? If not, what foundry do you currently use?

A: We're currently fabless in our three main technologies. Silicon Bipolar, IBM is our primary supplier, gallium arsenate MESFET, Triquin is our main supplier, and gallium arsenate HBT, TRW is our only supplier. They own the technology and that's the technology they've licensed to us and will be operational our fab next year.

Q: How have cellular and PC handset manufacturers embraced ICs versus the imbedded discrete solution.

A: For the most part, they've embraced RFICs tremendously. It allows them to shrink the phone and gets them out of a lot of the manufacturing tuning and testing issues. It lowers their costs over the discrete implementations. Our sense is that if an OEM can figure out how to get an RFIC to do the job, they'll do it. That will be their priority.

Q: What are prospects for non-communication solutions such as automotive?

A: Are you speaking of our company or wireless in general? I'm not really sure what you're trying to get to. RFMD is strictly a wireless company. We're 100% focused on that. Eventually, we expect some of the same technologies we use today for wireless will have applications in other areas where linearity and efficiencies at very high frequency is important. Certainly, fiber drivers, high speed data, automotive, undoubtedly this kind of technology will play into many, many different markets.

Q: What kinds of RFIC design tools are currently available today?

What is the need for these tools?

A: One of the things that makes RFIC design so difficult is that there is really not a turnkey equivalent to what you might see in the digital world. You don't go to Cadence and buy a complete suite that goes soup to nuts and you're done. In the RF world each engineer gets used to a certain set of tools. There's no womb to tomb kind of package. Seems like every engineer has his favorites, P-spice is a very common one for circuit analysis. There are probably three or four different layout tools people use.

I would say if this is from a computer-aided design company, they should take a hard look at RFs. It's really wanting in that area.

Chapter 11: IC Manufacturing 2000-2010: A Decade of Momentous Change

C.M. (Mark) Melliar-Smith
President and Chief Operating Officer
SEMATECH

Moderator: My name is Joe Grenier. I'm going to be your host this morning for the second day of our conference. Our first speaker is Dr. Melliar-Smith. He was named President and CEO of SEMATECH last year, November 1996. Before joining SEMATECH he was Executive Director of AT&T's Bell Labs, the IC Division and Chief Technical Officer for Lucent. In his 25-year career with AT&T he has worked in a wide variety of assignments, including fundamental research, electronic and photonic device development and manufacturing and business unit management. He is also chairman of Semiconductor Research Corporation Board of Directors. Dr. Melliar-Smith has a Ph.D. from Southampton University in England and an M.B.A. from Rockhurst College.

Dr. Melliar-Smith: Good morning. I have the rather unenviable task today of getting everyone warmed up this morning for today's session and maybe in a few cases actually waking people up. I took the liberty of choosing a somewhat hyperbolic title, namely that we face a decade of momentous change in the semiconductor industry. That's really what I want to chat about today. Momentous versus incremental depends on your perspective. If you are outside the fab, the changes may look like they're incremental. Believe me, if you're inside the fab, over the next 10 years they look pretty momentous to me.

Before I actually get started with my presentation, I wanted to provide a little bit of background as to who I am and what we're doing at SEMATECH. SEMATECH is a consortium of about 10 U.S. semiconductor companies that's engaged in manufacturing R&D. It started about 10 years ago. It's located in Austin, Texas. We have about 600 employees of whom about 160 are engineers on loan from the member companies. That's a very valuable important way we transfer technology back to our members. We run a $120 to $150 million a year budget, and our focus is very much on equipment and unit processes. Fully integrated process we leave to our member companies. They usually view that as being proprietary and strategic, so we tend to concentrate on individual tools and particular problems.

This is a curve that is fairly familiar to everyone in the room. We were regaled yesterday morning with

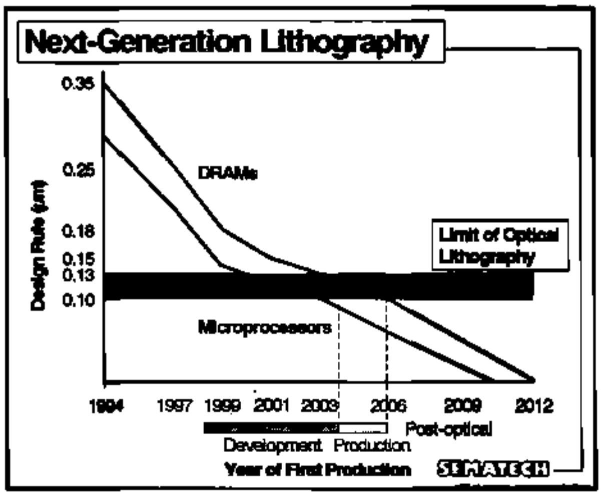

comments like 2.7 trillion e-mail messages helping drive this business to $150 billion a year. For those of you like me who have a hard time understanding 2.7 trillion e-mail messages, it works out to be several hundred e-mail messages for every man, woman and child in the world, and we were also told yesterday that half the people in the world had never made a telephone call. It's fairly clear to me that this is a concentrated business. If nothing else, this curve and this tremendous exponential growth for 30 years is a remarkable achievement in an industry. Since almost everyone in this room is dependent on this curve for their continued employment, I think it's worth looking at what are some of the things that could possibly cause us to come off this curve or have it turn over on us.

In the Andy Grove model of only the paranoid survive, I tried to look at four different things that might cause us not to be able to stay on our present growth curve of about 15% to 20% a year. First is obviously that we are blessed with what appears to be a completely elastic market. You lower the price, and the volume goes up even more. If you run out of new markets, that's not going to continue. A second problem we may find is that we run into insuperable technical problems. We've had a history of being able to solve anything that physics throws at us, but as we get down to 15 nanometers or less design rules, there are some very significant physics problems that we're going to have to solve. Another problem we have is managing complexity, whether it's complexity in design, complexity in software. We also have heard a lot of talk about systems-on-a-chip, and anyone who's been involved in systems R&D knows that the biggest problem there is not hardware, it's software. That's something that if this industry aspires to doing that, it's going to have to change very significantly into a software centric industry. Last but not least, and this is the thing I want to talk about today, is that we fall off the manufacturing productivity curve.

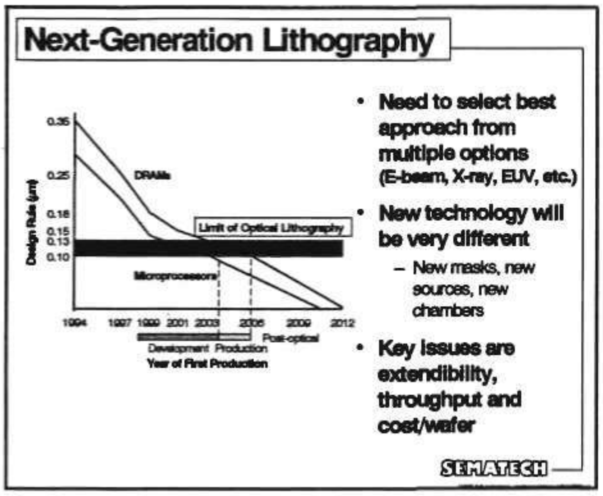

What do I mean by that? If you look at this industry for the past 30 years, we've been able to deliver to our customers an improvement in functionality, and that functionality could be the number of transistors, it could be the number of bits in memory, it could be millions of instructions per second if you're selling DSPs or microprocessors. All of those have improved approximately 25% to 30% compounded annually per dollar of price. That is truly a remarkable productivity improvement, and it's based both on cost, lowering the cost of the transistors by changing the design rules and doing other things. It's also based on performance, running the clock speeds up, making the individual transistors do more work for us.

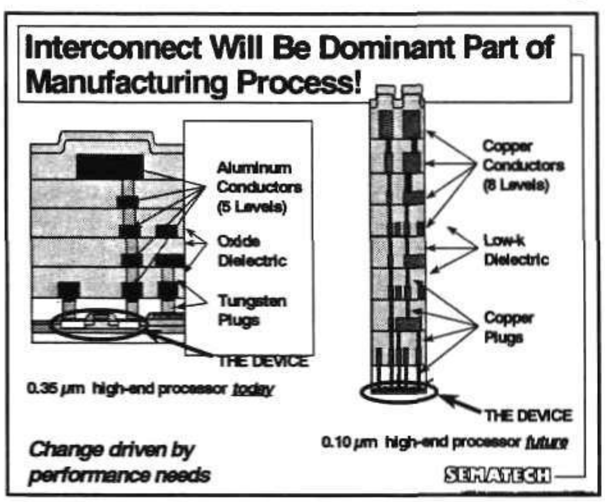

As I look at this curve, this is what concerns me. If we start to fall off of this curve, then obviously the things that have made our industry so valuable to our customers and so important to us will begin to change. You look at that curve and you say that's all right, you've done it for 30 years, just keep doing it for another 10 until I retire. Unfortunately, that's not actually going to be quite so easy to do in the future.

The problem is that in the past 30 years we've essentially had incremental change in our technology. What do I mean by that? Let me explain that anecdotally to match all the data. When the integrated circuit was invented about 40 years ago, Jack Kilby in Texas hooked up the individual components on a sliver of gimanium using wire bonds. Almost simultaneously, Bob Nice in California joined up the individual components on a silicon substrate using photo-defined aluminum. Both of those technologies are in almost universal use today. That's what I mean by incremental. What I mean by discontinuous is that both of those technologies, aluminum metalization and wire bonds, are coming to the end of their life, and we'll talk about that. We really do face a decade of non-incremental change. That's what I worry about on the curve.

What are some of the non-incremental changes? Next-generation lithography. We'll talk a little bit in more detail, but optical lithography is running out of steam. We've got to talk about copper replacing aluminum. Low-k dielectrics for interlevel dielectrics, no longer using

SIO2 based dielectrics. Even as we start to go below 100 nanometer design rules, we may start to change the fundamental materials we've used for many years in the gate junction. In fact, as you look at these last two, low-k dielectrics and deposited high-k dielectrics for the actual device, it's also worth remembering that the reason that silicon was chosen as the industry standard semiconductor material was not because of its semiconductor properties. There are, in fact, semiconductors with electronic properties at least as good if not better than silicon. It was chosen because it grew a very good oxide on the surface, and now we're beginning to say maybe we don't want to do that anymore. That represents another very significant change.

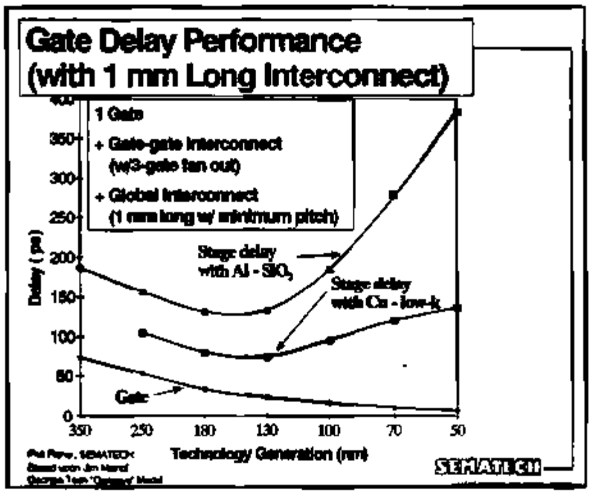

You've got all these non-incremental changes plus, of course, all of the things that I will now call incremental, like going to 300 millimeters, reducing the design rules, $300 billion fabs and the like. There are a lot of problems that we face, and I want to go through these in the next few minutes. Unfortunately, I don't have enough time this morning to go through them in detail. This is going to be a fairly high level dissertation just to give you a sense of where we're headed.

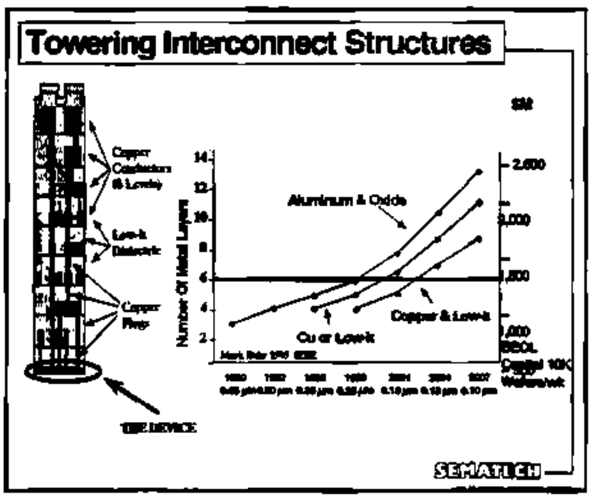

The first thing I did was try to break it down a little bit more quantitatively. This is the 25% to 30% curve, and it has varying different contributions to it. Obviously, the largest contribution is that feature size has been coming down. That's the classic scale of integration, make the transistor smaller, pack more on a wafer, get the cost down. If anything, that part of the issue in the future is beginning to accelerate. Partly that is because of product reasons. We can make devices faster if we make them smaller, and partly it is to make up some of the slack of where we're falling off. For example, some of the yield improvement—once you get the wafer yield up to 25%, it is just physically impossible to get much more out of it. What we'll do is we'll look at each of these in turn and see what the future holds for us.

The first one on feature size shrinking, it affects a number of different things, but let's start off with lithography. This particular graph shows the minimum design rule in beginning production as a function of time, and you can see it's been coming down very steeply with microprocessors actually in the last 10 years overtaking DRAMs. The reason for that is that the critical lab, gate level patent for microprocessors, controls in many respects the clock speed and clock speed controls the average selling price. There's a tremendous pressure here to keep this small. We haven't found quite the same way to drive average selling prices for DRAMs, much as we'd like to we heard yesterday for the past 18 months without much success. If we look at where we've been in photo lithography, the way we've come down this curve is to shorten the wave length of light, going from i-line down to 248, and we'll go from 248 down to 193 nanometers over the next couple of years.

The problem we face is that even with the best efforts in terms of mass making, it looks like it's difficult to make a patent more than perhaps about six-tenths of the wave length of light. That means that 193 nanometer lithography is going to run out of steam somewhere around 130 to 100 nanometer design rule. Just where it runs out is a matter of scientific debate, and it also will depend a little bit on your product.

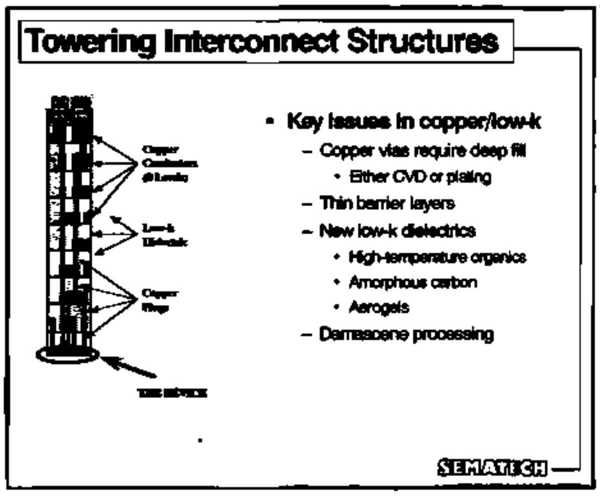

However, once we get to this point, we can't go further in our present UV technologies simply because once you go below 193 nanometers, nothing transmits ultraviolet light anymore, so you can't very easily get it through lenses. You can go to calcium fluoride, but I don't think there's enough calcium fluoride in the world to let us do that. Resists get to be extremely difficult, and even going through air gets to be a problem. We're looking at a limit around about 130 to 100 nanometers, and we need a new technology. If we need that to ramp manufacture around 2005, we need to start making decisions to give the lithography equipment companies a chance to develop it sometime in the next year or two.

We need to select the best approach from multiple options. There are about six different options, a couple of variations on electron beam and extreme ultraviolet technique going all the way down to 13 nanometers, X-ray technology, EUV and E-beam technology. Each of these technologies is capable of doing 70 nanometers. They all have their scanning electron microscopes. It's what we call in the industry their baby pictures, which all look very nice. The real issue here is trying to decide which is the technique which will serve us best, and it involves a number of different things, extendibility and cost, that sort of thing.

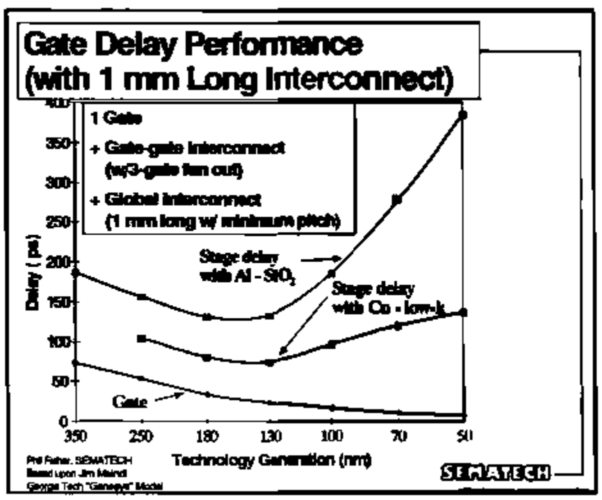

The industry as a whole needs to come out with a consensus because we really can't afford to go down three or four paths simultaneously. We need to reach a consensus, and that consensus obviously is going to be a free market decision. It's not going to be a decision made by a couple of SEMATECH directors in a smoke-filled room somewhere. It's made by people who write checks with eight zeroes after the first significant digit to buy the equipment. Nevertheless, SEMATECH is providing a forum where all this data can be put on the table in an open and clear way, and we're trying to build a scientific-driven consensus is the best technique.

The new technology will be very different. The masks are going to be very different, and I'm sure in 10 years we'll look back with nostalgia at the simplicity of chrome on glass masks. The sources will be very different. If you're going to be a syncotron instead of having a laser source, sort of half the size of an automobile, you're going to have a syncotron source the size of a parking garage. New chambers are going to be done in vacuum. That's a new experience for anybody making lithographic equipment.

In addition to that, the key issues are extendibility. Whatever we choose, we want it to be able to go down a long way. I don't particularly want to go through this transition again in my lifetime, so we want to make sure we can extend them.

The other issue is throughput and cost per wafer because fundamentally the real issue on this is not can the technology, the new technology, do 70 nanometers. The key question is how are we going to do 70 nanometers at $50 a level? I used to work for a manager who said his definition of a good engineer was somebody who could do for $1 which any fool could do for $2. That really is the name of the game. If we want to stay on our productivity curves, we've got to keep the costs under control.

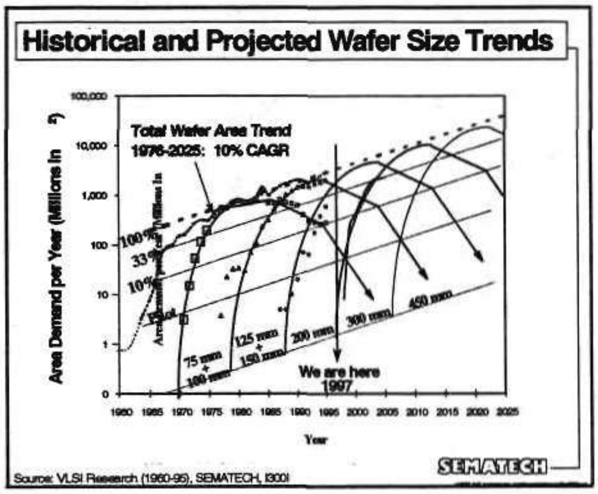

The next thing to look at is interconnection, and that deals both with the cost and also with performance. These are typically where we're at today, and we're looking to go to very complicated interconnection structures. You can see the device now becomes a small part of the integrated circuit. In fact, when I look at this, I kind of smile to myself at the slightly arrogant point of view that I see engineers take over their printed wiring board brethren in terms of high technology. The IC world is moving into the printed wiring board business. Essentially, we're putting micro printed wiring boards on top of our devices.

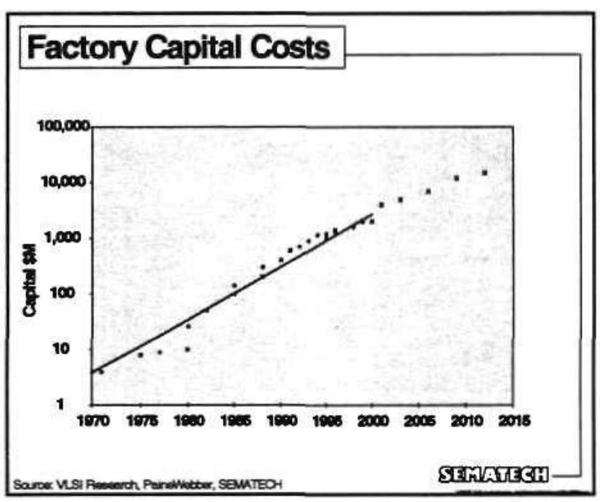

The reasons for that are well known by everyone in the room that essentially as we make the devices smaller, the gate delays become an increasingly smaller and smaller fraction of the total delay. The exact position of these curves depends on your particular choice of a simulation, but in this particular case we're looking at one-millimeter lines on minimum pitch and you can see that once we get down below 150 nanometers, even if we go to copper, you end up in a situation where the clock delays are starting to get dominated by interconnection. The reason that that is the case is that we chose some fairly severe requirements in terms of a millimeter length, line length and minimum pitch. If you keep the line length shorter and if you go to reverse scaling so the longer lines are wider, then obviously this has less of an impact. It comes at a price. If you want to go to reverse scaling and you want to make the upper levels of metal wider, then you end up with more levels of metal.

You can see on this chart that if you've got around eight levels of metal, then the capital cost for a large fab just for the back-end of the line, just for the interconnection, starts to get prohibitively expensive. If you're going to spend $1.5 billion on just the back-end, you may make our friends from Applied Materials and LAN very happy in doing that, but it starts to become a very significant burden on the cost of building a device.

You also face some technical problems in these interconnection situations. With copper and low-k dielectric, the copper is going to be required to fill deep, deep trenches. Typically, we're talking about nearly 100 nanometer veers with an aspect ratio of up to five to eight. That's going to require either some form of probably chemical vapor or some form of plating.

In addition to that, we need very thin barrier layers to keep the copper up here away from the device down here. As another vignette, in the early days when people were struggling with silicon transistors about 35 years ago, they knew there was something that was poisoning the diodes and destroying the diodes and they didn't know what it was so they christened this impurity deathnium because it would kill the diodes. It turned out the deathnium was copper. Now here we are going to put it all over the wafer, so we need barrier layers but they need to be thin. If you're going to drop that barrier layer down the inside of a 100-nanometer veer, you can't afford to have it very thick or all of a sudden it takes up the whole veer and doesn't make it conductive.

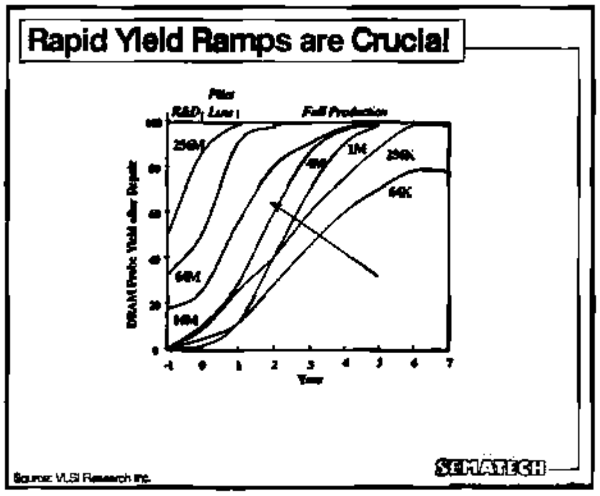

New low-k dielectrics. We're looking at things like fluorinated amorphous carbon, parolenes, but ultimately we're headed towards aerogels. If you can imagine trying to put down many layers of metal that's a rather different technology. It also gives us a challenge when we're thinking about aerogels, which are air, and copper, which is a conductor. When you get to copper and air, it's hard for me to think of any materials which are likely to be much better than that for interconnection. You start to run into one of those things that I call an insuperable problem. There will be plenty of people out there saying we'll do it with integrated optics. We just happen to be about five orders of magnitude in cost away from where we need to be, or perhaps we can do it with superconductivity.

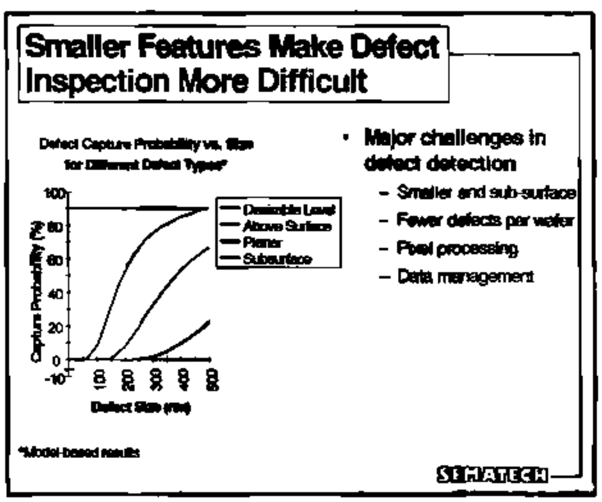

The problem is the superconductors we know todaydon't work at room temperature or anywhere close. Also, they don't work with the sorts of frequencies that we need for our ICs to work. We've got some very big material problems which are not going to be solved in the next decade here. That's the challenge we've got in interconnection.

In transistors, our challenge here is I've assumed that the gate was going to continue to speed up as we go down in dimension. The problem is once we start to get down around 17 nanometers, we begin to run into some fairly significant problems with the device physics. The leakage currents get to be a problem. It gets to be a problem. The oxide gets so thin. Once in that range in about the next 5 to 10 years we're going to be looking at a lot of change in the transistor structure to potentially metal gates to get away from depletion effects with polysilicon. New high-k dielectric constants. Different direction in terms of dielectric constant from where we're going in interlevel dielectrics. We're going to have to get much better threshold voltage control because we're headed down to one volt, but people will still want decent performance and dynamic range. Very shallow junctions and doping contacts. You're literally going to start to be able to count the number of doping atoms in a junction, and that's going to draw us down towards some basic statistical limitations that we can't get around. Also, potentially new sub like silicon. That's really covers what I want to say about feature size coming down.

Wafer size has given us a continuing cost reduction. It tends to come in big jumps, but on average it runs a few percent a year in productivity. The 300 mm transition will begin starting next year with the early pilot lines, but actually when it sort of comes up still faces a number of challenges.

The first is obviously the timely availability of a capable tool set. The biggest problem we have there is litho late. The DeepUV 300 mm scanners are not yet available even almost as alpha tools, so that's an issue for us. We've got some in repeat cameras which are available, but they're probably not suitable for production. We need the cost and footprint to be right. There was some discussion yesterday afternoon in the panel session about that, and I'd like to emphasize what Dale Harborson said from TI—we need the cost to be no more than 30% above that of 200 mm tools and we want the same footprint. We need cost-effective automation. The cassette carriers will be too heavy for the average people to walk around with, so we have to have full cassette automation. We're also looking at mini-environments around the tools.

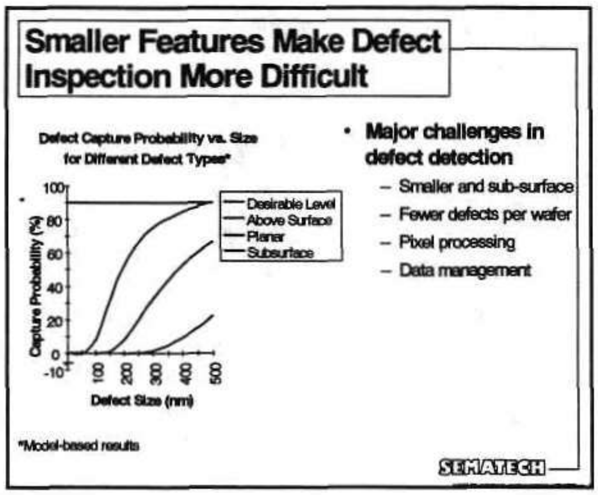

Here's another which is a challenge for us, and it comes back to the first problem—to develop the tools, we need silicon wafers to test them, and obviously it's difficult to get those without litho tools. We're having a hard time getting enoughsilicon wafers at quarter micron and 180 nanometer to give to the equipment vendors to test out their tools. Finally, and this is one that the semiconductor industry owns, we have to get out of the mode of preaching standards and buying specials. That is a particular problem in our industry, and we need to try to fix that.

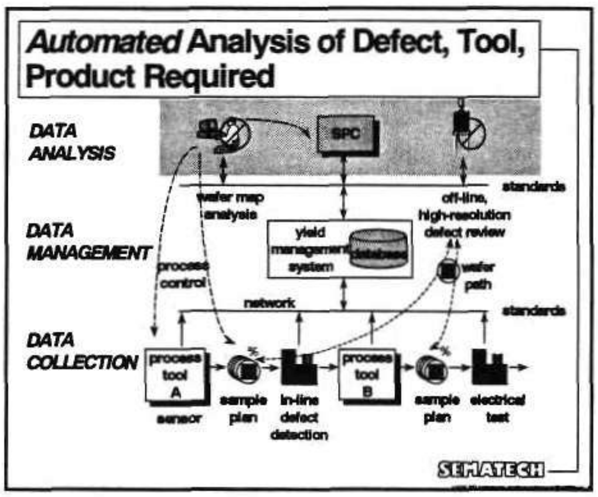

This shows you why I think this transition though will be somewhat incremental. What you can see along here in thousands of millions of square inches is the total increase in the area we process each year, and then underneath it you can see each of the share of the different wafer sizes. Three hundred mm will be providing 25% to 30% of the total production sometime early in the next century, and we're right sort of on the target for coming up this curve.

You may have wondered like I did yesterday—what is a billion square inches? It's hard for me to visualize that. For your information, 5 billion square inches is about a square mile. Right now, we are processing in our industry about a square mile of silicon every year.

Equipment productivity relates to a whole variety of things. Equipment costs I have already spoken about. We think there's some limits on what we can afford to pay for it. In addition to that, we also need to look at yield and defect detection and control.

If I go to equipment costs, as I looked through the attendee list a couple of days ago, I saw there are a large number of people from financial institutions. This is the chart for the financial institutions. It shows the cost of wafer fabs for the last couple of decades, and it's been on a fairly long linear curve going up. It's around 2 billion today.

If we carry on this curve, sometime early in the next century we'll go through $10 billion, at which point I think a wafer fab will be the most expensive industrial building or entity in the world, exceeding even nuclear power stations. The problem I have is I can't see any compelling reason how we're going to stop this continued exponential growth. We've taken a couple of runs at it in the past without much success, so I'd have to say that history tells us it's going to keep coming up this curve. I would say we're going to be looking for some pretty heavy funding requirements.

One of the problems, of course, is that if you are going with a $10 billion request to your board of directors, it's likely that your board and certainly the chairman of the board and CEO want to know that you're going to do a good job ramping. This shows the improvement that has been made in DRAM ramps. A lot of this is hard data. Obviously, the 256 megabit is data that suggests that we're being hopeful there. You can see there has been an improvement to the rate of which we can ramp the fabs and turn the tremendous expense into a profitable asset.

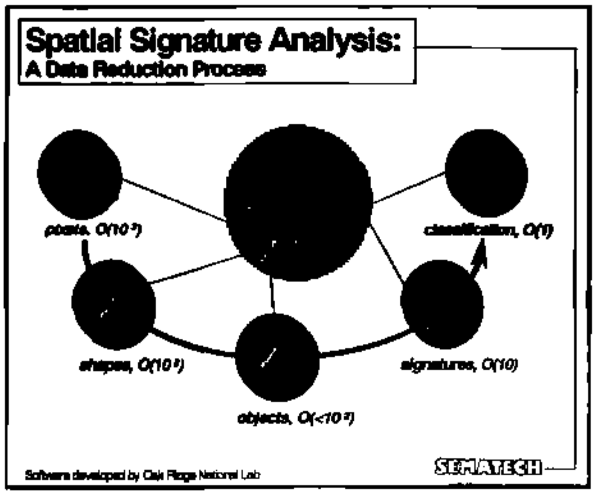

To do this, we're going to have to get better control over our defect detection. This is a pretty big challenge for us. The first problem we face is obviously the defects are getting smaller and very deeper in the wafer. One of the hidden secrets or less public things is that when we worry about optical lithography—everyone knows about optical lithography—but the defect detection people face many of the same problems. You can see as the defects get smaller, our present optical technology is getting much less effective in their ability to detect them.

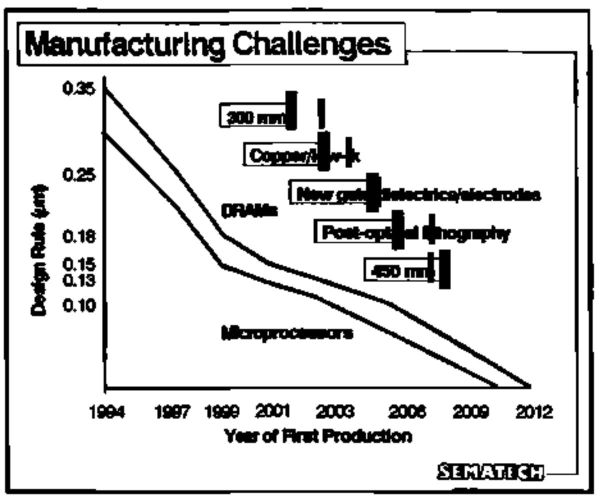

I said I had a financial analyst chart. This is my media chart for the media people in the audience. You can see it's a picture of the world and a golf ball and a satellite. People will say what on earth has this got to do with anything? If we're looking at 25 nanometer defects, which typically would be fatal at 100 nanometer design rules, looking for a 25 nanometer defect on a 300 mm wafer is equivalent to looking for a golf ball on the face of the earth. Not only that, because your detection equipment can't be at the wafer surface, it's equivalent to looking for that golf ball from about 100,000 miles out in space. That is a really big challenge for the industry.

In addition to that, we've got many fewer defects per wafer as we try to drive the yield up. In fact, we're moving to a phase over the next five years of mean number of wafers between defect, not the other way around. We literally are going to be hunting down needles in several haystacks. You can imagine that to chase down a defect on that sort of density level you've got to process literally an untold number of pixels.

Finally, we've got to data management. Let me show you some of the things I'm talking about in terms of that. We're going to have a huge data gathering activity down here along the process tools. There will be far too much data to be able to bring it up for human data analysis. Those are the red bars going through the people. We will not be able to handle people up there. We've got to fold this data directly into a yield management system which allows us to control where we're heading here.

An example of what we're trying to do is something we call spatial signature analysis. If you go in and look at a wafer, you pick out 100 pixels and it tells you there's a problem. Then the machine automatically converts that into a shape. From a shape it calculates an object. From the object positioning it calculates a signature. From that it will determine the classification of the defect which will allow the engineer to determine what to do. It reminds me in a way of an old IBM advertisement which I thought was rather good. It didn't last very long, but IBM said, "It's not just data, it's reality." That's what we're after. Data itself is of no use to us; it's reality that's important.

Let me just summarize where we're at. We have a huge number of manufacturing challenges coming at us in the future—300 mm, we've got copper and low-k. In fact, copper with oxide will probably start sooner than this. Low-k maybe a little later. We're going to be looking at new gate stacks sometime just early in the next century. Around 2005 we'll be ramping the next generation lithography. Ultimately, at about 2006 perhaps, we'll be looking at 450 mm wafers—really large wafers.

We need to do this to stay on our growth curve. We'll need help obviously from the 42 long marketing people. We're going to need help from our deep physicists who are going to solve some of these insuperable problems. We're going to need a lot of help from the EDA industry, and that's a lead in for Wally Rhines, who's going to follow me. He's going to save us in this design problem. Finally, we've got some challenges in the factory. If we didn't have challenges in this business, it wouldn't be fun. If it wouldn't be fun, we wouldn't all be sitting in this room today. I look forward to the next 10 years even though I think there are going to be some very, very significant non-incremental changes.

Q: Two questions. The first is: talk about the copper wiring program at SEMATECH and in the industry.

A:That could be a topic for a 45-minute talk alone. Basically, there's a very active program in the industry right now. There's been announcements made by IBM and Motorola and many others are following them. Those programs are largely directed towards copper on SLI2-based dielectrics, just simply a replacement of aluminum with copper. The work we're doing at SEMATECH is a little different than that. We're looking at combining copper and low-k dielectrics. That's a little bit more of a challenge and probably will be several years out in time. It's my expectation that we're looking at different sorts of equipment. We're looking at barrier layers as I mentioned, but I think copper will be here certainly within a year in preliminary production. We know that. Probably copper low-k a couple of years after that.

Q: When and where does silicon technology end in the future?

A: That's an important question for us. In fact, I'm not sure that silicon technology will end. It may be that even though we run out of the present productivity improvements and we can't continue shrinking the devices—and typically I certainly think it's good down to 15 nanometers. That's three or four more generations. Even when that happens, it's not clear to me that silicon will stop. It's a little bit like steel. There are many better metals in many ways than steel, but steel is still the most cost effective way, and I think silicon will remain the most cost effective approach as we go out. Even if we fall off the manufacturing curve somewhat, I'm convinced that silicon will continue to hang around as the material of choice. I haven't seen a material which I think will do better. It offers lots of flexibility and good deep physics like that. The problem is a 450 mm wafer would be a tremendous technical challenge. It's a very heavy material, a very brittle material, and it's a very expensive material. It's not clear to me, in fact, that it will be able to compete with silicon even if we can't make devices smaller than 15 nanometers.

Q: Any predictions for winning solutions, low-k, copper barrier, copper deposition?

A: I'm going to duck that with integrity. The low-k business will be an evolution and the various organic options, the various inorganic options, and eventually we'll end up with aerogels. My own personal favorite is going to be initially an inorganic solution. I think perhaps a fluorinated amorphous carbon may be the first step before we go to aerogels. Copper barriers are a number of different refractive metals that could be used. Again, that's something the companies consider to be very proprietary. At SEMATECH we're looking at several different refractory metals which have been talked about. As far as copper deposition is concerned, I'm sure that both the plating and CVD houses are hard at work. I've seen depositions from both of those technologies, and I don't think the PVD people are going to give up without a struggle too. They may very well be able to do the first couple of generations with PVD. I think it gets much more difficult as we head down to 100 nanometers and below.

Chapter 12: Intellectual Property: The Enabler of System-Level Integration

Dr. Walden C. Rhines
President and CEO
Mentor Graphics

Introduction: If a fab costs $10 billion, if we are able to build 10 of those per year, and that represents, in capital costs, about 20% of semiconductor revenue, those 10 FABs translate to a $500 billion semiconductor market.

Now that we know how to build all these 100 million transistors, Dr. Rhines is going to tell us how to hook them up. Dr. Rhines is President and CEO of Mentor Graphics. Prior to joining Mentor Graphics, he was Executive Vice President in charge of TI Semiconductor Group. During his more than 20 years at TI, he held various positions and was responsible for the development of the speech synthesizers in TI's first "Speak and Spell;" he was also instrumental in the development of the TMS 320 DSP family of processors. He was a Sematech board member and is currently on the boards of Cirrus Logic and Trident. Dr. Rhines has a Ph.D. from Stanford and an M.B.A. from Southern Methodist.

Dr. Rhines: Thank you. It is a good challenge to have — to keep those designers employed just like those process engineers. I do want to make an appeal, to all the marketing people, to make sure that more of that 50% of the world population, who haven't made a phone call yet, get a bigger share of those 2.7 trillion e-mails, because I can't read all these things. It's really getting tough.

Today, we're going to talk about the reusable design information. It's unfortunately referred to as intellectual property; unfortunately because that brings the image of lawyers to the mind's eye and, in fact, there are many lawyers involved in this concept. The reality of reusable design information is the one element that will make possible the utilization of the silicon that Mark just talked about.

We've talked about reuse for years — it's not a new topic. Everyone has thought for a long time that we should reuse what we've already designed, but it doesn't happen very much. I supervised design groups, back in the '70s, in which we would set up a project. It was going to run 24 months, and we would agree at the front-end that, this time, we're going to document the design information, and this time we're going to build a library, so that it can be reused in the next design, and yet at some moment during those 24 months someone would come to me and tell me that we're behind schedule. This is a tough design, what do you want — a well-documented library that can be reused, or the design completed on time? Being a pragmatist, I normally selected the latter, and usually I got neither one.

We're going to go forward — it's going to happen. It's going to happen for some very fundamental

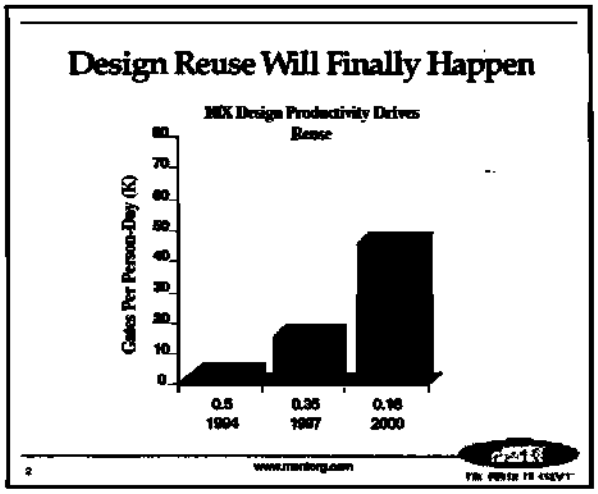

reasons. The basic driving force this time is simply the number of transistors or gates or bits available, and over the six-year period, mid-90s to early next decade, that's going to increase by 10x, simply by virtue of the capacity that's being built.

The actual design cycle time required for new products is not increasing. In fact, in industries like telecommunications, that formerly took years to develop a new system, the fact is that now they're being required to do wireless cell phones in six to nine months, simply to keep up with the product cycles. The number of designers per chip is not increasing appreciably, simply because of the difficulty of growing design teams in size and coordinating the activity.

If the time taken to design a chip doesn't increase, and if the number of designers doesn't increase appreciably, then the net result is that productivity has to increase, and the gates per person they design, or however you want to measure it, is going to go up in order of magnitude over a six-year period, and it's got to come from somewhere.

You might argue the designers should work harder; they're a bunch of prima donnas and simply need to speed up what they're doing. I don't think that will do it. You might argue the EDA vendors ought to get some tools that can fix this problem, but they've got their own problems, with how design tools have to handle all these deep submicron effects and do power analysis and reliability and things, so I don't think you're going to get it there. There's really only one place to easily get it, and that's from design reuse, using what's already been designed before.

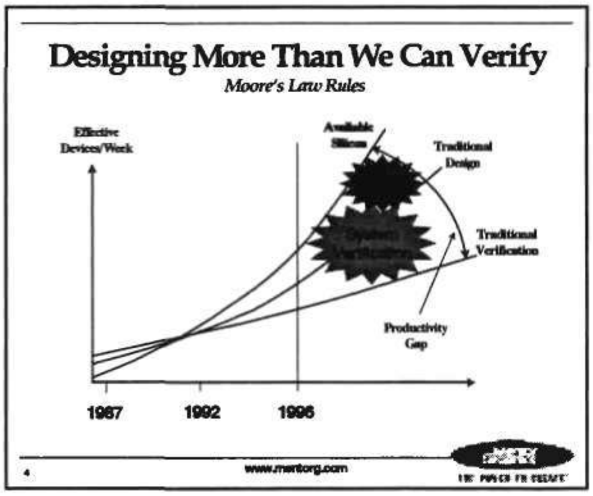

You could argue that even if the transistors go up 10x, I don't have to use all of them, and not every chip has to be to the limit, but this is one of these phenomena that falls into the category of "If you build it, they will come." The designers will use all of the available silicon, because they always have.

When I went to Michael Slater's microprocessor conferences, he'd give out notebooks with all the leading edge chips mounted on the cover. I always found it very interesting that they were all the same size and shape. Why is that? Obviously, a LAN controller isn't the same as a host microprocessor, and isn't the same as a DSP. The fact is that those chips were that size because that's as big as we make the sockets. If we made them bigger, the chips would be bigger. People will put the transistors in and take advantage of the additional real estate.

There's some other reasons why you would want to reuse what you've designed again and again. One reason is that much of the design today adds no potential value. It's the same companies designing the same functions again and again throughout the world — standard protocol blocks, other pieces of design information. The fact is that they don't add any value. If your design team wants to do another PCI bus controller, what's the advantage you're going to get? You can't run a 66 megahertz PCI bus at 67 megahertz and have a differentiated advantage. You're much better off to use the one everyone else uses and spend your time doing things that add value.

There's another reason. There's the deep submicron physical effects. When you get a block of physical design put in place, and it performs to a given level, and you've analyzed it for power and noise and reliability, in the future generations, the rule is going to be, "Don't fool with it; reuse the block." There's too much invested in getting it to work the way it works today.

At .35 micron we saw an interesting phenomenon, saw a little bit at half micron, but usually the design teams of the world want to stay away from the new technology as much as possible, because they've learned back at one micron or whenever they got their experience, a new design would be committed to a customer based upon a new process. As you got closer and closer to release of the process, then the design rules started to change. We can't make that spacing; we can't get that contact spacing the way we spec'd it, and so on.

Suddenly, you're in turmoil and the process comes out and doesn't yield and there's no wafers and you're late to market, so designers learn not to be the first one on a new process.

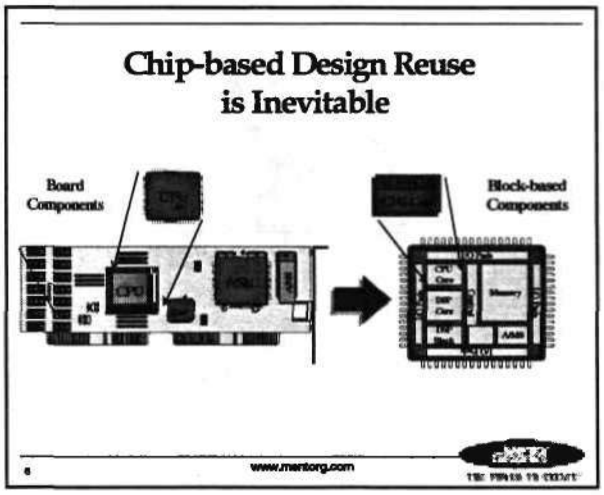

At .35 microns, when the capacity came on-line, it amazed the management of all the semiconductor companies because, in fact, it yielded. There was capacity. The big difference was statistical process control had made its way to the factory. We now could characterize the process capability of individual steps, and then the product of those steps, instead of being zero, was engineered to have a process yield, and we had wafers and we had no designs. Suddenly, we had to head off into a shrink the designs we had and run them, and not take advantage of the new capabilities because we hadn't gotten anything designed, and semiconductor companies started re-assessing their whole design process. This caused a boom in the electronic design automation industry.

Looking forward, those companies are now focused on two things that have to occur. One is that design reuse has to happen in a big way, both within companies and across companies. Two, we can design more than we can verify, so system verification tools become the key. Many companies said, "I've got to get so I can design everything and reuse it across the company, because I want to be self-sufficient and design my own chips." That's not enough. I can remember that we had years at Texas Instruments when we thought we could do everything better than everyone else, so we continued to enter every possible market segment, developing products, memory to microprocessors and IO.

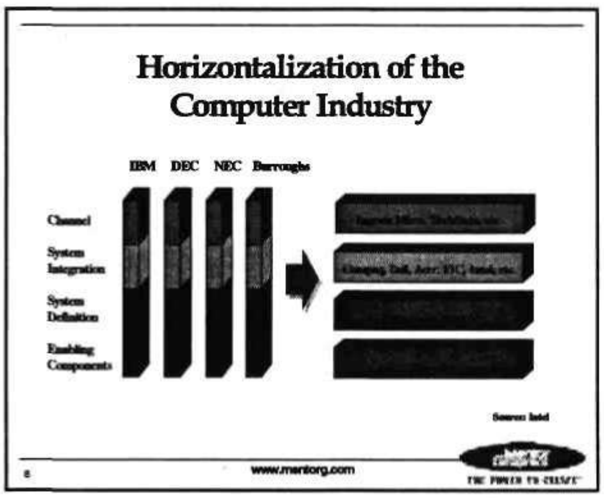

We always hoped that some day we would see one of these customers that would bring us a printed circuit board that looked like this, the wonderfulness of it – CPU's, RAMs, all from a single vendor. It can't happen. It couldn't happen. It never could have happened. The fact is no one company can be best at everything. If you're not going to be best at everything, then you need to use the things from people who are best at what they do. That indeed is what has to occur in the reusable IP business.

The chips of the future will, as Mark noted, increasingly look like the printed circuit boards of the past. You're already seeing many embedded core DSPs and embedded micro-controllers and microprocessors, but the steps ahead are even more dramatic in the reuse of functional blocks. Those of you who were around for the TTL era or the RTL or DTL eras (and I see fewer of you at every meeting), what happened was that people had to decide what voltages should we work at and what should the pin-outs be, because companies had to sell integrated circuits into system companies that used ICs from different sources. Those things were agreed upon fairly rapidly because they had to be. It was needed for survival. The same thing is happening here — defining standards associated with the reuse of intellectual properties.

One of the groups that's doing that is the VSI (Virtual Socket Interface) Alliance. Mentor and other EDA companies were founders of this group, which now has over 150 members, that tries to drive general standards. Rapid is an organization that is made up of independent suppliers of intellectual property who principally are worried about the common rules for marketing and selling their intellectual property. One of the most recent ones was one where Synopsis and Mentor Graphics joined forces, because it became clear that VSI was going to give high level specs, but people wanted a process, a methodology, a how-to, "what do I do today, so that what I did in design will be reusable tomorrow," and they needed a cookbook, a "reuse methodology" manual. We developed that cookbook in the consulting services to help people reuse or develop reusable IP.

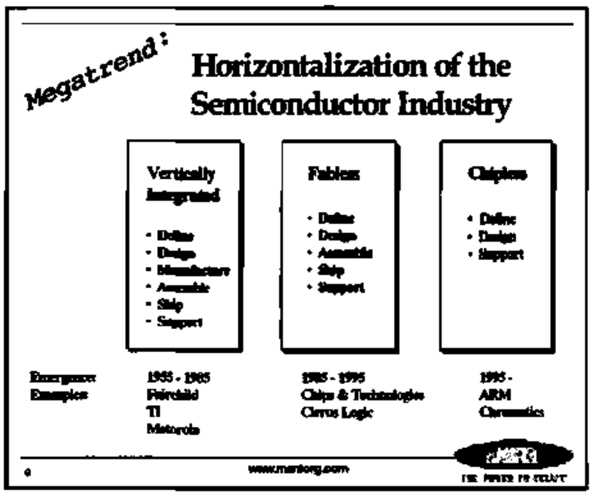

There is another phenomenon going on. We saw, over the last 15 years (it was pretty obvious to everyone), the horizontalization of the computer industry. The individual segments of the industry became big enough that it justified a horizontal attack on the markets, so 15 years ago most of the disk drives were built by IBM, DEC, Fujitsu. Today most of the disk drives are built by Seagate, Quantum, Western Digital, whomever. Each segment has gone horizontal. Here's a slide from Intel, so you can see Intel puts itself in each of the horizontal segments. It's a bit of a stretch, but the basic principle is there. The operating system is separate from the hardware.

What's happening that's an equivalent in the semiconductor industry? Basically, the model is shifting. Traditionally, vertically-integrated companies defined the product; they designed it, developed it, manufactured it, assembled it, shipped it, supported it, and that was necessary because the process was tightly tied to the design. That will always be the case for differentiated processes, like linear and like power, even like DRAMs. This model will probably stay around forever, but an increasing percent of the market will go to what we've seen in the 80s: the fabless model, where you drop out that line that says manufacturer, and you define the product, and then you ship it to a customer and support it. This is driven by companies like Chips & Technologies.

What's the next wave? It's very clear. It's the chipless company. It's the company like ARM or DSP Group or Chromatics, that simply provides design information, and that type of company is becoming the fastest growing segment of the new product business world today, as new companies are formed every week to attack that chipless market.

It's not hard to see from the data, the percentage share of wafers built by foundries, versus wafers built by captive suppliers, is increasing, doubling about every six or seven years. There's data to support the trend, so you have to ask what should a vertically-integrated semiconductor company do? I would offer this advice.

I believe it's a big opportunity for the vertically-integrated companies, but you need to be sure that you're not subsidizing inefficient manufacturing with good intellectual property. If your manufacturing costs are bad, form a relationship with a foundry designed to the same design rules, and shift your capacity back and forth as needed. Don't subsidize one with the other, and the same thing on the other side. If you've got great manufacturing costs, go out and acquire designs, acquire intellectual property. Every company needs specialized IP that is shared within the company, and reused in the company, but not made available elsewhere. At the same time every company needs to get access to all the standard blocks of IP or design information from third parties, and to put design flows in place to make it usable and reusable in your own design flows.

Lastly, why not take advantage of a key strength. That is, if you've got wafer fabs, emphasize the part of the business that utilizes differentiated processes where it's relevant, because no one is going to design IP that is easily portable for non-standard processes, things other than standard digital CMOS, .5, .35, .25 micron.

Where is all this intellectual property going to come from, that everyone's going to reuse? Mostly from where it's always come, from semiconductor companies and from systems companies, but the new entry into this age is the independent provider of IP, and the new support challenge is for the EDA vendors to provide the tools and the infrastructure to support reusability, for designs created by semiconductor and systems companies, as well as by third parties.

As we look at the roles ahead, I would expect that the semiconductor companies and systems companies will continue to be the dominant creators of IP, but increasingly there will be third parties that create intellectual property that's used. They will own it; they can market it where it's appropriate, but that won't be the primary thing. The primary thing is using it and doing system integration to develop system chips. That's where the greatest value is added, by semiconductor companies and by systems companies, and that will be their primary role. Of course, the semiconductor companies will shift the production units and so forth. There are secondary roles for all of these, but the third parties mostly will concentrate on creation.

What about the EDA suppliers? There are actually four models that have been discussed at EDA industry meetings. Most of them support the EDA company providing the tools to make reuse possible, and the consulting services and the things needed so that you could verify and use IP. There are several different approaches in the industry for the other roles.

The one Mentor pursued is the distribution and support role, because support is such a key challenge. That is working directly with the designer of chips and providing whatever standard building blocks that designer needs. This means that we must have the expertise of whoever created the block, and we can do that by working with third parties, or we can do it by acquiring IP or whatever, but the main focus is on the system designer helping designers design more with less.

Another one, the system integration model, is one that Cadence has popularized; in this one you provide the actual design services and do the actual chip design, or in some cases handle outsourcing of the design process, so the focus is on use and system integration and, as Joe Costello, the former President of Cadence, and I have discussed at meetings — Cadence becomes a customer for Mentor in that model and, in fact, has become a customer. That has quite a bit of potential too, but it involves EDA companies taking over the design process, so Mentor has shied away from it, but it certainly has been successful financially for Cadence.

The third one, Synopsis, has driven portability tools, their cell-based array architecture and so forth.

Lastly, there is one represented by quite a few companies, Aspect and Mentor and Compass and others, and that is simply doing the physical libraries. Most of the ASIC libraries in the world today are designed by EDA companies, simply because there's such a great economy of scale in reuse of libraries.

What should system companies and semiconductor companies do to take advantage of what's going on? Every system company or every system designer ought to have a pool of proprietary IP. They ought to have things they get from third parties; they ought to have their ASIC libraries that they get from the ASIC companies that have their own building blocks. A company like Mentor is in the business of providing virtual libraries for the standard building block functions, and the reason this all works well is that the EDA companies are typically compensated per design, and the semiconductor companies are compensated per unit, so there is a synergism. The lower volume designs tend to get much more support from EDA companies.

The other part is that many system companies and semiconductor companies need to get their own IP taken to customers, or made public, simply to cause standardization, and they need someone to help drive that. That's one thing EDA companies can do through the libraries they provide. In our case, it's a virtual library we install at the design company, and they use it to do their designs, being able to call up those building blocks as needed.

The IP itself is plentiful. The problem is it's not reusable. Reusable intellectual property is scarce because it's not verifiable, it's not testable, it's not supportable in multiple design environments. At semiconductor companies there are tons of designs done, but reusing them is the challenge.

There are quite a few reasons behind that, including the fact that they weren't designed to be that way, but supportable IP is very expensive. For us, over the last 10 years, RTL level design information source level — we spend about the same money every 18 months to support it as we did to develop it. Jean-Louis Beret at LSI Logic says it sometimes takes up to 400 views to support a piece of intellectual property and make it reusable. It's a big task and the people who design it really don't want to take the time to translate VHDL to Varilog or do other functions like that.

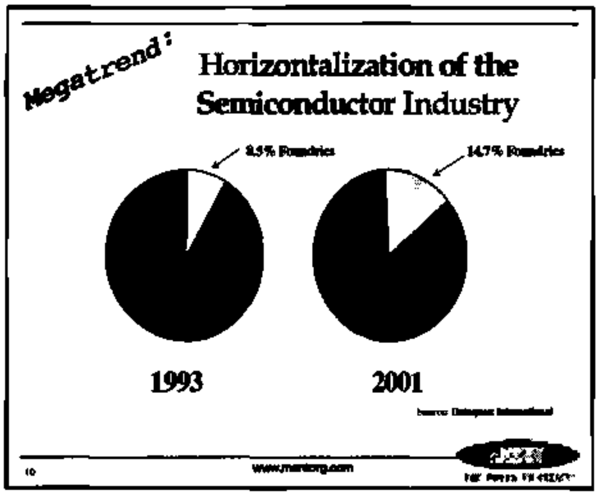

It's inevitable that reusable IP will be supported by EDA companies, because they can sign long-term contracts, because it's simply like the model for tools, where they provide worldwide coverage and, if necessary, on-site applications support and hotlines. I've talked to many designers, and they don't want to simply pick it off the Internet and use it, because they don't know if it's stolen, they don't know what happens when the next generation comes along and they need support, or they need it ported to a different process, or they get in design trouble and they need some help and they need some leverage against someone. EDA companies are ideal. You buy tools from them, so why not pressure them to support the IP? In addition, if you really get in trouble, you need a company that you can sue. Who other than an EDA company? EDA companies have money, they've got contracts, it's someone you can go after and get the support you need.

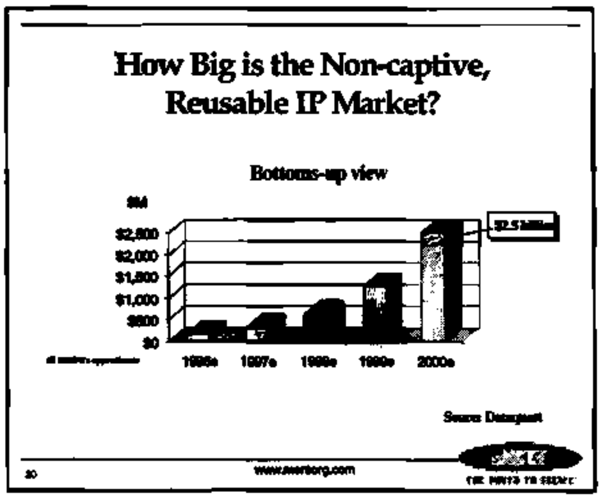

There's also the problem of verification. The EDA industry needs a major focus on this. In our case, we've put most of our resources on just the verification tasks of simulation and design-for-test, and hardware/software co-verification, and all of the physical verification required. You will never find a designer who says, "I have simulated it enough, I'm done, take it away, it will work first pass." You will always need to do more system verification, and that's the challenge, even after reuses, that we have to face.

Lastly, how big is the market for this reusable IP? The numbers are all over the map, but this one uses a Dataquest curve with a different starting point.

Whatever it is, it says it's very big, billions of dollars, but not that big unless you look at it the other way and say how much IP is there really around or what could it be. Here we do an estimate where we say: "If I simply look at what a commodity integrated circuit sells for, and compare it to one that has substantial design value in it, like a microprocessor?" Even if you take out Intel's 90% and simply look at the 55% that normal people run on their microprocessors, you still find out that there's about a 15-point gross margin differential between commodity IP, where there is no value given for the design, and the average chip in the industry.

That means that today there's about $25 billion worth of IP being sold. It's simply being sold tightly coupled to the silicon. Over the next 10 years or so, an increasing share of that will be decoupled, and this becomes a very significant market or, if you're buying it, a very significant investment, or something that companies should consider taking advantage of; that is, to sell what they've already designed.

To summarize, I believe that reusable IP, and the ability to separate design information from the silicon, will cause an acceleration in innovation in the industry because it will involve thousands of designers, who were not part of the process before,

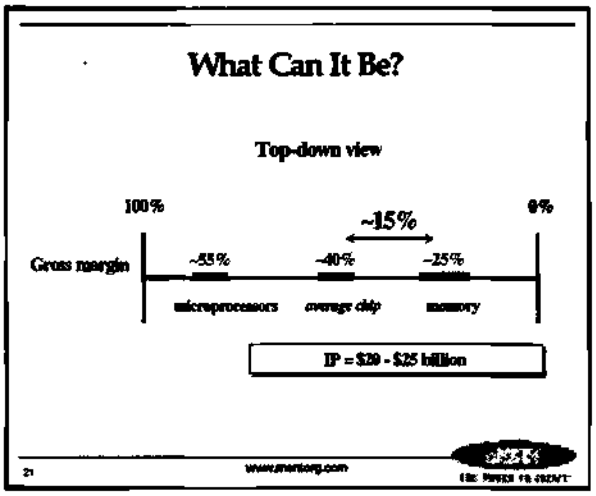

and who are able to independently create intellectual property. To do that they'll need that IP to be supported to customers all over the world. They'll need people who can add all of the design views they need, to put in testability or take out testability or support a particular new tool, and that's what the EDA industry will do.

As we move to a greater mix of designers, you'll see a greater leverage of design innovation. It's not unlikely that for something like Universal Serial Bus (USB), you're going to want to buy the copy that has already been used in 50 or 60 designs and is debugged, instead of having your designers tell you: "I can design it better, why don't you let me do my own, and we'll figure it out from there?" The task is making that impossible.

You can't redesign million-gate blocks in your own company and still stay competitive, when you need to spend your time on the differentiated blocks. Spending more time adding value, providing more specialization... all this adds up to a much more productive industry, an industry that can create system chips that have the best of everything available, and achieve the end result in the most efficient way. It's going to be a major discontinuity and will have a large impact on our total industry.

Q: How do you view the required patent licenses in light of reusable IP? When I buy a reused cell from Mentor, do the patent rights come with it?

A: The issue is with regard to patent rights — first, let me say there are more issues than simply patent rights. The big issue — or one of the big issues — is copyrights. All the IP that we do and other reputable suppliers do, if it's developed from a data sheet, if it's done in a clean room, there are inspections by lawyers regularly. You'll also take what's on the data sheet and create VHDL or Varilog that exactly reproduces it and is free of any copyright violation. That's step number one.

If you get IP, you need to be sure it's defendable against copyrights. On patents, no company that I know — certainly no rational company that I know — provides indemnification for IP. It's impossible. There's no way you can indemnify people for patents, and I might point out that simply because you take a reusable piece of IP from a Mentor Graphics or some other supplier, they are unaware what patents might affect it in the future, what might be issued that isn't existing, and even if you use an alternative method and say, "If you won't indemnify it, I'll design it myself?" You've still got the patent problem.

Basically, the suppliers of reusable IP are not solving the patent problems except where we have relationships with companies who say, "We'll simply give you a package deal, whenever you sell this please charge $10 more and we'll give you a patent license for our patents." That happens in a few cases. Actually, with many companies we perform a service. We take the low volume customers and support them so they'll stay with the company's architecture, and in those cases the companies like to provide a package deal. That's fairly rare.

Q: Will EDA companies, via acquisition and/or alliances with foundries, also become a supplier of wafers and ICs, so as to become a one-stop enabler of system level integration?

A: The answer, of course, is absolutely not. Some of the EDA companies talked about that. Mentor stepped forth very early and said "we will never, ever own silicon." The reason some of the EDA companies talked about it is those people have never managed semiconductor businesses and they think it's real easy to own silicon. I personally have written off more than $1 billion worth of silicon, and I don't think that owning chips is any fun, and I don't think it's what the EDA industry does well. Our focus as an industry has to be on helping designers at other companies design better chips with better tools and the building blocks they need, and if EDA companies try to get into the silicon business, I predict they will fail and fail very badly.

Q: Who will be the primary group to support the customer when the IP chip made with the third party or EDA IP doesn't work?

A: In every case of IP that we re-market from a third party company, we set up an exclusive agreement where we get full-time dedicated support from the people who created it, so that you can always get back to the prime source and deal with it. In most cases we've got over 130 cores in our library, and in most of those cases we have in-house expertise — but not always. Whatever it is, you need the direct expertise of the creators to support customers, and simply being a distributor in the arm's length sense is not an adequate way to handle the business.

Q: How important/feasible is design-for-test at the core level?

A: It's not only feasible, it's necessary. You used to be able to get by without designing in testability. Now it's not simply an issue of getting more reliable products; it's simply being able to debug the ones you have. Design-for-test needs to start at the beginning of a design. You need to use available tools to assess the testability and then design it in with complete testability; not simply scanned but looking at built-in self test,or you're not going to be able to afford the testers, you're not going to be able to debug the parts, and you're not going to have efficient trade-off of manufacturing costs versus capability.

Chapter 13: 2010: A Vision of the Future

Gordon Bell
Senior Researcher
Telepresence
Microsoft Corporation

Gordon Bell will now give us some visions on how those highly integrated designs will be used. Before I introduce Gordon, I will mention that, at the breakout session this afternoon, the Dataquest people will also talk about intellectual property issues, so if you're interested in pursuing what Wally has talked about, you might come to that session.

Mr. Bell is a Senior Researcher in the Telepresence Group at Microsoft. He is also a computer industry consultant at large, working at computer-related projects, especially parallel processing. Mr. Bell spent 22 years at Digital Equipment Corporation as a Vice President of Research and Development. He was the architect of various mini- and time-sharing computers and led the development of DEC's VAX computer environment. He was also Professor of Computer Science and Electrical Engineering at Carnegie-Mellon University. He's the author of several books, many papers, serves on the boards of many companies, and has received numerous awards, including the 1991 National Medal of Technology. Mr. Bell is also a Fulbright Scholar to Australia. Please welcome Gordon Bell.

Gordon Bell: I'm going to give my version of why computers are like they are, and then look at that from a historical point of view and see just what we can glean. Then we will look ahead, basing our views on Moore's Law. If we look back at why things are the way they are, we discover that they came, first off, from two great inventions, the computer and the integrated circuit (IC), which those of us in this room will agree, are the two great ones. Maybe you won't agree that the computer is the greatest one, but certainly one of the great inventions of this century.

In addition to that, there's a quest for cyberization and I'll tell you what that is. I want to talk about why classes form and then the markets that make that whole thing go. Now we will look ahead into some of the technology and what's going to be inevitable.

One of the two great inventions is the computer, which is a 1946 invention. The reason why

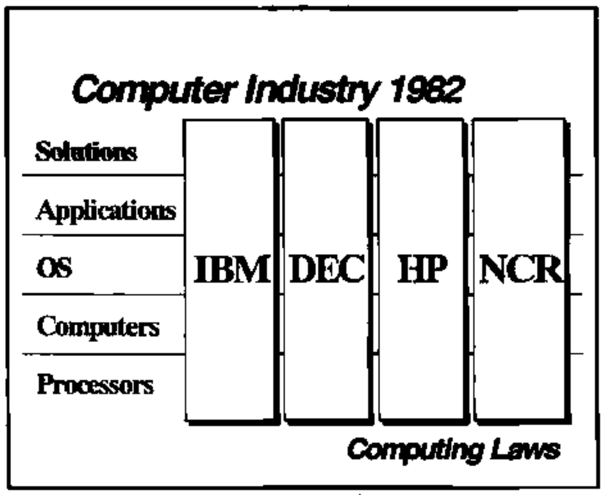

computers are so important is that they fundamentally are capable of supplementing, and in many cases, substituting for or supplanting other information processors, including humans. There is a sub-theory that computers are so powerful, in themselves, that they recurse. One builds computers on top of other computers all the time. That's why Sun wants to build Java on top of our PC's. Then there will be a JavaPlus or a HotJava or a SuperJava on top of that, so that we have these layers and layers of architecture taking all of these processing cycles and interpreting everything. We'll just keep it up and keep the users happy interpreting all of these machines that have to run on top of other machines.

The other reason is that computers are built in a well-defined way, on a component structure, in terms of processing memories, having the ability to switch and the ability to transduce. They get bits from the physical world into the computing environment. The second great invention, which occurred at the same time, was the transistor and, subsequently, the IC.

The other thing that's happening is that everything that is cyberizable will go into cyberspace. That's really driving the use of computing. It's the cyberization of the world. I view this as fractal-like, in that we can think of this as a series of universal networks into the worldnet, onto the continents and down into cars and homes. Everything will be connected, and have an IP address, all the way down to body-area nets. Even our dogs will have two or three IP addresses on them, to see where they are, to call them and know their GPS coordinates.

What is cyberization? Basically, it is the encoding of all of the information into this one universal network. It's really a coupling of all information. To information processors, it's pure bits like printed matter. It's bit tokens, like money. It's the state of everything—places, things; just knowing where everything is. It includes the state of physical networks, knowing everything about highways and where cars are and things like that.

Did we come up with this? Well, in 1945, Bush made a statement that there will always be plenty of things to compute or to get information from. With millions, and now billions, of people doing complicated things, we've got a potential market here of five times 10^9. I call it five gigapeople. That's a nice number. Always divide by that when you're thinking of all these, in terms of how much is available. He described a thing called MEMEX, which was going to store everyone's books, records and information, and that could be consulted with tremendous speed and flexibility.

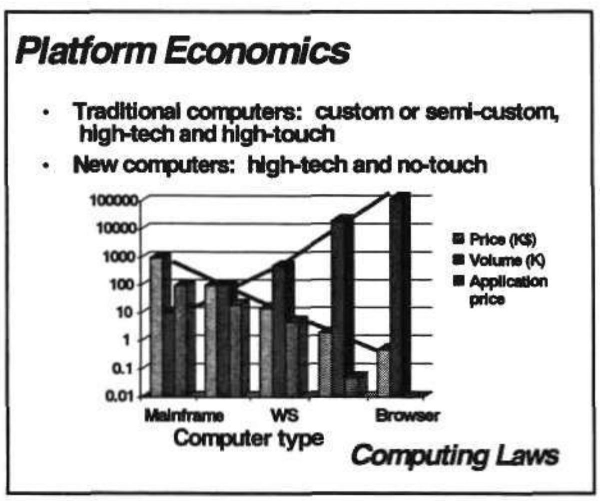

He also described many other things, one of which was a matchbox-sized, $.05 encyclopedia which, in fact, was available about five years ago in the form of the CD. He talked about speech to text. That's still to come. I got out of the speech business in the 1960s, having declared it a 20-year problem that I didn't want to work on for 20 years before seeing any results. It turns out I was wrong on that, as usual. It turns out it was a 40-year problem. The other thing he mentioned concerned a head-mounted camera, and I just came back from Japan with a couple of cameras that I've succeeded in attaching to various computers, one of which is certainly head-mountable. It doesn't have a radio or IP address on it now, but it will in the future.

Let's look at the collection of laws that makes all of this possible. We've got the two inventions. We've got this forest to fiberize the world, to encode everything, put it into a network form and give it an IP address. What we have to look for now are transducers to make this possible. Moore's First Law really governs what's happening. You all know that law, so I won't go into that. You should remember that, with this type of exponential growth, the past doesn't matter. So, 10X here, 10X there, and soon you've got something real. PC platforms have declined faster, due to volume, learning curves and the demand curves.

Also, computer components must all evolve at the same rate. Amdahl's Law says that for every instruction per second you need a byte of memory and a bit a second of I/O. Processor speed has evolved at 60%, if you ignore the fact that you can't expect to get that 60%. People haven't learned about processor speed at the application level, as opposed to the benchmark level. Storage is evolving at 60%, and it is substantially more important than the semiconductor part of all of this.

Last year, the computer museum, of which I'm one of the founders, put on a show, with Intel, on inventions of the PC. We had the first floppy drive and the first hard drive there, but asked ourselves what the floppy and hard drives had to do with the PC. That shows a funny centric view. I just got a small PC, 700 Mbytes and it turns out that it's already full. Wide Area Networks evolve at the same speed; and then there are the depressing Local Area Networks.

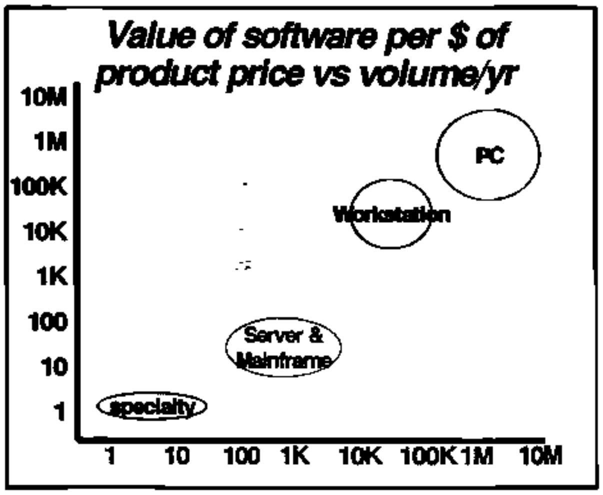

Local Area Networks had been dormant for awhile, but now are evolving at a high data rate. Grove's Law is saying that the telephone service asked to thwart all the gains everywhere else. Over the last 15 years, telephone service growth has evolved at only 14% and certainly, if there's a big, black cloud on the horizon, it's the lack of telephone service to our homes.

A law that I have is actually a corollary from Moore's Law; however, it's a law that governs computer classes. The computer platforms emerge based on chip density. They require three factors: the platform itself, the form factor and cost. They require a unique network and they require some form of cyberization, the connection of the computer to something, whether it's a car, a human or anything else. Then applications follow that particular class and subsequently each class becomes a vertically disintegrated industry, that we saw before, based on the hardware and software standards. I observed that

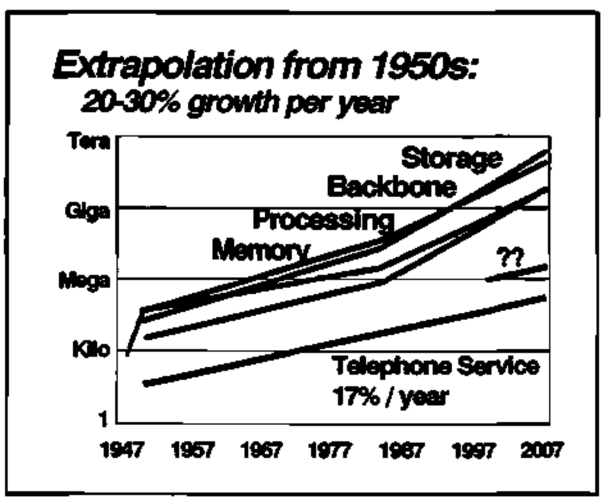

computer disintegration in the early 80s, when I saw the UNIX people buying commodity UNIX as all the UNIX vendors started up.

Essentially, this explains how computers evolved. How these classes formed— mainframes at 10^6, 10^7 price range, mini's that came in at the $100,000 level, then their friends, the workstations, came in and PC's, more importantly, came in at the $1,000 to $10,000 level. So, these new things you can take more as law in two ways. One is that you can go the way of the Intel strategy, which is to put more and more on a chip and try to keep the price constant for those chips. The other thing is to take the system, integrate all of that on a single chip, and then take the resulting cost-reduced chip form.

With the 256 Mbit chip and the 32 Mbyte chip coming out, why not simply put a processor and some I/O on there and have a single chip system. You know what the cost is, so you decide what the price is, based on what operations are done. With this theory of platform interface, cyberization and network, we can observe how computers have evolved.

We have the computer form, which was a mainframe in the very beginning. It was a tube core/drum tape/batch operating system. The interface was either direct or batch. It had no network. With the minicomputer and time-sharing, we built those out of SSI-MSI. Disks were important in their time-sharing operating systems. There was a terminal command interface, and we were content with using the plain old telephone service for the interaction there.

With the PC and the workstation, the micro, the floppy, the disk, the bitmap and all of that was critical. The thing that really made this grow was the WIMP interface, the coupling of this with people. That is Windows, Icons, Mouse and Pull-down menus. Then there is the LAN that connected all of those computers together. Now we're looking at another generation forming, which is the web browser, telecomputers which are telephones and computers connected together that access the web, and the TV computers. That convergence is already pretty clear. That form is really PC's and scalable servers. The interface there is really the Web and HTML and, of course, the Internet is what's driving that. We've got the Internet here to save us from what would have been a period of time with nothing really happening.

National Semiconductor Technology Roadmap (size)

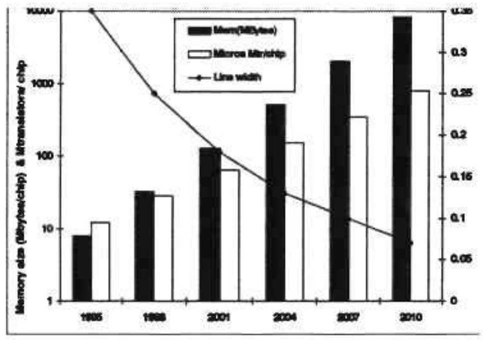

Computing Laws

With that theory of how things form, you can look at nine tiers of computer prices. If you look at $1 chips, those are the embeddables. Things like chips that play in greeting cards. Wrist watch and wallet computers and money. Portables in the $1,000 range. All we have to do is get rid of these LCD's, which we can easily do by having LCD glasses. That's going to change a lot of the demand for a machine when we get low cost displays. There's no reason for having to carry a 13-pound monster around if one could get rid of this display problem. Personal computers, anywhere from $1,000 to $10,000. Departmental computers, the old minicomputers that sit in the closet, and that Sun tried to knock off with the workstation and—now it's saving them. That's at the $100,000 to $1 million level. Site computers in the glass houses, regional computers and now the other thing is that we've broken the $100 million barrier for computers. It used to be that supercomputers only cost $30 million but now the government pays as much as $100 million for that. It has its three centers in Los Alamos, Livermore and Sandia. That was done by lashing about 9,000 PC's together which is now computing at over a terabit level.

We've already heard about how this disintegrated industry was originally an integrated industry where computer companies made everything from processors and chips to solutions. I heard that AT&T, at one point, had its own beach for sand. What we see today is an industry that is a completely disintegrated level of integration, and all of that has happened by standards. All of these levels can be done by itself. This has to be updated a little bit. I need another layer there that says we got the Justice Department in there working on the architecture of the operating system.

National Storage Technology Roadmap (size, density, speed)

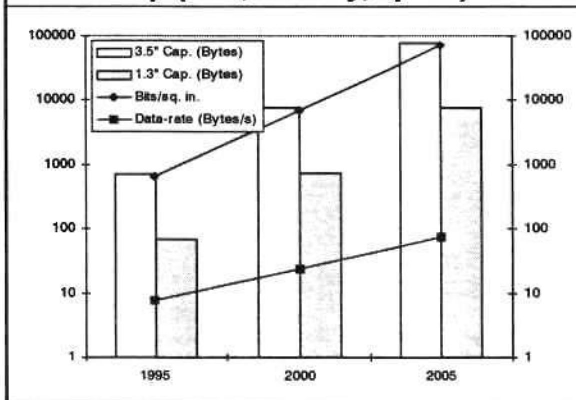

There's some good news. I just got this Porsche last year and it's got these round dials, and I wanted a color interface and a different steering wheel. So next year I suppose I can trade that in and not have any steering wheel or any dials and then buy the dials and interface for all instrumentation from somebody else. Of course, not from the same vendor. So, I'll have the option of having a totally different interface. It makes no sense at all. I contacted the Justice Department to see how you get this new interface. I really want a color panel, not this crazy old analog system. That's all analog, I don't like that, anyway.

The economics is the other factor that drives everything in that the demand, which is — as the price declines, we get a doubling of demand. We have the learning curve that says every time the cumulative volume decline increases by 2X, you get a unit cost. Bill Gate's Law defines the economics of software; it says you have to do things at a large scale. Nathan Myhrvold's Law defines the virtuous circle. Metcalfe's Law defines the value of a network that allows us to sustain the market. These all work.

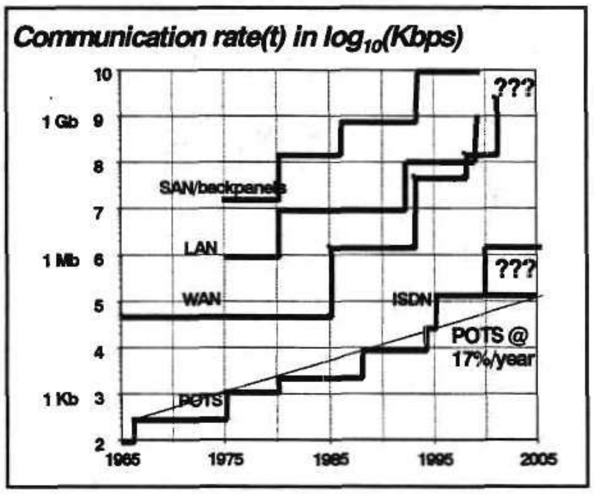

Bill's Law says the price of things is really related to the unit cost. Bill Joy initially said you can't afford to write software for under 100,000 platforms. That is, if you want a $1,000 price. That all comes out of marking up the engineering expense and looking at the various lines. Bill Gate's Law is that you can't afford to write software for under a million platforms. I'd say that's probably 5 to 10 million platforms right now.

You've got a fixed engineering cost or you've got a fixed cost and you divide by this number of units. If we look at the marginal cost of distribution, that's gone from a cost of $1 or so for CDs, and all the cardboard and air that's shipped with these boxes, to zero or near zero, as you distribute software on the Internet. We see what happens when you look at UNIX versus NT; a massive economy of scale, just because of the units. If you look at how the price of Oracle software has come down, this gives you some clue as to why Larry Ellison wants to get rid of PC's, because for some reason now the SQL-Server that used to cost $100,000 on a Sun box suddenly costs $6,000 on a PC. That will eventually will be shipped with NT and there will, of course, be another lawsuit.

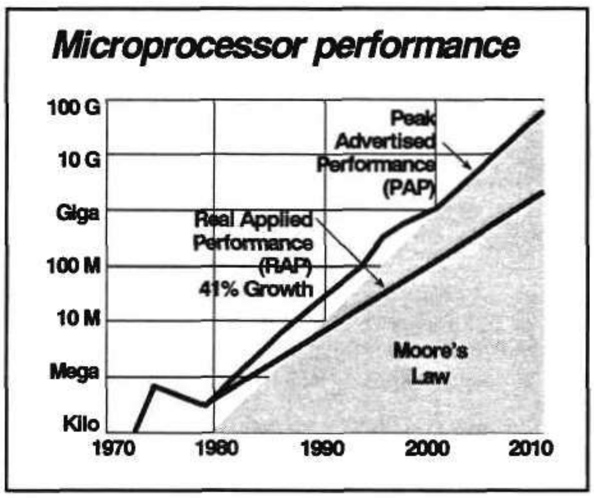

It's unfair to ship a SQL-Server to a non-database person. It's just a file system. Do you want a file system included in your computer or not? I happen to want one, and I'd like to have a relational database. That should be the standard. It also explains why there are no spreadsheets or presentation packages on UNIX, VMS, MDS or those other machines. You just can't afford to design for that.

Nathan Myhrvold the Chief Technical Officer of Microsoft, describes his laws of software that makes it all happen. He's right, software is a gas and it expands to fill whatever container that it's in. Software then grows until it is limited by Moore's Law, the amount of memory. That's why we can put so many features in. Software growth than makes Moore's Law possible because, in fact, if you didn't have it, you wouldn't have to have these memories and there wouldn't be other applications to fill them. Of course, software is really only limited by human ambition and expectation. I would add to that our ability to cyberize, to encode things into the computer.

That whole thing is really a combined vicious circle that starts with an innovation, which presumably has some utility and value; otherwise, people don't buy it. That leads to volume, and volume leads to competitors, which then goes back to innovation. All of this is not possible without standards.

Standards are the things that make that whole circle work. If it weren't for them, everybody would be able to operate in their own stratosphere, as the UNIX vendors have been, by having their own little standard cycles. However, they are, in fact, being driven back to a single UNIX standard. That will eventually happen as the PC standards drive them out of business or they finally adopt a single standard. We are seeing consolidation in the UNIX space and we are seeing people leaving the microprocessor. We see HP leaving. DEC will be next, followed by MIPS, with Sparc and IBM as the last to go in terms of the standard for high volume manufacturers.

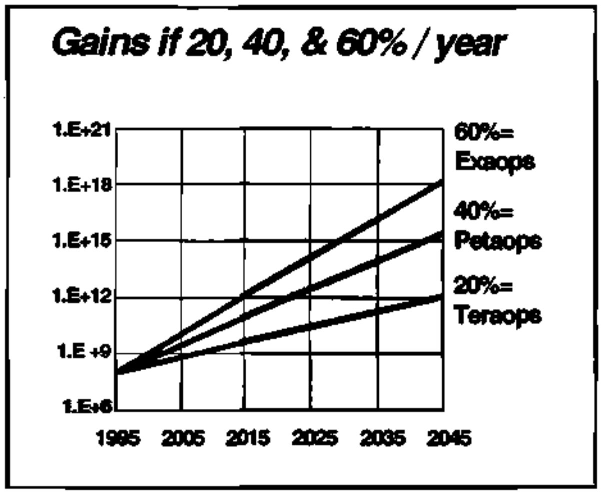

The final thing is Metcalfe's Law, which says that utility of a network grows as the number of users squared. The utility to an individual is the number of other people it connects to, but it's actually the sum of all users, and that's how you get that square term in there. This indicates why you've got to have everybody on the Internet, or why you want everybody on a single telephone system; because it grows exponentially.

Let's look at the hardware technology that's making it all happen. We've got many factors. Number one is that we get more. I was supposed to talk long-term, I'm sorry, this is near-end. There is a 256 Mbit chip that will happen in the next few years. The 32 Mbyte computers and 80 Mbyte computers will all shrink down to single chip systems. LSI Logic has proclaimed itself the system-on-a-chip company, with a large number of gates. Megabit bandwidth by then will be as easy to get as ISDN is today. That's a fairly small number of users. There are about a million ISDN subscribers in the U.S. today. Incidentally, there are 3 million subscribers in Japan. What we've got is a collection of networks that are all forming. These new networks, like the pager networks, the phone FAX, and all of these have to be connected together to form this cyberspace that I talk about. This is enabled by having IPs on everything which is really critical.

If we go back 50 years or so to see what the trends were, we've been evolving rapidly in the last 10 or 15 years. We've been on a Moore's Law curve for processing memory backbone and storage. Those trunk lines growing about 60% a year, contrasted with telephony growth, which has been about 14% a year.

The question is, will that get up into the megabit level, which is what's critical by this time. Semiconductors, you know what those are. The densities are increasing. The communications disaster. Processor performance is going up. If you run these out 50 years, you've got these curves. I've added another set of curves: these are processing and memory curves. A factor of ten up from that is mass storage, which is the critical factor for the applications that we're looking at. New overtakes old and comes in and wipes out bipolar, and it also wiped out mainframes and supercomputers. Things get cheaper. These were curves I did in 1975. I didn't believe those curves. We just couldn't comprehend those kinds of change. New and cheaper

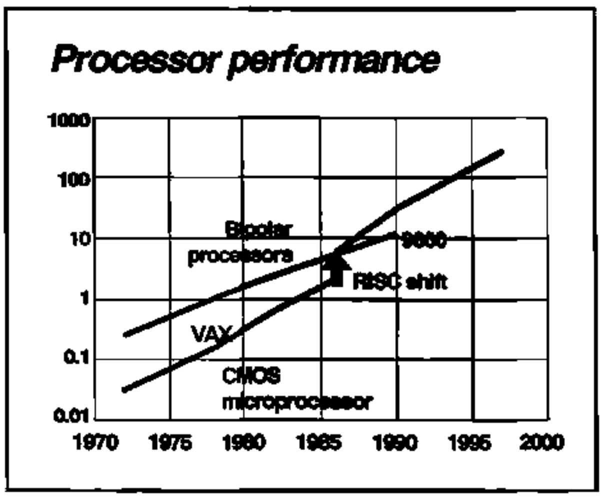

always wins, if it weren't for the Law of Inertia. The Law of Inertia says that data and programs sustain the platforms. That's why the mainframe will still be in operation in another 100 years. It will probably be working on the 2100 year problem at that point.

You have to understand the goals of the hardware suppliers are uniqueness and to differentiate and lock in. The goal of the software vendors is to differentiate and lock in but to operate on as many platforms as possible.

This is something that Nathan Brookwood and I talked about last night.

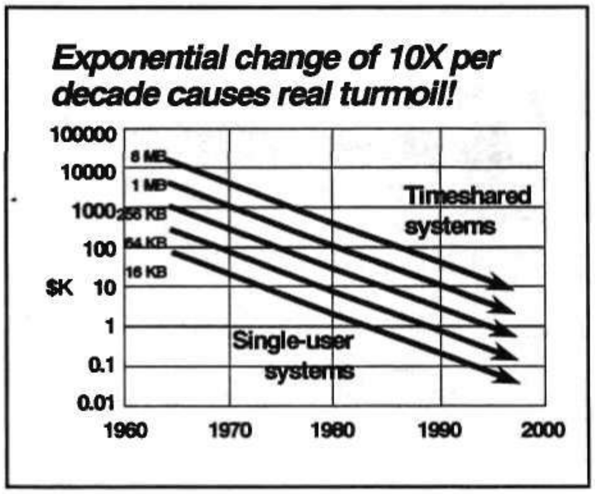

Will there be a need for high volume, higher performance micros if the PC's continue? I don't know. I can't type any faster than I ever typed, so what am I going to do with it. Where's that need going to come from? So far, speech and video seem to be providing this motivation. However, we have no idea what's really going to happen with video. So far, video is just coming into the computers and we're streaming them out.

Once we start to understand that, the need for cycles is going to go up enormously as we process video and compute video. The big need is that everyone will have video servers that will be served from their TV's or whatever. Those will be Web-based servers. The area that I'm personally working in is the explosion of storing everything — photos, voice, video. That's going to require significantly more processing and memory.

I'm not going to take any questions. I have some new data which I want to give you. You can e-mail me the question: gbell@microsoft.com

This is something that I couldn't help but do last night. I gave this talk once before, and it was boring to me — I wanted to add something to it. I was in Japan last week and got a couple of cameras. Here I was, as a road warrior. I had a Toshiba computer about that big. Of course, you couldn't look at it and my fingers were too big to type on it so a friend gave me a nice Compaq LCD terminal, mouse and keyboard. Every hotel room should have one, along with this fax thing. I don't know why anyone would ever have a fax in a hotel room.

Let's look at the kind of information that we have. Business cards, pages, a fax don't require very much space. My world can sit in a portable like this with 4 gigabytes. You can take all the business cards you've ever seen, all the pages, all the snapshots, all the books. Well, maybe not all the books because you can only store between 40 and 500 of those per gigabyte.

Now it gets interesting in terms of the project I'm working on. I want to store everything I've read and heard; not seen because that would take a lot more storage. If you look at the middle column, we've got everything you've read and a few pictures you've seen each day for the last four years stored in 2-10 gigabytes. Everything you've heard each day for the last four years stored in 40 gigabytes.

If you look at how much you've seen each day for the last four years and put that into what I call video light, which is VHS Light, the new encoding WEB-TV uses for video, that's only one terabyte. This is for four years of recording encoded motion. I'm not advocating that, but it's not out of this world. It's only a gigabyte a day if you wanted to carry around everything you saw in a day.

How does this all map into... [copy down this address: www.lesk. Michael Lesk is at BellCore.] This is where the thought gets interesting, and I have to quit, unfortunately. He looks at how much is going to be shipped in '98. There's quite a bit of storage in the worldwide base. It's on the order of 10-15 exobytes, that's 10^{18}, and then you can do the division back from that.

Look at all the bits around. The Library of Congress has a pedabyte of the pages there. Thirteen million photos, he did the arithmetic 13 terabytes. The maps and 3-1/2 million recordings, that's 2 exobytes or 2,000 terabytes. In terms of the Library of Congress,

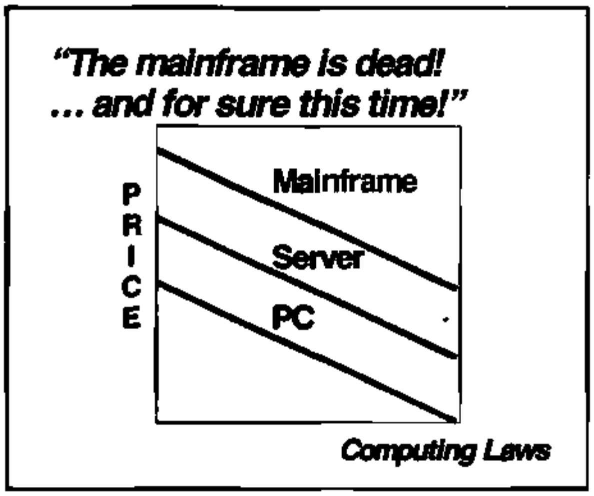

certainly we're producing enough to store all of that. If you look at other bits per year, cinemas take up only a couple of hundred terabytes. Images get scaled back by a factor of about 100. That's probably only a pedabyte. Capturing all the broadcasts, that's probably only a small number of pedabytes. All the recordings that are made, that's only 60 terabytes.

If you wanted to listen to everyone in the US, that's 500 billion minutes and that's 400 pedabytes. And then, of course, there's video. Lots of predictable computers, including the non-predictable computers.

This is how I see things going. I have quite a few bets that there will be a non-predictable computer that none of us see today. The nice things about this are surprises. Larry Ellison believes NC's will outsell PC's by 2000. I haven't been able to get him to take this bet. He's off by a factor of 10, which is the Larry Ellison height factor, at least.

This is why he and Scott are so irritated. This is really the whole business, scalable network platforms. Some of my friends believe that you can make arbitrarily large computers from PC's. That's based on commodity hardware and commodity networks, whether it's Ethernets or other kinds of things. With that, you can put them together in arbitrary ways. We put 50 of these together on May 20th and demonstrated the world's largest transaction processing system in New York. For us, scalability means reliability — always upscaling in terms of number of nodes. Scaling in locations, putting it anywhere you want. Scaling it with machine generations. That means not having these multi-processors, like the vendors are trying to push. How many microprocessors can I put in a big box. That's not scaling. You can put 10 to 20 of those in a box. To us, scaling is if you can put more than a thousand in a box.

The model is that various kinds of networks will be put together and integrated to perform as a single system. I don't think the vendors see this change coming because they've got these great businesses that wouldn't have been there without the Web. The Web has changed the whole computing server market but, in fact, that's going to decline. I don't have time to talk about bi-computers but certainly they're happening now. The wearables, I like these — the implants. Particularly since I've had a couple of heart attacks. Thank you. Send me E-mail.

Moderator: As Gordon predicted, we are out of time. When we come back from break, all of the morning's speakers will come back to the stage. We'll have an opportunity to have Gordon answer the questions that you provided just a moment ago. One thing I would like to mention is that this series of three talks this morning was absolutely outstanding. I'd like to thank all the three speakers for providing us such insight. Thank you.

Chapter 14: Inside the Mind of the End User: What Semiconductor Users Really Care About

Jonathan L. Yarmis
Vice President and Manager
Special Projects
Gartner Group

Moderator: Datquest Analysts focus on the vendor side of the industry. The Gartner analysts focus on the users' side of our business. We've been talking throughout this conference, on the vendor's side, of how to make the chips and what are the technologies? Our next speaker is an invited Gartner analyst, Jonathan Yarmis. He's going to give us the user perspective for all of this technology we've been talking about.

After Jonathan gives his perspective, he will conduct the panel, consisting of the four speakers that have preceded him. They will come up to the stage, and after that we will break for lunch. Let me introduce Jonathan Yarmis. Jonathan is the Vice President and Manager of Special Projects at Gartner. He has been at Gartner for ten years, most recently, as Director of the PC Service. Previously, he was with General Instrument, where he was a PC analyst and Corporate MIS Manager responsible for all aspects of corporate PC implementation, purchasing and issues. He was also at Touche Ross, where he was Director of Microcomputer Services. Please welcome Jonathan Yarmis.

Jonathan L. Yarmis: Thanks and good morning. I'm not a semiconductor person and, after having listened to a day and a half of this stuff, what's inside the mind of this end user is a little bit of mush. I also figured that, in the spirit of Gordon Bell, I'd start with my Larry Ellison joke. Some of you have probably heard it. The Internet makes being topical on humor really difficult because jokes spread with the speed of light now. I'd love to ask Larry how it's changed his business. It used to be that my broker would always call me up under the pretense of telling me the latest joke and only then, would get into flogging stocks. Now that I've heard them all before, he's got to come up with another pretense. The joke goes: What's the difference between Larry Ellison and God? God doesn't wake up in the morning thinking he's Larry Ellison.

I'll try to make that the last cheap shot I take although, when you talk about what's inside the mind of the end user, that will be a challenge for the next half hour. It's a fascinating market dealing with end users, as we do on a daily basis. One of the reasons why our business continues to grow so robustly is that our clients are being overwhelmed. Here is an industry where we used to think the pace of change was electrifying, and then along came the Internet, which increased things by about an order of magnitude. Every once in a while, we'll be talking to a client, and I hear a screaming sound, which I suspect is some person running out of their office looking for a South Pacific island to move to. There are quite a few challenges out there.

When you talk about getting inside the mind of the end users, — well, we see two kinds of users out there. Number one is the individual. The individual

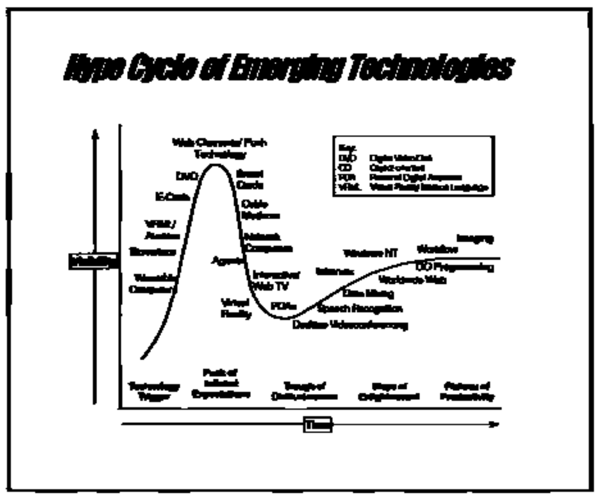

just loves what's going on. He likes this notion of getting more for less. There's a whole bunch of them out there, waiting for Microsoft to reduce Internet Explorer before they take the plunge. Given that the product price is zero, then again given Bill Gates' bank account, they might have a prudent investment strategy. On the other hand, we've got corporate users who, these days, are getting overwhelmed by change.

I loved all the curves we've seen this morning about how Moore's Law is going to continue forever. God knows it may even accelerate with some of these advances. Users have gotten to a point where we're going too fast for them. That is going to be one of the challenges to growth as we see it in this industry. The users are saying give me less. I'll even pay more for less. "Your product cycle, it's almost a sick joke."

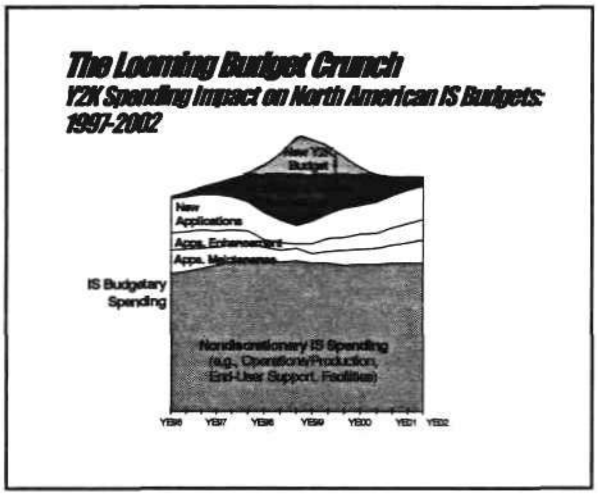

We've got a client that evaluates a product by taking it in and exhaustively testing it. They test it with all their standards and networks, and anything else they can think of. They put it through a rigorous six month evaluation process before they declare a product ready to be purchased. The average life cycle of the products they're buying right now is about five months. Therefore, they qualify a product for purchase the month after it's withdrawn from marketing. It's a ridiculous problem in the end user community. It's one of those things the industry will have to wrestle with.

When we look at what's inside the mind of the end user, let's also understand that, in many instances, the things that they're most excited about are not necessarily your best selling opportunities. We see what's called the "hype cycle" of any new emerging technology. What happens is that when it's first introduced, it's the most amazing thing ever to strike the face of the earth. It's going to cure cancer cells, global warming, render El Nino a non-issue and reinvigorate economy. Then three weeks after these massive pronouncements are made, people say that it has failed to meet expectations.

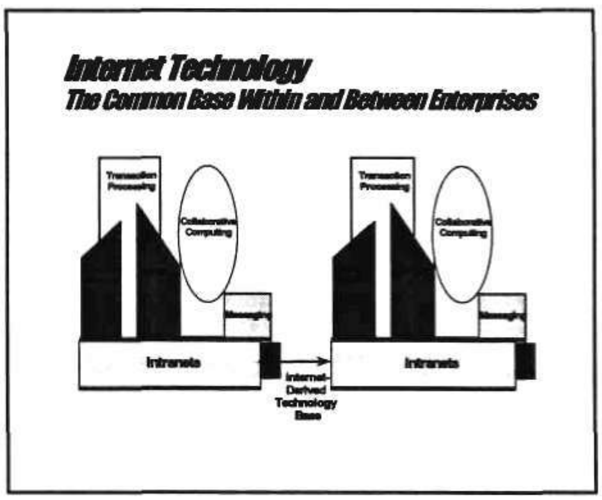

We go into this trough of disillusionment that says that this product is never going to work. "How could I have been so stupid to ever think this was going to work. This is a dead technology." This is usually about the time that people are going to start buying this stuff. We go into the period of reality, in which we tell ourselves we have this inflated sense of expectation about what this thing could do, and it couldn't do any of that, but when I went to dismiss it, there actually was something to that technology that's interesting to me, and maybe now it's time, with realistic expectations, to go into that space. We see it time and time again.

We have a list of technologies out here. This is probably a chart we ought to have in real time. At the peak, we've got Push Technology. The amazing thing with the Internet is things go from the hype cycle to disillusionment to reality in about a four week period. When we did this chart for the production deadline about a month ago, Push was the darling that was going to save the industry.

Now, you see the stories about how Microsoft is backing off from Push. PointCast should have sold out to Rupert Murdock because they'll never be able to have an IPO because no one will want technology. The truth of the matter is that this is a long term, essential, important technology. The good news is that now that we're going into that disillusionment period, we can begin to get real about the capabilities of the technology.

When you're looking at users, understand that hype cycle. Many times when they're speaking most vibrantly about a technology, you have zero selling opportunity because you can't deliver what they're looking for. When they're the most disillusioned, you've got the best selling opportunity. It's a lot easier to exceed modest expectations than it is to exceed hyper-inflated expectations.

This is something we see continuously in the market. It's great to look at, but understand the challenge in anticipating this stuff. With some of the technologies, you go from the technology triggered to the plateau of productivity in about a six month period. Then there are things like voice recognition or handwriting recognition. Anyone remember the Newton era? The period Gordon Bell was talking about 20 years being an overly optimistic timeframe for the voice stuff? Understanding the timeframes are important, but also understand that the peak is generally not the best selling opportunity into the user community.

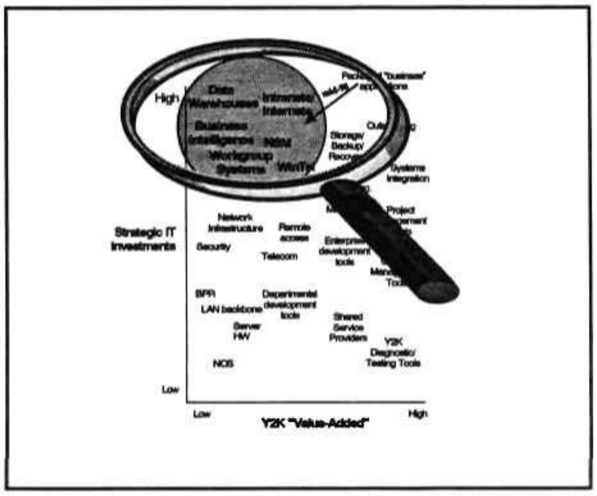

What do users want? Every once in a while, we try a novel approach and actually ask them what they want. We've got a program called ITEP, standing for the IT Executive Program, made up of a couple of hundred CIO's. Every year we ask them what they want. There are two interesting and important things here that have some relevance to the transformation that's going on in this industry. The North American commercial market turned out to be 1 and 2. That is aligning IT and business goals and IT for competitive breakthrough, competitive advantage. What's new about that? We've been talking about these things for the past 10 years in this industry. It masks an industry transformation.

Number one, what did it mean to align IT and business goals? It was the way our end users made themselves feel important. What it really meant is that we should figure out what the users are doing, then back off and automate that stuff. IT people were like children, best seen but not heard. IT was best when it was invisible.

Now there's this whole notion of aligning IT and business goals. The IT person is in a fascinating place in the organization these days. They've seen the future. They get the Internet and what it lets them do at a visceral level. The business people have yet to figure it out. You mean, I can put my catalog up on-line? We'll talk a little bit about that in the panel session.

There are some profound business transformation issues that are interesting in the user community. This enables IT as a competitive differentiator. We echo some of the comments Larry made about long term sustainable growth in the IT industry. This competitive advantage concept is something we've talked about forever. In most instances in this industry, competitive advantage was a fleeting thing. You'd gain competitive advantage for about 72 hours and then everyone would see what you were doing and they'd copy it. All of a sudden, competitive advantage was lost.

What we did was redefine the level playing field. We believe that the next 10 years will be the greatest opportunity for competitive transformation that any of us are going to see in our lifetime. The way people use information technology to reach new customers and new markets. The way they transform their businesses and more tightly integrate themselves with their customers and their suppliers. There are massive opportunities to transform industries, for example, Amazon.com: maybe we'll talk about that a little later. Just one example of someone who's changed the competitive ground rules of an industry. We see chapter and verse of that taking place.

There are many technology issues in the user community. What do you look for in real estate? Location, location, location. In Technology it is Internet, Internet, Internet. You cannot talk enough about how companies are seeking to transform the way they do business, given Internet technology. It's replacing much of what they've already done. The fascinating thing is that it's also creating incredible new opportunities. There's a side of me that says, "Let's be honest about the Internet. This is the biggest step backwards this industry has ever taken. What type of breakthrough is this? We can display text and graphics on a computer screen — and this is what we are hailing as the greatest breakthrough?"

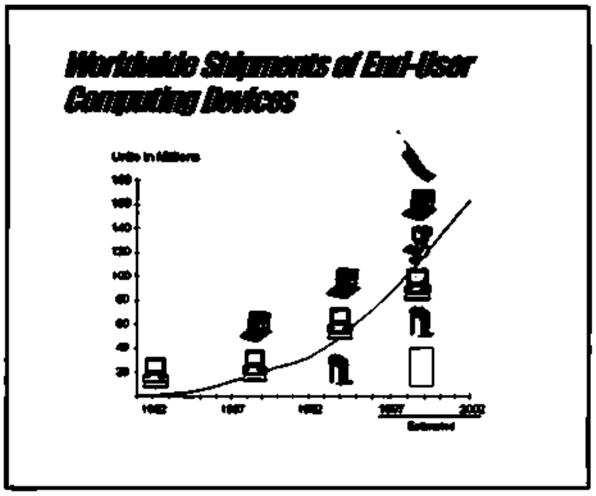

Somewhere along the way in obtaining this ability to display text and graphics, we decided that sub-second response time wasn't really all that important. Maybe users will tolerate sub-hour response time. We've taken major steps backward in an opportunity to achieve a global connectivity. That is such a radically transforming idea that it is worth all of the pain and suffering that we'll be going through. From a technology side "Internet, Internet, Internet" is what's on the users' mind. Year 2000 is also. They wish it wasn't on their mind. It's certainly anticipated to have a big impact. However, when they're looking at strategic directional technology issues, we can't talk enough about the Internet.

There's this looming problem, the Year 2000 problem. There's a side of me that I'll tell to this audience. You'll never see me do it for publication, because I don't want to precipitate a global banking crisis. This whole issue terrifies me. The bottom line is whether you are going to be Year 2000 compliant when it comes time to be Year 2000 compliant. We estimate it's about 50/50. About half of the companies out there will not be fully year 2K compliant. We define full compliance as not only your company being compliant, but all your business trading partners, and anyone else in your sphere.

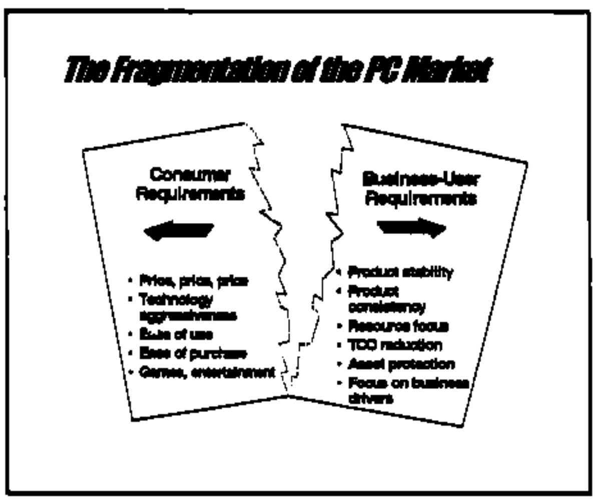

I was talking with a regional bank in upstate New York a few months ago. They were congratulating themselves on how ready they were for the Year 2000 problem and I asked them if it would trouble them to know that the Federal Reserve Board, for instance, hadn't begun their Year 2000 initiative. Or that the State of New York had budgeted as much money for year 2000 compliance as the State of Mississippi. Either Mississippi is going to be very ready (and some of us would argue that they haven't made it to the year 1900 yet), or New York will be in a world of hurt. This is an issue. When I say that I don't want to be alarmist, there's a side of me that says, "Let's ignore Larry Bowman's advice and put everything we've got under our mattress. All hell is about to break loose in the banking industry.

I've said this to a couple of people in financial institutions and they said that I was being a bit alarmist. They said they understood. One of them said that they would take 15% of their wealth and put it under a mattress; the other said 20%. I told them to think about it for a second. If what they'd just done becomes generalized behavior, that's called a run on the bank.

The very fear of Year 2000 non-compliance becomes a self-fulfilling prophecy, causing catastrophe in the global financial market. I'm the type of kid who liked to pick the wings off of flies. I have fun doing this with people in financial institutions. This is a truly terrifying issue. There are many places where I walk in and they talk about Internet terminals and transformation being interesting and enterprise applications being fascinating. Maybe I could save my business but all of my resources are going into Year 2000 stuff.

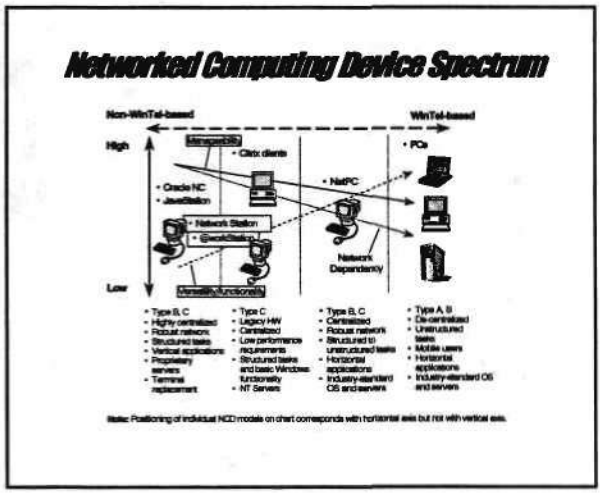

In a corporate marketplace, this is a big and growing issue. Anyone that is trying to figure out what's on the minds of corporate users, figure out a Year 2000 angle. It's a fascinating one. This industry hasn't done enough talking about Year 2000 issues. Since I am the rabble-rouser that I am, I met with them a couple of weeks back and asked them if the traffic grid system was going to work and they asked if that was an issue. Well, many of those streetlight systems are run by a calendar, and the calendar isn't Year 2000 compliant. When we get there, the system is not going to know whether it's a weekday or a weekend, and your traffic light system may be on the wrong pattern. There are many places in the semiconductor marketplace where there are non-intuitive date issues. The whole issue of what the lurking liability is and what the impact on the industry is, is certainly something that's very much on the mind of end users, and is growing more every day.

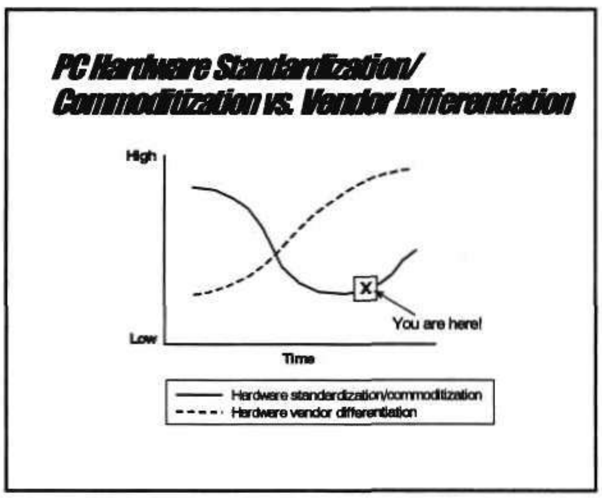

I can't say enough about the Internet. There's a side of me that regrets that every new technology that's been introduced in the last 50 years has been deemed a paradigm shift. This is the paradigm shift. This is the single biggest business transformation opportunity any of us are going to see in our lifetime. The technology has enormous impact on everything we do. I'm a walking Internet access device. This cell phone isn't Internet-enabled yet, but probably some of you in the room have Internet addresses on your cell phone. This is 2128595@skytel.com.

We're going to see the Internet pop up in so many different things. I'm going to talk about someone that you probably haven't thought of when we get to the panel. Just be thinking about what happens when the Internet meets your gas station. Think about how that's going to change the way you buy goods and services in the next couple of years but, even more profoundly, projecting out 5 or 10 years.

You cannot say enough about how the Internet is changing business. IT organizations are spending quite a bit of time looking at this. The business community is about to get captivated in a way that would support Larry's notion that IT is going to become more and more prominent in global economies. It's not just the IT organizations — we are capturing the imagination of business users around the world, in a way that's going to mean great things for those of us in this industry.

Where are people investing? There are only three that matter. Number one, obviously, is the Internet. Number two, Year 2000-related. If you have something that solves the Year 2000 problem, you're very close to having the ability to write your own ticket. This is one of those things where, if your companies haven't jumped on the Year 2000 bandwagon already, then how are you going to solve the problems, I'd love to know who you are because I can write you a really big invoice to help you get out of an enormous problem.

There's quite a bit of money going into this space. The amount of resources that are available are going down. The demand is going up. It is a great solution place to be in. If you could just start labeling semiconductors Year 2K compliant, you could probably price them 10% higher. That's probably being conservative.

Lastly, the whole notion of business value-add. These enterprise resource planning packages that people saw some bonds and SAP's. There's quite a bit of money going into that for two reasons. One is that the vendors have a bit of an inflated sense of self — yet they are going in because they are a great solution to some profound business problems in terms of how the business will be managed and integrated, and put on a common planning framework. There are many people buying them for very tactical reasons.

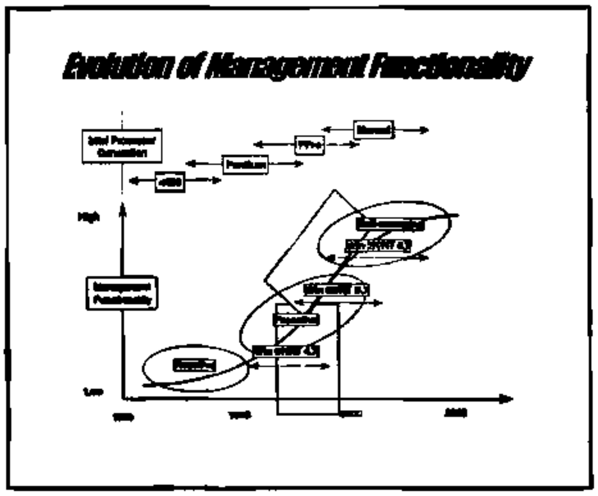

For example, Year 2000 compliance — I can make it or buy it. I can shift the burden to someone else and sue them if it doesn't work. That's another terrifying statistic for Year 2000. I don't know how many of you saw Lloyds of London. If there are any publicly traded law firms, go invest in them. Lloyds of London estimates that the legal claims arising out of Year 2000 non-compliance, in the United States alone, will be a trillion dollars. I look at that number and think it's a ridiculous number, so maybe they're off by an order of magnitude of only $100 billion.

This is a profound business issue that is very much on the minds of users. It is not just an IT issue anymore, when you start talking about a trillion zeros. I'm not as conversant with the number of zeros as the semiconductor people are, but it's the one with a whole lot of zeros at the end. Enough to put some of us out of business. This is very much on the minds of end users.

We see accelerating growth in the PC market — not only in the units but also in the variety of things. We've done ourselves a dis-service in the last day and a half, as we've been focusing on PC's as desktop devices. The variety of things we'll be carrying around: the Palm Pilots, the credit card device that is being sold, my pager, my cell phone, my eyeglasses, my wristwatch. The number of things we're going to have that are Internet-enabled — communicating and computing devices — are going to grow dramatically in variety.

Think about where that microprocessor is going. The user is not thinking about just putting it in a desktop or portable computer anymore. With the price points continuing to plummet and the functionality continuing to accelerate, we're going to see these things everywhere. The average high-end automobile today has about 40 or 50 processors in it. I just built a new house that has the home environmental system. You name it, we'll start to expand the notion of what we think of as a communicating and computing device.

We've also seen a fascinating change in the marketplace in the past few years. It used to be that the corporate markets, because they have the money to spend, led the market in the adoption of new technology. Starting with the CD-ROM, it's really the consumer marketplace that's leading the adoption of new technologies. We expect that to continue for some period of time.

There are certain innovations that consumers will adopt much more aggressively because they see personal business benefits. Why did we all buy CD's? It made sense. Buying the CD drive and then getting an encyclopedia for $25 was cheaper than buying Encyclopedia Britannica. There is a company that missed the market. A fascinating case study of someone that really blew it. The whole notion of the consumer marketplace leading some of these new technologies is something that those of us who follow more closely with the business marketplace are just coming to grips with.

If I had to align myself with either Gordon Bell or Larry Bowman, I would take the position that there is a market in the world for these NetPC's but it is probably not at the expense of personal computers. There is room in the world for both devices. Many people are going to want more functionality and power on the desktop.

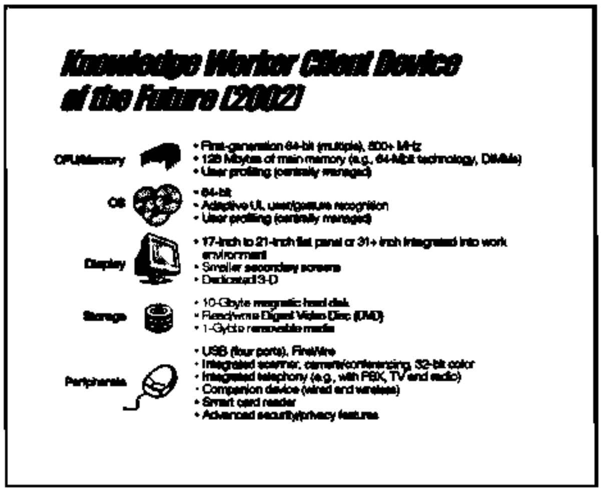

That said, there's a category of users for whom the PC is overkill and there's room for a large and vibrant NetPC market. After 10 years of harping about it, users have finally gotten on this total cost-of-ownership bandwagon. This says that in deploying this technology, the truth of the matter is that the acquisition price is only about 15% of your 5 year life cycle cost of the technology. We finally started to see users say this is going to be an active decision criterion.

Contributing to reduce total cost-of-ownership is a big selling proposition and the user market is galvanizing very quickly. For 10 years, we've lived and breathed that stuff. Now, it's rolling like wildfire into the end user community. I give Sun, Netscape and Oracle quite a bit of credit for trying to use that, for obvious marketing reasons, but it's caused the users to wonder what it really costs them to deploy technology on the desktop.

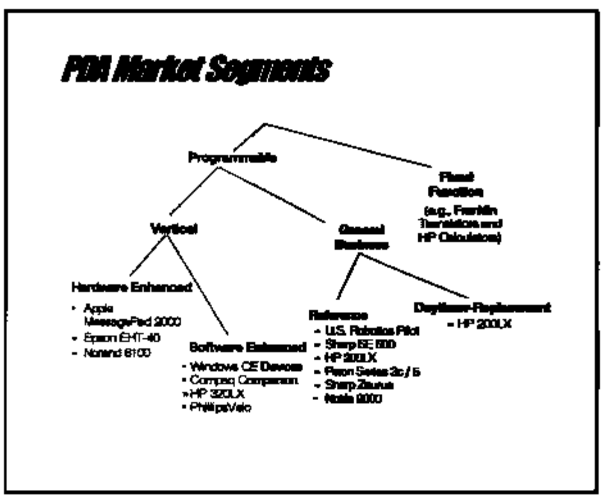

We've gotten faster in introducing new technology to the market, but users have not acclimated to the situations. We keep talking about how you've got to shorten the cycle times, but we believe there's a big issue in the user community on how they're going to adjust as we keep ratcheting up. As we heard this morning, it's going to get faster and faster and faster, cheaper and cheaper and cheaper. Let's tighten everything up and make it quicker. We haven't done anything inherently to increase the user's ability to assimilate that technology. There's a very great backlash potential out there.

I would concur with one of Nathan Myrakravold's Laws, in that I am amazed at the ability of software developers to take any hardware we put on the desktop and bring it to its knees. And do it with meaningful functionality. We've just begun to scratch the surface. Anyone who believes the graphical user interface is the be-all, end-all user interface is missing the point. There is so much more we can do to make the system easier to use. Natural language interfaces are on the horizon; real-time translation, that is enormously processor-intensive. You give me a high-end machine, and I guarantee you we will fill it up with meaningful things very quickly.

The PDA market is not just one market. Cases in Point: Apple's notion that the Newton was going to be the answer, or Microsoft's first generation Windows CE abomination. We're still wrestling to find the right form factors. We're going to see oodles and oodles of these things, to a point where everything you've thought is big enough to incorporate a microprocessor will likely do so in the next several years. It makes this an incredible growth opportunity market and challenges the creativity of people. How are we going to enhance functionality with the technology?

On the software side, what users care about is organizational productivity. We've spent 15 years in this industry solving the personal productivity problem, only to discover we don't have a personal productivity problem. We've got an organizational productivity problem, and the flow of information goods and services is very much what people are focused on. Don't give me yet another solution to a personal productivity problem, give me Groupware, give me Enterpriseware, give me InterEnterpriseware. That's where all the focus is in the software market.

In conclusion, users want to change their lives on the Internet while surviving the Year 2000 problem. They want to learn how to cope with the increasing complexity that we're delivering out there. How do we, at the same time, deliver an incredibly rich technology solution, yet make it easier for them to manage. Lastly, the fragmentation. I hesitate to call it the desktop because that's perhaps the least interesting piece of real estate we're fighting over. It's the mobile top. It's how much technology are we using today. There's a side of me that's a representative of a technology company. It embarrasses me that we give you these really fat binders. And the only reason why I'm not more embarrassed is most of you, even if we gave you a CD-ROM, aren't there yet with the technology. Think what it would be like in a couple of years when you've got your small handheld, handwritten, annotated device that has high resolution display graphics capabilities. We would beam the presentation to you and all throughout it, in real time, you've got the information available to you. That, to me, is the fascinating transformation process. Apple referred to it as anytime, anywhere computing and we're really on the cusp of that — that the prices, the form factors, the supporting infrastructure and all of those things are taking place that enable incredible business change. It makes this remain the vibrant industry in our economy and a great place to work. Thanks.

Moderator: Jonathan, we have a couple of questions.

Q: Manny Fernadez open the conferences with his comments about the Y2K problem and that it would greatly impact IT budgets of the users. Why haven't your IT surveys shown this?

A: In fact, our IT surveys have shown that it will greatly impact budgets of the users... the problem is in the way it's impacted budgets right now. Most of them are saying, "I'm going with my budget the way it is and oh, by the way, I've got this shortfall based on Year 2K compliance, and I'm using that as a way of increasing the budget outlay from the organization." What's happened is that we haven't seen too many slashed expenditures based on Year 2K compliance, but we've seen some interesting expansion of budgets. That said, it differs greatly from company to company. There are those that haven't done a good job of defining the business problem, and IT doesn't have a great track record. We really see two kinds. One is where they've expanded the budget, so they've gotten incremental money; and one where, just now, they're starting to run into the wall and we're starting to see discretionary programs slashed. You'll see that turn up much more in budget surveys next year.

Q: Who are the companies who will benefit from the percentage of IT budgets not spent on new applications?

A: The people who are making quite a bit of money. Number one, outsourcing vendors. If you're not spending it on new applications, how are you going to manage the old applications? Increasingly, the solution being used is outsourcing. This is a trend that we see is growing, and long-term, and durable. It's not a cost savings phenomenon. Rather, if you've got scarce resources, where do you allocate them? The biggest allocation that isn't going into the new application space is going to the outsource vendors.

Chapter 15: Special Speaker Panel: Bringing It All Together

Jonathan L. Yarmis
Vice President and Manager
Special Projects
Gartner Group

Larry Bowman
Founder
Bowman Capital Management

Mark Mellier-Smith
President & COO
SEMATECH

Walden C. Rhines
President & CEO
Mentor Garphics Corporation

Gordon Bell
Senior Reasearcher
Telepresence Reasearch Group
Microsoft

Before we get into specific questions, I wanted to outline some visions of the future. Perhaps also to see whether you concur with the visions and what you see as the business opportunities arising out of that vision. Some of this is related to several consulting engagements that I have had in the last couple of weeks. One of them was with British Airways where we're working with BA on their Terminal 5 that they're planning to build at Heathrow Airport for completion in the year 2004. This is an interesting exercise.

Also, the gas pump experience when I was meeting with Arco. BA's vision of the future is that technology will be so cheap and so broadly available that given the price of an airline ticket, they will be able to give you, at your point of initial contact with them, a little card device that has a GPS system and a dimensional sound speaker on it. As you walk through the airport, they will know where you are and be able to give you customized personal announcements. "Wally, your plane's leaving in 10 minutes, you ought to be heading to the gate right now." Transforming and hopefully getting you entering the airport to the airplane quicker. If you're one of those lunatics who actually listens to them and you get to the airport two hours before check-in, as you walk around the airport and walk by a retail store, it may tell you to go in now and they'll give you 10% off a purchase. That's a business transformation.

I'd like to comment on that. If you look at my homepage, I've been working on telepresence and telepresentations. A couple of weeks ago, I was in Heathrow. Every year I go to Heathrow and I can't imagine Terminal 5. The thing gets worse and worse. It is the worst experience ever. Having a thing stuck in my ear or anywhere else in me, would actually be a benefit. Something that would put me out of my misery as I go through the place. Personally, I don't want to travel in those kinds of environments. Sure, the technology is trivial to do that kind of thing, but I have a personal bet out, that's still got four years to go, on video telephony and the impact on that. Already, we're using it a lot in a collaborative thing. We had a conference, the ACM 97 Conference, the 50 year conference of the Association of Computing Machinery. There are 2000 users and so far there have been 20,000 people on-line looking at those presentations. There are people out there who don't really want to travel and, in fact, this conference would be a great one to telepresent. I don't see any reason to be here, frankly.

Q: Are there any bandwidth benefits today in the corporate environment?

A: In the corporate environment you have adequate bandwidth and then you can cache this stuff down. I see 200Kbit, certainly VHS quality, that's what the new WEBTV stuff is. It's perfectly all right for that. If you're interested in just the content of this, as opposed to the networking, you can get it all that way and a lot more efficient. I'm an advocate of that.

Q: Are we going to be able to build all this stuff?

A: No problem.

Actually, there's a lot of cost demand tradeoff. On the GPS side, it might be nice to have the GPS chips and have them tell you where to go. Any of you who've used GPS systems in cars, they're phenomenal. You drive along and it says you should turn in 200 feet and so on. The problem is, it's $2000 or $3000 added to the price of the car and the average person can't afford it. We have to do an awful lot of cost reduction to adjust not only the chips, but the displays, the interface, the speakers and everything.

I was at the Akihabara Japan last week and they were down to $500. They're in there to do all that stuff. You get rid of the screen and make it all voice actuated. The GPS will be a throw away device.

Moderator: What breaks that chicken and egg phenomenon with GPS? We understand that with the PC model, if you make enough of them, it becomes virtually free. How do we get down that learning curve? What are the drivers that are going to accelerate certain of these technologies?

A: The economies of scale that can be achieved are going to happen naturally. With some resistance because they require cooperation. We've seen it on the manufacturing side. Years ago, everyone did their own manufacturing equipment and everyone did their own process flows. Now, you buy a piece of manufacturing equipment and process recipes can be included. Everyone uses one of a few alternatives for each process step. The same thing has to happen on the design side. As long as 50 different companies design 1394 interfaces, everyone's got to debug them and everyone's got to put effort into it. The big leverage point is going to come when one design for 1394 is the one that everyone uses. When it's been in a hundred designs, it's debugged and it's high quality. When you just pull it up on your computer, use it and spend your time doing the things that are unique. Big economies of scale-end design yet to be realized. That you haven't felt yet. The manufacturing economies- we already have the mechanism and it works to cause more specialization and more economy.

For the system-on-a-chip, there has to be an integration at the same level as the PC. That is the original bus structure so the 1394 interfaces to something. You need the layering and architecture for system on a chip. I predict there will be three or four of those particular things. There will maybe be four 1394 interfaces for the various varieties. There will be an Intel chip. There's certainly a chip that has all of the integration out there that works with the software. There are handlers there so people can write software to those various system on a chip platforms.

Panelist: You will always have only a few alternatives. It's so much better than everyone doing his or her own thing.

Panelist: You can't do it the other way and make the plethora of stuff that we want. The stuff that fits in the pockets or the appliances that we use.

Moderator: Larry, how do we pay for all this stuff? You're talking about the attractiveness of small and midcap stock. I heard these guys talking about how we're going to be spending $10 billion to build a fab. Are the capital markets sufficiently available for this business to fuel the growth?

Panelist: I think I'm at the wrong conference. I haven't understood a word you guys have said for 15 minutes.

Moderator: That's why I asked you the question.

Panelist: I'd like to talk about slams. This one I can deal with. Capital markets are really neat. If you remember the movie, greed is good. Where there's money to be made, there's capital available. Period. Where there's money to be lost, there's money to be made. The money's there. Period. It will be there. You talk about $10 billion fabs having to be built. The neat thing about high technology is that the real money makers are the people early on. The small design houses that develop innovative technologies, patent them, get licensing breakthroughs, get royalty streams off the intellectual capital. Rambus right now has a $2 billion market cap and $25 million are revenues. The revenue stream, the royalty stream doesn't even really kick in until '99, 2000, 2001, 2002. It's the start-up companies of today that will give you the homeruns and grandslams, the return on your invested capital, that will allow the designs and implementations to achieve whatever it is you guys were talking about.

Moderator: : Every three months they say the market is overvalued. Do you think there's an appreciation in the market of the transformation and role that information technology is going to take so that maybe those old rules of valuation, at least the ones that apply to this sector, don't apply anymore?

Panelist: The rules have already changed. The highest quality growth stocks in the world, typically ones we'd look at, whether it's Home Depot and Coke, trade at their growth rates. If they're going to grow earnings at 30% a year for the next 10 years, they trade at 30 multiples.

It used to be that technology companies, had low multiples because they were technology companies. You never got over 15 multiple for a PC company. You never got over an 8 multiple for a disk drive company. Intel never got over 15 or 16 multiple up until one or two years ago. Now, the markets are starting to realize that maybe they can count on PC's continuing to be around in 3 to 6 years from now. Whereas, before, we thought they were just a fad. Three or four years after PC's started up, in '85, they hit a wall because we ran out of all the single user app's. We had a downturn in units in the PC industry. Now, there's more long term confidence that technology is here to stay and again it is a growth industry. Confidence and multiples are certainly going up and people are much more willing not to just puke out all of their contact because they had a bad quarter. It's a $26 billion company, growing at 30% a year and people are starting to recognize more and more that there are ups and downs within the cycle. Even with the short term Intel is up, almost 100% in the last year. People are realizing what real value-added these companies have.

Moderator: Who is puking out?

Panelist: Puke out is selling slightly slower than blowing it out. Which is what we should have done this morning.

Panelist: Let's get at what may cause you to puke out a quarter or something. When you're talking about these intellectual property issues, we've always known it was out there but there was a tradeoff between time-to-market and intellectual property. It sounds to me as if there's going to be a short term pickup in this space. You're saying for the longer term good, to survive we need to get to the reuse but in the short term, it's going to slow me down. What's going to cause someone to act? If I were running a company, I'd want to be the second, third or fourth one to do that. Let you fall on your sword and then after you've imploded, I'll learn from your mistakes and the market will start to be more receptive. Is that pickup opportunity going to retard the market?

Panelist: I don't think there's the kind of pickup you have when you have to tear up everything you've got and start over again which is what happens in these technology waves. In this particular one, you can keep what you've got. It may not be that useful to you moving forward, but at least you don't have to throw it away and start all over.

What you have to do is start designing with the methodology that allows you to reuse it. That requires investment, which means capital, training, learning and so on. The ones who make the investment earliest will get there quickest and the ones who make it later will get it later. The good news is that the IP that ends up being well supported will be a benefit to everyone because it will be available to everyone as you get your methodology in place.

Moderator: Where are the discontinuities? That's something that as market forecasters, we're as guilty as anyone. We can draw these great extrapolations and then along comes something like the Internet. I'd love to sit here and tell you we predicted that. There were people out there who predicted it. You generally thought of them as lunatics before they predicted it and even though they were proved to be right, you still believe them to be lunatics. Where does the discontinuity come in? The things that say that Moore's Law is conservative and we're going to kick that curve up or that Moore's Law is over and we're going to ratchet it back 30% a year. Do you see something like that on the horizon?

Panelist: I actually believe Moore's Law is unconstitutional because, actually, Moore's Law is not the phenomenon. It turns out that it's unconstitutional to shrink feature sizes and grow wafer diameters and that has caused Moore's Law to be true. The real underlying principle is the learning curve- the cost per bit or the cost per function. That's what drives it. Looking out in the future, it may be that going to all these exotic technologies that Mark talks about isn't the lowest cost way to get there.

Panelist: Modeled chip solutions and going vertically instead of horizontally, shrinking and that kind of a thing becomes a more viable mechanism. If I looked to the design side and say where is there a discontinuity as we move forward, the one that's on the horizon the soonest to arrive seems to be at about the quarter micron point, maybe a little sooner for some, a little later for others. You just can't use the same toolset to verify a full chip solution. You get to the end of the design or somewhere in the middle of the design and you put it on the computer. You network every workstation and server that you've got and it runs for days and days and it never gets there. That's the breaking point in which you need a new tool, a new methodology.

You need to reuse blocks so that you don't have to verify the individual things they're built from. You need to get a capability that does, for example, hierarchical verification instead of flat verification, which means you throw out all your old verification tools and buy a whole new set and a new methodology. That's going to occur in the next two years.

Moderator: Wally, that makes an assumption about architecture and design that, in fact, I hope is not true and that people are not going to do that. That, in fact, when you get ten to the ninth transistors on a chip, that doesn't mean you've got all these designers each trying to design something like that that you could never possibly verify. Their secret there is it's all memory. Give us as much memory as you have and we'll fill it with software. Don't make them so complicated. When do you bind the complexity? Bind it after the fact so that you don't bind it in the factory. People used to ask what we were going to do with all those chips. There was a theory for awhile that you'd put more microprocessors per chip. One of the most efficient designs out there is the arm chip. At one point, it was 32,000 gates or so. It's grown now. Fundamentally, it's a small chip. It's a small design which is known to take every architectural trick in the book and to add new books. I don't think people are going to go to this design complexity. Look at the design. Look at how much of that is RAM and how much is cache. You don't have to verify at that level. I don't think the real complexity is growing like that. It shouldn't grow. There's no need for it to grow at the level you're saying. That stuff has got to be memory based.

Panelist: By definition, it will grow only as fast as the tools let it grow.

Panelist: Fine.

Panelist: For example, chip designs of the past have basically assumed that the embedded software I put in ROM will work with the hardware and that if it doesn't, I can go back and make a quick change to the ROM and fix it. Or I can buy devices that have E prompts so I can change it. Now, you put that embedded arm core, MPEG core or other core on the chip and there's a certain amount of software that has to be embedded on chip. It has to be firm, hard to change, so you need to verify software in advance. The simulators of today don't even verify enough hardware instructions to boot the operating system much less look for subtleties of function. The tool developers have to get tools that, in a reasonable amount of time, do hardware/software co-verification for that embedded processor. And they're doing it. It's one of the faster growing areas of the EDA industry. A whole set of things like that is required that wasn't required before. That's what the discontinuity will drive.

Moderator: Is there a scenario here? I listen to all the talk of complexity and I'm sorry Mark wasn't up here because he talked about some of the things that could be barriers to industry growth. Is there a credible scenario that says heads are going to start exploding, that the barrier to growth isn't what you can do in technology but what we can do with people? The number of available people aren't growing as fast as our ability to grow this technology and that's about to be an inhibitor to industry growth.

Panelist: A lot of people would argue that is the inhibitor today. It's so hard to find people that people are going all over the world setting up design groups and other things just because of labor shortage. I wouldn't say at any one point in time that any one of these dimensions will be the sole stopping force. There will be times when people are the shortage, times when there aren't enough designs or designers. There will be times when the wafer fab capacity is insufficient. It will go back and forth. As one becomes the critical limiter, then we'll get more innovative in those areas.

Moderator: Larry, is it a capital market issue that you and I are probably the only two people in this room who have no idea about 80% of what they're talking about yet they're going to come to you and ask for a $10 billion check. The ability to make money attracts money but part of the reason why every Internet venture under the sun got funded two years ago is that at least you were able to understand the business plan. Is the complexity of the market going to be an issue to capital attraction?

Panelist: I'll take your 80% comment as a compliment. In answer to your earlier question: Are people a barrier to implementation in technology? Absolutely not. I'm a simple person in a simple world. We need to get to the point where a person can do something easily. They want to talk to someone in Hong Kong so they do a teleconferencing. It needs to be that you walk in and turn on the television and there's all the data you need on a worldwide web. You say open the door, let out the dog, start the car, turn on the hot water, what's going on at work? What technology has done greatly today is provide horsepower, CPU cycles. What it has not done is provide band width, which I see as the great theory to implementation of all these technologies. Software certainly hasn't kept up with the hardware. When we come full circle, when Larry Bowman can pick up a big red crayon and accomplish everything he has to do in a day by drawing a couple little pictures on a piece of paper, in 100, 200, 300 years, then we're starting to come full circle again. Technology right now is the biggest barrier to technology implementation. As far as capital markets, were you asking are the people that invest smart enough to invest in technology companies?

Panelist: I guess you told us these people are smarter than the experts making the pick. At what point of complexity does that become a problem?

Moderator: The neat thing about the industry I'm in is, it's crass, but money attracts brains. Every one of these conferences I go to, I get a dozen resumes from bright doctor level engineers that want to work in this business. They see there's a reward and people are excited about the stock market. It's the world's biggest poker game. This is a lot of fun. There's nothing more chaotic. There's nothing more competitive. There's nothing farther out there on the risk/reward curve. Most entrepreneurial people, the brightest people, do two things. They go into venture capital, the stock market, or they start the next leading edge companies.

Panelist: We in industry want them to stay inside companies. We don't want them going off and raising money. That's a waste of engineering talent. If you can engineer, stay in there. The challenges are far beyond any of that. There's nothing new about market. People have been buying and selling for a few thousand years. They did it with stones, now they just do it and then we'll just do it with bits in a few years. Don't send your resumes. Stay in engineering. The challenge here is so much greater.

Q: Is it really?

Panelist: It really is.

Panelist: The challenge here is greater?

Panelist: Yeah, you're just making money. Hey, I help start companies, probably about once a month or so. I help put money in startups. To me, that's the challenge. Not going off and recycling money. That's the recycling business.

Panelist: I don't really recycle money. I try and grow money.

Panelist: I'm going to stay on Larry's side. Next to his name, it says Microsoft. Here's a company that were it not for that ability to grow capital, Microsoft pays below market wages and the reason why they can do that is one, Seattle's cost of living is less than the Valley's, and two, you go to Microsoft knowing full well that when your stock appreciates at the rate it's been appreciating, your salary is irrelevant compared to your total compensation package.

Panelist: I don't think we do it for that at all. We have fun.

Panelist: We have no interest in that at all. I'm in research by the way.

Panelist: We had noticed.

Moderator: I want to turn to your notion of the cyberization of everything because I share that belief. I probably don't get to research it as much as you. I want to talk about one scenario. I worked with a client in the last couple of weeks and we talked about the gas station experience 5 to 10 years out. Just some of the implications for this.

I don't know how many of you have seen this yet but Mobil is doing two interesting, seemingly unrelated, things that has profound transformation capabilities. Number one, if you buy gas at a Mobil station in Florida these days, a lot of them have pumps with *CNN Headline News* playing on a CRT right at the pump. What do most of us do with our five minutes at the gas station right now? We look at the car, we clean the windshield, we stand around looking at the other people standing around looking at us. So, Mobil has now given me something to do with my five minutes.

Also, in Connecticut now, I believe it's in other parts of the nation too, they give you what I call an easypass device. It's one of those things like an automated toll thing where you can drive straight through the toll booth and it just debits your account. Mobil's giving you one of those devices so when you drive up to the gas pump, you don't have to take out your credit card. You can start to combine those two things. When you drive up to the pump in a year or so, in the Internet enabled gas station experience of the future, you drive up, your Yahoo! screen will appear on the pump so you've got a personalized, customized information experience plus a retail experience. They'll remind you that the last time you were there you bought a Coke and they'll give you $.25 off if you buy a case of Coke this time. Press here, right now, and they'll deliver the Coke to your car.

There was a story last week in *The Wall Street Journal* talking about how gas stations and fast food places have stopped fighting over corner real estate and decided co-location makes a lot of sense. The gas pump experience in the nearterm transforms itself into the five minute retail experience. This is where you'll get your Coke, beer, diapers and lunch when you fill up your car. The five years beyond that phenomenon is when our cars get more intelligent. I may well program a set of rules in my car that when I'm below a quarter of a tank and there's premium gas on sale then the car will send out a notification that I want to fill up my car. Then the roving gas tanks will say, "Hey, Jon's car needs gas." I'll come out from work at the end of the day and my car will have been filled up. These, to me, are things we can do with technology today. This is not science fiction. This transforms the way we do business. This is one manifestation, Gordon, of your cyberization. Comment?

Gordon: I don't get too excited about a television set on a gas pump. In fact, Shell has a thing now that they're testing. They put a sticker on your car so when you drive up in your car, it reads your car, knows where the gas tank is and puts the gas in. When it's all deployed, I suspect we've all got a screen, hopefully not a screen, but a voice thing that's sitting there reading whatever we want from the Net. I believe that's going to be there. That scenario doesn't get to me. People all need to understand that the big phenomena here is the Internet and how that's going to transform. It's really a bits to atoms transformation that we all do. If I haven't said it already, I don't believe books are destined to be forever. That's a matter of screen. That's not necessarily the best interface ever.

One of the projects I'm on is how do we cope with the paper kind of thing. Paper is a wonderful media in a lot of ways. It's substantially better than a pilot. I don't happen to own a pilot. I use paper for those functions. I can see lots of things. What is the appropriate technology there? I just came back from a week in Japan and was thinking that we really have to solve that interface problem otherwise the growth can't happen the way we want it. We've at least got it so the computer understands Kanji. We kept Japan at no growth because of ASCII. Mainly, you try to cram 20,000 characters into an 8 bit word and it doesn't work very well.

Moderator: Let's talk about that interface problem.

Panelist: I personally believe that it is significantly a paper interface. I'm trying to trigger a lot of work that was aimed at products so that we can get better use. There's a tremendous amount of faxing going on within an American corporation and there's about 60 e-mails. People are either doing Hiragana or Katakana depending on what kind of a keyboard they have. For me, that's a bear to do. They'd rather scribble. Computers can't understand the scribbling. I say, don't worry, the computer doesn't need to understand that. Use the fax, use interfaces there. Paper is a wonderful interface to buy, you should just never have to store it or transmit it. Once you get that idea that it's really a screen and it's a temporary computing media, then good things can happen. There will be more paper in the world.

Moderator: Let's talk about the interface problem. We're in a world where most users are stumped by managing the information environment on their standalone PC. We are now making available to them the collected wisdom and ignorance of the world. Statistics say that with the current rates of growth, every man, woman and child in the world will have a Home Page within three and half years.

Panelist: I don't think it's that fast. Maybe another decade or so.

Panelist: Yes, but the bottom line is that there are people who can't navigate their desktops and we're giving them the world. How are we going to cope?

Panelist: Frankly, it's going to go via TV. To me, WEBTV is the kind of answer that if all of you who communicate or want to communicate with your parents or PC challenged friends, we'd just go buy them WEBTV. You'll find that it's significantly easier than those VCR's.

Panelist: The number of people who want to watch *Baywatch* demand interactivity. No, not that kind of interaction.

Panelist: No, it's not interaction. I don't regard getting the news when I want the news or a movie when I want the news or being able to get the weather as interaction. I get all the news now. I stopped taking the newspaper a year or two ago simply because I prefer it on-line. I get it at the font size I want, the speed I want, I read exactly all the things that I need. I put myself a little bit ahead, possibly a little leading edge there. But I get what I want. Basically, everyone wants that. They don't want to be bothered with these computers. They just want to get to a room or make a reservation or order a book, until books go out of style, which will probably be in another 20 or 30 years until we get a good display.

Panelist: One of the ways you can ease this transition is to build in intelligence in the actual appliances as they are used. So, a telephone that can look up phone numbers on the Internet becomes more useful than one that can't, or even one that can dial a number and ask someone. Those kinds of things are a driving set of entrepreneurial ventures going on doing the Internet interfaces that are sort of transparent to the user. And use one form or another, either Direct Internet Addressing or Java apps or whatever. Build it into the electronics.

Panelist: When do you think my dog and my dishwasher are going to have IP addresses? Is that something that is a 5 to 10 year phenomenon? Or sooner?

Panelist: They probably do, in the sense that they are probably in databases that have IP addresses already.

Panelist: But you have to have a need for them to communicate. Need has to drive this thing. You have to be able to buy this thing. You have to be able to install it. The reason why I say I am the typical user is that my patience for this stuff is very low. I got two cameras which one has already been given away because I've got a non-monogamous PC. Plugging that one in happened to blow the machine so I said it had to go away.

Panelist: I can define the need very easily. My advice is get married. When I was single, my clothing was probably 30% more expensive because I would mix colors and they would bleed and I would shrink things. I'd love to be able to throw it all in the washing machine and have it figure out what color clothes I have and what temperature it needs to be. It's to a point now that some of the jeans manufacturers believe, merely on step prevention alone, they can justify putting a microchip in every pair of jeans they ship. We're getting very close to this kind of stuff.

Panelist: Usually this stuff comes as things transition from wants to needs. The things that are very want-oriented tend to take a long time because people won't pay for them. The ones that are needs-oriented. When the PC became a needed part of business, it really took off. As long as it was a fun thing to have, it wasn't as big. The things you're talking about are in the same category. You'll probably live without the washing machine that can sort out your clothes for you for awhile, but you may get to the point where it's almost impossible to get by without being able to get phone numbers looked up by your cell phone.

Panelist: I just went through some home modifications. Actually, Dave Cutler, a friend of mine who built NT, just built a house. I asked him what he did with his house. He told me there were no square windows in it. I asked him if he did electronic controls and he said, "switches." No central control. I looked at this thing and this thing has got a programmable thing you can connect to the PC and he said, "switches." I'm not going to have another interface in there. I have another pass at the Motorola PCS cell phone manual and I have yet to get my PCS Motorola cell phone going. I hold this up as the worst human interface. It dwarfs anything you can ever say bad about those VCR things. I have yet to program it to get a number to do anything in any suitable way. I hope those guys have to program if they go to heaven or hell, whatever. I hope they have to read the manual they've created to get numbers, get anything into the thing. This thing is beyond belief.

Panelist: Try the Ericsson one.

Panelist: I'd be happy to try anything.

Panelist: If you know Dave Cutler, you can confirm or deny the story for me. All of you know *2001: A Space Odyssey*? You know Hal was one letter removed backwards from IBM. Take VMS, roll it one letter forward and you get WNT. That's how NT got named.

Panelist: Is that right?

Panelist: It makes a good story.

Moderator: I'd like to thank the panelists and you the audience for attending this session.

Thanks.

Graphics not available at time of publication.

Chapter 16: Panel Discussion — Internet Appliances: Are You Ready for the Cyber-Consumer?

Moderator:

Dale L. Ford
Principal Analyst
Semiconductor Application Markets Group
Dataquest

Panelists

Kevin Fielding
StrongARM Product Line Manager
Digital Semiconductor

David Limp
Vice President of Consumer Marketing
Network Computer Inc.

Raj Parekh
Vice President and General Manager
Volume Products Group
Sun Microelectronics

Jay Freeland
Vice President of Strategic Marketing
Spyglass Inc.

Van Baker
Principal Analyst
Dataquest

[listed in book: Farid Dibachi-Diba,]

Moderator: My name's Dale Ford. I'm an analyst at Dataquest, and I'll be moderating this first session of our afternoon panels. I'd like to welcome all of you back for this session — those of you who haven't rushed out to the telephone to close your positions in the market.

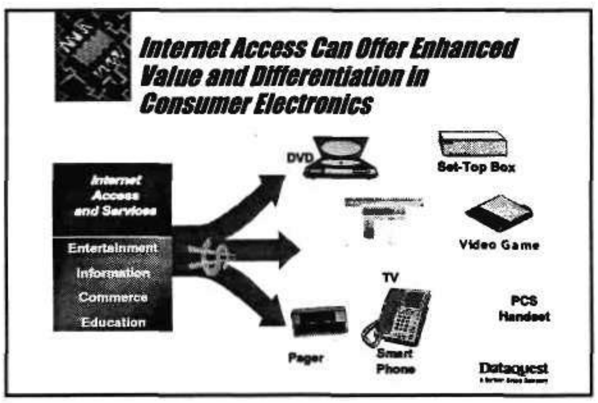

The topic of our discussion today has permeated this conference, and it's probably permeated many of the discussions you have had: in your work, and in your business decisions that you are making, and the plans you are trying to make for the future. I'm very pleased that we have the panel members that we do today; they represent a spectrum of key players in this industry and will share their insights and knowledge with us at this conference.

One of the key issues in making this panel a success in your mind will be for you to take an active part in the discussion and ask questions. The format we will follow is this: each speaker will take 10 minutes to present the key issues, as they see them, in this market. After each speaker's presentation, I've asked our highly trained professionals to go up and down the rows to collect your questions; this way you'll have five opportunities to ask questions and you can ask them when they're fresh on your mind after each speaker has given the presentation.

Let me quickly frame the topic of this panel. As so many have discussed prior to this session, the Internet is becoming a major driver in the high tech industry, presenting opportunities for software developers, for telecommunications suppliers, for Internet service providers, for equipment manufacturers and for semiconductor manufacturers. We've already seen the impact of the Internet in the PC, which is where the predominant impact has been today. Now we're looking forward to the possibility of leveraging the Internet and what it offers into other consumer-oriented products.

Back in January, at the winter CES show, executive after executive paraded to the stand, first lamenting the slump that the industry had experienced during 1996, and then discussing their plans for reinvigorating their industry and their companies as they moved into the future. The Internet sits at the core of all of those strategies. In fact, we've already seen many of the early execution plans as these companies try to move into this area. Many of the consumer electronics manufacturers, who have derived their revenues and profits solely from hardware, now recognize and are attempting to move into a different realm, and derive profits from the software and services segment, and not rely solely on the hardware side.

We're already seeing some very interesting dynamics in the consumer electronics industry, similar to the razor and the razor blade, where you essentially give away the razor and you make your money on the razor blades. We've seen incredibly cost-reduced products in the market as these companies try to shift their business models and tap into the services and revenue aspect of the Internet.

Indeed, the Internet can deliver value to many platforms beyond the PC. This provides a symbolic representation of the potential that exists in the market. We can range from products that are solely dedicated, like a separate and distinct box such as we've seen with the Web TV device, that would sit on top of your television, to a number of devices which we call Internet appliances that are not obviously an Internet appliance. Rather, this may be the television that now has Internet access

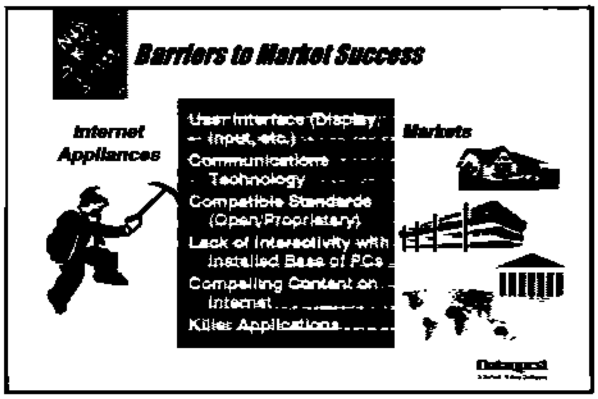

capabilities embedded in it, or it may be the telephone in your home that now has a screen and the ability to interact and derive value from the Internet.

This creates an interesting challenge for measuring the success of this market in hardware terms, because we can't go out and count separate, distinct Internet appliance boxes per se, and say this is the Internet market. Rather, it will come into many other products that we already have in our homes today.

In fact, Dataquest, in an attempt to create some type of structure for measuring this market, has created three separate categories: Internet televisions, which include either a set-top box such as the Web TV product, or a television with Internet access embedded in it; Internet telephones that include both wired and wireless telephones, such as the Nokio 9000 communicator; and Internet data terminal devices, which can also be either wired or mobile, and are separate from the prior category, in the sense that they do not have voice capability associated with them. We do not include the Net PC in our Internet appliance category.

Those who are seeking to develop the market for Internet appliances have clearly identified many of the key markets they are pursuing, ranging from the home, to businesses, to government and education, and penetrating into the international scene. However, there are significant barriers between where we stand today, and where we want to go tomorrow, with important issues related to the display and the input and output of information.

The entire user interface issue remains a very open question, as well as developing solutions that will be consumer friendly and adequate communications technology for delivering the content and providing sufficient interactivity.

Standards continue to emerge and play a critical role in shaping all of the new emerging markets that we see in the electronic marketplace. Certainly the Internet, as well as the underlying semiconductor technologies, are no exception.

The issue of lack of interactivity, with many devices today with the installed base of PCs, or lack of interactivity with other Internet-oriented devices, or lack of interactivity with more than one Internet service provider, stands as a challenge for this space.

Fundamentally, we need to look at the need for compelling content and the search for the killer application that will drive this market forward.

Having quickly framed the topic for this panel, we'd like now to turn to the presentations from each of our panelists, who will provide us insight on issues ranging from the underlying semiconductor technologies to the embedded software to the systems in this market.

Our first speaker today will be Kevin Fielding from Digital Semiconductor. Mr. Fielding is the Group Manager for the Digital Semiconductor StrongARM product family. They have marketing operations in Palo Alto, California, Austin, Texas, Hudson, Massachusetts, and in Reading in the United Kingdom. Mr. Fielding joined Digital Equipment in 1989 as an IC designer and later became the Product Marketing Manager responsible for Digital's Alpha microprocessor chips.

Kevin Fielding: I'm not here to talk about Alpha in the embedded space, you can rest assured. I'm going to attack this topic from the point of view of the

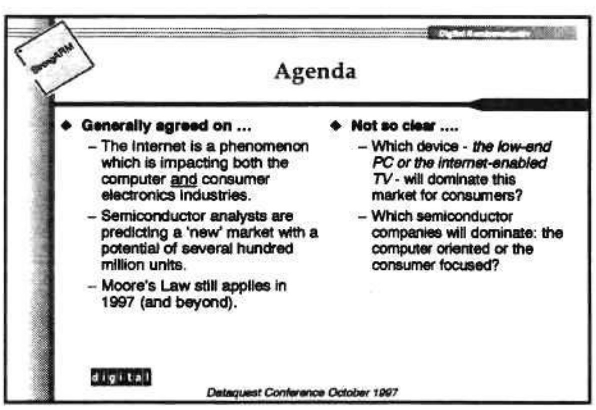

underlying silicon, and the implications of Internet access appliances to silicon suppliers. I'm going to try and hit on a couple of topics. Let me first hit a couple of bullets that I believe we are all in fairly general agreement on.

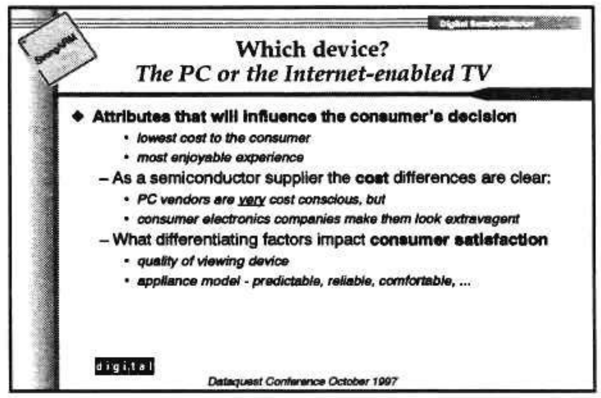

First of all, the Internet itself is a phenomenon that is here, it's there, it's everywhere. It's real, it's very useful, it's providing real utility. The second point is that companies like Dataquest, and several of the other analysts, are pointing to the huge opportunities that the emergence of Internet access devices will provide to semiconductor suppliers. Again, that's broadly agreed upon. The third point is that Moore's Law still applies within this market, be it consumer- or computer-oriented.

The two points that I would like to address a little further are, first of all, there's a raging debate about which of the appliances, an Internet-enabled TV, or set-top box-type device, or a NetPC, will dominate this market. The second point that I'd like to go into is taking that down a level, to the semiconductor suppliers who will again do best in this emergence of Internet appliances. Will it be the computer-centric, more traditional suppliers or the people who have traditionally been very strong in the consumer area? Those two points are still very debatable. I'm going to take a stance on them and leave myself wide open for questions later on.

First of all, in terms of considering whether it is going to be the PC or the Internet appliance, we within the StrongARM Group have been engaged over the past 18 months, almost two years now, with a number of different companies, be they computer-centric, PC-oriented companies or consumer electronics companies. We're firmly convinced at this stage that the consumer electronics, the TV type approach to this, the appliance approach, will certainly dominate. There are a number of reasons.

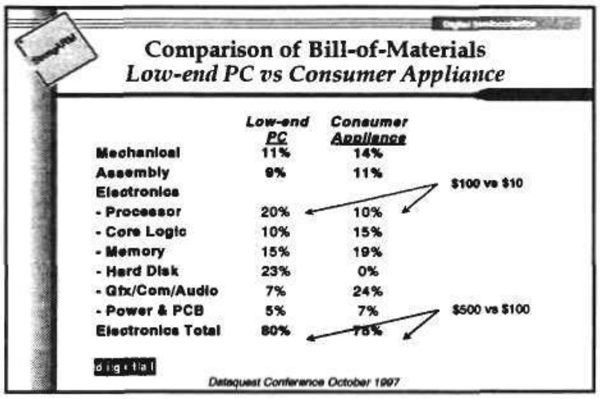

Let me select two of what I believe are critical attributes of system level products in this market, and critical attributes with the consumer, the mass market buyer in mind. The first of those is cost, and no surprise to us here, the cost is a huge, most important attribute when it comes to consumer products. The second I'd like to address is the user experience.

Let's talk about cost. Certainly, as a semiconductor supplier and a supplier of CPUs to both the computer industry and the consumer electronics industry, I can attest to the PC industry being very, very frugal in terms of how much they'd like to pay you for the various components that they need to put together a PC. Obviously, their margins are pretty thin as well. If you consider the PC vendors are frugal, then in comparison, the consumer electronics companies don't bear description.

Let me make a couple of observations. What I've done is compare a couple of low-end PCs, where I've aggregated the bill-of-materials for the two. On the right-hand column I put together a bill-of-materials of three consumer-oriented companions to a TV. One of those is a Web TV box, one of them is a video game console, and the third one is actually a cheap video conferencing solution targeted at consumers. What first strikes you is the similarity in terms of how the money's been spent in the bill-of-materials. Overall, to be honest, there isn't a huge difference between how much those two types of firms pay for the various types — assembly, enclosures, packaging, the CPU itself, the memory, so on.

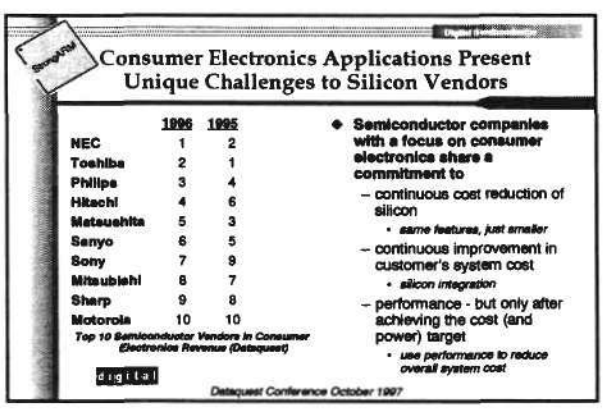

Two things do stand out. One is the bottom line bill-of-materials for a PC. The most aggressive kind of bill-of-materials is going to come up to about $500, whereas the consumer appliance will hit a bill-of-materials more in the $100 or below range. The second thing that's quite striking is the amount of money, as a percentage, that the PC vendor allocates to his CPU. That generally ranges from 20% to 25% of the overall bill-of-materials. On the other hand, the consumer electronics companies are more willing to spend about 10% or less.

What this translates to, and this is very disheartening for suppliers of CPUs to the consumer electronics industry, is a difference between a price of about $100 for the cheapest PC CPU to an average price of about $10 for the CPU that goes into a consumer appliance. Certainly an independent observer might say that's fine. If you can build a chip for $10 and sell it for $100, you can make as much money if you build the chip for $1 and sell it for $10, provided you have an order of magnitude better volume.

Certainly there's no question that volume is there. The markets in the consumer space obviously are several orders of magnitude greater than the PC industry. However, the chips that are going into these appliances are not built for $1. In fact, the actual build cost, the complexity, the performance, the power and attributes that go with these processors, be they in the right-hand column or left-hand column, are actually fairly comparable.

Let me give you a couple of anecdotal examples that'll back that up. First of all, let's look at the Sega Saturn, obviously a huge success in the embedded space as a video game console. That particular chip doesn't give you only one 32-bit

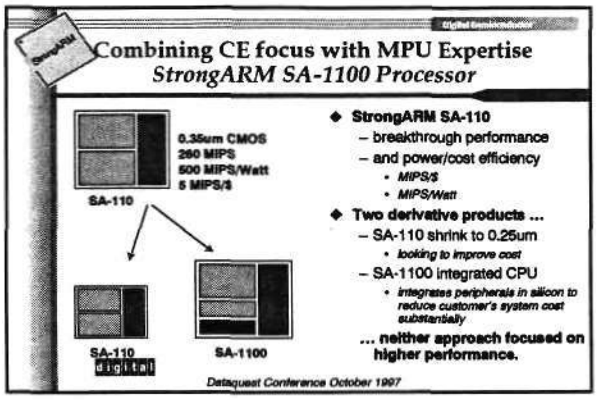

microprocessor. It actually contains four very high performance 32-bit processors.

Let's take another example. About a year ago, when Apple launched its Newton product — Newton is one of the few products that Apple targets at the broad consumer market. At the time when they launched that product — let me ask people — how would you say it ranks against all of the other computers, from a performance point of view, that Apple shipped at that time? Do you think it would have been half the performance of the others? Maybe as high as 70% of the performance of the other systems that Apple sold at the time? It might surprise you to hear that when it shipped, the Newton was the highest performance machine that Apple shipped, bar none. You can see this consumer market turns things on its head.

Let me give you third example and this one relates more to the consumer side of things. Intel three weeks ago launched its most aggressive Pentium-type product for the mobile, the notebook market, the Tillamook. The Tillamook comes in at about a 200 megahertz processor and is fairly extraordinary in performance, and it does that with a reduced 1.8-volt battery power source which gets it to about 3.9 watts, which is under four watts of power for a very beefy microprocessor.

A lot of people have talked about these hand products, like the Pilot. Let me point to, as an example, the consumer product like a Pilot or HP 320LX, versus a notebook computer. The chips that are going into the Pilot and the 320LX, and that class of device, Windows CE devices, come in at about the same performance as the Tillamook. Take an example from my own camp; StrongARM comes in about 200 or 233 megahertz. That same product, with the same performance, comes in a factor of 16 lower in price and a factor of 16 lower in power consumption than that state-of-the-art PC chip. A factor of 16 in power means you halve the power, you halve the power, you halve the power, and you halve the power, and only then are the consumer system people happy.

Let me talk about a second aspect, which is less on the cost side and more regarding the user experience. Let me talk about two things. One is in terms of the viewing quality of these devices, PC versus a TV. For $100 today you can buy either a 15-inch VGA monitor or a 25-inch color TV, and you'll see, if you compare 17-inch versus 27-inch, from an equivalent price point view, dollar for dollar, you're getting three times the viewing area on a TV as you are on a PC.

The last points, in terms of appliance model, that are going to be fundamentally important to the PC model versus the appliance model. There are three things that I find are fundamentally important, that make something accessible, usable and enjoyable to the consumer. It must be instant on and always available, it must be intuitive and easy to use, and it must be predictable. Every one of them must be fairly similar.

That model doesn't apply to a PC today. It does apply to a phone. If you walk to the hall and eventually get through the line there, you pick up the phone. First of all, you hear the sounds as it twirls away for awhile and then it says we're making your connection to your long distance carrier and it twirls away for awhile. By the time it comes back and tells you it's about to load Windows 95, you've walked away from that. That is not a usable model.

The last thing is, who's going to win from the semiconductor supplier point of view? Is it going to be the traditional suppliers of consumer electronics products? I've listed the top 10 on the left. Is it going to be more the computer-oriented people, the Intel's, the AMD's?

In fact, we believe that it's going to require a mix of best practices from both of these companies in order to be successful in the consumer space. If you look at the emergence of consumer electronics products, DVD players, satellite and cable set-top boxes, digital cameras, digital camcorders, all of these products have, predominantly, 32-bit high performance processors in them. Certainly, there is a place for the performance-oriented computer people. Equally, it's going to require a extremely strong look at cost reduction.

Last point, and this is an example, to give you a balanced view, of how it's going to require both computer and consumer focus. When the original StrongARM product — and indulge me if I use an example again from our own stable — when the original StrongARM was launched a little over 18 months ago, it delivered fairly significant performance, best in class performance at that time. In fact, it achieved about the performance of the original Cray I supercomputer, all embedded in one chip, consistent architecture top to bottom, 32-bit architecture, very much what you'd expect from a computer-oriented company.

About 18 months from then, according to Moore's Law, normally what you'd expect of a company — like a Digital or an HP or an Intel — is twice the performance, doubling the caches, adding better branch prediction, adding out-of-order execution and so on. In fact, what we've done and what we've been forced to do by consumers is not quite ignore performance but make cost reduction the primary issue. What we have done is made two different chips, one of which only shrinks the device and halves the task, so you still get the performance of a Cray I, but now you get it for $10.

The second approach is keep the cost roughly the same, in terms of the device cost, but help to substantially reduce the cost to the system vendor, by integrating quite a bit of functionality on the chip. Again, extremely different from a more traditional computer approach. It's not going to be either the computer semiconductor companies or the consumer-oriented people, but rather best practices in both of those firms, that will eventually be successful here. Thank you.

Moderator: Our next panel member is Mr. David Limp. Mr. Limp is the Vice President of Consumer Marketing at Network Computer — NCI. He's responsible for marketing products to the consumer retail, cable and satellite marketplaces. Additionally, he manages all consumer content partnerships for NCI. Before merging with NCI, Mr. Limp was the Director of Marketing for Navio Communications.

David Limp: About 14 months ago I sat behind a mirror in Minneapolis, looking at 12 angry-looking consumers trying to figure out how to work a brand spanking new Netscape Navigator that had been ported to a television. The results of that focus group were, "Don't produce it; we won't buy it." About 12 months ago I watched another focus group in Tampa trying to work with a brand spanking new version of Netscape Navigator ported to a very cool prototype of a digital phone. The consumers looked at it, and they said, "Don't make it; we're not going to buy it." Why, after the write-off of a couple of weeks ago, did Web TV get bought for $700 million, and why did Navio get bought for quite a bit of money, when none of these devices that we're talking about here, none of these appliances, are fundamentally going to work in the market?

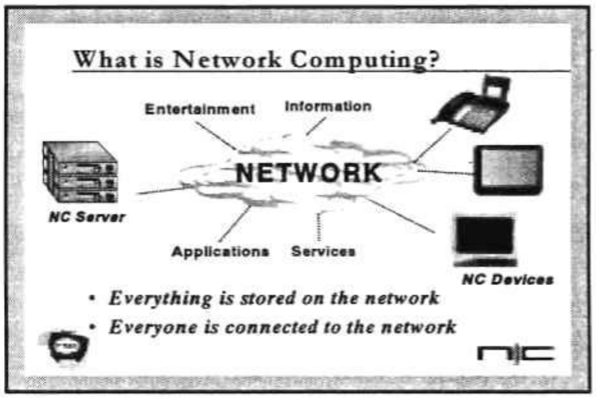

The issue that I want to talk about today is why I believe there is a shift happening and why we still think that there's an opening. I'm going to use as a case study television, but I believe you can move this to any other device, a phone, a game console, a cable set-top box. They all work similarly, when you start thinking about them correctly. The thing you need to think about is that technology only matters to the OEM. In the end, NCI is an OEM company. I want to have a brand like Intel inside. I want to have a brand like Dolby.

The fact of the matter is that we sell to the list consumer electronics manufacturers that were listed up on the screen a second ago. The Sony's, the Matshusita's, those types of people. They care about technology. They care about cost. They care about what goes in the box. They care about the chips that you're worrying about producing, the microcode, the embedded OSs.

The consumer could care less. The consumer cares about simplicity, cares about content, and that's it. Every time a device has been successful in the computer realm, it has had content driving it, plain and simple. The telephone — the content was being able to connect to your friends and family easier than ever before. The radio — it was Sunday night talkies. The VCR content was rental videos. It's always about content. That's truly what we've begun to look for. Our newest products and what we do, and what I believe you'll see from people, like Web TV and others on the panel, is that the focus is going to move away from a consumer focus on technology and what is WYSI in the box, to making the device more simple, making the device more intuitive, with better access and cool innovative content.

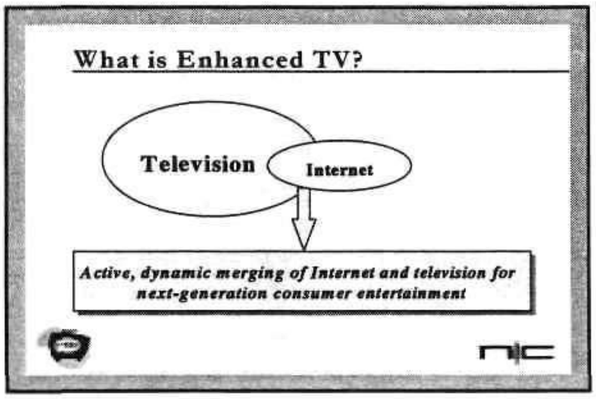

That said, I also want to say that the technology is still important. We're still producing browsers that fit into 500 kilobytes. We still want to make sure they're portable to every kind of processor that's out there because, make no mistake, the industry that we are all in has quickly become a subsidized industry. Subsidized industries, as any of you who come from the cellular phone market know, are tough businesses.

There are set-top boxes that are being sold at retail right now that are selling for under cost. It's a very, very difficult market, and we're beginning to have to drive that. Where's the money going to come from? The money's not coming from people who are simply throwing it in. The manufacturers have to show a business model; they have to show shareholder value, to get that money. They begin to show it through things like content, added-on applications and ways to get into some of the premium services as we move forward. Those are all content-based, so what's going to pay for this is content.

I'll use TV as a case study. I said that we showed to a focus group 14 months ago a browser bolted on top of a television. They yawned. They quickly left the room. They ate a few free M&Ms, and that was about it. The end, though, was that they said if you could somehow mix TV and mix the Internet and use the Internet instead as a vehicle to transport information, that might be interesting because "I could still get my mail, but then I could get interesting kinds of enhanced TV content."

It's important to understand why this is important in the TV space, as well as in many consumer devices. A consumer device like a TV is an active device. It's the user that's passive. The user is a couch potato. This is exactly opposite of the computer that's sitting here. This computer is a completely passive device. It will do nothing until I click on it. Proof of that: when suddenly we got screen savers and fish started running across the screen, everyone got excited. The computer's suddenly active. TVs do that every second they're on, every day, every night — and that's important.

The question becomes, "How do we make TV active with the Internet content even more enhancing?" TV's great. It's one of the biggest industries in the U.S. The content behind it is worth multi-billions of dollars. The telephone industry — multi-billions of dollars of industry. What we want to do is figure out a way to sell a few more features into those industries, and then you can justify the subsidized cost of the value of whatever technology it takes to do this

In the phone industry, if I can sell an additional phone line to a telco, if I can sell call waiting, that pays for the business. In the TV industry, if I can sell someone some EPG services, a program guide, or I can get them to order one more pay-per-view movie because it's a little bit easier to do, that pays for the business. That's what enhanced TV is all about. That's what this software becomes? We begin to figure out ways that you can write applications that are seamless with a television. They blend in with TV, as some of these graphics show. They wrap around TV, and they begin to make the user feel that this is still an active device, it's still a familiar device, but it's something much more approachable and it's something that they can use much more every day.

When I took this to a focus group about six months ago, the difference was dramatic. People were saying, "Is this digital TV?" No, it's only analog TV with some cool Internet content. Suddenly, they're thinking it's digital TV because it's something new, it's something compelling, and they're used to things like overlay. When you watch the World Series tonight or tomorrow night, you'll see they put scores above, they overlay them on top. That could just as easily be an Internet content, and that becomes important as we look forward. Many people quickly tell me, when I show them this case study, "we went through this already." I know... I was around. In 1991 I did interactive TV; I lost my shirt on interactive TV, as many of us did.

There are many differences between what's going on right now and what happened with interactive TV. First and foremost is that the TV companies are interested. Consumer electronics companies had no vested interest in moving their technology at all for interactive TV. Now they're very interested, because they're being mandated, in many cases by law, to move to digital signaling and digital TVs. As was mentioned in the opening statement, the consumer

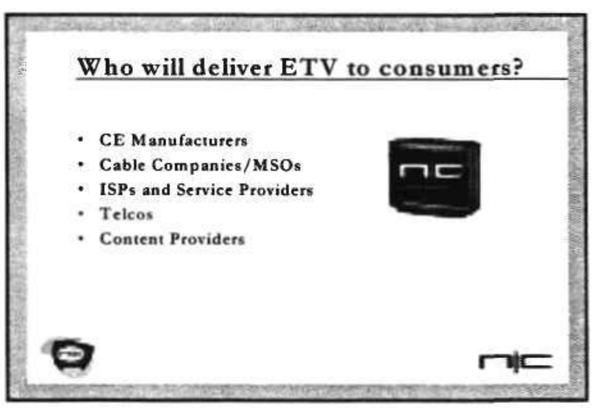

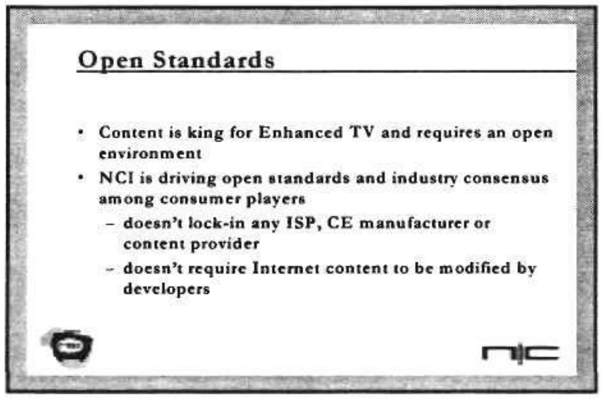

electronics companies are experiencing a bit of a lull in their sales right now, so a new product is something that's exciting. We have the weight of the industry behind us.

The second thing is that when we all experienced interactive TV in the past, any time you wanted to get the content (that thing I just told you about that is going to make or break this business), every time you wanted to make a content relationship with a Disney or a Time-Warner or a Viacom, it cost you about $1 million, $1 million per telco, $1 million per cable company, $1 million per CE company.

There's no ubiquitous standard that made this possible. This has happened in all of the different content areas. You can't imagine if your telephone couldn't communicate to another vendor's telephone. You couldn't imagine if we still had two versions of videotapes. It just wouldn't work. The industry wouldn't have grown as big as it is today. The same thing is true about the content that's going to come across the wires. For a Disney or a Viacom to author something — they want the ability to author once and deliver many.

Our technology's all based on open standards, whether it scales up or down, and this is going to make the industry a success, because the Internet, as a content or an application environment, makes quite a bit of sense on TVs and telephones and other things. People aren't going to be writing, or word processing, or using Microsoft Excel or Power Point on their television. Make no mistake. Those are passive-based applications. They don't work well on TVs, and people that try to market them for TVs will fail.

However, people are going to be interested in chat along with television. They're going to be interested in cool interfaces for pay-per-view, and those can all be written in HTML, Java Script, in Java, languages that are out there today, simple to access and it's easy to get talent out there that understands how to author it. The great thing about going down to a place where animation or other types of things are going on, as at Disney, is that when you go down there you see these people, you see that they know how to use their tools. They're great at making compelling, interesting content. They can do it now on the Internet, they're doing it for their Website anyway. It carries into television and into these other devices very seamlessly.

Who's going to deliver this? The other big problem about interactive television was that there wasn't necessarily a delivery mechanism. For the first time what's happening is we have another dynamic going on in the industry right now. Content becomes, if not free, much less expensive because of the Internet, but competition has taken hold. In many cases, these people are competitors of each other, and they all want to get closer to their customer. What better way to get closer to your customer than have something in the house that you can, for the first time in your life, address?

Until now, it's been the telephone, but now the Internet gives you a second address into the home. You can directly identify the person that's on that Internet connection. You can't do that with a television; you can't do that with a VCR; you can't do that with a DVD player. Everyone's excited about getting that tie to the customer just a little bit closer.

Satellite customers are saying. "If I could just get interactive experiences across my satellite, we'd have all the ability to be able to get closer to our customer, and we could push the telco and we could push the cable company out." The cable company's saying, "I've got this high speed Ethernet-like wire that goes into the household. I can use that and get closer to the customer and get very high speed bandwidth." Meanwhile, the telcos are saying, "We've figured this out; it's XDSL, it's ADSL, it's all these other technologies that we can bring into the house and take this over." As each one of these constituents starts seeing the differences in the competitors, they're driving forward.

The thesis is that this will happen. Consumer electronics companies are very interested, because digital TV is one of the things they must and will do over the next 10 years, and there are 600 million devices to be replaced. Phone companies and everyone else will do this because they see competitors on their heels, and they're looking for that way to get that much closer to their customers. Content companies for the first time see a business model that can conceivably work and pay for this whole infrastructure, pay for the hardware that was just talked about, because they have competition breathing around them in every space, and they want to make sure they can write once and deliver many, so that their content keeps up with all the other content vendors.

Moderator: Our next speaker is Mr. Raj Parekh from Sun Microsystems. Mr. Parekh brings more than 20 years of engineering and management experience to Sun. He oversees the division's research programs, international R&D, architectural innovation and Internet-enabling component to support Java and other networking technologies. As Vice President and General Manager of the division's Volume Products group, Mr. Parekh oversees the proliferation of 32-bit SPARC and Java processors into the emerging embedded network marketplace.

Raj Parekh: Thank you very much. Sun is not in the consumer business. We have not made a single chip, single system, single piece of anything for the consumer market yet, but we believe that the whole dynamics of the market is changing. Internet came around and that provided connectivity everywhere. Unfortunately, pure connectivity is not as useful. You don't know which computer is sitting on the other end of the line, which equipment is sitting on the other end of the line, which operating system it is

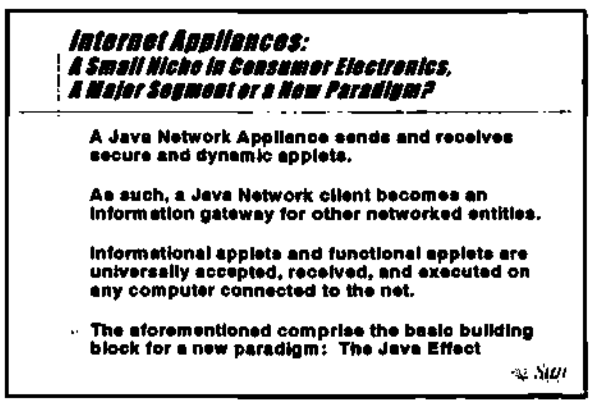

What Is the Java Effect In Technical Terms?

While Java is a monumental technology, it offers a way to change the dynamics of the fundamental way in which business is transacted.

All informational transactions between suppliers and consumers will accelerate.

Boundaries between customers and suppliers, customers and customers, and suppliers and suppliers will blur as they find economic advantages to sharing informational and functional Java applets.

To describe the Java Effect in functional terms, an example is required.

running, which application it is running, which version of software it is running, which version of hardware it is running, have they loaded the latest ECO or not. You don't know. How can you communicate? That's where Java came in. It came from Sun. It was purely accidental. Someone had to do it, and Sun happened to be the company who did it.

Java allows you to communicate from anywhere to anywhere, from anything to anything, and you don't have to worry about which processor it is, which operating system it is, is it of yesterday or is it of tomorrow, you don't have to worry about it. It is all standard, completely open, available to all, including Microsoft. The important thing is everyone has to conform to the spec. By the way, during my speech, if someone feels offended, it is not intentional and I apologize in advance.

I have this presentation. It is already in your book. I will use that presentation, it lasts one minute, but let me just talk about why this paradigm is so powerful and so important and, if you want to play in this market, what are the key issues you need to figure it out. I would say don't look at yesterday. Only study yesterday to understand the history. What is happening in the future versus what has happened in the past and to be very careful. Don't just repeat the history because that is not the formula for success.

Let me give you a simple example of the telephone. I have a little telephone which I carry around in the U.S. It works quite well, but when I take this phone to Japan and open it up and turn it on, it doesn't work. In Europe it doesn't work. Why? Is the technology the problem, or is the problem one of politics?

Actually, the problem is politics, because it is giving a message in Japanese in the first place, which I don't understand. Think about it. If I had this phone and it has Java in it and when I go over there and I turn it on, the first thing the phone will do is look for the nearest site. The nearest site will download a little applet into my phone, find out who am I through Internet, go to my home base, check out whether I am a good customer and paying my bills on time or not. If I am, it downloads another little applet so that, temporarily, it is enabled to work in the local environment. Now I can make the phone call with my same phone. By the way, it will figure it out that this is an American customer, so it will say "Please dial zero before you dial your number," instead of giving the message in Japanese. Very simple.

Now that's the power of Java. There is no consumer device like that, and how much extra would you pay for such things? People will first say, "I can pay a little extra for that." Then when you say it doesn't cost anything extra, people get very excited. The whole dynamics here now is that the semiconductor companies' responsibility is different suddenly. They understand what their customer wants, and their traditional method of charging is to ask for 5% for one value-add and 10% for another value-add, but they also should now try to explain to them the value, because the customer at the end of the day will pay for the value. You can build a chip which either you can sell by weight or you can sell by value. If you sell by value, that's when multiples of stock take place in the industry; that does not necessarily happen when selling by weight.

If you take this simple phone and add a smart card, a simple smart card, to the operation, so that when I put my smart card in, then it becomes my number. I

The Java Vending Machine

The Java Vending machine can periodically broadcast data to suppliers over the net by way of a secure dynamic Java applet.

Conversely, suppliers and equipment manufacturers can query the Java Vending machine by sending it functional specific Java applets.

The company concerned with replenishing supplies in the Java Vending machine executes an applet which updates it's delivery schedules, inventory, and order placement.

Managing supply and demand "to order" in near real time, over the net, represents huge costs saving to the supplier since this paradigm affects nearly every OpCo (e.g, Operations, Sales, Accounting, Inventory, Delivery, etc.)

take that smart card out and give it to someone else. They put their smart card in; it becomes their number. Beautiful concept. When I take that smart card out, by the way, and go home and plug this into the home telephone, now the telephone will ring there. The same thing would happens if I'm in a hotel, and the phone there has a little slot. I plug my card in, and that way it authenticates me, and now my telephone will ring there.

I go to the computer at my job and plug my card in, and it will bring up my environment, but when I'm traveling in Japan, to my local office there, or someplace in the hotel, I plug the card in and in a few seconds it will bring my environment up. That's the power of the Internet. That's the power of Java. That's the world of tomorrow.

What we did here is we did not make funny telephones, or make the telephone more useful. Suddenly, we can have an equipment to equipment conversation. Right now when I call someone, I can call someone's equipment. If that person happens to be there, then it is person to person conversation. Otherwise, it is connectivity between the equipment. By virtue of this smart card, you change the paradigm, and now you are actually going from person to person, and the numbers and all the rest of the things will just track with it.

This paradigm is so powerful and the applications are so large, if you look at all this big equipment, whether it is switches or a piano, or whether it is an industrial production line or a simple thing like a copier or a printer. All of them can have a small processor inside and a telephone line going to it. What will happen is remotely you can get into the equipment, download a little applet, say "run this diagnostic for me and tell me what's happening." Are you running out of paper? "Call me so that I can send more paper there. Measure some temperatures for me. Measure some physical parameters. Do something for me and report back to me." If you look at it that way, we can have this whole service industry done differently. Again, the Internet brings you that. It's not only for the pretty pictures, nor is it only for looking at stock on your PC. The application ranges so widely it is limited by imagination.

The Java Vending Machine
(Continued)

The vending machine manufacturer can query each machine individually and test it for functionality, reliability, etc. Hence, it can view any machine as if in a virtual test lab which allows the analysis of their design in actual usage rather than just in a manufacturing test lab.

When tampered with, the Java Vending machine broadcasts a Java security hazard applet and it is received and executed by a security contractor's Java computer.

The Java Vending Machine's insurance company offers a substantial rate reduction to the company managing the Java Vending machine, because it now has a tangible mechanism to track theft and vandalism. In fact, the insurance company now uses the Java security applets launched by the Java Vending machine to dynamically compute rates moment to moment. A brand new paradigm!

Theory and Practice

While the above example may seem futuristic or far fetched, the Java Effect is already taking place.

CSX uses Java to manage logistic tracking.

AVEX, an electronics contract manufacturer, uses Java to allow their customers to access manufacturing data over the Internet.

Home Depot is using Java to minimize paperwork for interdepartmental transactions.

The list goes on and on.

In the future the companies and organizations will be successful who, on one hand, look at the history and not create $500 processors for the phone but, at the same time, look at what value can be provided and what kind of alliances we must make with others. We should do this so that we can build on their knowledge, their experience, their intellectual property, and combine our strength with their strength, and come up with a uniquely different solution. At this time, there is no existing company today, in my opinion, perfectly suited for the world of the Internet. We have to migrate towards that. How do we migrate? We can either learn ourselves, or we can make the partnerships with others. That is one of the key considerations for the people in the future.

Let's take an example of a Coke machine and see how that works. Again, this is a new paradigm. This is not just a minor improvement of something, and this whole issue of client/server. We can talk about it at a very high level or the server can be anything. Any kind of device can be a server. Once any kind of device is a server, on one hand, this momentum of Java, which is very strong, is taking place. Right now every company who is a consumer electronics company or a computer company or wannabe has

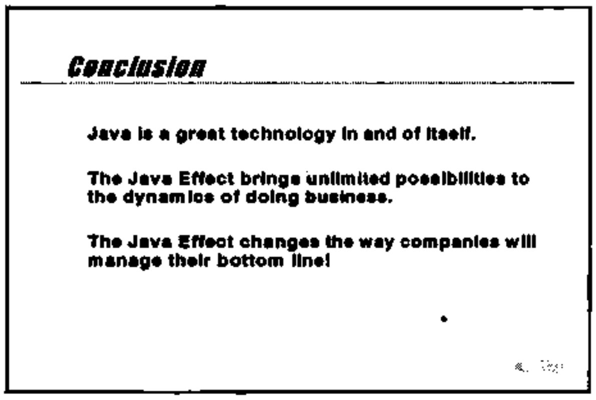

already licensed or is thinking about licensing Java. It is available for pocket change, so it is going to continue to happen. As the transaction is happening. Now the knowledge is flowing at such a rapid pace that one has to keep up with it. That's where the magic comes into play.

We did this Java vending machine. The machine is just sitting there in a mom and pop shop in the middle of nowhere, and periodically just broadcasts the data on the Net in a secure way in a little tiny applet so that the person who owns the machine can look at it and say is this working properly, looks like this connection path is working and everything is working all right.

If I am a supplier, then I will say, "Mr. Machine, tell me something — is the temperature okay?" The machine will tell you. "Are you running out of Coke or are you running out of Pepsi? How many are left?" Before it runs out of it, it sends the signal and sends the command and says "I'm getting ready to be refilled," or "There is too much money here at this time; send an armored car because we've got all this money to be transported." This machine will tell you.

This applet will eliminate all the unnecessary runs of those people. Right now, they open up the machine and see that everything looks okay, they put in some Coke and then leave. They don't have to do that. With this Java applet, they only go to the machine when it is necessary.

The people who are providing that machine can also do the diagnostics, make sure that the situation is okay. If the temperature is actually rising slowly day after day after day, they know it is time to send a service person before it breaks down. Very simple. It eliminates breakdown, it eliminates the cost of non-service. Not only that, but this approach can order more not only from the supplier but all the way to the manufacturer, all the way to the aluminum can supplier or the ore supplier because it is all electronically done. It is done online.

This machine can also tell you whether it is being tampered with. If someone is putting in a dollar bill which is phony, they can call the police right away for you and stop the crime as it is happening, or before the major damage is done. That way the insurance company will feel that for this kind of machine they will provide insurance at a lower rate, simply because it can actually prevent thefts and other hazards to proper operation.

There are many examples like this, and there are many people who are already using this. These are not all the consumer companies, but here are some places where some of the examples are very interesting, like Home Depot for example. Home Depot has so many programmers that are doing so many projects. They have stores all over the place, and they have decided that every line of code they will write from now onwards will be written in Java, and every single piece of equipment they have will not run Windows but will run Java.

With that kind of approach what is the semiconductor piece? The semiconductor piece is we can use any processor and we at Sun are also coming up with a very powerful processor which is purely a Java processor. It's a DSP-like approach. It does only one thing, and one thing well, and that is called Java. Later on we went in and provided C execution as well, because it makes sense so all the communication will always be done in Java but the drivers and things like that underneath can still be utilizing C. We are going for a very powerful approach in this area, and that way we believe, although Sun is not in consumer electronics today, we will be able to enter that area, and this whole paradigm is done very openly so every company has a chance to get into this particular environment, think about the approaches and applications, and make it happen.

Moderator: Thank you, Raj. I confess I have not heard of the smart card application in pianos so I'll be talking to you about that later.

We are particularly grateful to our next panelist. He stepped in here to help us at the last moment, so in your binders you'll see this presentation under Randy Littleson. Joining us today in Randy's place is Jay Freeland from Spyglass. Jay is responsible for initiating new strategic marketing activities at Spyglass. He is the strategic marketing manager and Vice President. He was also, prior to being employed at Spyglass, a co-founder of Surfwatch and a number of other high tech startups, and he also spent five years at Sun Microsystems.

Jay Freeland: I'm reminded of the famous adage: "It was the best of times, it was the worst of times," because I believe we're at a series of intersections in the industry today. As several of the panelists have mentioned, we're running into a situation now where we're seeing, for example, the major consumer electronics companies trying to come up with compelling reasons to provide consumers with new functionality, and yet we're in a bit of a lull.

The move to digital television and the whole concept behind that have delayed a whole set of things in terms of the marketplace. We're in one of those moments wherein people start to worry. That's the worst of times. In fact, the other thing I believe that's happening is, if you look at the PC industry, there are many people that are very proud of the fact that the PC has 40% penetration, but it looks like it's not going to get much more than that. Sixty percent of the population aren't going to use PCs, and various panelists have talked about why that may be true.

What we see from our perspective is the other side of it, which is that it is also at the same time the best of times. In the last year more new content has been created than in the entire previous history of humanity. What's driving this? It's a combination of the Internet and more.

Let's stop for a second and think about how many hours of television content, television programming, are created each and every day. Unfortunately, you could also argue that much of that content is not particularly valuable. I certainly would agree with that, but we see that there is this massive dynamic that's happening, which is that there is indeed new content, and that if you can convince consumers that there is some way that they will either be able to save time in their lives or that they will be able to be provided new information or, also very importantly, that their children will have some kind of advantage, and think about each of you in this room as a consumer, that's what starts to make this seem compelling.

I'm going to talk a little bit about what Spyglass has been doing, and also a little bit about the industry dynamics, in terms of the various places that you might start looking at content and how you might start doing it, and I'll also talk about a couple of case studies of partners of ours and where they're being successful today, not successful in 2002, which is when many of these technologies will have a chance of being deployed.

Everyone has talked about the wide variety of places, such as Home Depot, that are using these new technologies. We've heard about office buildings, a little bit of talk about telephones, televisions, and I'm sure everyone here is familiar with what's happening in automobiles today. If you go to the San Diego Airport, you can rent a Hertz car and it has Never Lost in it. What that gives you is the capability of finding out where you are and where you want to go. That's a compelling application for a consumer who wants to avoid being lost.

Let's talk a little bit about the market characteristics, and this also leads to where we think that this is heading. The world is becoming more and more embedded. Several of the speakers have talked about you having a PC on your desk. That's not an embedded application necessarily. The kinds of applications that are happening are using chips, and are appearing everywhere. Today's average microwave oven has more compute power in it than computers of 10 or 15 years ago. That's a certainty. What's happening is more and more devices are being created, but how do these devices all communicate with each other.

The other dynamic that is very important is that no operating system, no chip — we don't have the same Wintel world being recreated that was created for the PC. Some people argue that this means there will be no critical mass. I would argue just the opposite — that the critical mass will come from a whole variety of applications. The consumer stops thinking about the fact that they're interacting with some kind of computing device. That's the wrong paradigm. Consumers need information, whether it's how long the potato needs to continue to cook or whether it's a stock quote on a cellular phone — from the entire spectrum of human experience. We see that in order for this market to be successful, it's about getting consumers information and getting them compelling content, for whatever they happen to be looking at, at that particular moment.

The other thing that we think is so important is that much of this work is going to be done in the infrastructure where you will never, ever see how the content is created, how the content is delivered. To the consumer, it will appear to be seamless.

The best example that I like to give is if you look at how the industry has existed in the past. I was at the Internet Explorer launch and heard Bill Gates expounding about the Internet lifestyle. His argument was that 10 years from now — and he was very strong about the fact he was going to be around 10 years from now and that he had to make his prediction accurate — we will all be living what he calls the Web lifestyle.

What does a Web lifestyle mean? You don't think of today as living a telephone lifestyle, yet every person in this room is critically dependent on the telephone. In fact, you can argue that, without the telephone, our lives would be potentially substantially better in certain ways, and substantially worse in others. If you look at the telephone, the television, radio, VCRs and personal digital assistance — I happen to have a Palm Pilot. We have a Web browser for the Palm Pilot. I wouldn't say that having a Web browser for the Palm Pilot is a particularly important thing. In fact, I used to swear I would never use one of these things. I was a person who used a paper calendar, and I would refuse to use this because you can smash it on the ground and lose all your data. Most of the problems have been solved.

The next big problem for this is getting connectivity. It's connectivity for information that you want at the right time. I don't see many people using this in their automobiles, with little maps on them showing what the traffic patterns are. That's not probably going to be a use for it. On the other hand, to get information about a particular telephone number or a live update of a package delivery status — let's say you have to trace a package if you're a salesperson. Those are the situations where small amounts of data getting to this device at the right time, and in the right place, could make a big difference.

In terms of what the success factors are, and this is trying to look forward and say how are these things going to play together — we would absolutely agree that standards are critical. The fact that tremendous amounts of content are being created in HTML, or delivered using HTTP, and also delivered using all these other different kinds of Internet-based protocols, does mean that you can start to integrate some of these concepts.

The other thing is that we have to address the consumers' issues. For example, if you go out and ask women on the Web what their biggest concern is, they'll tell you it's privacy. If you ask a parent what their biggest concern is, they'll tell you it's content being available for their children. These are many of the same issues, politics and the like, that have emerged out of other environments. I can't imagine, for example, the world of television, as it moves into the future, without things like the V chip, or imagine not having 976 blocking on a telephone.

Think about the past as it applies to the future, and you'll start to see where consumers have said these are critical issues to me. I'm not going to put my credit card in, I'm not going to give you information unless I know this information is secure. It has to go back to the question of why I care. I'm a consumer. Why do I care about this?

One of the things we were talking about in advance of this particular session was the question of why, for example, is someone going to watch a combination of television with the Internet. There's some entertainment value, there's some service value to it, they can get more information. The World Series example was a very good one. There was a major debate at NBC. They were not going to put that little box of information in the corner. That little box of information has become very important to them. It tells you what the score is. We see the whole layer of services that can go on top of this as being driven by these Internet standards.

Finally, the other thing that's very important is this has got to fit into the technology pricing curve of the consumer electronics vendors. Kevin put it extremely well. We don't believe that the PC TV is the paradigm. You're not going to boot your television. It's not going to happen. What's going to happen is you're going to start getting capability. Suddenly, it's merged. There's much work that has to get done behind that. Probably the best example I can give here involves a company we work with called Worldgate. They use some of our technology to deliver Internet, combined with cable television.

The particularly important feature of this system, and I'll call it a system, is that it uses existing cable systems, both analog and digital, so it leverages the existing infrastructure; to connect to the Internet takes less than three seconds. In that particular case, you can have an environment in which you're watching something on television, and you see a URL at the bottom of the screen — www.toyota.com. You can click the button on your remote and you can go to the Toyota site to get more information. That's compelling, not necessarily to the consumers as much as it is to the advertisers who want to draw the consumers in. What we've got to do is look at the entire buying cycle.

In terms of Spyglass' approach, what we try to do is integrate core technologies and provide value to our partners, and we work with everyone. We announced a relationship with Nokia for cell phones on Tuesday. We work with people like RCA Thompson and, as I mentioned, Worldgate. The whole idea is, as consumer electronics companies move into this space, they're struggling with the question of how do they get — everyone says I want the Internet on my TV or I want the Internet in "x" device.

What they mean, it turns out, is that they (the consumers) want more information, better, faster, an easier way to live their lives, so they don't have to go to a telephone book. The usage of telephone directories has gone down dramatically since Directory Assistance came along. You can imagine how you can go even farther than that if you have an Internet screen phone. What we've done is approached this from a partner-centric model. What we try and do is provide that technology leverage to our consumer electronics partners.

Moderator: Thank you very much. Our concluding speaker for this session will be Mr. Van Baker. Van is a principal analyst at Dataquest and is leading in the launch of a new service, digital consumer, looking at the consumer markets in the United States and opportunities for these types of devices.

Van Baker: I don't have the luxury of slides, unfortunately. Dale and I got together a couple of days ago and decided it would be a good idea for me to come down and participate in this.

Let me react to some of the information I heard this morning. I'd like to know how many people in the audience think that a consumer is going to get excited about seeing their personalized Yahoo! screen on a gas pump when they go into the gas station? I don't see too many hands coming up. I believe we ventured into the front yard of Fantasy Land this morning, in terms of carrying this thing forward, and that's the watchword that I'll have — not that I'm going to pretend to come up here and be a nay-sayer to all the visions that these gentlemen have outlined. We're very supportive of it. I just would caution you to keep in mind who you're selling to, how they use it, how they feel about technology.

The consumer is overwhelmed by technology at this point in time. I'll give you an example of some research that we did, and then later corrected, a while back when we first were kicking off some of the research in this program. Last fall we performed a broad survey, asking our sample population if they owned a digital camera. We had over six million households respond" "I've got a digital camera." How many people think there were six million digital cameras in the market last fall? Certainly I don't.

To measure the level of confusion, we went back to all of those same people and we said, "Just to reiterate,. a digital camera is a camera that does not use film but captures images electronically for downloading later to a PC. How many of you have digital cameras?" About 5.3 million of that 6 million went away, saying, no, I don't have that. There's huge potential for confusion in the consumer market. They think it's a digital camera if it's got an LCD display on it, because those are digits.

We do measure some standard things, like PC penetration in the marketplace. When we first started this, in mid-'95, we had a little over 27% penetration. Last fall we got to 36% penetration. We are interviewing, as I speak, for this fall's data, but preliminary indication looks like we're going to get to 41%, 42% penetration of PCs into households. Last fall we had about 17 million, of the 35 or 36 million households that were out there, accessing the Internet on an ongoing basis, so about 17% of the households in the U.S. — we know that's going up dramatically. I don't have any preliminary numbers on that at this point in time.

Certainly technology is becoming much more pervasive in the consumer space, and maybe one of the most telling things — we went out and asked consumers — we presented them with a statement that said that a household that does not have access to a PC will be at a significant disadvantage in the future, and asked them to agree or strongly agree or strongly disagree or disagree with that on a 5-point scale. If you look at the people who agree or strongly agree with the statement, the responses include 75% of the total asked the question. Theoretically, penetration could be 75% of the market as of last fall. It may go up again as we interview.

Certainly the benefits of technology are there. The penetration is increasing and the consumers are interested in this, but they're confused by it. They're overwhelmed by it. Most of them are nervous about it. They also tell us they're concerned about the impact technology has on their lives going forward, and that's true not only in the non-PC households, but in PC households, as well.

What we're talking about with these devices is the convergence market, if you will, the coming together of computing technology and consumer electronics technology. Some of the things that we think are important to pay attention to in these markets is that what we've got are two different markets coming at each other. One of these is very data-centric; that is the PC and computing base which, even in the home the majority of this is far and away a data-centric, a work-oriented set of tasks going on with this device Versus this data-centric market is an entertainment-centric set of devices, which is a much less interactive environment, a much more passive environment versus a highly interactive environment on the data side. That needs to be taken into consideration.

I was pleased to hear some of the comments from these people relative to the importance of paying attention to the environment in which these devices are going to be used. Social context is hugely important in this environment. People use devices very differently in the living room versus the home office versus the car versus the workplace, and you need to take the social context of each one of these devices into consideration when you do that.

We've got two different industries, if you will, a consumer electronics industry and a PC industry, that have hugely different business models that are approaching this market. We've got the consumer electronics industry that introduces new technology at a relatively high price point (in the CE industry approaching $1,000 is a high price point), and they ride the price point down and have long life cycles, and they don't typically hit very high volumes until they get underneath the $500 price point. The PC industry, on the other hand, has had a history until just recently of keeping that price point constant and refreshing the technology on an ongoing basis to keep the price point high. Very different approaches to business. That has to be reconciled in some form or fashion in order for consumers to be able to realize the benefits of this convergence.

There are some accelerators and decelerators taking place in this. A big decelerator is bandwidth into the home, from the curb to the home. There's a huge problem with that. We've got two industries that are competing or vying to provide this, the telcos and the cable industry. What we've got is first the telcos saying "we do voice, don't worry about data, voice is our business, that's what we care about." Second, the cable companies and they go "we do video, don't care about data, that's not our business."

Now we've got two infrastructure alternatives here, neither of which seems to be very interested in data delivery to the home. Admittedly, people are trying to skew that in the marketplace. Microsoft runs around and drops a billion dollars here and a billion dollars there and gets the valuation of the stocks up and maybe the cable companies are a little better prepared to invest in the infrastructure, and they're motivated for a different reason.

In order for the cable companies to compete with the satellite companies, in terms of delivering numbers of channels and quality of broadcast, they've got to upgrade their infrastructure. As long as you're doing that, you might as well upgrade to two-way and be able to provide data delivery infrastructure in place. Cable looks like the most likely at this point in time, but it's a very slow process. We're not going to get big increases in bandwidth from the curb to the home for quite some time, for a few years yet.

A potential accelerator that's not in place right now, but will be in place sooner than we'll get bandwidth from the curb to the home, is the home network. However, it's an issue of having the infrastructure in place and being able to do it. Everyone in this room is from this industry. How many people in the room have their home wired for a network? There's about four, five people out of this group, and we're in this business.

I would suggest to you that networking via wired connections is probably not very likely in the home. It's going to have to be via power grid, which is, in a sense, a wired connection, but there's a real issue of bandwidth on power grid, or it's going to have to be RF, which is doable at good bandwidths but it's very expensive now, so the price has to come down. That has the potential to completely stand this on its head and to completely change the convergence pattern, but it's not here yet. We do think it will come, but it's still a couple years out.

Thinking about all of this, let's have a little look at Web TV and their evolution, and how some of this applies to that. When Web TV first came out, we were very negative toward it. We had survey data that showed that hardly anyone even knew what an Internet set-top box was. Very few people were interested in implementing Internet on their television, and we were very skeptical on it. The main issue we had with Web TV was their positioning. Would you like to get the Internet on your TV? "If it costs anything, probably not," is what most people said, and Web TV had some serious issues of getting people to adopt the product, but they learned.

They learned partially with some help from competitors, in the form of Net Channel and the folks from Navio and Oracle. They repositioned their product as an enhanced television experience. The user doesn't care whether something comes over the Internet or not. Yes, it's a buzz word. Yes, there's much interest around it. Yes, many people access it. Do they care how the Internet works or do they care that it is the Internet? No, they care about the content that they get. Web TV repositioned themselves as an enhanced television experience.

They also attacked the bandwidth issue by putting a hard disk in the Web TV Plus product. Suddenly you could send quite a bit of data down, cache it on a hard disk, and now as you switch from page to page to page on the Internet, it's delivered very fast — a satisfying experience for the user. They accomplished this through not only the hard disk but also implementation of vid modem technology, which lets them send data down using the overscan portion of the broadcast signal over cable and, in essence, refresh this gigabyte worth of information every night. It's not a huge pipe; but it's significantly better than a modem. It's four or five times faster than modems out there, but the fact that this trickle is going on constantly lets them refresh an awful lot of content, so they're getting at the issues that the consumers care about, which are responsiveness and content. They are positioning this as an enhanced television product.

That's how we look at these products as they come out and how the consumer reacts to them.

Relative to prospects in the future, in the near term we're relatively bullish on Internet telephones. The reason why we're bullish on them is that we think it's going to take a cell phone model. We think the telcos will use the screen phones to get service revenues from the consumer and, in essence, give them the phone for free. Most of the folks working on the phones are at least in negotiations with the telcos about that, and we think they'll wake up to that, look at it as a revenue opportunity, and implement those that way.

In the intermediate term, as we find home networking possible, certainly home servers and Internet televisions become much more viable products.

In the longer term, as we get bandwidth from the curb to the home, and then with the network computers being scattered and proliferated throughout the house, with all of them accessing things over a fat pipe coming straight into the home, that lets you customize your services and everything. We think that is very viable.

Q: One for you, David. When will non-proprietary set-top boxes be available for smaller ISPs to provide to their customers?

David Limp: I guess you'd have to say what the definition of small non-proprietary is. If you mean that you can put in your own DNS address and put in your own phone number, I know that, regarding my own company, they'll be available early next year from a major consumer electronics manufacturer. Many of the telcos are very excited about this — just like Internet phones taking set-top devices and trying to sell on added services as well.

Q: For you, Raj. Someone's basically saying, "show me the money." After your presentation they're saying how does Sun make money on Java if you're giving all this away for free?

Raj Parekh: Thank you very much for being concerned about Sun's well-being. Think of it as the Ford Motor Company putting the brake on the left side of the accelerator pedal. Is it something to make money? No, but it is to unify all the cars in the same way, so the user experience will be a lot better. Keeping fenders at the same height on all the cars and licensing that — should it make money? No, that simply prevents an accident, with one car going underneath the other and having a serious injury occur.

Java licensing is done for one very simple reason, that it will enhance the user experience. We believe that for our core business, which are the servers and workstations, that's where we make our money. The tremendous migration taking place or the hit coming from Wintel offended us very much. We want to enable the consumer electronics connected into this, and that's where Java comes into play. Sure, we get a little royalty here and there, and we'll make some money on Java as well, but if you wonder, "which is Java going to do, make us famous or rich," I would say it will make us famous. As a result of it, every device needs a server on the other end. That will make us rich.

Q: Why don't we ask you that same question with Digital Semiconductor?

Kevin Fielding: How do we make money? With much difficulty. As you saw from my earlier presentation, the kind of margins that are possible within the PC industry, especially as a CPU supplier, are unheard of in the consumer electronics market. The way to make money here is to be able to obtain the volumes that completely wash away your investment in engineering, and you make your money back through manufacturing only.

Q: Let's take a moment here with you, Jay. One of the key things that Spyglass is trying to do is make Internet content available on a variety of platforms. As you're dealing with that, we've had a brief mention earlier of Java as a critical standard in this environment. Are there other standards that are critical to make this an open environment, a non-proprietary environment? Are those standards issues all settled, or are there some key issues that still need to be worked out?

Jay Freeland: We absolutely think there are some key issues that still need to be worked out. I don't know if most people know this, but today almost every Internet browser passes a string back that still says Mozilla, because many of the Internet services out there want to know that they're talking to a Netscape browser. In fact, it's a little known secret that Microsoft actually passes that string back to lots of different servers.

There are many issues, and one of the things that we've been trying to do — I like to say we now have an application that's the Robin Hood of the content world. It robs from the rich and gives to the poor. If you have a rich PC-oriented device and you've got a Windows CE device or some other small handheld device, what you can do is deliver content, that was originally designed for that extremely rich environment, down to something that's much, much simpler.

We think this is one of the compelling things that will help drive the industry, but clearly there are still issues around standards, ranging from "how do I identify who I am and what kind of content I can take," all the way through "how does HTML and Java and all of the extensions that are going on — how does that play in this world," when you may have devices that are much less sophisticated.

Van Baker: Let me add something to that. It's important to keep in mind that while I believe there's quite a bit of value in being able to convert content on the fly, and make it viewable on lots of different devices, it's important to understand that different devices are going to get used in different ways.

The Internet screen phone is going to be an enhanced telephone. It's going to have a telephone use model. You don't sit down in front of a telephone for two and a half hours and look at screens of stuff. What you do with a telephone is, you pick it up, you punch a couple of buttons, you get your task completed, and then you end the transaction. Internet content that comes to the telephone is going to have to keep that in mind, that the consumer is going to want to get on this, get it done simply and quickly, and have the device be in essence what we call inspection-usable, meaning you don't have to read anything to figure out how to use it, you can just look at it and figure out what to do. They're going to push a couple of buttons, get the answer that they want, and then hang up. You need to keep that in mind as you look at how content gets delivered to these devices.

Q: Let's follow that train of thought into another question out of the audience. One of the audience members mentions that in the presentation it was mentioned that in an Internet appliance you won't need to take time to boot up. They say, however, given the enormous bottleneck caused by bandwidth constraints and requirements of loading, even multiple loading of Java, won't these appliances be frustrating? When do they become friendly enough to the consumer to be pervasive? We're telling people that these are instant-on devices. Are you saying then that the overhead from Java or whatever software won't be a serious bottleneck, won't get these devices bogged down?

Raj Parekh: Absolutely. If you look at things like e-mail, it takes the same amount of bandwidth whether it's done on a PC or on a Java-based anything. It is all the same. The booting piece — most of the Java-based consumer electronics items we believe will boot from the RAM. It won't need a disk drive, it won't need anything, it'll just boot from the memory and practically instantaneously, and then the applet and what you want to do comes down from the wires.

If you look at the application software and see how big the world is versus what piece you use, only the thing you use comes down and it is always fresh, not the entire application all the time. That particular model we think does not put significant stress on the bandwidth itself. We, of course, believe the bandwidth will increase over time and more and more sophisticated work will be done, but to begin with, it is not a big deal. When you arrive at the airport and have this Palm Pilot that tells you where the Hertz rental car agency is so you can go there directly, it is not requiring huge bandwidth, but it's still very useful.

Q: Do you agree with that?

Kevin Fielding: The bandwidth issue is — I agree certainly with what Raj is saying. The issue of performance running Java is a extremely important issue right now for the embedded market. It's not just Java; it's any interpretive language. The exact same issue holds for Inferno. The same issue will hold for Newton. It's basically applying about a 50X brake on the performance of your processor. Good news-bad news. Good news for us in StrongARM. We happen to be taking the high road in terms of performance right now, so it certainly is an entree for us, and to many consumer electronics companies, because we and they are very intrigued by Java, and the Internet-enabled TV people, the Internet-enabled phone people are very intrigued with the use of Java. The good news for us is that StrongARM is one of the few processors that makes it tolerable right now.

Raj Parekh: That's why I believe it will be a very interesting idea for Digital to license the Java chip license. That way you don't have to interpret. You execute and the screen problem disappears.

A: Raj, the problem with that is, of course, Sun itself is bringing out version two of Java while the rest of your licensees are still on version one. If I thought I'd be kept on an even playing field maybe we could talk.

Q: Van, someone wants to know if enhanced TV and/or network computer forecast takes off as we're talking about here, how does this impact the 15% to 20% PC unit growth forecast?

Van Baker: It certainly has potential to impact it because as we move towards digital broadcast and there are so few constraints put on the broadcasters in terms of how they use this six times the bandwidth that they're getting for free — and you combine this broadcast bandwidth with push technology — theoretically, the broadcasters, if they woke up to the opportunity here, and there is much doubt as to whether or not they have their eyes open enough to wake up to this — but theoretically they could use push technology in broadcast to steal the focal point of the Internet, which is what most of the excitement is around today, away from the PC and put it on the TV.

If that scenario unfolded, it could have a large, negative impact on the PC penetration in the household and Internet usage in the household. Also, you simply have to look at the situation as we see these new consumer electronics devices emerge to the market, with the consumers having disposable income. The disposable income that they use for PCs and the disposable income they use for consumer electronics is completely interchangeable, so there's an definite potential to impact the growth in the PC market in the home in a significant way.

Q: One questioner wants to know about the opportunity for the Internet market beyond the U.S. We've talked about penetrations in the U.S. What's the Internet penetration beyond the U.S., and how does that market model differ from the U.S. market model? Let's hear from Spyglass.

Jay Freeland: I believe one of the very interesting things we're seeing is that the rest of the world is playing a bit of catch-up, but it has the potential in many cases to leapfrog the U.S., because much of the infrastructure issues that we're worried about here — for example, curb to house connectivity and the like — would be bypassed because new infrastructure can go in. We see the rest of the world outside the U.S. as being a especially phenomenal growth area and potentially a tremendous market.

Q: David, how do you see the difference?

David Limp: I believe you'll see the biggest TV-based applications be either a satellite like the BSKYB project that's going on in the U.K. right now or cable projects. We're doing one with Viewnet in Japan — will actually install quicker than the U.S. We're finding it's faster overseas and that they, in fact, are leapfrogging other things and deploying right to it. I know for sure our first major installations will be overseas.

Q: One final quick question. Just a very brief answer from each of the panel members. In 10 to 15 years from now, when we look back and the Internet appliance market's taken off, what will speakers in presentations like this be identifying as the killer app that made it happen?

David Limp: The Disney Channel. It will be a different kind of Disney Channel.

Q: Raj, what's the killer app?

Raj Parekh: The killer app, in my opinion, will be a state of mind: you won't even know that you are living in the Web Age or the Internet Age. It will be all around you, but you won't consciously think about it, you'll just use it.

Jay Freeland: I'm going to agree with Raj on that one. As much as I hate to agree with Bill Gates, I believe this Web lifestyle concept — well, you don't think today about making an 800 number call as being the telephone lifestyle. I believe it's going to be an integrated part of our lives and maybe watch out for video mail. I believe that could be another real killer app.

Q: Kevin?

Kevin Fielding: Unfortunately, I can take the high ground on this one. I have to be a bit crass. I believe the killer application will be the same as for VCRs: games and game-related things.

Q: Van?

Van Baker: Video conferencing. Absolutely. The notion of being able to put the family in front of the television or some other kind of device connected to the TV and talk to Grandma on the other side of the country is hugely compelling, but it takes bandwidth to be able to do it.

Chapter 17: Panel Discussion—Bandwidth: Who, What, When and Where?

Moderator:

Gregory L. Sheppard
Chief Analyst
Semiconductor Application Markets Program
Semiconductors Group
Dataquest

Panelists

Brett Azuma
Director and Principal Analyst
Public Network Equipment and Services Program
Telecommunications Group
Dataquest

John Armstrong
Dataquest

Don Miller
Assante

Moderator: This next topic does not need any further introduction. This topic is to all of us an opportunity as plastics were to Dustin Hoffman in *The Graduate*. This is the future. It's a topic that's full of arcane terminology, hundreds of billions of dollars in revenues, particularly on the service side as well as equipment and religious wars, shall we say, in terms of the different technologies and standards going after the different market spaces.

We have done a bit of shakeup on our panel today. We had a couple health problems in that John is going to be sitting in for Steve Diamond and we have Don sitting in for Bobbi Murphy.

We're going to be looking at the whole food chain. Brett's expertise is in the services area, public telecom services. Don is coming from a quite varied background. He's coming from a router and WAN perspective. He's going to be looking at the access marketplace issues. John is going to be the within-the-enterprise guy looking at the enterprise bandwidth issues.

John Armstrong will be our first speaker today. John comes to Dataquest with a wide variety of experience. He did a stopover at Synoptics as well. He's been with the company for about six months and comes fresh out of the trenches.

John Armstrong: Out of the trenches is a good way to put it. As Greg said, there are a lot of interesting and innovative things happening in the networking space and we're going to look primarily at the backbone here. I know that Steve Diamond was supposed to give this presentation and I had the presentation placed on my desk the other day. We're going back to the source so it's not a bad thing.

Let's step back a little bit and take a look at what's happening in terms of all the technologies that are converging into the networking marketplace. That's what makes this such an interesting space to follow because there's so much happening at one time. You can see that there are some standards activities that

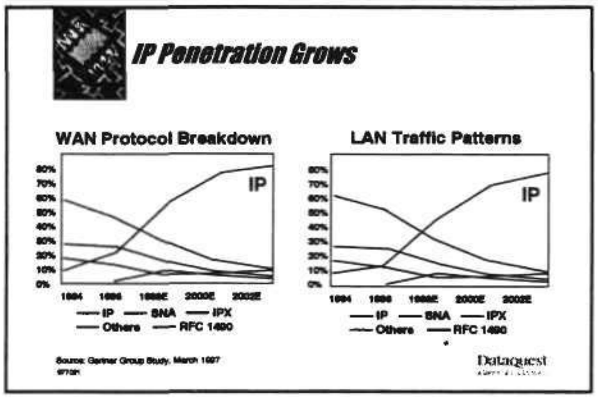

are taking place from a number of different bodies that are dealing with IP issues, Multicast, RSVP which is a reservation protocol which is a quality of service protocol. You see a number of video standards that have been proposed and are in the process of being implemented by vendors that will allow video to be transferred over the local area network.

You hear vendors talking about their new data, voice and video strategies and that's where all of this is coming from. Now they want to add more value to the networking solutions that they provide by allowing you to put video and voice on the network besides data. ATM development has had a lot to do with that. Asynchronous transfer mode is a networking access method that's been around for a while that was actually designed from the ground up to support multimedia. We have enhanced local area networks now. We have a Gigabit Ethernet, which is 100 times faster than the old Ethernet, 10-megabit Ethernet, 10 times faster than Fast Ethernet. We have wide area network capabilities now that are very complementary to the local network, xDSL and ATM. ATM's actually a LAN and WAN technology. All these things are happening at once.

We also have another phenomenon that's occurring in the networking space and that is the predominance of IP or TCP IP as the protocol of choice in the network. You can see our forecast over the next few years, both in the wide and local area, is that the traffic patterns are going to support primarily IP. What's significant about this? If you're designing networking products and you only have to design a product to run on one protocol primarily, it makes it a whole lot easier for you to create a very efficient design to push packets quickly across the network.

Ten years ago when traditional routers were being developed, they had to be able to support multiple protocols. New products coming out the chute, such as Layer 3 switches, don't have to do that necessarily. Most of them support TCP IP or IPX.

As a matter of fact, if we take a look at these growth rates here, we will see a steeper growth rate for IP. We've shown that through a recent user wants and needs survey that we've done.

What's happening in the Intranet here? Web server growth is increasing both on the Internet and Intranet

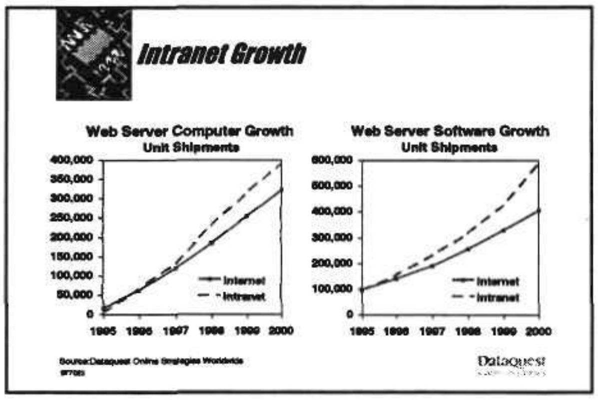

and the software growth is increasing almost in a hockey stick effect.

That puts a lot of pressure on the local area network backbone. We have a lot of demand for bandwidth now coming from sources that didn't really exist in any significant fashion a few years ago. A lot of push technology, for example, is also entering the LAN backbone. What we see here is an inversion of the old 80-20 rule. The 80-20 rule was that 80% of the traffic occurred out in the workgroup, out into the user areas, out of the wiring closets in the LAN space and only 20% of the traffic was concentrated

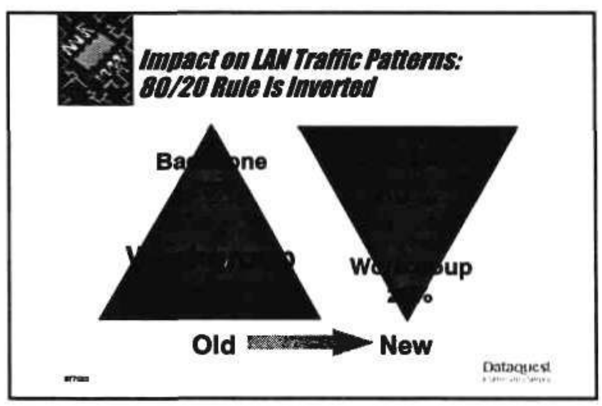

in the backbone. That has reversed itself pretty quickly and now we have the opposite situation. A tremendous amount of pressure on the backbone in the network and that's the network center, the confluence, the convergence of all the traffic that comes into the network.

If you're a betting person, maybe you want to put some money on some of the vendors who are going to make it or maybe just break down in terms of the new technologies. Older technologies, such as FDDI, are proving to be unsuitable for a lot of the new generation networks, not that FDDI didn't have its place but its actual market share has peaked and is

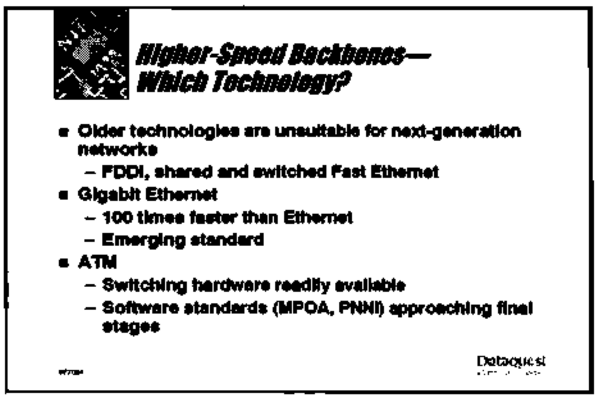

declining over time. One of the reasons was cost and complexity.

Now we have Gigabit Ethernet as I mentioned earlier, which is an extremely high-speed networking protocol. The standards should be completed by mid 1998 and we have ATM which has been available now for about five years commercially, which has actually matured quite nicely in technology and is being adopted fairly widely in the backbone now by customers who are looking for technology that has

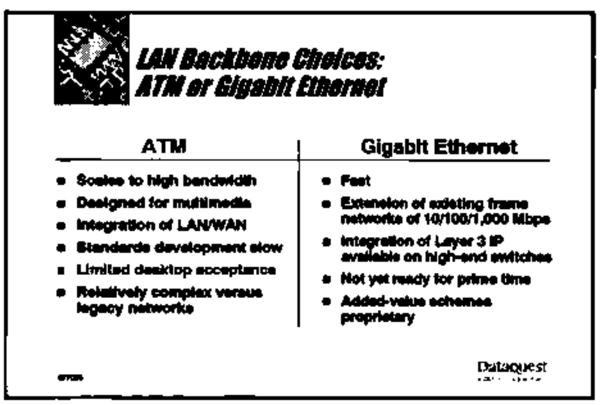

had a presence in the market for some time. Gigabit Ethernet is not quite ready for prime time yet.

We have a comparison chart here. ATM's pluses and minuses compared to Gigabit Ethernet because there is quite a debate raging here in some quarters. Greg mentioned technology religion. Smart vendors are not getting religious about technology. Most of the big guys are providing both technologies. ATM was really the only technology that was designed from the ground up by engineering design committees to support asynchronous traffic, that is a combination of voice, video and data together. It really was designed for "multimedia." It's used in the local and the wide area networks with relatively equal ease and actually will probably be more predominant in the wide area than the local area over the long term. However, what's hampered ATM is the slow development of standards, the inability for ATM to really get any critical mass out on the desktop and its perceived relative complexity compared to the Ethernet standards, Ethernet Fast and Gigabit.

For the Gigabit Ethernet the pros and cons are that Gigabit Ethernet is really fast. We're talking 1,000 megabits per second and it's viewed as an extension of the existing 10 and 100 Ethernet. Same frame size fits in there. Customers have the perception it's

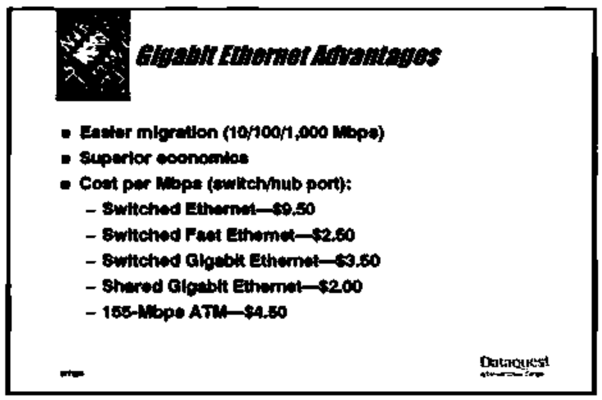

relatively easy to upgrade, although that may not be the case with early implementations but that exists nevertheless. A lot of the Gigabit Ethernet vendors are integrating Layer 3 switching which is actually fast routing into their products so they're getting added value in there that makes the products much more attractive than ATM but, as I mentioned earlier, not quite ready for prime time. The big guys are not out with a full set of products to support gigabit.

We're seeing a lot of announcements. We do have a number of start-ups who are selling product to the customers who are not risk averse. Also in Gigabit Ethernet we see added value schemes such as quality of service capabilities still being implemented on a proprietary level.

Nevertheless, lots of advantages to Gigabit Ethernet. That is the migration, superior economics. If we take a look at the numbers here, this is what's really compelling to a lot of people who are taking a look at Gig Ethernet. Even as a nascent technology, just coming out the chute, if you compare the cost of 1 million bits per second on a switch or hub port to existing technologies, Gigabit Ethernet looks very attractive. It's still cheaper than 155 ATM, which has been around now for four or five years commercially. It's a very attractive technology from an economics standpoint. We know that those numbers are going to decline over time relatively steeply I would expect. Right now a port on a Gigabit Ethernet switch is around $2,000. We should see that go down by about half within 18 months.

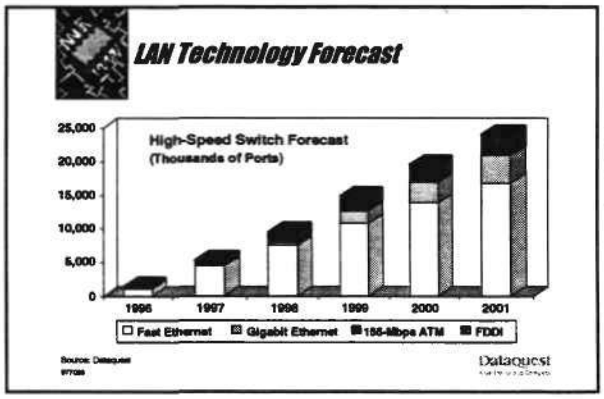

Gigabit Ethernet is also absorbing ATM-like features. Quality of service capabilities—things that were previously the domain of the ATM product area are now being adopted by Gigabit Ethernet vendors. The challenge here is for these vendors to actually make these added value capabilities standard across the board so that there's interoperability between

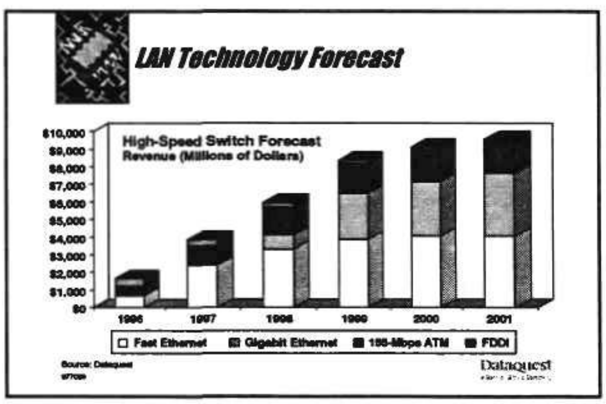

products. That doesn't exist on a lot of levels yet, especially on the QOS side. We see things like link redundancy and link aggregation being implemented by some vendors but not by others. Not everyone is doing QOS. Some QOS schemes are much more complex than others. Some vendors simply have a QOS which is high priority/low priority, pretty straightforward stuff. As I mentioned, Layer 3 stuff is being implemented on these switches which allows the switches to do fast routing, typically only with the IP or maybe IP and IPX protocols. Nevertheless, these products can be used to augment routes that exist today which, as we all know, are fairly costly and complex.

Here's our technology forecast. One of our fortes at Dataquest is to take a look at the numbers and the technology and make forecasts. As you can see, we're pretty bullish on Fast Ethernet. In terms of the actual total number of ports, that's going to be the predominant technology but we also see some pretty strong growth in Gig and ATM, all through 2001. FDDI will decline.

In terms of dollars, you can see where the big bucks are. The big bucks are in the Gigabit Ethernet side. In terms of a per-port cost, it will be a premium price technology. As I said earlier, if we take a look at 1 million bits per second and the cost of moving 1 million bits per second out of a switch port on Gigabit Ethernet, it's actually more cost effective than the other technologies.

Here is Gigabit Ethernet versus ATM in terms of revenue and port shipments. As you can see, they're both going to experience good strong growth, Gigabit Ethernet stronger than ATM.

Gregory: Now I'd like to invite up Don Miller and

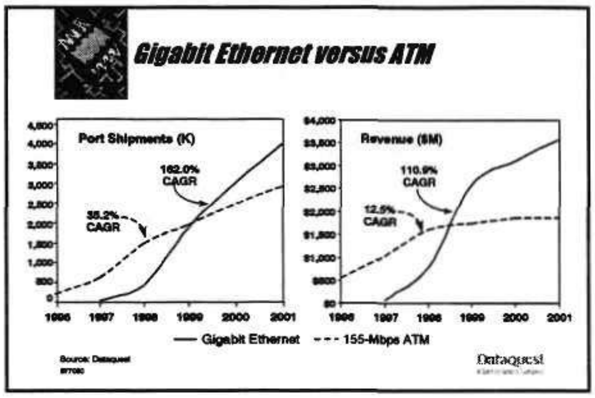

he's going to be looking at the access area for us. I think you have worked just about everywhere, mostly recently at Asante.

Don Miller: This presentation was originally developed by Bobbi Murphy, who's my counterpart in what we call the remote LAN and Internet access service and they focus specifically on technologies that are used for access into the Internet. This includes things like DSL, dial modems, ISDN, cable modems, a variety of different technologies. We're going to take a quick look at some observations that we made in this particular area and some patterns we see emerging. There ought to be a couple significant things that will jump out at you as we go through this fairly quickly.

One thing we've seen is that in the remote access market it's tough. It isn't easy anywhere in the technology sector but it's extremely competitive

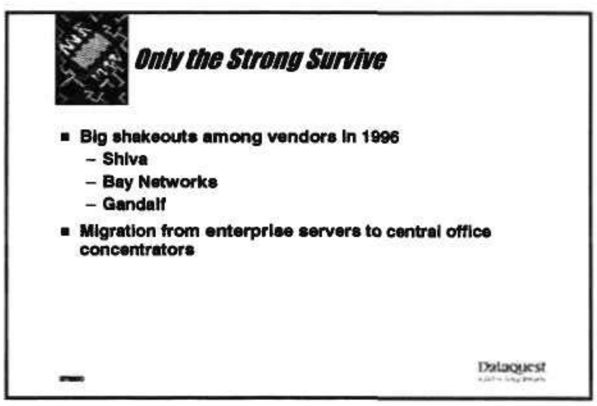

here. In 1996 we saw some shakeups in terms of a couple of small players as well as large players. Shiva and Gandalf may have been dealt the terminal blows in terms of what happened to their business. Bay Networks, a large player, was well established as a leader in the enterprise space but has also been impacted by a shift that's taking place in the industry from access servers to access concentrators.

The next couple slides give you the idea of who's doing what in terms of the access concentration. You notice at the end where you see U.S. Robotics and 3Com, we've combined their two revenue streams to represent the combined company after their merger.

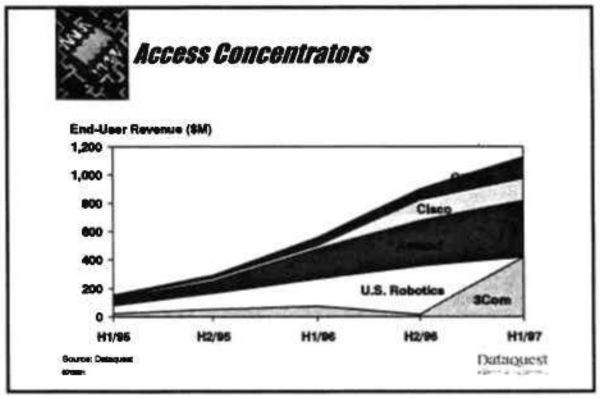

These are the top tier players you see in the access concentration market.

To give you an idea of what's happening in terms of servers, these are a different set of players. You see Bay and Shiva. You see how they're taking gas basically as the market has shifted because the overall line is going down, not up. In our business we're all used to seeing that hockey stick, so something like this is a significant foreteller of the future.

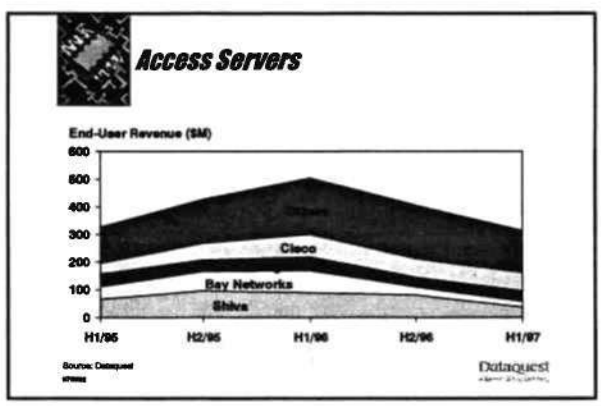

Here's another comparison that we've taken a look at. You can see that the access servers are extremely flat. This is not good news, so there's a clear signal that concentration is the way to go. If you're going to survive as a vendor in this space, you have to have a very competitive product. The markets that are going

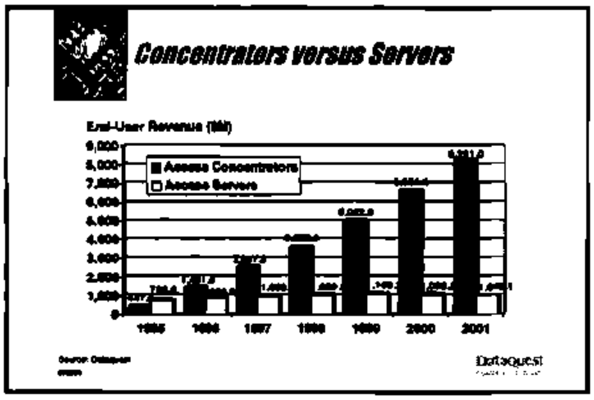

to be assimilating this are primarily the carriers, the Internet service providers.

What we've seen here is this shift in terms of the sales process. It was for a while an enterprise sale. It's now moved to a service provider sale for a couple of reasons. Service providers are now building out their infrastructures and they're looking to seek new business opportunities. They're looking at the enterprise networks as a nice juicy plum that they can go after. Enterprise customers are receptive to this because they've got this huge headache that they'd just as soon get rid of if they can find a good partner.

With the explosion of the Internet, the small office home office channel, residential users dialing into the Internet is also accelerating the growth and interest in this market.

If you're not in the top five because of the consolidation in this industry, you're going to be relegated to the walking dead. There's a certain amount of companies in the industry that achieve typically under $100 million in revenues and they plateau. I call them the walking dead. They never get beyond that state. Consolidation continues to drive the very top level companies, so you'll see a handful of very large players at the top that are going to be the big dogs and that are going to drive what the future is going to look like.

Outsourcing is viewed as a very big opportunity. When you look at remote access, if you're an enterprise customer, you're trying to support a network of branch offices that is scattered around or a very large community of telecommuters. This remote access stuff is voodoo in many cases. It's not straightforward, it's very labor intensive, you have to deal with a large user community and it's an IS organization's nightmare. Instead of being able to walk across the hall to go to someone's office to ask them whether or not they've got they're PC turned on to determine whether or not they're on the LAN, this person's across the city somewhere, maybe across the country. You have to be able to communicate with them over a long distance and provide support. Cockpit errors, as we know—some of us are guilty of this—are difficult to deal with. People are looking for 8 x 24, seven days, 24 hours of service. You have global organizations that never close and they've got to have service around the clock. People can't typically staff. It's the old pig and the python. How do you handle the peak load? If you size your network for the peak load and if you try to staff to that, you'll go out of business because 90% of the time either the resources remain unused, sitting idle or you'll be paying for excess capacity when you don't really need it. At the same time you suffer dearly if you have a peak and you can't deal with these spikes in capacity and usage.

You're seeing a WAN cost change with shifts in the footprint. Also, if service providers stick to their knitting and don't take their eye off the ball, they're going to be very well positioned in the long term as the people who can take advantage of these problems that exist within the enterprise. It's a business opportunity. That's why you see a lot of the consolidation in the industry right now.

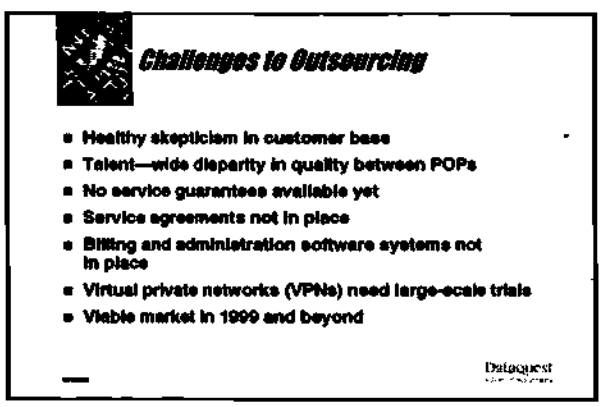

This all sounds great, right? You're just not going to flip the switch and go into the outsourcing business. This is not an easy thing to do. The same headaches that the enterprise customers have had initially in this space are going to basically move over. You're

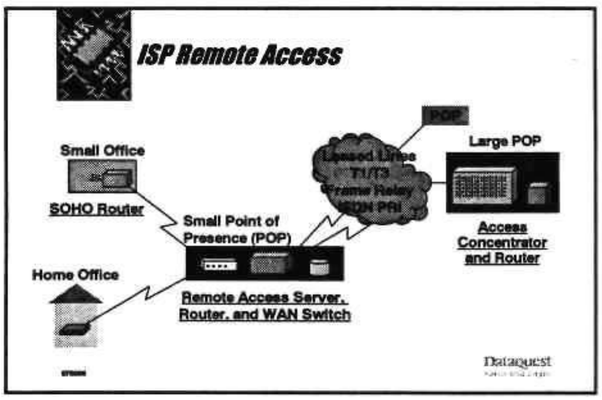

going to have to hand the bottle of Excedrin off to the next guy, so they've got to be very skilled. There's still skepticism within the enterprise customer out there because they've tried it and they've broken their pick on it. They've had plenty of headaches and they're thinking you're going to do a better job than me? There's a certain amount of skepticism out there.

In terms of the ISPs that are coming to the table saying I can take care of this for you, there's a wide disparity in terms of the quality of services they provide within their network. Some POPs are overloaded. Look at *Boardwatch* magazine which is a book that reports on the Internet. Among other things, they've started recording statistical data on performance of various Internet carriers, ISPs and they graph the response times and other data by location. For any given ISP in there, it will vary dramatically from city to city. One thing I noticed, as I looked through all these ISPs to see who had the best rating in the Bay area, without a doubt the access was the worst in San Jose and the Bay area. Primarily, it's because there's a bazillion of them trying to get to the net, so no one's got a real good solution there. That's the kind of disparity in terms of performance that you have to deal with. As an ISP, how do you build a network that is going to be resilient enough to deal with these traffic patterns that are changing with demographics and various geography? Customers want consistent service across their network. They don't want hot spots where they've got a lot of problems, especially if they've got a high concentration of users. There's other mechanical things. Service agreements need to be hammered out. How do you bill for this stuff? Some of the technologies to enable the deployment of these outsourcing agreements are referred to as virtual private networks. Most of this hasn't been deployed on a large scale and it's coming. There's a lot of focus here. We feel this is going to emerge as a fairly viable market but not probably until 1999 and beyond, so there are going to be some pioneers here. Hopefully, there won't be a lot of arrows in the back.

As I mentioned earlier, here are some bullets that you have to deal with. If you're an ISP, you have to be able to scale your network and scale it rapidly because the growth in the Internet and the traffic from these Intranets that will be outsourced are growing dramatically. Along with that, you have to be able to scale your operations, so this is a capital-intensive aspect here. This is ongoing cash flow. You have to have skills at multi vendor integration because there's no one vendor out there—regardless of claims of some of the larger vendors—that's going to provide all the solutions for you. At the same time

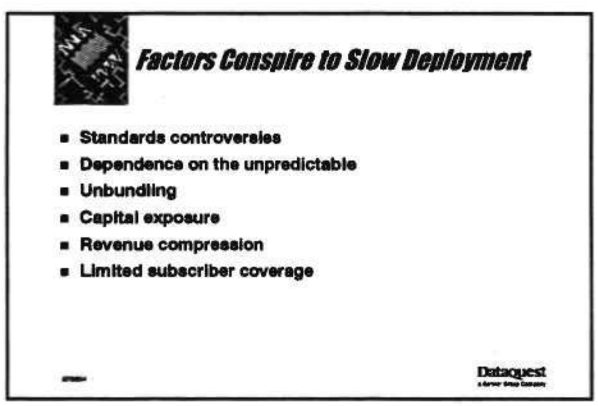

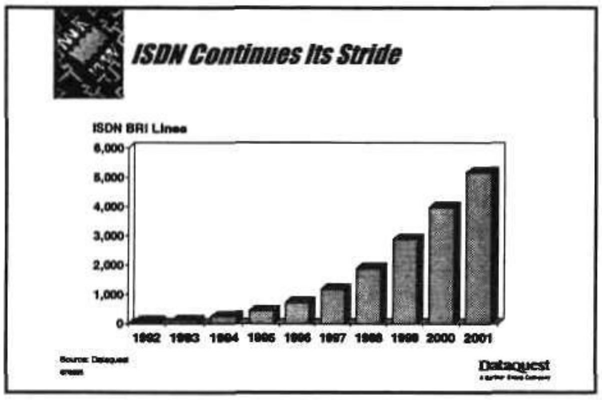

you're fighting this rapid commoditization of services, so it's the 1995 all-you-can-eat. It's hard to make a business out of that and it's a chicken and egg situation. You have to be able to build and deploy and bill and deliver truly differentiated services that you can add value to and charge more for. At the same time the cash flow that's feeding you is the 1995 all-you-can-eat thing. Our pat comment in this area is, looking forward, the 1995 buffet will no longer include meat. There's a variety of new competitors coming in, so just managing the growth is a real change. Many of these folks have not got a lot of experience in doing that.

Here are some changes you see in the deployment of architecture. Today a lot of ISPs have built networks that look like this where they have small POPs or points of presence, so you have users dialing in. This is you at home with a dial-up connection or maybe in a small office location, a remote office, where they lease line, dialing into a local point of presence. This is not a long distance call here. This is a local exchange call. With a server, a switch, probably a router here, connecting into the cloud, the Internet or the backbone network, aggregating this traffic in these large POPs where you have access concentrators and routers. These are then networked out to network access points within the Internet.

What you see happening is this evolution where these little stars represent the small POPs. The economics are changing and you're starting to see an evolution that looks like this, fewer and fewer POPs as people begin to back-haul traffic and aggregate concentration in fewer points. Why? It's easier to manager for one thing. Another reason—these lines are becoming less expensive. There are other technologies that are allowing people to aggregate this network and these networks are built around support for dial.

One thing I want to point out before we leave that. How many of you are familiar with the xDSL types of technologies? What it amounts to for xDSL to be deployed widely it depends upon a footprint. This looks like a nice footprint. Who's going to deploy this service? ISPs? Watch this. The footprint's shrinking, so we've got this diametrically opposed model going on out there in terms of this is being driven by economics and you have people in the DSL side of the house trying to figure out how to deploy this or get people to deploy it in a wide footprint. This is making for some interesting observations going forward here.

What we're seeing on the dial side is this mega-POP or POP in a box. They are very dense sites. For example, the America Online network has a total of 30 POPs. This is huge. What we're actually seeing is that within the ISP community more than likely you're going to have the smaller ISPs subcontracting the services, the dial services, to reduce their costs. They're actually going to go to larger carriers who are going to have excess capacity and negotiate agreements with them to carry the dial traffic for them. Much of this comes from this—economies of scale. It's like laws of large numbers.

What do these POPs look like? They're highly dense. They take advantage of DSP types of technology for either high-density modems, ISDN types of terminations with the capacity to terminate large numbers of telephone calls.

If you take a look at this chart, basically it's saying that prices are going down on access concentrators and that's a function of port density. The denser you get, the lower the price per port. Acquisition cost in

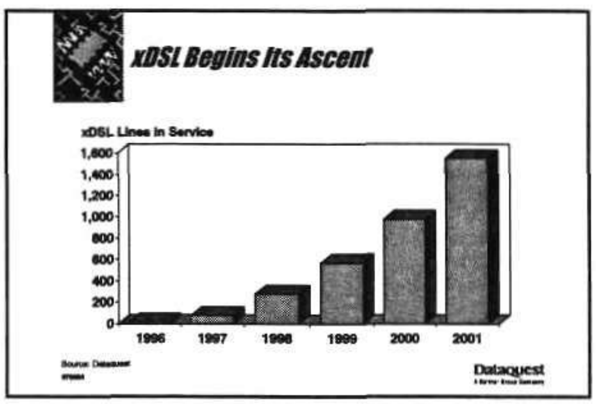

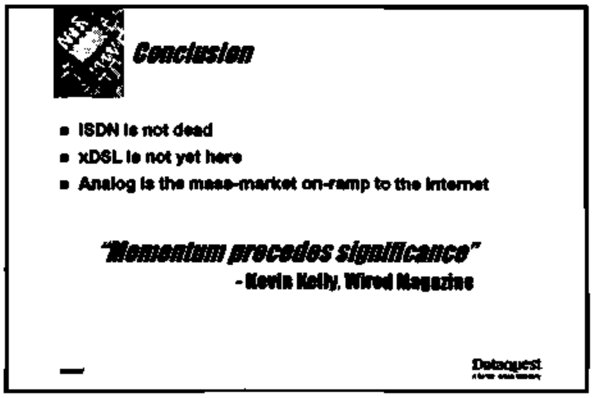

terms of capital expense is a big issue here, as well as cost of ownership. Space— these things have to live in a central office that has been where you're leasing space from a carrier and some of the rates I've heard here are pretty usurious compared to what you pay for square footage of regular office space.

This is Bobbi's poke at the 56K dial modem standard. We give this the most high product of the year award. It's a problem. Is anyone using a 56K modem out there? I am. Anyone gotten a 56K connection yet? From my house I can't. There's no way. I got online with the tech support people and

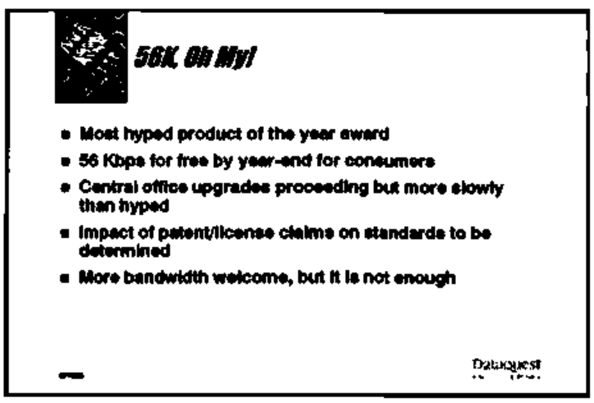

U.S. Robotics actually has a test site somewhere in the world that you can literally dial into from your home or wherever you're dialing in and it will test the characteristics of your line to determine whether or not you'll ever be able to run it over 28.8. 33.4 is the fastest I've ever made a connection at. This is fairly broadly-hyped. The standards issues are still working their way out. Some of the central office-related work is going slow. It's still a pretty dicey situation but it's going to muddle through and we'll eventually get here. There are some vendors, for example—I believe Bay has technology in their new access concentrators that are DSP-based that support both of the 56K standards or the current technologies. By and large, in spite of the hype behind 56K, we expect it to be the dominant technology going forward.

One of the problems is if we've got the ISPs building these POP in a box, super concentrated, high density nodes. What's the last mile going to be? There's a lot of debate and arm wrestling here and dollars being thrown. The target markets that we see here for business is going to be ISDN, fractional T1 symmetric service. We estimate some numbers here that Brett is going to comment a little bit more on later. In the residential market we still see it being primarily analog with some ISDN. The monthly line revenues— people are honing in on that number. We have it figured out.

To steal a line from our present government, we're on an analog bridge to the 21st century, so the predominant amount of connections out there for remote access are going to remain analog modem of some kind or another, mostly the higher speed one. We'll see low speed DSL deployed primarily in the small office locations, a lot of remote offices and folks that are subsidizing telecommuting connections. There will be options out there. We're already aware of some companies that are going to offer DSL types of connections to corporations. Like in the Bay area, they'll go to 3Com, Cisco or Oracle, any of the large purveyors there that have a large population of people who telecommute. Standard Oil will be another one. Instead of offering them ISDN, they'll offer them a higher performance line at a more competitive rate. You're going to see this type of thing emerge in the not-too-distant future.

Cable modems are still going to be out there. One of the issues here is the capital required to deploy this. A lot of the cable companies are under severe attack by the satellite guys. I love my DSS, thank you. By the way, TCI doesn't get hardly any money from me anymore. That's happened a lot around the country, so they don't have a lot of excess cash in order to finance the upgrades that are needed for their infrastructures to deploy the capabilities for Internet access. Some of them are doing a better job than others but we do see this coming on. Then there's a variety of satellite options that are being considered.

What's ahead? Deal-making. We see more consolidation and there is going to be a new form of carrier class, remote access company. This will be a new form of competitive local exchange carrier that will emerge to take advantage of some of this. They'll have some deep pockets and they'll need them for a while but they'll have a very broad presence.

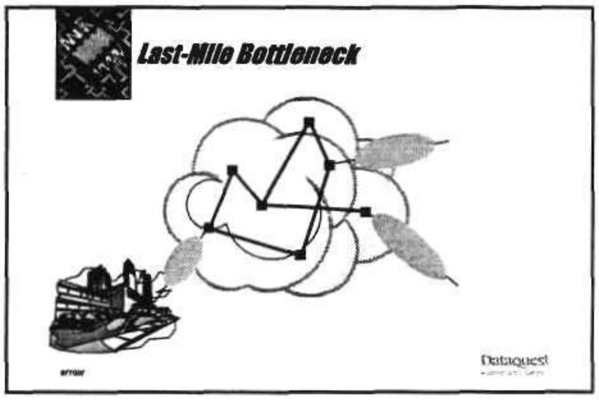

Gregory: Next up is Brett Azuma and he tracks the public services area for the North America market for Dataquest. He's out of Pacific Bell and he told me once upon a time he worked with Scott Adams.

Brett Azuma: I'm here to provide you a carrier perspective on things. The one way to try to characterize the carrier market, if you try to look at things like ISDN and DSL over the last several years, is that things don't move nearly as quickly as everyone expects. In fact, one way to perhaps characterize the market is that it's a somewhat inert market. Not inert from the standpoint that nothing happens, it's just that the market continues along the same trajectory and at about the same speed over a long period of time. It requires a fair amount of effort to create momentum to move the market forward.

I'd like to start off with some factors that conspire to slow-roll the deployment of these services. Last year I started seeing articles about the death of ISDN. One of the pieces of advice I've been giving the clients is look for stuff that people say is dead and circle that because that's probably going to be an area of boom. xDSL is kind of the favorite new child of access technologies. It's not moving along as quickly as everyone had anticipated. There are some very interesting reasons for its slow-roll and ascent. Then I'll provide you with some conclusions.

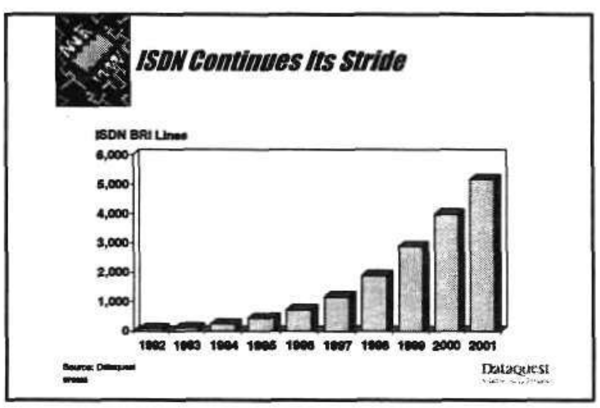

A couple of things about the services market and what happens with the carriers. There's a law of three actually that I use to describe the primary driver for massive carrier deployment of any technology. There are only three things that move the market forward. Item number one is regulatory fiat, some type of new regulation that comes out. Item number two is some type of operational cost savings and item number three is market opportunity. I would almost challenge you to find number three in gross abundance. The first two and arguably the first one, regulatory fiat, have been the biggest drivers of evolution within this market.

Factors that conspire to slow deployment. We talk about holy wars. We've had the standards controversies. The standards holy war in the past was around ISDN, custom ISDN versus national ISDN versions. The holy war of the last year has been DMT versus CAP for xDSL. The holy war of tomorrow is going to be 2B1Q versus the winner of the CAP or DMT controversy. What you're starting to see is the emergence of new xDSL technologies as opposed to ADSL that are based on a relatively familiar line coding scheme, 2B1Q.

Number two, dependence on the unpredictable. Everyone likes the idea of laying fiber to everyone's home. It sounds like a great idea. The problem is it costs a lot of money to do that and it takes a lot of time. For the foreseeable future, you're looking at a copper-based world. That copper-based world has been installed and upgraded over the last 70 or 80 years and it wasn't designed to handle digital services. In fact, it was designed to handle voice grade services and it's done that very well for a number of years. The problem is shortcuts have been made to optimize cost, which was the intent of the carriers and the intent of the regulatory schemes to milk every voice connection they could out of that plant. They put in things like load coils and bridges taps and all sorts of other interesting telco type terms. The definition of that for me is stuff that makes digital services not work terribly well. The fact is that copper is not designed to handle advanced data services and we're asking it and we're modifying it on a one by each basis to provide these advanced services.

One of the other interesting things that's happened is if you look at the plant having been installed over the last 70 years, record keeping wasn't done very well. Even though you think of the telephone company as being terribly well organized, the fact is there are a lot of manual records that were entered in over a number of years. People did things quickly, restored services, things like that and they forgot to update the records. One of the problems with inaccurate records is it makes it very difficult to understand where you can and cannot provide service. I had a panel discussion about two or three months ago where one of the ISPs expressed concern that the Telcos would give them all the bad loops and give themselves all the best loops. The ugly fact is the Telcos don't know where the good loops and the bad loops are. I went out with a construction crew a number of years ago to watch them install T1 circuits. They had three orders for T1 circuits so I expected us to actually process three orders for T1 circuits. The fact is they were only able to install one of those orders. The other two had record inaccuracies that prevented them from being able to install it. Those are some of the problems that are facing these guys. They've automated the record keeping but realize that the source for that record keeping was those old inaccurate manual records. Garbage in-garbage out.

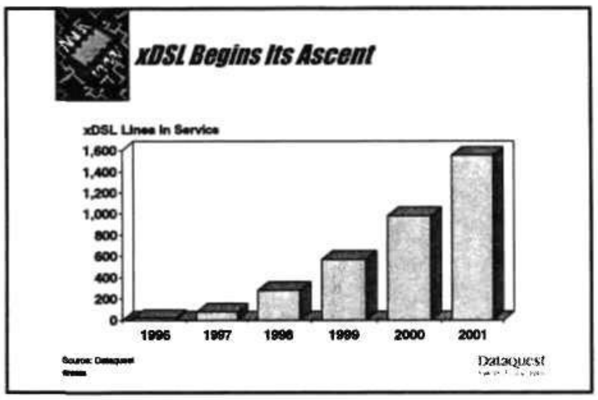

The other interesting thing in the installation process is that every time you have to go out and install something like ISDN or xDSL, it becomes an adventure and adventure is the thing you don't want when you want to mass produce something.

Unbundling is what people have been talking about doing within the public network to enable further competition, to enable new service providers to enter the market and utilize the same assets that the carriers have used for years. In theory, it ought to accelerate deployment to new services and in theory it should move things along. We're finding just the opposite. The reason for that is now instead of having one group placing new technology into the network, you have several, so things like spectral incompatibility issues, which we're starting to see with high speed data services like ADSL, are starting to show up. When you can't necessarily predict what's going to go in the plant, it becomes much more challenging—in fact, it becomes virtually impossible—to maintain service integrity to your existing subscriber base. Realize the carriers are still responsible for providing 7 x 24 bulletproof, five nines, type of service reliability in the public network.

There are two pieces to capital exposure. When we start talking about cost, the cost of being wrong in this market is extremely high. When I talk to people that manufacture products, they have a hard sense of understanding deployment. If you ship product, like a PC, to a store that has relatively low traffic, you can actually redeploy that PC to a store that has high traffic at some point and the cost in doing so is really transportation only. If you place fiber to the wrong neighborhood and your take rates are real low, there is no way to recover from that. You can't redeploy the fiber terribly well because most of your costs are labor costs.

I was looking at what it would cost to deploy xDSL on a ubiquitous basis throughout the U.S. There are 23,000 central office switches in the United States. That's a lot of locations to have to drop this equipment. If you were to assume that each DSL line was going to cost about $1,000 in capital deployment and that's a pretty good approximate number and you had to deploy a minimum of 10 lines per DSL node, then you're talking about an outlay of about $250 million of capital before you see dime one of revenue.

One of the other issues that you see associated with capital exposure is back to the unbundling issue I brought up earlier. What the carriers are going to be required to do over time is offer to their competitors at their same unit cost access to all of their facilities, including things like DSLAMs. That means that pride of ownership is about all the carriers will enjoy for that capital exposure that they're going to be facing. It'd be like some of you who have chip manufacturing plants having to offer capacity at your unit cost to any and all of your competitors. It's an ugly dynamic from a business standpoint.

Clearly the carriers are going to see a fair amount of revenue compression as it starts shoving a lot more bits down the same lines. What they're facing is a revenue compression that ranges between 10 and 15 to 1. Whereas before they could sell a T1 circuit that provides 1.5 meg of data, they're looking at having to sell that down in the $50 to $60 range. In fact, some companies over the next month or two will be announcing prices even lower than that. What was surprising in talking to the ISPs is the revenue compression they face is even more severe. Their current charges for a T1 access into their ISP cloud is about $1,800 a month and they're looking at a forced compression to get them down to the $15 to $20 a month range as well. For both the ISPs and the local exchange companies, we're seeing a fair amount of blood being spilled in terms of revenue compression.

Distance is always a problem. Any of these new technologies—you go back five years and you think about ISDN—distance is a huge problem. Right now if you take a look at the technologies that are being offered, you have about 15 kilofeet that go from the central office. There is about 30% of the subscribers that you would try to target that won't actually be able to get your service. It's difficult marketing a service when you can't offer it on a wide scale basis. People complain about that and you end up coming up with relatively expensive engineering solutions to be able to provide that type of coverage. In addition, if you're advertising—you can't advertise terribly well if you can't offer blanket market coverage. The challenge in deployment, as we look at these things, is to provide for wide geographic coverage with relatively low density.

Let me give you some figures to think about. If you look at custom calling features, things like call waiting, roughly one-fifth of the telephone subscribers subscribe to a lot of these advanced services. If you look at the more nichey products, like call forwarding or speed dialing or some of the other custom calling things you see from the telephone company, about one out of every 20 people buy those services. It's really difficult to develop scale when your penetration rates run that low.

Enough about the ugly stuff. ISDN—last year's dead product is this year's booming product. A 40% growth rate over the forecasted period. Why is ISDN moving? It's got momentum. It's got product in the channel. It's got a large footprint of deployment. At the moment, it's the only access service that's available that can give you anything above 28.8 on a somewhat reliable basis. The ISDN people are actually doing some smart things to enable their service. They're packaging CPE together. They're developing technologies that allow you to get to take e-mail and push technologies without running a four-figure usage bill every month. We believe ISDN has a lot of legs.

What happens? ISDN's original root and, in fact, the original application for ISDN was not data; it was actually voice; it was actually focused at providing feature button phones within large enterprises. We believe that the shift will occur back. As ISDN starts to get displaced by xDSL and the data market, ISDN will become attractive once again as a voice product. It enables all sorts of neat call control features and it's well suited for telecommute applications.

xDSL begins its ascent. I've got to warn you about these numbers. These look really nice and large. Realize that a big chunk of these is IDSL, which relies on the same chip sets, relies on the same equipment as ISDN. xDSL on these charts is about a third of the bars that you see. By 2001 we anticipate a little under 500,000 lines. You make a lot of money off 500,000 lines so I shouldn't belittle it. We're looking at about a $3.3 billion annual revenue stream for the service providers by selling 500,000 million lines of xDSL but it's going to happen slower than people anticipate.

ISDN is not dead. ISDN has got a lot of legs to it. It's got momentum. In fact, we see ISDN growing for a very long time.

xDSL is an exciting technology and over the long term it's the ideal data product. However, it's not there yet. You can't order it. You'll see a couple of announcements for deployment this year but the fact is you're not going to see anything that looks like serious deployment until about mid to late next year.

For the time being, as Don was indicating, analog is the bridge into the 21st century.

This is one of my favorite quotes—"Momentum precedes significance." It takes time to develop momentum in this market but we believe that significance will follow but it follows momentum, not precedes it.

Q: What are ATM's compelling features in the LAN for the long term and will Gigabit Ethernet eventually replace ATM in the LAN as it's rolled out? Start with John.

A: There are some good reasons to have ATM in your LAN backbone. First of all, it's mature and it will maintain some leadership as a mature technology for at least another year or so until we see all of the major vendors roll out gigabit products and integrate necessary network management capabilities to the same extent that they're available in ATM now. The other thing you have to remember is that ATM, those same cells that are used in the ATM backbone in the LAN, are used in the ATM in the wide area network. You don't have a similar technology on the Ethernet side that can make that crossover as easily, so large enterprises that want to have cell-based traffic throughout their network and between campuses and out on the wide area are going to lean more heavily on ATM than they will on gigabit.

We've forecast that by late '99 the actual number of Gigabit Ethernet shipments will exceed ATM shipments into the local area network backbone, so at that time the lines cross and gigabit will take over.

Q: We talked about Internet protocol, IP, as opposed to intellectual property which we've had a lot of running around the meeting today. At one point there was IP squared, intellectual property as related to Internet protocol but I won't get into that. This IP thing is the great leveler it seems. How is it going to impact LAN and WAN equipment? Are we looking at redesign? We have the chip guys out here that are looking for a window of opportunity.

A: One of the reasons IP is emerging as the predominant protocol is the support of a lot of major vendors, such as Microsoft, large application developers who are building large scale applications that are now being ported to IP environments. The Internet is another reason for that. It's becoming the lingua franca. What it's going to mean is hopefully less of a chaotic environment that our enterprise clients are going to be dealing with. The shift is away from these proprietary protocols or these one-offs like the litany of things that you find in a network today. Easily you can find IPX, Decnet, X&S, Apple Talk, Net Buoy, lots of different things floating around. IP managers are trying to kill this stuff off because they'd rather support one protocol. In terms of opportunities for people that are working in silicon, a lot of recent products we've seen on the systems level have been targeted at handling IP traffic in a very efficient manner, especially routing. These Layer 3 switches are optimized to provide wire speed or the line speed support of routed packets using IP within the LAN at very high speeds. If we're going to take the excess protocols and say you don't need to support those anymore, it's easier to apply focus and technology development and silicon to a single protocol suite.

Q: One for Brett. The recent spate of merger and acquisition activity in the services area has created a lot of ill feeling about the outcome of the Telecom Reform Act. Where is it going to go from here? It seems like this is a big quagmire from regulatory and M&A perspective. Any thoughts?

A: It's a huge quagmire. When I talk to people about what consumers see as a result of telecom reform, I have a rule of three—torn up streets, from trenchers dropping fiber into the ground; clogged courthouses and higher prices for consumers. AT&T's minute rate for long distance has actually increased since telecom reform was passed. We're starting to see an upward trend for rates at the moment. What consumers are going to see over the long term are going to be higher rates. They're going to see a greater variety of services and business customers will probably be the big winners in all of this as competition starts to drive down prices and increase the diversity of services. The cross subsidization goes away. What ends up happening are more opportunities on the business side. Things happen really slowly and the fact is this stuff is going to get locked up into the courts. As soon as you get an attorney involved in things of this type—and that's what has happened—you can immediately multiply the amount of time you expect something to happen by about three. That will give you a sense of when stuff will happen in the market.

Q: One of the hot things at Inter Op this time around was voice over IP, voice over Internet. There's obviously opportunities for the chip guys here, particularly on the compression side, with DSP processors. Likewise, as Don mentioned, silicon acceleration, at the Layer 3 level. Is this market going to go anywhere? Are we going to have free phone calls and things like that?

A: We've recently given a number of speeches on this topic and one thing that needs to be clarified right up front is there's a difference between voice over IP and voice over the Internet. Two different things. Voice over IP happens to be voice over the Internet. It's packetized voice. What you have to look at are the technologies in the deployment of packetized voice. I'll try and give you a short answer here. The current architecture of the Internet and the latency problems that are involved in delivering toll quality voice on a consistent fashion, we see that the initial implementation of this technology will be predominantly in enterprise networks where you can control the environment and deliver predictable performance from end-to-end. We also see some of the predominant applications that will be deployed first will be the non-real time types of voice applications, like audio broadcast, bulk movement of voice mail messages from point A to point B and not so much the Internet telephony types of applications.

Q: Is the quality there?

A: It's getting there. One of the problems is you have compression algorithms and other variables. Aside from all the things that the semiconductor guys are doing, there's this latency delay in the network and that has to do with network design. There's no single component that's going to come in and be nirvana and solve everything. You have to have a good solid network that gets under 100 milliseconds before. I think what we found is it's roughly 200 milliseconds today. We've got a long way to go in terms of getting latency down in the public network.

A: Our perspective is we're in the process of actually doing extensive study work on voice over IP. It's a real interesting market. There's an awful lot of energy, an awful lot of investment. Along with that, of course, is an awful lot of hype. What you're seeing right now is a number of vendors— we've identified seven or eight distinct market segments within the voice over packet networks if you want to genericize this market. Best in class performance right now is something that sounds like a satellite connection from a latency standpoint and that's the delay between when I say something and you hear it on the other end. Voice quality we believe is actually going to become indistinguishable from toll quality. The fact is you have some pretty decent codex, you have some very decent DSPs that enable you to squeeze the bits through. The problem that you can't control, however, is the performance across the network. Right now the PSTN for an international call is just a little bit over 100 milliseconds, so if you're trying to provide comparable type service, you want to target something that looks like that. Best in class, as I said, is about 250 milliseconds, maybe 200 milliseconds over a controlled network. What's the attractiveness? For international calling, even though it's a little bit cludgy and a little bit clunky, it's really cheap. They're able to eliminate between 80% and 90% of the cost basis of an international long distance call because they're able to avoid settlements and access charges. Settlements and access charges for international calls are 80% to 90%. That means they can make a potload of money and still reduce the price by about 50%. Do these arbitrage opportunities last forever? No. They last for four to five years in a majority of the countries and in the developing world that arbitrage opportunity extends for maybe a 10-year period. If you take a look at these voice over the Internet companies, they're focusing primarily on the international marketplace and very few of them encourage domestic long distance calling mainly because the performance delta is more noticeable and the price delta is less noticeable.

Q: Working our way back to the enterprise for a moment. Does fiber channel have a place in the need for bandwidth quest? John?

A: You need to think of fiber channel as a network that's distinct from the local area network. The term that's been used for fiber channel is a storage area network. They are distinct networks because there's really only one common device that links the existing local area network as we know it today, which we would call a real true, more ubiquitous networking technology and fiber channel and that would be a server that would be common to both the fiber channel and the local area network. There really aren't any switches available, at least being widely deployed, that support fiber channel and Gigabit Ethernet interfaces. For the time being, for the foreseeable future, those networks are not going to be linked closely together. They're going to be quite separate. Fiber channel has its place as a medium for accessing storage devices and shunting information very quickly into storage from a switch of some kind, whereas the local area network is really a ubiquitous distribution system for interconnecting multiple users together into the network and ultimately to the wide area. The convergence just hasn't occurred yet in those network systems.

Q: A question I've been getting lately concerning gigabit—could we see a repeat of the Fast Ethernet scenario where we have 10/100 today, could we see 10/100/1000, assuming the price is right?

A: I don't see a demand. That capability is going to depend on the desktop acceptance of Gigabit Ethernet because what you have in the network are pretty much distinct areas where you have a backbone—an example, gigabit backbone feeding to Fast Ethernet and maybe Ethernet connections out to wiring closets and to desktops. If Gigabit Ethernet was to gain for some reason very high acceptance levels at the desktop, yes, something like that could occur but we don't see gig as a desktop technology certainly in the near term. There's going to have to be some type of tremendous growth in applications that require that amount of bandwidth to the desktop. You'll see small pockets of Gig Ethernet to the desktop for certain applications, digital pre-press, medical imaging, stuff like that.

Q: Here's a question to get the gray matter going. What about security of transactions over the Internet? No one has addressed this yet at our conference.

A: We didn't intend to. Security is a whole different animal. Our focus is more on the plumbing pieces of this stuff. We have other people within Dataquest that look at these types of applications. Certainly it is a lynchpin to broad deployment of commerce within the Internet, extensions of the Internet using corporate Intranets and another term that we hear now, Extranets. Extranets are like the Internet equivalent of EDI and a lot of the security aspects need to be in place. There are some standards going on. There's IPSEC and there's a variety of other encryption capabilities, all the standards that are out there that can be deployed in these environments already. It's not been our focus in the scope of this forum to talk about that.

A: Let me try to cover a piece of it, however. If you take a look at ADSL technologies, one of the primary competitors people see are cable modems. One of the trade-offs you're going to see between ADSL or xDSL and cable modem is going to be the issue of security. Cable modems run over a bus type of architecture versus their lower cost to provide service. Being over a bus technology, cable modems will be able to provide high performance at relatively low deployment cost because it's over a shared medium, whereas xDSL technologies typically require one-to-one relationship between the subscriber and the central office type equipment. The one thing that you are buying, however, by having that one-to-one relationship is absolute security because you don't even need to worry about encoding to the network. It's a private connection.

A: If you're looking at the enterprise LAN, security is going to take place where the gatekeeper is and the gatekeeper is the traditional router and the Layer 3 switch. That's the link between the subnets and other parts of the network. If you're designing security systems, they have to be designed to fit into those specific infrastructures in the LAN.

Q: A question in the wireless area. Wireless is another way to do the last mile. What about MMDS and LMDS? Any chances for it in the future, Brett?

A: I'll be honest with you. I hate the copper so I look at things like—I like things like radio because I don't know as much about it. It's like to know me is to hate me. Some of the challenges with technologies like LMDS and MMDS is the way that you need to allocate bandwidth. One of the difficulties is that you can't allocate on a statistical basis; you have to do it on a time division basis. That produces a constraint almost right off the bat, number one.

Number two, you typically have to transmit in the clear, so things like encoding, encryption, that kind of stuff, become a major factor. Number three, these technologies don't work terribly well when it rains. Those are some of the challenges.

Will there be opportunities? Yes, you'll probably see these technologies be most prevalent in the outlying areas. In fact, I recently saw a study about xDSL which I neglected to bring up in my talk, where if you take a look at all the copper loops in the U.S. and all the subscribers, even with today's technology you can only hit about 85% of those with ADSL services which sounds pretty good but the problem is when your objective is 100% coverage, it falls really short.

When you look at LMDS, MMDS and some of the other wireless packet type technologies, that may be a way to close your deployment net so that you can get the other 15% and actually advertise for service and not create a bifurcated society where you have people that are in range and people that are out of range.

Q: Going back to the enterprise again or, in this case, the SOHO market, can you comment on the role of 1394 versus 10 or 10/100? What are people going to connect up in a small office environment?

A: Whatever is cheap. Seriously, that's what's going to win. Simple and cheap. Right now I'd say Ethernet is—you have a new generation of PCs coming with built-in Ethernet adapters, plug it right in. Very low cost; very low cost hubs. You're getting Fast Ethernet hubs at $50 per port list. It's a real no-brainer.

A: We figure if the trends continue the way they are that the vendors will be paying users to take the equipment.

A: One note on that from the chip side, I know that particularly Intel is working to build 1394 directly into their main chip sets, so it might be they're free at some point.

A: If that occurs, then you'll have some choices that are going to be out there. Simplicity really counts a lot down there and price is another big issue and interoperability. Can it plug you into a lot of different devices? Can I do it without an engineering degree? Do I have to take out a loan? Can I get it at Frye's? Can I order it in a catalog?

Q: Back to cable modems. There's some clarifying questions here. Basically you have all of these alternatives, xDSL, ISDN, so forth, Don talked a little bit about cable modem. It's off and running. We've got the @ Home service up and going. We have $39 a month service. What do we think? You mentioned the financial wherewithal of the cable MSOs versus the other guys. Where is it going?

A: It depends on the area of coverage. Can you get it? I was talking with an associate who lives in Fremont and he has a cable modem and he loves the service. It's only available there. I'd love to have it. I can't get it because I live in Saratoga. How broadly deployed is it? In general, many of these technologies are going to be adopted if they're available. The issue is making them available. Why does Brett have ISDN jumping off the map here? If you're a stationary telecommuter, after years and years and years of waiting around, the RBOCs have finally put this stuff in and you can get it in a lot of areas. Can you get an xDSL loop? No. Can you get a cable drop from a lot of the cable companies? No. A lot of it is the demand there and what's available at the time. Eventually all of these technologies will probably merge and have some share of the market.

A: I'll have to admit I was a bit of a naysayer with cable modems initially. Sharing bandwidth, it's going to choke the network down, you're going to get crummy service. The fact is if you take a look at the numbers, cable modem—when I talked about wide geographic coverage and low density—because I'm not talking about everyone having computers at home that they want to hook up to the Internet. I'm talking about a relatively small number of people that will want to do this at really fast speeds. Cable modems do that reasonably well. The concern is the build-out process. I actually had a panel discussion at Interop where we talked about this with someone from Cox Cable and the cable companies are moving to modernize and update their plant for operational cost savings. In many respects their plant is being modernized to enable these types of technologies and it makes me a little bit more bullish about what they're going to be able to do. The challenge that everyone points out about them and the reason people throw darts at them is the concern that what if a bunch of people log on. If a bunch of people buy cable modem service, that's a really good problem to have if you're the cable company. The price points they're going to hit are relatively decent; their capital exposure is relatively low. The one piece in the whole argument that has me a bit concerned is the motivation for updating their plant and it's being driven primarily through operational cost savings. Being in a carrier company, I've seen projects that were meant to reduce operational costs disappear off the priority list when budget crunch time comes. That would be my only concern about the deployment.

Q: A last question I'll direct at our two hardware guys. Basically the framework is how can the chip industry play? We have people out there adding value to what they do. I'm hearing even chip guys doing Layer 4 capability and what they can do, so it's an arm wrestle for value. From an equipment vendor's point of view, where are they going on their road maps in terms of features into the future, what are they willing to surrender perhaps to a semiconductor supplier in terms of subsystems and where are they going? Where is their value-added space moving?

A: Any well established mature technology that can be integrated—something that's not in flux so there's not a standard that's emerging and we've seen this already. You guys have already adopted this in terms of higher levels of integration in terms of max and Ethernet interfaces. You see more integration, fewer discreet components, fewer discreet chips that ride along with the development in terms of semiconductor technology. Probably the biggest opportunities are those that are going to be well documented, fairly fixed in terms of the maturity of that particular technology. One of these things that people are doing now is a this Layer 3 switching we talked about. It's still a dicey situation. IP routing and using protocols like OSPF and Rip and some of the other protocols that are out there take time to mature and get to a solid code base where you can fully implement those types of software applications in silicon. Vendors are a little risk-averse in terms of their design approaches because they need to have the flexibility that they need in order to modify the code, in order to tweak it. If they run into a problem downstream in order to be able to correct that without having a spin at ASIC, so you need to understand what they're going to and try to focus at the things that are pretty solid that you can move. Another editorial comment—Layer 4 isn't cutting it. That's just some marketing hype. It basically is routing, so let's just call a spade a spade here.

A: I'd say follow what's happening in the TCP IP space very closely on the LAN side because that's going to be the predominant protocol, that's the area that vendors are going to focus on in terms of their support down the road with new products and that's a more focused space for chip makers to be involved in in terms of reaching a larger market. Beyond that, it's hard to predict what technologies are going to achieve standardization down the road but certainly Gigabit Ethernet is an excellent one to follow. On the Layer 3 space, once vendors and standards committees start to get their act together regarding things like QOS on Gigabit Ethernet, that'll be another opportunity for ASIC's integration and added value on chips.

A: I'll put on my voice over IP hat for a moment. If you take a look at all of these voice over packet network type services and technologies, they gobble MIPs like crazy. If you take a look at what's going to solve the latency issue—I'm a software guy so it really pains me to say this—but it's going to be a hardware solution, not software. The only way to get performance in latency down is to embed it into the hardware. It's going to be a hardware-oriented type solution and there are some opportunities to tweak some things down, perhaps look at other transport layer technologies like frame relay, which may be able to provide the performance characteristics that will actually provide a commercially-viable grade of voice telephony service over a packet network.

Chapter 18: Panel Discussion - Major Device Trends and Forecasts: Looking Ahead

Panelists

Nathan Brookwood

Principal Analyst
Personal Computer Semiconductors & Applications Program
Semiconductors Group
Dataquest

Bryan Lewis

Director & Principal Analyst
ASICs Worldwide Program
Semiconductor Group
Dataquest

Jordan Selburn

Principal Analyst
ASICs Worldwide Program
Semiconductors Group
Dataquest

Jim Handy

Director and Principal Analyst
Semiconductor Memories Worldwide Program
Semiconductors Group
Dataquest

Here's what we want to do with this session. We want to talk about three hot areas in the device area—embedded DRAMs, SLI and microprocessors. We're going to have four Dataquest analysts come up and make three 25-minute presentations.

The first will be Nathan Brookwood, our principal analyst for computer microprocessors. He will be followed by Bryan Lewis and Jordan Selburn, the principal analysts in the ASIC SLI area. Then Jim Handy, a principal analyst in the memory area, will talk about embedded DRAMs. We're going to run all three of these one after another. Please hold your questions. They all hang together so after they're all done, you'll see the synergy between these three talks. After the end of the third talk, we'll call for questions. Write down your questions. Someone will pick them up and bring them up here. When we're all done, we'll have a 15-minute session in which all the speakers will sit down, we'll ask them the questions and then take it from there. Nathan Brookwood, our principal analyst for computer microprocessors will tell us the latest happenings in microprocessors.

Nathan Brookwood: Good afternoon. We're going to be talking about the X86 market and the impact that the clones—AMD, Cyrix, IDT—may have on that market. We'll look at how we anticipate Intel will respond to this, how the clones will respond to what Intel does and finally, we'll finish up with a discussion of whether or not this time is going to be any different than the last times, or if we'll simply see the dead bodies of a few more companies lying in Intel's path.

Before I got to Dataquest, I used to be very interested in the RISC market, and everyone would say, "Well, why do you waste your time on that? It doesn't really matter." It wasn't until I came to Dataquest and started looking at numbers that I finally got it. This is a chart of the computer microprocessor market for the next five years. As you can see, out of the total market of about 83 million units, something like 77 million units were X86. Then the PowerPC, SPARC, MIPS, et cetera, split up the rest. If you're sitting in the back of the room, you may not be able to see the area that's above the green but let me assure you, there is some there.

We made this forecast before Apple eliminated the clones, its cloning policies, so next time I suspect we may see a slight shift where the PowerPC, which is the yellow stripe, goes down a little bit. I don't think you'll be able to notice it on the chart.

Looking at 1996, you can see that the market basically consisted of Intel and the seven dwarfs here. Intel actually gained market share last year. IBM, Motorola, and AMD all lost a little bit. We'll see that going forward. Here, I show the entire computer microprocessor market. We've got the SPARCs and everything in there and you can see that was about a $16.5 billion market. Of that, Intel

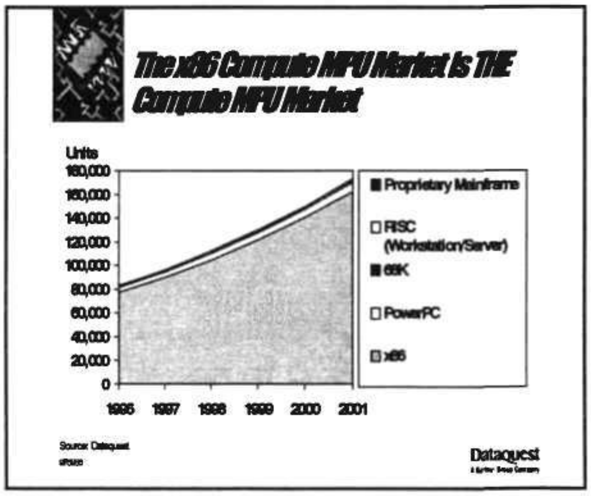

had about $14.5 billion. This is not one of the slides they intend to use with the FTC, I'm pretty sure.

Nor will they use this one, where we take the X86 market alone and get rid of all the peripherals like SPARCs and Alphas, which probably will happen anyway. Here, we can see that the market was about $15.5 billion. Of that, Intel had about $14.5 billion. Everyone else split up $1 billion. It was not a year anyone wants to think about, if you weren't Intel.

As a matter of fact, if you look at it from a market share standpoint, Intel walked away with 96 cents out of every dollar that was spent on microprocessors that go into PCs, workstations and Intel-based servers—96 cents—and everyone else argued over the four pennies that remained. On this slide, it obviously looks like Intel gained market share, from '95 to '96, but that's a little deceiving. You'll notice that we started this chart at 80%.

Basically, Intel simply increased their market share by a little bit. They went from about 91% to 95%.

Looking forward, we see where the X86 computer microprocessor marketplace is going to increase by a little bit less than 2-1/2 times from 1996, where it was 77 million units, to the year 2002, when it will be something in the neighborhood of 180 million units. That's going into PCs, workstations, servers—all based on X86 architecture—and excludes things like SPARCs. That's pretty good growth. As you see here in this chart, the majority of the growth, over the next few years, will be in the red area, which is what we characterize as 6th generation processors. The sixth generation processors include the Pentium Pro, the Pentium II and AMD's K6 and Cyrix' 6X86-MX. You can no longer tell whether a processor is a fifth generation or sixth generation processor—whatever that means anyway—by its socket. That's

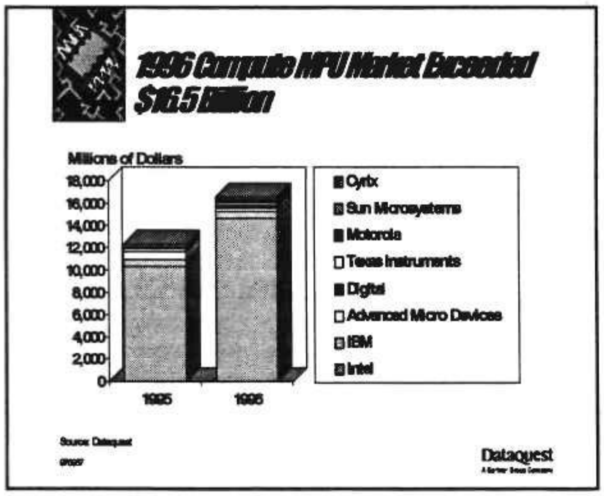

an important distinction we make going forward.

Revenues over the next six years are going to increase by a factor of three, from a little bit under $15 billion in 1996 to almost $44 billion in the year 2002. If you do the arithmetic and say, if computer microprocessors are going to increase by a factor of three in revenues and a factor of 2.4 in units, what does that mean? It means that the average computer microprocessor is going to go up in price. That's a counter-intuitive thing, since we all know that the prices of these things comes down on a regular basis.

Next week, Intel's going to announce that they lowered prices again and everyone is going to say it must be a competitive response to AMD. Tom

Kurlak from Merrill Lynch will lower his estimate on Intel. Everyone will sell. Merrill Lynch will pick up a little money on commissions. Then, about six weeks later, Intel will announce they had a pretty good quarter after all. Tom Kurlak will say "buy," the stock will go up, and Merrill Lynch will make quite a bit of money on commissions.

This has been going on for a long time, and I'm amazed that people haven't broken the pattern. *The Wall Street Journal* printed it on page one a couple of months ago, but Kurlak continues. There is a document, that we have coming out at the end of the month, that goes into the mechanism that allows

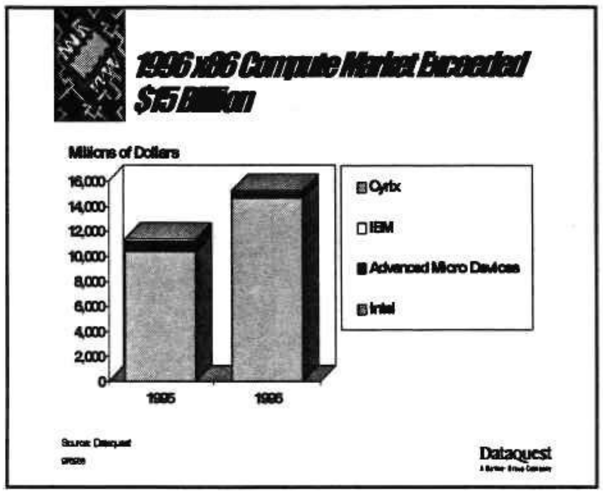

Intel to keep lowering its prices while people pay more for what they're buying.

The non-Intel portion of that market—this is a split based on our projections of what Intel will be selling, which is the solid area and what the clones—AMD, Cyrix, IDT, and others to be named at a later date—will be selling during the same time period. What you can see here is that the clones will be taking more and more of the revenue. They will go from about $770 million a year ago to $7.7 billion in the year 2002. That's not a bad business. It's not quite what Intel will be doing at that point but still $7.7 billion is enough for people to want to chase that market.

The non-Intel computer microprocessor shipments will increase by about a factor of five during that same time. The reason, that they can increase by a factor of five while their units are going up at a slower rate, is because their ASP is going up from where it was last year, which is easy, because last year their ASP were abysmal—I'll show you that in a minute. Again, a reasonable amount of unit shipments and revenues demonstrates that you can play in the X86 market even if you're not Intel.

As I mentioned, the ASP continues to increase. It's been bouncing around $200, for the market as a whole. It does that even though we see that the Pentium products have come down in price. They're about $200 this year. They'll be way under $200 next year. The Pentium Pro or 6th generation products in 1996 were $600. Probably not too many of you bought them in 1996 because they were so high. This year, they got a little bit more reasonable. Next year, 1998, they'll be phenomenally reasonable. You'll go into the computer store and you'll be able to buy a $1,200 or $1,300 machine with a Pentium II in it or some other sixth generation processor like an AMD K6. Why not? For $1,200, might as well go with a 6th generation part. That's what drives the volume here.

Somewhere in the Year 2000, Intel will introduce the

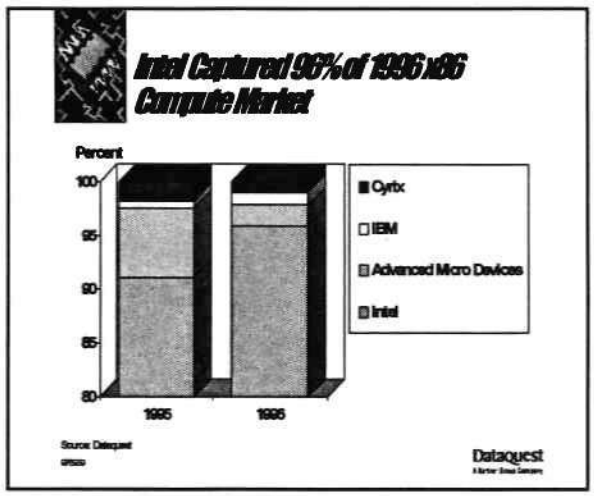

Merced products—the seventh generation. Those will be very pricey. You probably won't be interested in running out to buy one of those when they first come out. As a matter of fact, we show it coming down to $600 by the year 2002. Even then, you probably won't see them in mainstream PCs. They will continue to be laboratory curiosities, and things that power users buy, rather than things that consumers buy. The seventh generation or Merced products—IA-64 — will have a relatively slow penetration into the market because of this pricing characteristic.

Last year, Cyrix tried a disastrous experiment early in the year. They decided that since they now had Pentium level performance, they could charge Pentium level prices. The market told them, "ah-ah-ah, doesn't work." They have now learned their lesson. They are charging this very noticeable discount from what Intel charges. We would expect that "Intel Inside" will continue to command a reasonable premium over what the cloners can charge.

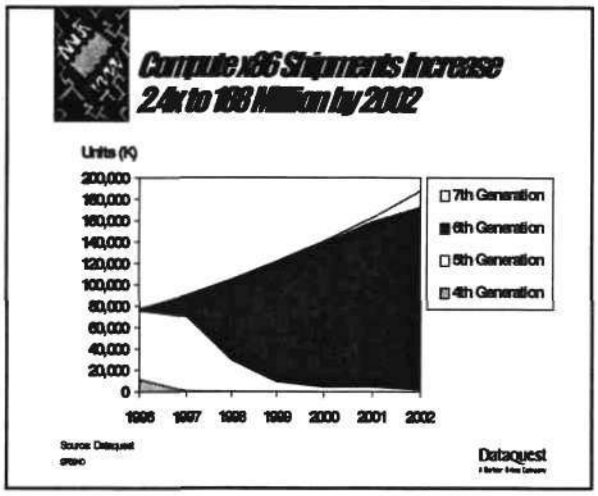

In fact, the point of stability in the market is that companies like AMD and Cyrix—if they can sell out their production at average selling prices of $100, they're ecstatic; whereas, if Intel sells out its production at $100, they're very unhappy, the stock market is very unhappy, revenue growth goes all to hell and that would be a disaster for them. The stasis here is that the cloners pick up the low end of the market below $100 and they're ecstatic, and Intel continues to focus on products that average way above $100 and go over $200 and they're ecstatic, and everyone lives happily ever after.

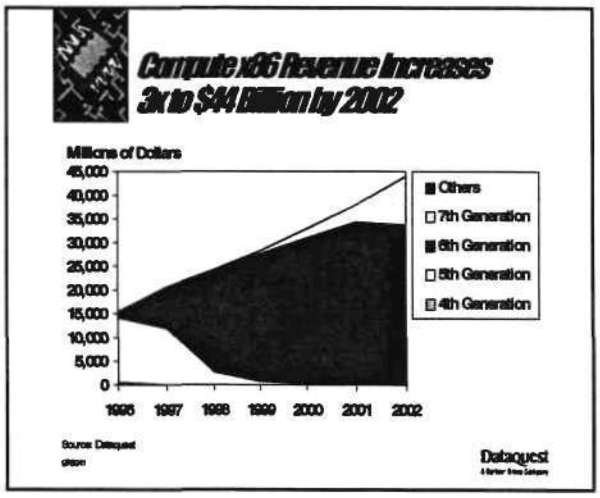

Let's take a look at what Intel is doing and what they'll be doing over the next 12 months. Fundamentally, we think of Intel as a semiconductor company, but they're really a printing company. They've learned how to print dollar bills using .35 micron lithography. Of course, you can't print that many dollars when you're using .35 micron lithography so they're going to be going to .25 micron lithography— already doing that in terms of their mobile products, and soon their Pentium II products, so they can get more dollar bills per wafer and increase their revenues that way.

Even so, making the products smaller, they don't have enough manufacturing capacity. Throw up a few more fabs here and there—another $4.3 billion—the estimate for next year is actually up to $4.5 billion now, in new equipment—and turn out more dollar bills. They introduced MMX at the beginning of this year. Last year, we were really worried. Everyone knew MMX was coming in January. We were afraid buyers would hold off, and not buy machines in the fourth quarter, knowing MMX was coming. That didn't happen by and large. It was a reasonable fourth quarter.

However, MMX has had a pretty turbulent effect on the market this year, in that those guys in the dancing bunny suits convinced consumers that this is what they wanted—not only consumers but even corporate customers—and it became harder and harder to sell non-MMX Pentiums or K5's or M1's, to the point where they almost became a glut on the market. You couldn't move them at normal pricing, and Intel has had to continually discount its prices in order to be able to move the product. For a while, they hadn't really geared up to have them cross the product line that way.

You started to see all sorts of dislocations. Intel has addressed that issue to a large extent by accelerating the rate at which it brings MMX from the high-end and mid-range, which is where it started out with

$400 or $500 processors, down to the low-end—$100, $150 processors. To blunt the initiative of the network computer—that $500 thing that really cost $700—they introduced what they call the network personal computer, or NetPC, which is really a PC that they took the floppy off of, put on a lock on the box so you couldn't open it up and then put in your own add-in cards and said, "Now organizations can use this."

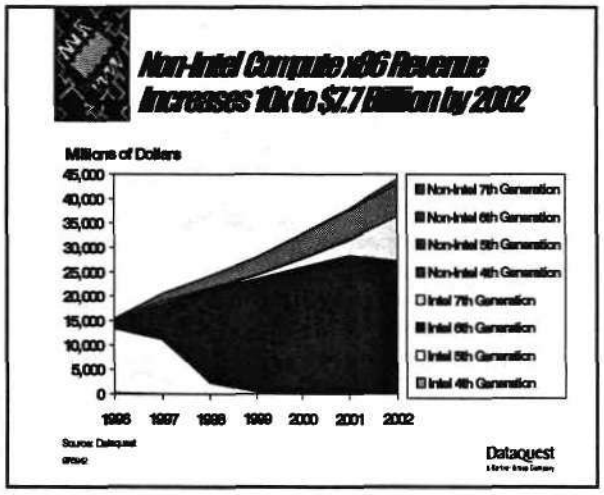

Heaven help the poor guy at his desk or girl at her desk who tries to install some software that MIS hasn't authorized, because that's where quite a bit of the support costs come from in PCs: people get stuff in the mail and they install it on their machine at the office. Then the PC doesn't work any more and MIS has to come in and untangle it; that takes hours and everyone complains about how unreliable PCs are.

Intel said, "If you want to sacrifice flexibility, we can sacrifice flexibility for you. Here's a sealed box, now go away." That's let quite a bit of wind out of the sails of the net computer movement. Intel, this year, has been pushing their Mobile Module concept, where they put the processor, the cache and the north bridge of the chipset all on a little module that can be dropped into a laptop computer.

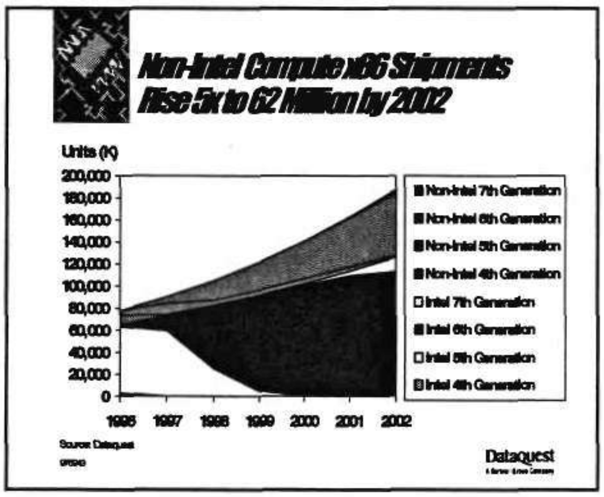

This is not intended to allow end-users to upgrade the product, but it certainly makes it easier for system assemblers to adapt to newer versions of the processor as they come out. The ultimate test of that will be next year, when they put a Pentium II on that mobile module and tell the vendors, who have been geared up to support the current Pentium laptops, to go ahead and move to Pentium II. You'll see a very quick transition for high-end laptops to Pentium II early next year.

Regarding AMD—last year, we were all waiting to see how the K6 would turn out. AMD had been suggesting that everything was under control, the design was complete, and performance looked good. They launched the product in April. It really did look like a very competitive product. The die is smaller than the Pentium II, the performance was arguably in the range of the Pentium II.

I say arguably because clearly, on quite a few benchmarks, AMD's chip comes up on the slow end but in fact, for quite a few things that most people do in their offices, even in their homes, the K6 was certainly as fast as most of the Intel chips and a whole lot cheaper, and you could build an argument that it was comparable. That's what AMD was doing. It looked like they had enough fab capacity to produce millions of them and we were really optimistic.

As the year has progressed, however, our optimism for them in this near term has somewhat abated, primarily because of yield issues: they shipped 350 thousand units in the second quarter. That's an average selling price of $287, which is more than five times what AMD averaged last year. Wow, what a deal. They were going to do 1.5 million units in the third quarter. They ended up doing about a million, throwing away half a million units—literally. They didn't work at all, so they put them in the trash; probably will sell them as earrings and trinkets and paperweights, and give them to Mike Slater to put in portfolios and things like that.

That translates to $100 million in scrap as opposed to bottom line. Really disappointing. They're trying to do 2 million this year. The jury is still out on whether they have conquered their yield problems. This is a real problem for AMD, and a real wildcard for how the year will turn out, in terms of clone shipments.

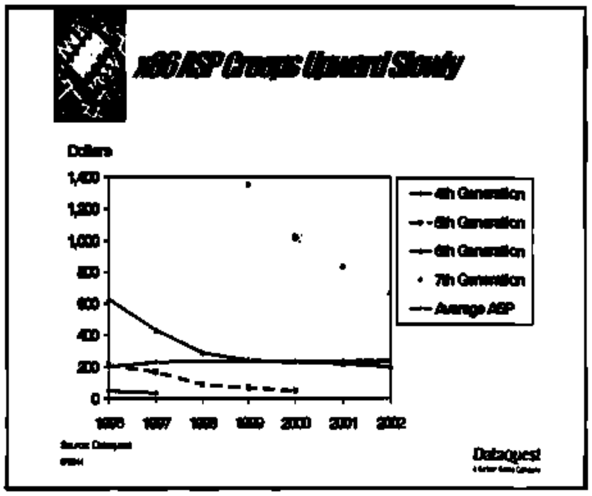

Cyrix and IBM have been a bit more successful. They've had yield problems, but they haven't been as visible. Compaq introduced the Presario 2100, based on the Cyrix Media GX, in March and it has been an incredibly high-volume product for Compaq—their second highest volume product across that product line in the previous quarter. Very, very successful. It has enabled Compaq to really tackle the low-end market, and still make some money at it, which is otherwise tough to do with $100 microprocessors from Intel.

They'll be adding MMX to that processor next year, and moving it to a higher performance course. They've got a whole strategy for high-performance integrated processors going out a couple of years. They also launched their M2 product, the 6X86-MX. Very aggressive pricing. They learned their lesson last year. Don't try and charge prices on a par with Intel; give customers a huge discount so they can't resist. That has been fairly successful.

You've got Acer using it in some products; you've got IBM using it in a variety of its PC products. Now that National is about to acquire Cyrix, that does a couple of things for the company. First, it will eventually give them a lower cost of manufacturing than they've had in the past, when they've worked on a foundry basis with IBM. Also, National intends to really try and focus quite a bit of the effort onto Internet appliances. Now, those don't have ASPs of computer processors. They'll try and stay in the computer processor business as long as they can.

What's Intel going to do? Intel is going to do what it's been doing for the last several years—drive performance. They want to drive performance as high as they can to put some competitive distance between themselves and AMD and Cyrix, to encourage all the people who bought a machine a year or two ago to say this machine is too slow, I need another one.

I expect that by this time next year, we'll see 450 MHz Pentium II's become available in mainstream PC product lines. They're going to try to use their overall slot 1 or Pentium Pro-bus architecture with AGP (accelerated graphics port) to differentiate their products. Their newest chipsets support the AGP, and they want to convince everyone that if you want AGP, you're going to have to get a Pentium II—and that will drive the market their way.

They'll also broaden their "Intel Inside" campaign to cover the entire world; in some parts of Asia, they still haven't heard "Intel Inside." That's two things that they have to do in these less-developed countries. They have to learn how to make a phone call and they have to learn about "Intel Inside." They'll both probably happen next year. Intel is not really interested in a price war. They're still capacity-limited. When you're manufacturing capacity-limited, it doesn't pay to lower your prices to increase the market because you can't ship additional product. Fortunately, Intel went to the same business school as AMD on that. I really don't anticipate that there will be a price war here.

AMD's goal is to scale the K6 performance to match what Intel is doing as best they can. They also better address their yield problem so they can ship something, or else no one will care about their performance. They're going to increase their performance by adding a level two cache onto their chips. That eliminates some of the problems that led Intel to go to the slot one architecture with cache chips mounted next to the CPU.

AMD also plans to add 3-D capability and fast vector floating point operation, to speed up the 3-D transforms for games and other 3-D operations. Intel would like the world to move to slot one and, therefore, is not adding AGP to any of its Pentium chipsets. AMD and Cyrix will be killed if they don't have AGP in that Socket 7, the current Pentium socket, arrangement. They're working on their own and with all the chipset vendors to provide AGP for the current Pentium-style bus structures. How well that works and when that shows up will be an important issue for how well AMD does next year. AMD is also committed to providing a 25% price advantage—that's on the microprocessor, not on the system—to keep people coming to them to buy their products. They're not interested in price competition.

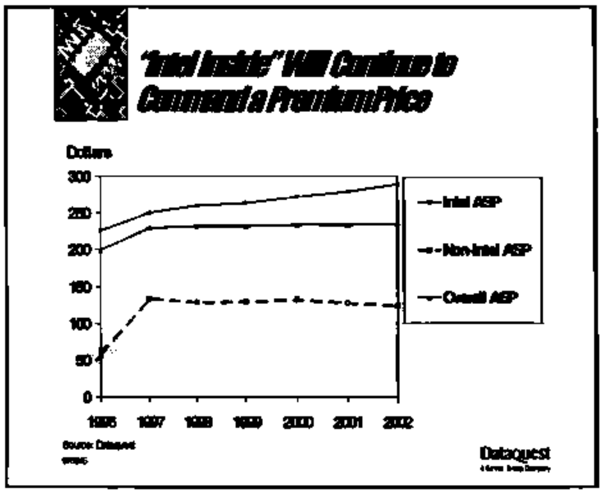

Cyrix wants to scale 6X86 performance. They'll also be working on improving their integrated cores. They need to have AGP for the same reasons AMD does. They're all going to be increasing the speed of the Socket 7 interface to 100 MHz to overcome some of the performance bottlenecks and the access to the main memory. Finally, they'll make sure that there are some Socket 7 motherboards out there. They'll also continue the GX.

The question is why is it going to be different this time than in the past? What's really different? I would argue that, this time, AMD has really coalesced their manufacturing capacity, their process technology—they're close to coalescing their process technology—and their product design, to be able to compete more effectively with Intel. They can also have two development teams and can develop two processors simultaneously via the NexGen acquisition of almost two years ago. You don't have to have them both on the market at the same time, but this allows them to stagger their developments, and more or less match what Intel's doing.

Of course, with the two teams, the assumption is that they're both going to work, that both teams are going to hit on every shot. When one of those teams develops a product and it's late or doesn't hit its performance target, they'll have a problem, as Intel does, but Intel has three separate groups doing this, and Intel can cover a little bit better than the cloners.

Intel's performance advantages in the current Pentium II system, in the areas of floating point and so forth, really only help them in certain niches. They're not applicable to the broad set of office productivity environments. That means that it's possible to buy clone-based systems, that don't have quite the performance of Intel, and still not suffer from it in terms of their usage models.

Finally, because of the disparate market shares and antitrust considerations—of course, the Federal Trade Commission hadn't yet rung Intel's phone at the time I put these slides together—that simply heightens the case—I can't really see or imagine that Intel's going to be terribly aggressive on price or try to sell things below cost simply to unhook its competitors.

Is this a sure bet? Not by any means. Execution remains a major concern. AMD's execution in the last quarter has been less than perfect, and that's a problem. Intel continues to execute with very few hiccups. The near-term concern is whether the AGP performance in Socket 7 is close to what it will be with slot one. If it is close, then at least from a performance standpoint and product spec standpoint, the cloners still have quite a bit of running room.

If AGP doesn't show up on Socket 7, or if it runs poorly on Socket 7 , that's going to limit their opportunities. We also need to see how Intel comes out with their next generation 32-bit designs, and how AMD stacks up against that. In the longer term, we'll need to see what Intel does with its 64-bit Merced architecture and how the cloners compete with that.

The new ingredient in the microprocessor scene for desktop computer microprocessors—competition.

It's there; it's struggling but it's there—more than we could say last year. AMD and Cyrix have a window of opportunity. The degree to which they succeed will be very much a function of how well they can execute over the next year or so.

Moderator: Thank you. Let me now introduce Bryan Lewis. Bryan is our principal analyst and director of the ASIC/SLI service. We've simply changed its name to be more representative of the marketplace. In any event, Bryan's been with Dataquest for so long that he actually used to work on applications-specific vacuum tubes before taking over here.

Bryan Lewis: Thank you, Nathan. It's a pleasure to be here today. As Nathan mentioned, I have been here for awhile. It's interesting to reflect, being back here in San Diego. This is my 14th straight Semiconductor Conference.

System level integration. You've heard quite a bit of talk about it over the last couple of days. Why is that? What's the key ingredient to what's making this market start to take off?

The concepts of putting an entire system on a chip have been there for years but the market hasn't taken off. Now, one of the key ingredients in the manufacturing technology. We now have.5 micron,.35 micron and we're really moving into the .25 micron area strongly, and we'll shortly have the .18's. With the .25's, we can really get into having over 3 million gates on a chip. That gives us quite a bit of real estate, so now it is practical to put the microprocessor on the chip, put the memory on the chip and the logic and really have the entire system on a chip. The manufacturing technology is the key enabler that came around recently.

There are quite a few system-level macro houses or slams or however you want to talk about it, but the macros are key enablers again. If you don't have design reuse, which Wally Rhines clearly went through very heavily, you can't do this type of product. The macros are out there. We have quite a few third companies out there, EDA companies out there, system suppliers out there and ASIC suppliers out there developing these macros and that's really enabled the technology to take off.

Applications—it's clear that there are plenty of applications out there. When it comes down to time-to-market, if you don't have design reuse, you really don't play. The emerging markets that are really starting to take off right now are set-top boxes, quite a bit of the wireless communication stuff, the DVD, DVC (digital video camcorders); they are starting to adopt this technology very rapidly, because they must shrink the system, get the performance up. System on a chip is here.

We have a "great news and very poor news" scenario. The great news is that not everyone's moving into this intellectual property area if they can start to get some of the system profits, because that's where the value is. The bad news is that everyone in the entire industry has figured this out, and is trying to jump in this market in one angle or another.

A key example is Intel, and looking at the IP value chain. Nathan went through quite a bit of the Intel story here. A key example is seen in the PC market. We see that the PCs remain relatively constant at about $2,000. Meanwhile, Intel's ASPs have risen from about $177 up to $230. As you can see, they have captured a higher percentage of the revenue, as well as the profits, and that's why they're so profitable out there. He who controls the IP is in the driver's seat.

This graphically depicts it. You can see that their

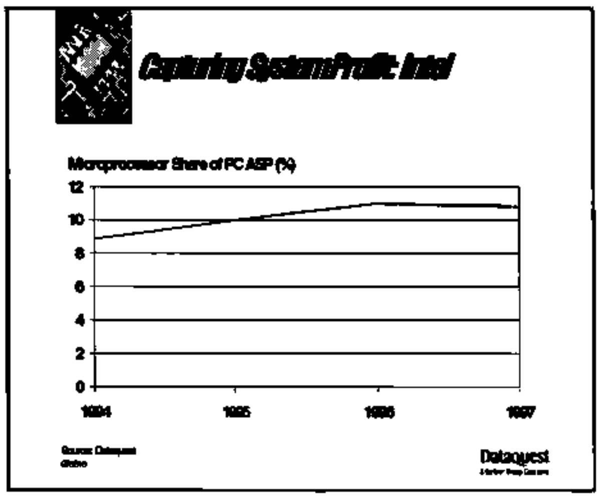

market share is growing as a percentage of the entire system.

Today, we're going to cover the system level trends from, more or less, a high-level view. Then we're going to look at this battlefield which we talked about, and see how everyone's entering this market and whether everyone can stay in there? What are the implications five years down the road? There are quite a few changes that are going to happen in this marketplace. Then we will give some strategies on how to be a clear winner in the long haul.

You've heard what system level integration is, but let me run through it real quickly. It's conceptually taking large functional blocks—such as the microprocessor core, PCI cores, a variety of other cores—and integrating them all onto one single chip. One of the key ingredients here is the microprocessor core.

Looking at the forecast here, this shows you in the yellow bars what's non-system level. Those are basically vanilla ASICs. You can see that they're already starting to peak out and will start to decline. In the red bars you can see, in 1996, that the market was about a $2 billion market and that moves up to approximately $4 billion in 1997. When you look at the year 2000, you see that system level integration is over half the marketplace. This is in terms of revenues. If you're to look out in terms of designs, we'll see the crossover closer to 1998.

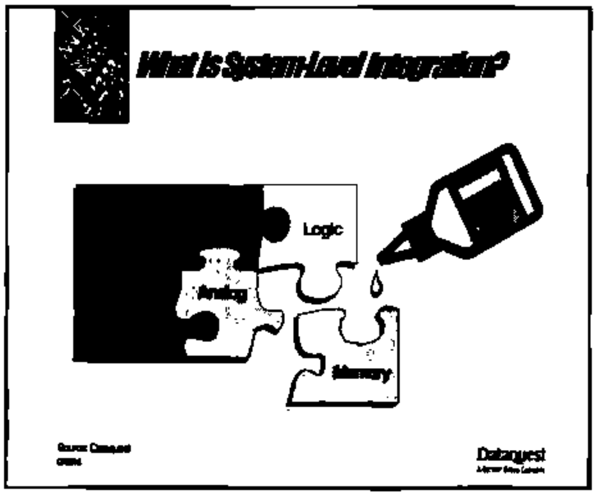

This market is taking off very rapidly. One thing I want to caution you on, system level integration, there are ASICs—which this forecast is—and then there are ASSPs, application-specific standard products. Now, you have to add another, almost equal to what revenue is out there for SLI ASSPs, in that forecast to look at the entire system level market. We are in the process of analyzing these markets in quite a bit of detail, looking at the ASSPs and the ASICs and the tradeoff by each application market. We'll have some new data out within the next six months on this marketplace. As you can see, it's taking off very rapidly.

It's not all rosy though. There have been some stumbling blocks as people have run into these things. Gary Smith talked quite a bit about methodology. Methodology is critical, going forward here. The key people that got this stuff out on the street first were the people who had been in it, had large system business and had quite a few investments in it. IBM is clearly one of the winners in this area. They have been pushing very hard. They've learned quite a few things internally. They have a very good design flow, which helps them get these products out the door. Methodology—there will be a large focus on methodology, as Gary Smith talked about.

The core standardization—Wally Rhines also talked about VSI, Rapid and then there is the ASIC Council that's out there as well. They're all working on standards. If you have different intellectual property out there from different vendors, in order to mix and match, you really have to have some standards, so that they can communicate with each other.

There is clearly the IP protection issue. How do you continue to get revenues back from your product? How do you protect that? How do you protect it on the Internet? How do you protect it so you get these revenues in the future? These things are being conquered. They're being looked at very quickly, and this market is poised for some rapid growth over the next years, and will hit the mainstream shortly.

As we talked about, there is this battle going on. It's among the system suppliers, the OEMs, the semiconductor manufacturers and the third party suppliers. Some of these can even be EDA suppliers as well, as Wally Rhines went into. The bottom line is, "he who offers the values will get the profits." Jordan is going to take you into this battle zone and explain it more.

Jordan Selburn: Good afternoon. The concept of value in system level integration is really driving a fundamental change in the structure of the ASIC industry. If you go back a number of years, the value an ASIC company would bring would be in their service, the applications engineers, people working on layout—because the product itself was actually fairly generic, even though the name is application-specific. All the parts were five volts. If you go back, they were gate arrays. There was no such thing as an embedded function. The value came from the people at the company helping you do your design. Now, things are changing quite a bit as we go into the area of system level integrations.

Let's talk a little bit more about value. Value is a great thing, right? Value is great if you've got it. It's not so good if you have to buy it from someone else. Taking the PC, if you were to ask Compaq or Dell what they thought of Intel's value, privately, they're not as pleased as Intel is.

We've recently completed our 1997 user survey. We received a record number of responses. One of the questions we asked is, "In your system level designs, what makes up the core area of the die?" For the current year, we see it's a split—about half random logic, about one-third memory and one-sixth cores or system level macros.

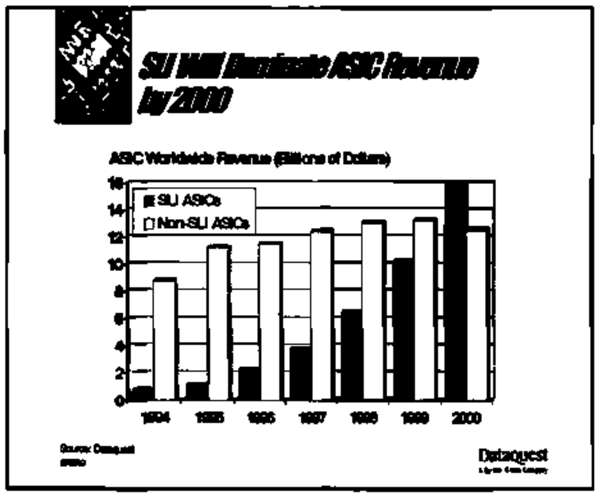

Going into 1998, the memory stays fairly constant, but we're seeing a shift from random logic to cores. Cores are going up substantially—actually about 40%, relative to where they were in 1997. This is a significant space change, because the entity that adds the value in each of these areas can be different. As an observer, Dataquest can see this as really more of a warfare in which everyone's trying to get everyone else's profit and the customer-supplier relationship doesn't always have to be quite so benign.

Let's a take a look at each of these areas, and establish who's adding the value in those various areas within the core. If you look at memory, it's a function of the type of memory—embedded DRAM, embedded flash—those are relatively new. They're not available from many suppliers. It's the semiconductor manufacturer that adds the value in those areas.

For more commodity SRAM and ROM, those could have value-added, if you take a look at a company like Artisan Components. They develop very high-speed memory compilers and, as a result, they bring some value to the party which they share with the semiconductor manufacturer. Most memory doesn't bring much value to the table, because it's available so commonly.

The logic—this is where the system designers apply their expertise and turn it into silicon. They're the people adding the most value there. The semiconductor manufacturer can also add some value here, if they've got a very high-speed library or low-power library. In this case, the value would be shared between both the system designer and the semiconductor manufacturer. The system designer wants to be able to retain this value—as was discussed in great detail this morning—by taking this intellectual property and turning it into reusable cores. This is something that they can maintain control of. The control of intellectual properties is really a key.

If we look at the system level macros—here, I refer to them as core—they can come from any one of a number of parties. That can be something designed by the system designer, from a previous design and reused, it can come from system level macros developed by either the semiconductor manufacturer or by a third party, or from an EDA house as well. The person who owns it and controls it brings the most value.

In general, there is some value for the "distributor" of this, which would be the semiconductor manufacturer. That's something that changes over time. We've talked about cores that have value, we've talked about cores that are really generic and don't have value—but it's not all digital, it's not black and white. System level macros and other forms of intellectual property have a shelf life and a sell-by date over the course of this lifetime, and the person who commands the value changes.

Again, from our user survey, if you take a look at the owner of the system level macro, the place where it originated, there are some dynamics there also. Even today, the majority of system level macros are developed by and ownership retained by the system designer. We're seeing greater and greater shares going to the semiconductor manufacturer and also to dedicated IP developers.

It's not necessarily so rosy a world. There is a battle going on for the control of this intellectual property, and the value that it brings. The OEMs want to use the intellectual property to sell more boxes. The semiconductor companies want to retain ownership of their IP, and use it to get more of the dollars that the system designer was going to get.

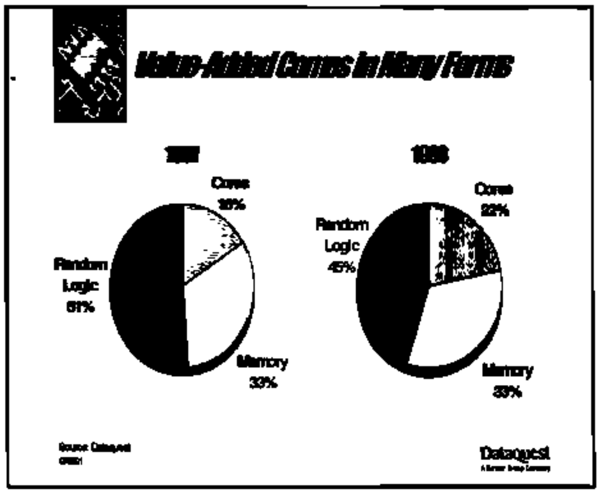

Going farther back in the food chain, the third parties want everyone to have their IP. They don't care which semiconductor manufacturer or which system designer. They want to distribute it as widely as possible so they can capture the most value. VSI is moving in this direction; it's becoming a marketplace where people can buy and sell system level macros, particularly. It's mainly going to be the semiconductor vendors and system house doing the buying, with the third parties doing the selling.

I've talked a little bit about how the ownership of this is changing over time, and I apologize for putting an eye chart in here for the people in the back. I'll go over it a little bit, because this especially details how the value of intellectual property changes over time, and I've provided some examples as well.

Going across the top of the chart, we've got the availability because Economics 101 does rule in this marketplace. It's pure supply and demand. The number of suppliers is the availability. Usage, that's the demand for the system level macro; and the owner and the distributor, the silicon vendor, are fighting to control the revenues from this supply/demand balance.

At the introduction, typically, there is only one owner. This is the entity that came up with the idea for the system level macro, and he's generally only made it available for limited distribution. A semiconductor manufacturer is only distributing it—in that they benefit from having it embedded. An example of this is embedded flash memory which today is only available from two or three sources. Not too many applications benefit from onboard flash but that's starting to grow.

There is much value to be captured by both the owner and the semiconductor manufacturer. In this case, it's the same companies. As system level macros get wider and wider usage, the ownership is still limited but they're available through more and more sources. An example here is the Oak digital signal processor from the DSP Group.

The DSP Group owns this but there are a number of semiconductor distributors or manufacturers that you can get this from. There are more and more applications that are moving towards using on-chip DSP. There is quite a bit of value for the DSP Group. It's starting to become less so for the people who distribute or manufacture this core in silicon.

Taking that one step farther, to a more mature phase, if we take a look at the ARM processor. Again, there is one owner of it, although there is competition for a number of other cores such as SH, MIPS, X86, et cetera; but there is only one owner of ARM, while there are over 20 licensees. All the value of the ARM processor—all the incremental value—goes to ARM and none of it stays with the semiconductor manufacturer.

I've gotten a number of calls from semiconductor manufacturers asking what price premium can they charge by virtue of putting an ARM core on-chip. The depressing answer is basically zero. It all goes to ARM, because the customer can always go to one of 20 other sources and get the same processor, or they can go to a different processor architecture if they so choose.

Similar to the mature phase are products which become standard. An example here is PCI. This is something that anyone could go out and develop, but there are also sources such as SAND where you can buy it and bring it into your process. Again, very little value for the semiconductor manufacturer. Most of it—although there is not necessarily much—goes to the owner or the distributor of the standard.

Looking at this graphically, over time, I show here how the value is split between the silicon manufacturer and the owner of the intellectual property. At the far left-hand side, you'd see things that are brand new, hot pieces of intellectual property. There is value to be shared by everyone. Over time, if you look in the middle of that, maybe you'd see the OAK DSP. Towards the far right, you'd see something such as the ARM processor where the only benefit to semiconductor manufacturers is that this allows them to sell the rest of their silicon.

Let's take a look at the potential outcome of this in five years, 10 years and paint a picture for you for what this fight for value could lead to in the ASIC industry. Almost all of the ASICs are going to become system level. That's almost inevitable. As the size of these components and availability becomes more widespread, it becomes more and more compelling to bring them on-chip. Manufacturing becomes a commodity and is something that is competitive strictly on a price basis for the most part. Certainly, you'll have the occasional manufacturer taking a quick step out ahead of everyone else, such as IBM did with copper.

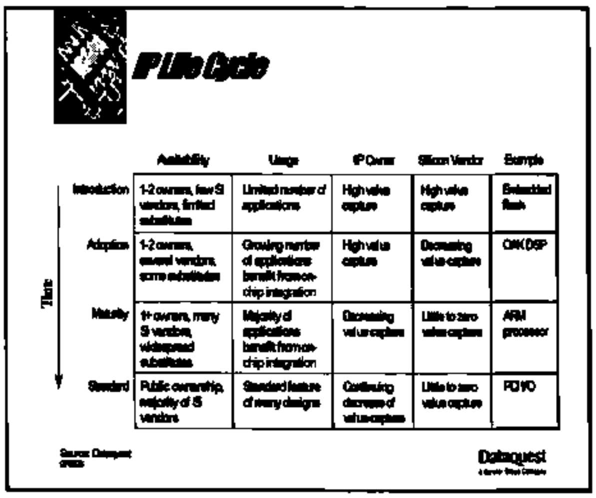

In general, however, manufacturing of these system level chips is going to be a commodity with the competition being strictly on price. Large foundries are going to be able to compete in this. Vertically integrated large companies like NEC, Samsung, and IBM are going to be able to compete in manufacturing. It's going to be a very tough road for the mid-size ASIC manufacturers such as VLSI Technology or LSI Logic to continue to compete as a silicon manufacturer. They've got some very interesting decisions to make over the next five to 10 years.

Along with these very few silicon manufacturers, there are going to be many places all fighting to control intellectual property, all hoping to be the next person out with that hot processor. There is going to be quite a bit of industry consolidation as the brokerages try to acquire more and more of the hot IP and another two-person garage shop springs up and has the next hot piece before they too, get bought out. I'll turn it back to Bryan Lewis.

Bryan Lewis: Let's pause for a second. As Jordan mentioned, we're talking about consolidation being rampant. Many of the vendors that are out there today may not be here later. There are over 100 ASIC vendors, can they all survive? Will they all have manufacturing? There are many implications in this, and it's a thing that everyone should be thinking about now— how am I going to get my manufacturing in the future, and what is this industry going to look like, and what are the tricks to succeeding in this marketplace?

The way to stay in this marketplace, the number one thing, is to really know your application market. System knowledge is king, really understanding the entire system. We're going to look at each one of the vendors in this marketplace and look at a strategy. The semiconductor manufacturer—this is a heavyweight battle. There is no room for people who don't have deep pockets. It's going to be key to have very large economies to scale and you're going to have to be very, very efficient. This is a very cost-competitive market and if you don't have the cheapest silicon out there, then your customers will go look elsewhere.

They're going back to this having an application focus. This even applies to the semiconductor manufacturer, even if he's a pure manufacturer and doesn't do design. The reason is because each one of these application markets has a different element requirement. If you have DRAM on there, that'll have a certain set of masks and you have to be able to do that. If you have flash on there, that's another set of masks. If you have analog on there, that has incremental masks sets. You really have to be flexible to move with this marketplace, and really understand where your customer's needs are.

Moving into the IP provider—again, you've got to avoid jumping into the heavyweight battle unless you're a heavyweight. Avoid the manufacturing temptation because it's going to be more and more expensive. That's not to say that you can't sink some money into a foundry relationship to guarantee supply, but to own your own fab... It's pretty much been proven, even at $1 billion, it's very difficult to own your own fab in the ASIC industry. LSI logic has proved that to some degree. They're struggling a little bit because they're trying to open Gresham and they don't have enough products to fill this factory.

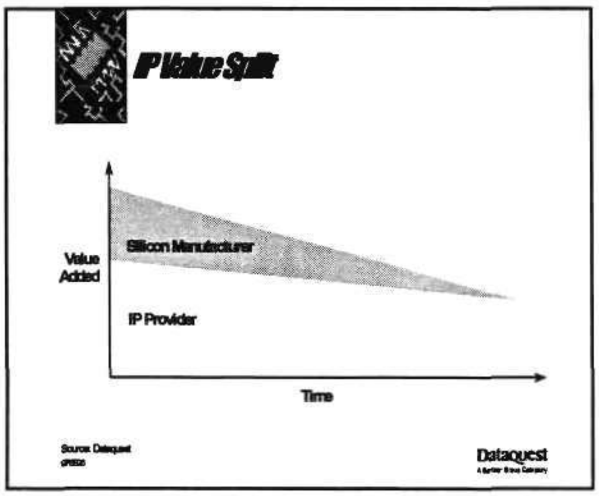

It gets very, very difficult going forward and it's already been proven by the number one ASIC supplier. The IP person has to really focus on "exactly what do they bring to the party." Get in early, develop this product and know when to get out. Distribution will be critical in moving these products in and out. Do you want to team up with an EDA vendor? Do you want to team up with an ASIC vendor? Do you want to team up with a system vendor? There are a variety of ways to distribute your IP but it has to be well thought out to really get your value out on the street to get your maximum returns.

Providing a comprehensive portfolio is not something everyone can do. That's not to say that a small five-man shop can't develop some intellectual property and make money off of it. The larger and larger that you get as an intellectual property supplier, you should really focus on what are all the pieces required for that given application. That is going to be critical as people grow in size, and if they don't get swallowed by another company.

The system vendor or OEM vendor has one of the bigger challenges in the marketplace. Everyone is attacking—the intellectual property, the value. They must focus on only the intellectual property or value that differentiates their product. They need to buy all the rest on the street and not devote any extra resource to it, except for integrating it into their own products.

The focus has got to be for all their internal people to really differentiate their product. Innovation is going to be key, because they're going to have more of a commodity-type of product. They're going to have to find innovative ways to market that, as well as try to get more features in the marketplace. Distribution channels, again, can be a good way to try to find new avenues of revenue out there. Bundling with other complementary products targeted at that application can be helpful too. That doesn't work for everyone. I think Microsoft has proven that bundling can be a little risky but bundling, in many cases, can really help many people out. Innovation and being flexible and watching out for garden hoses is going to be critical for some of these system suppliers.

The landscape is clearly changing. The SLI storm is on its way. It's moving on in. It's time to batten down the hatches. Really look at what value do you bring to the marketplace and where will you be in five years, because the landscape is very much changing. Thank you.

Moderator: Next up, we have Jim Handy. He's our memory analyst, principal analyst and director. He's been with us many years and he will come up here and give us a talk on embedded DRAM.

Jim Handy: What I'm going to be talking about is DRAM market, and whether or not it makes sense to look at it as an embedded market that's going to disappear. I have had the misfortune of watching over the DRAM business during a time of probably the biggest slide in history.

A lot of people are staring at their business and all the fab investment that they've made and the say: (a) are we going to be able to hang onto this market or is it going to become completely embedded? (Many trends indicate certain parts are becoming embedded); and (b) is there any way that I can capitalize on embedded DRAM to take advantage of the fact that I've got all this excess capacity?

In this presentation, I'll be talking about some embedded memory markets, some of the alternatives to using embedded DRAM and why, or why not to use embedded DRAM. Then I'll review a couple of the candidates who are already using this technology, and then what we think you should do as a DRAM manufacturer or user.

There are numerous embedded memory markets that have existed for quite some time. Since 1974, there have been microcontrollers and more recently, there have been DSP chips that use embedded memory. They're not DRAM; however, there has been a question lately—is Dataquest going to be covering embedded DRAM as a memory market? We're not. This is my explanation why.

We don't cover the ROM or the SRAM in microcontrollers. Look at a microcontroller chip and say how am I going to approximate what the size is of that market? One way of doing it is to look at the die size of the microcontroller and decide what portion of the microcontroller is Mask ROM, what portion of it is SRAM? On a microcontroller, about 50% of the chip is a Mask ROM and maybe about 10% of the chip is usually an SRAM.

Take the total of the 1996 microcontroller market of $10 billion and you end up allocating around $2 billion worth of SRAM and about $5 billion worth of Mask ROM. Likewise, look at microprocessor chips. Half of that thing is the SRAM that's used as the cache inside the microprocessor.

When I was at IDT before I came to Dataquest, we used to say that Intel sold the highest price per bit SRAM in the industry and they happen to throw in a microprocessor for free. A microprocessor is maybe 50% SRAM, probably about 10% Mask ROM. With the compute microprocessor market being around $16.5 billion last year, what that amounts to is an $8.25 billion SRAM market, which is about double of what the total SRAM market was last year—and a $1.7 billion Mask ROM market.

Finally, look at ASICs. There is some SRAM inside ASICs; little tiny ones—four kilobits or so, and some Mask ROM in them. Out of the total $16 billion worth of ASIC market, maybe about $800 million could be attributed to SRAM, maybe $200 million to Mask ROM. You add together all of these numbers and you end up with this phenomenal market. You've got an SRAM market of $11 billion that's embedded inside all of these products, in comparison with a $4.7 billion-$4.8 billion discrete SRAM market last year. You've got a Mask ROM market where the discrete Mask ROM market was only $1.3 billion. You can combine all of these applications into a $7 billion market.

These are things we don't count. We're going to be treating DRAM the same way as we have historically counted other embedded memory. We're going to be counting them inside the ASIC market. I don't have to worry about it. We'll put all that onto Bryan and Jordan's shoulders.

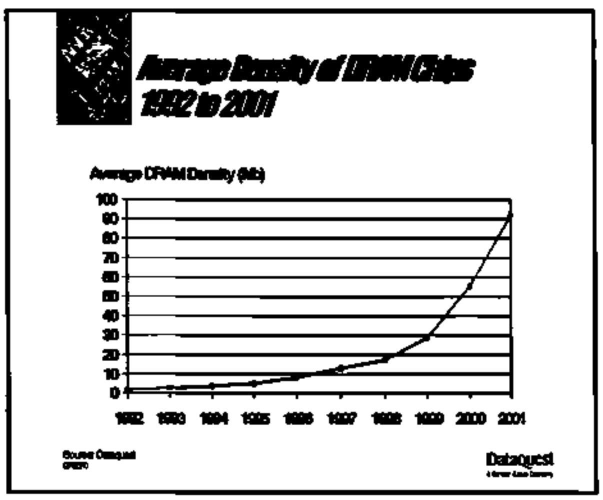

The embedded SRAM is a good alternative to DRAM. There are many good alternatives to DRAM. Embedded SRAM is extremely well understood. Anyone who can make a logic process can make an SRAM. It can be manufactured on an ASIC process using a standard ASIC six-transistor cell and it's much faster than DRAM. If you use an external SRAM as a competitor in competition with internal embedded DRAM, it still might be cheaper than using the internal DRAM. It is very fast.

You might not be able to get the bandwidth that you have available inside an ASIC chip, but it does give you much more bandwidth than you get by using an external DRAM. If you don't need that bandwidth, then an external DRAM is always going to be extremely cheap, because DRAMs are made in far more massive quantities than any individual run of an ASIC is.

Something that has been an issue is how to refresh that DRAM? It takes so many logic gates. Well, the amount of logic gates that it takes to refresh your DRAM hasn't changed over time. You add maybe one gate to a counter every time that a new DRAM

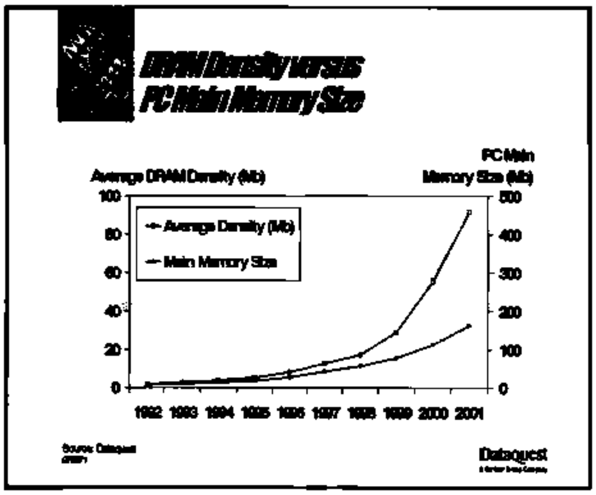

density appears on the scene. It ends up that, in the greater scheme of things, that's a very small and still-shrinking part of an overall ASIC. Nowadays, I don't think anyone really cares whether they put a DRAM or an SRAM outside of an ASIC, because it takes so little logic to cover that.

There are many good reasons to embed your memory. You can cut the system costs, because DRAM vendors don't want to support the old low-density DRAMs. No one makes 256 kilobit DRAMs anymore. Most companies don't want to make a one megabit part and they'll charge the same amount for a one megabit part that they're getting for a four megabit part. If you don't need four or 16 megabits, which are the current generations for DRAM, then it might make sense to use a one megabit internal DRAM inside your chip.

Also, it might be cheaper—although this is a stretch—to use one megabit DRAM inside an ASIC than to use one megabit SRAM outside the chip. One megabit SRAMs are now going for under $3. That ends up being a pretty cheap chip.

You can reduce your power consumption in systems. I'll show you some examples of how people have taken advantage of DRAM on ASIC technology to reduce power consumption. You can reduce your chip count. That's good for quality and power consumption as well as size constraint standpoints. It's like motherhood and apple pie. You could also reduce the number of I/O pins you have devoted to memory I/O on your ASIC. You can increase the bandwidth that you get out of the DRAM significantly in that type of an application.

There is an equal number of reasons not to embed your memory. One is that it's very expensive to put DRAM onto ASIC. A standard ASIC process is a metal intensive process. A standard DRAM process is focused on getting a good dielectric and not so worried about the metal. What you end up having to do is run two processes on the same chip to get any type of a small-sized DRAM.

There are DRAMs that are constructed using three transistors, on a standard ASIC process, but when you get to a three-transistor DRAM versus a six-transistor SRAM, the difference in size between those two—even though it is two to one—isn't all that great.

With standard DRAMs, you've got very competitive sourcing, high run rates. With embedded DRAM, you have limited sourcing, you have very low run rates. The run rates actually are much of what drives DRAM pricing down to the low level that it's at.

Finally, your design and test are extremely complicated. The design tools to put DRAM onto an ASIC chip are not as good as they could be. And your testing—you have to combine a logic test with the memory test and that tends to double your test time.

Here's a die photo of a picture-in-picture chip from Seimens[not shown]. I am showing you this because many people think embedded DRAM is something that's happening in the near future, so we have time

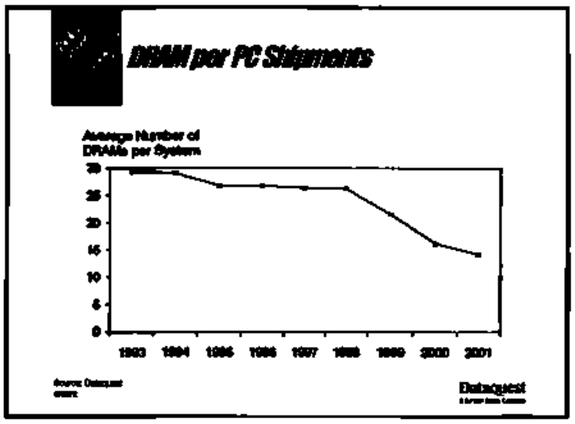

to look at this problem a little bit longer and think about it. Well, they're wrong and we can't.

This chip has been shipping in volume since 1991. It uses a one micron process. It has about a megabit of DRAM on it. PIP is picture-in-picture. They're using it in high volume consumer applications for putting a smaller picture inside a larger television picture. This is an example of one chip. There are some others from TI, from Philips Signetics and others that have been shipping for quite some time.

Embedded DRAM technology is here now. It's not a future thing. Now, it's simply a question of which of your competitors are going to be using that technology to their advantage before you. This is not a complete list of companies that are producing such products. They're simply different examples of different chips that are being made by different manufacturers—some of which we'll see in this program.

For example, here's the Mitsubishi 32-bit RISC processor in DRAM[not shown]. You can see the RISC processor in the center there, but this chip is mainly DRAM. I've gotten some pictures of Motorola's and Mitsubishi's ColdFire CPU chip; it seems like it's turned inside out from that. You've got very small DRAMs on a large processor chip. It simply depends on what you need out of the application as to what type of resources you throw at it.

The average density of DRAM chips increases 62% per year. Moore's Law allows you a "two times" density increase every 18 months or so. DRAMs have been on that forever. They seem like they're likely to stay on that track forever and our forecast is built on this assumption. If anything, it looks like there is a possibility that things could accelerate, because usually every three years at a remier design conference called ISSCC, R&D labs show the next generation DRAM. NEC waited only two years after the one-gigabit DRAM was introduced before they did their four-gigabit introduction. There is a possibility that we could be accelerating rather than simply going with a straight 62% per year density increase.

What are people doing with that? They're putting DRAM onto chips. These are things that sell in volume right now—Neomagic, Silicon Magic and a couple of other companies have versions of this. They're graphics controller chips. On the left-hand side of the photograph [not shown], you see the ASIC. That is the typical graphics accelerator. On the right-hand side, you've got the DRAM. That DRAM is being ported over into the graphics accelerator in a 64-bit wide signal path. The signal path could be increased in width; however, that would swell the size of the ASIC to the point where it would probably be no longer be cost-effective. Ironically enough, this is only a 32-bit graphics accelerator, even though it has got the 64-bit path into it.

If we look at that increasing DRAM density on the

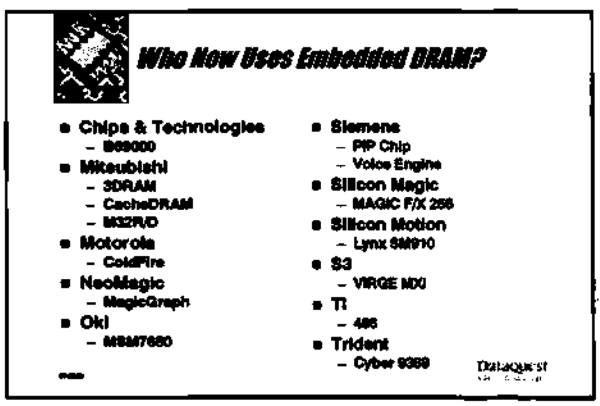

yellow line, and we compare that with what's marked on the right axis of this chart, of the increase in PC main memory size, you can see the main memory size in PCs is not increasing at quite the same rate. We're expecting the main memory in PCs to consume a shrinking number of chips. A phenomenon that was triggered by this was the advent of a need for a 16-bit wide DRAM in the transition from the four-megabit DRAM to the 16-megabit DRAM. That was part of the reason why we had the DRAM under-supply last as long as it did—from 1992 through 1995. DRAM manufacturers were unused to making wide chips and as they attempted to make the wide chips, they ran into noise problems they had never encountered before. That caused almost a year's delay in the adoption of the 16-megabit chip, and it caused a shortage for an extra year on top of the other shortage mechanisms that were in there.

If you take the DRAM growth that I showed you and the megabytes per PC and you divide the megabytes per PC by the average density of the DRAM, you end up getting the size of the main memory in an average PC going from 30 chips back in 1993—and this is including many extras that goes around the processor—down to somewhere under half that by the end of our forecast period. We are expecting PCs to go down in the number of chips. However, you notice I'm not predicting that we're going to have an embedded DRAM in PCs. The question becomes then where are we going to see embedded DRAM chips?

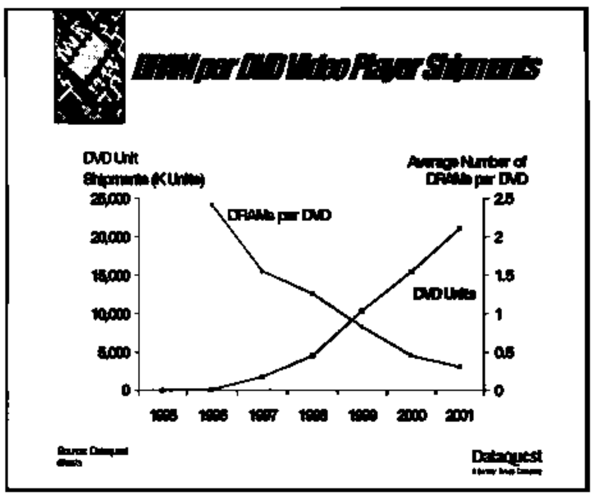

You do see them in graphic accelerators, but we're also expecting to see it go into the DVD market. Digital video discs require MPEG decoders. MPEG decoders require two megabytes of DRAM, and pretty much nothing other than two megabytes. We expect that to be the case now, we expect that to be the case in the future. Now, DVD units shipments are going up phenomenally. That could be a market in units for 16-megabit DRAMs. However, as the DRAM average density increases, then we can look at the average number of DRAMs that are required in the DVD. It ends up being a significantly lower number. Any time that declining line goes under the number one on the right-hand axis, that indicates that there is a need for embedded DRAM in that application. We believe that this will be a very good application for embedded DRAM as early as 1998.

That type of application could be suited for something like Toshiba's DRAM and ASIC mixture, in which they've actually taken their trench cell process, that they developed with IBM and Seimens, and they've used that to their advantage. They get a more powerful, higher gain transistor by using a trench cell process than their competitors are able to get using a stacked capacitor cell. As a result, they can have a faster ASIC. This type of chip can be used on those DVD players.

Another similar market is the direct broadcast satellite market. Once again, this is one that uses an MPEG decoder. You can see that the direct broadcast satellite shipments were expecting to be up around 30 million units by 2001, the end of the forecast period. Once again, we see on this right-hand axis how many chips are required using increasing density of DRAMs in DBS boxes. We end up with something that suggests, once again, as early as next year there could be a very good market for embedded DRAM.

There are other markets that I won't be going into, like the graphics controller market. That market is very rapidly going to have a number of alternatives that you can use—with embedded DRAM technology to get your graphics accelerator plus your graphics DRAM all on the same chip.

For these other markets that we're looking at—DVD, DBS but also other set-top boxes for wired networks—hard disk drives are being looked at very hard by Seimens and some other companies.

Graphics—I showed you all of those graphics chips. There is a possibility that, because of the complexity of the process, as the average density of SRAM moves into the domain of the static two-megabyte size required by these devices —that the technology could move over to SRAM. That's something to watch out for.

If SRAM does displace embedded DRAM in these applications, you can be guaranteed that there will be other, larger applications that will want to use embedded DRAM. There are many difficult issues that have to be confronted, though. I'm very glad that I'm no longer a system designer who has to make an economical decision about whether to use embedded or to use a non-embedded design. The guy down the street could have chosen the alternative, and the dynamics of the DRAM market are so crazy that it could very easily be that my competitor would end up making a system that was a dollar less expensive to manufacture than I could, simply because of the decision I made. Still, the things that effected this are the difficult process, the run rates are low (and so you're in competition with high run rate DRAMs that can squeak out every last cost), and that the die-size balloons that I showed you were all at least 50% DRAM. There are many cheap alternatives out there. We believe that for a few volume applications, embedded DRAM makes quite a bit of sense.

For the majority of the market, and certainly for the main memory market in PCs, embedded DRAM isn't where things are going to go. We're continuing to watch over the DRAM market as simply a standard market and we will watch over the standard part. We'll gladly let Bryan and Jordan try to track embedded DRAMs, and we'll watch as the processes in technologies mature. Now, I'll turn it back over to the moderator.

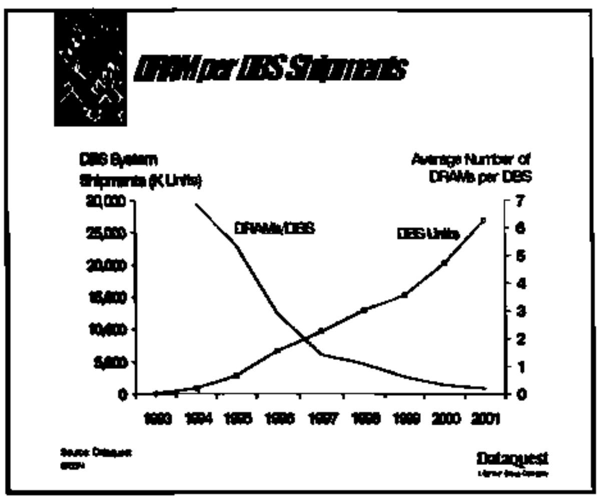

Moderator: We have a whole fistful of questions here. Nathan, we'll start off with you since you've been very patient sitting there for nearly an hour.

Q: Why will Intel's migration to Slot One not force motherboard manufacturers to redesign boards, cutting off competitors "at the pass" as they have every other time?

Nathan Brookwood: That certainly is an open issue for 1998. What's going to happen next year is all the motherboard vendors who have been squeezed out of the market or have suffered because they have been competing with Intel, and the chipset vendors who have been squeezed out of the market or suffered because they've been competing with Intel, see their ability to sell Socket 7 designs which are inherently lower in cost than the Slot One designs—at least for the foreseeable future at lower prices. Therefore, they see this as being a way that they can actually regain market share. We're seeing some of those alliances. It certainly is one of the last hopes for companies like SIS, VIA and ALI, who are the remaining principal participants in the chipset market.

Q: Why would Intel not develop a low-cost CPU?

The answer there is relatively straightforward. As long as you are manufacturing limited and can sell everything you make for prices over $100, why would you want to sell anything for less than $100. That's really the simple answer.

Q: Bryan or Jordan, how does the trend to SLI impact programmable logic vendors?

Jordan Selburn: Well, the programmable logic vendors are trying to get into this market as well. It's difficult, of course, because they don't have the logic density that the ASIC vendors have. The ASIC vendors starting to get into this market are actually having more of an effect on the ASIC vendors than the other way around, in that as the ASIC vendors become focused on system level integration, they move away from gate arrays and towards cell based.

A lot of companies have either de-emphasized or stopped making gate arrays entirely. The gate array/ PLD arena is where the PLDs and ASIC vendors butt heads. With system level integration, the ASIC vendors are becoming less interested in that competition, and that is providing some openings for the PLD people to expand in what was the lower-end of the ASIC market.

Bryan Lewis: As Jordan mentioned, there are really two angles on it. The ASIC vendors will attempt to merge some programmable logic into the ASIC solutions, and some of the PLD people will also attempt to enter this market and put some of the ASIC pieces onto their silicon. Clearly, there will be a combination strategy going after this marketplace going forward here.

Q: For Bryan or Jordan, which microprocessor's cores will be most popular in SLI?

Jeff Selburn: Some of that is going to be a function of the type of application. If you take a look, the processors are general purpose. There is some market differentiation. However, the wireless market seems to have adopted the ARM microprocessor. If you take a look at the consumer market, particularly the games market, that's very strongly in MIPS camp. Digital video seems to be up for grabs. A lot depends on the application. You'll also see processors such as the SH, PowerPC, X86—really all—getting a toehold in the market.

Q: Jim, can a DRAM company afford to ignore the embedded DRAM market?

Jim Handy: Yes, definitely. We're still expecting the DRAM market to be the mainstay of where all DRAMs are shipping. I've forgotten what number it was that Jordan and Bryan were bandying around. It was something like a $6 billion—$4 billion market in the year 2001 for ASICs that included DRAM. That's nice and some DRAM manufacturers are going to get into that market, but other DRAM manufacturers will be quite content to ignore that and simply participate in the other $67 billion or so standard DRAM market that's going to exit then. There is quite a bit of room for DRAM vendors.

However, once again, we're predicting that there will be a slight cyclical downturn in the year 2001 in the DRAM market. Along with that, there will probably be more manufacturers in the DRAM market. Many of these DRAM manufacturers will say, "well, it would have been nice if those guys hadn't joined in"— like they're saying now about Taiwanese firms.

Bryan Lewis: That's a bit of a standard product view, though. If you really look at the DRAM manufacturers, you see that Samsung, the leading DRAM manufacturer, clearly has targeted that embedded market. Then you go right down the list. Almost everyone of them, from our perspective, is entering the combination DRAM and logic market. The marketplace has voted and is saying that, in order to diversify a little bit, we need to start pulling in a bit more of the logic functionality. In looking at this whole system trend. I would have to disagree.

Jim Handy: Although there are manufacturers who are going to execute that well, there are others who are not. Your typical DRAM manufacturer doesn't understand how to sell to engineers like an ASIC manufacturer does and tries to win the business at the purchaser's desk, instead.

As long as they keep their sales channels lined up for that type of selling approach, then they're probably going to have some problems getting into the embedded DRAM market. Those who are willing to put in the extra effort to provide good support tools and good engineering support—and the design and development process and everything—will probably do pretty well. I would guess that a number of other manufacturers will fail at that attempt and they'll end up going back to standard DRAM—but they'll still make money.

Q: What is the premium for embedding DRAMs for similar die-size ASIC, versus ASIC plus DRAM in the same package, for volumes over 100,000?

Jordan Selburn: That's the same type of question as "How long is a piece of string?" There are no fixed numbers for that. You need to take a look at what the embedded DRAM is bringing to the solution, which is going to be very application-specific. If it's a consumer marketplace, by definition it's almost zero, because the only reason you're bringing it on board is to lower the overall system cost

The ASIC may be a bit higher-priced, you can charge a little premium there but it certainly can't be as much as a standalone DRAM solution would be. These are applications that you're not putting DRAM on because of size or power. These are sitting on top of your desk or the TV set you plug into the wall. The only reason you're adding the DRAM is a cost basis.

On the other hand, some people are going to embedded DRAM for performance reasons. There, it's a function of what can you do to your system clock speed by virtue of bringing the DRAM on-chip. If you've got a workstation and you can turn up the clock speed 25%, then increase your price from $20,000 to $25,000, you're certainly willing to share some of that with the person providing the embedded DRAM chip and they'll be more than happy to take it.

Jim Handy: One of the companies that is using embedded right now is Silicon Graphics. They have a proprietary DRAM and graphics accelerator chip that they're using simply so that they can get that extra 25% or however much performance Jordan was referring to. Likewise, Mitsubishi developed something that's a custom chip for Sun Microsystems that has some of Sun's graphics primitives in the DRAM, and it's something Mitsubishi calls the 3-DRAM. This is licensed technology from Sun Micro and I believe anyone who chooses to buy that from Mitsubishi automatically gets a Sun license from Sun to be able to get that extra bandwidth out of there.

Q: Nathan, I have a couple of questions for you. First of all, what is the impact of the IDT Centaur on a microprocessor market and how can ASPs hold for X86's if they are sub-$1,000 PCs? Won't Intel's ASPs have to decline?

Nathan Brookwood: First, with regard to the IDT Centaur announcement recently, IDT entered the X86 market with a proprietary design that came out of a small group they have in Austin. Their goals are very modest in terms of volume. They're really trying to cherry-pick the market and, from IDT's perspective, given the price that you can get for SRAMs these days, selling low-cost microprocessors is attractive compared to selling lower-cost SRAMs on a revenue-per-wafer basis. Their product is targeted at the low end of the Pentium market, which is the $100, $150 market. The savings they offer are relatively modest compared with the low end of Intel's lines. They're going to have a bit of a struggle gaining a toehold there.

Next will Intel's ASPs have to fall? Of course, there were a bunch of financial analysts who said precisely that. What Intel has going for it is that even if they do participate more at the low end of the market with $100 chips or below—and it will be tough for Intel to go much below $100—they really are making a major push now in workstation and server products where the ASPs are going to be $2,000, $3,000 and up.

These are not intended for you to try to use in your home. No one's going to want to buy a $3,000 PC with a $2,500 chip in it. That does have the effect, even though the volumes are relatively modest compared with the mask volume, mask market, of raising Intel's ASP in a noticeable way; a lot of the projection that I showed for the "Intel Inside" ASP increasing, as a result of Intel participating in server and workstation markets, where I don't think AMD and Cyrix are going to focus at all.

Q: Bryan or Jordan, what companies are best positioned to succeed or excel in SLI market?

Bryan Lewis: I hate to pick out some because I'm sure I'd leave a few on the side. Today, if you look at the premier ones, you see some of the vertically integrated ones—the ones that have been around a long time. I will throw out a couple of names—IBM, Lucent. They're vertically integrated. They know the systems. They know the technologies that need to be there. They know the design flows that need to be there. They have the packaging expertise. It's a fairly complicated market. Others vendors will buy these other pieces out there but the ones that are out there first are the ones that have explored this area from their internal system business over the years.

Jordan Selburn: I'll take a look at the flip side of that. The companies that won't succeed are the ones that are out there thinking that they can charge more than their competition because they have for example, an ARM processor. The people who won't succeed are the ones who don't recognize where they're adding value. They will pour money into development of commodity intellectual property, and manufacturing, where they're not going to achieve any differentiation, instead of the areas they can top their competition in.

We're seeing a decoupling of the components that make up an ASIC. If you're pouring money into resources and areas where you're no better than your competition, that's a recipe for failure. You've got to put your resources into the area where you've got the edge—to name it again, not to leave out others—Lucent—"we make the things that make communications work;" they're pouring their resources into process areas that are germane to communication, system level macros that are germane to communication, I/O's that are germane to communication. They know where they add value and where they can compete; and that's where they're putting their resources.

Q: Jim, embedding DRAM is also becoming a popular concept with network vendors for high function switchers. Any thoughts why your market prediction excluded this segment?

Jim Handy: There are many section segments to look at. There are enough of them that we could do a rigorous analysis of several different markets and probably still leave many of them out. I personally don't know enough about what the memory needs are of the networking market and where there might be small chunks that could be replaced by the DRAM.

However, I have heard, for example, at this conference about certain SRAM technologies that are in peril of being absorbed into the ASIC as we speak. I believe that there are probably many applications that I left out, not simply networking. If you folks would like, you can approach our consulting group about our putting together a multi-client study to look into that.

Q: We can't let a session go by without an allusion to copper. What is the effect of copper on SLI?

Jordan Selburn: Right now, very little. Copper right now is going to be geared towards high performance applications. Those tend to not be the ones that will adapt system level integration first. It's costly and you can do two things with copper. You can run faster, and you can make your chips smaller. Making your chips smaller is a good thing, and can get you some performance as well, but if the net result is a higher cost chip, then it doesn't really buy you anything. The focus, until it becomes a mainstream technology, is going to be in the performance end. That's really somewhat decoupled from the system level world.

Bryan Lewis: If you look at it, we're just seeing the announcements and some of them are involved with .25. Most of the announcements will be .18 and below. We're looking at, probably, at least three years before we see any major impact, but it will start being shipped in some degree late next year probably.

Q: This is up for anyone to answer. With the move to consumer products, i.e. the Sony PlayStation, WebTV, PalmPilot, are we moving to an application-specific CPU or SLI processor? If so, what about the microcontroller players developing products, like Microchip, Zylog or Motorola?

Jordan Selburn: I'll take the first part of that—are we moving towards application-specific CPUs? We're already there and we've been there for a while. It's simply that it tends to not go by the name processor. If you take a look at an MPEG decoder, that is a processor that has one goal and that's to decode an MPEG video stream. It does everything in hardware, as opposed to a general-purpose CPU, which does things with software.

I would say that we're already there. We've been there and we'll continue to be there because, despite the success the folks at Intel have, when you've got an application orientation, you can always do it smaller, faster, and cheaper. By implementing it in hardware, you give up much of the flexibility and versatility.

A: Let me add to that, that if you pry apart a Nintendo 64, as we've done in some of our tear-down work, there are actually two MIP processors in there. One of them is an off-the-shelf discreet MIP processor that NEC manufactures. The second one uses a MIPS core and adds in all sorts of graphics primitives, and is called the reality engine. Again, amplifying what Jordan says, we have application-specific processors already showing up in the consumer world.

Chapter 19: Panel Discussion - Will Big or Better Win the Flash Market?

Moderator:

Bruce Bonner
Principal Analyst
Semiconductor Memories Worldwide Program
Semiconductors Group
Dataquest

Panelists

Bill Howe
Vice President & General Manager
Memory Components Division
Intel Corporation

Walid Maghribi
Group Vice President
Memory Group
AMD

Miin Wu
Founder and President
Macronix International Co. Ltd.

Bruno Beverina
Vice President, Memory Product Group
General Manager
FLASH Memory Division
SGS-Thompson Microelectronics

Dan Auclair
Senior Vice President
Operations and Technology
SanDisk

Hyung-Kyu Lim
General Manager
Memory Division
Samsung Electronics Co.

Bruce Bonner: Good afternoon. This is the final session of the Dataquest Semiconductor Conference. I'm Bruce Bonner. I'm a memory analyst with Dataquest. Welcome to our flash food fight which is our final session of the event. With the purpose of this really is to expose you to some different ideas, different points of view in the flash memory industry. The flash memory industry is very, very exciting. It's not like any other memory segment where there's a lot of diversity in terms of applications and different types of technologies. I'm hoping that you'll find this very informative. Recently, we've reissued our flash memory forecast

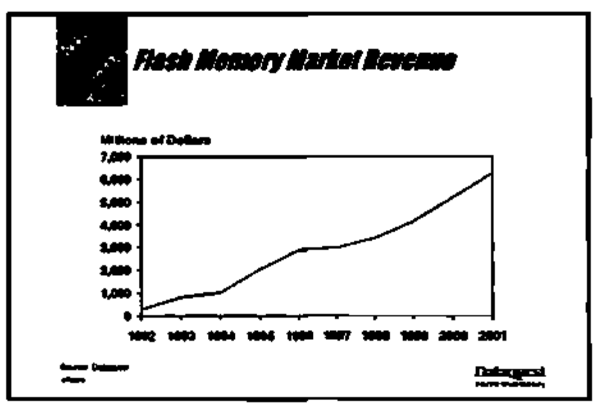

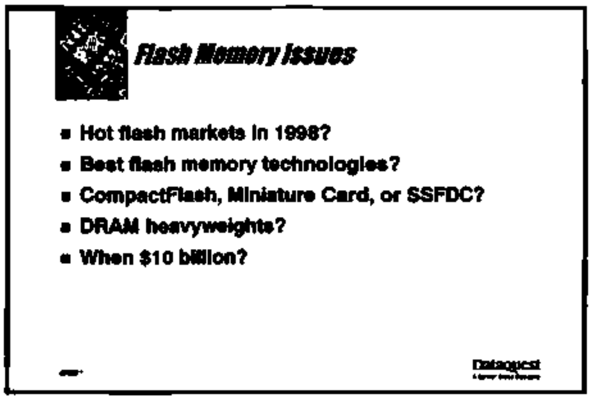

and we've actually had to take it down from where it was before. In the year 2000, 2001, it actually was around $10 billion and we've actually had to decrease it because of pricing down to about $6.3 billion. Also you'll notice that in the years 1996 and 1997, it's essentially flat. That is again, not because of unit shipments or bit shipments. That's because of very, very aggressive price reductions that have been happening in the market.

What we're going to be discussing are these questions. First of all, what are the hot flash markets next year and what are the different types of applications that are going to be driving this? There's basically two divergent application types in the flash memory market that are driving speculation.

First of all, the existing applications tend to be code storage. Everyone understands this fairly well in code storage. In data storage types, however, like digital film, those things are new, digital cameras use those and that really is fairly speculative that's new and there's a lot of question marks associated with that market. Also with those different applications, different types of flash memory technology are being used. There's a lot of divergence in the flash memory industry on what type of flash technology is the best. You're going to be hearing some points of view from the speakers today on what technology is best for those.

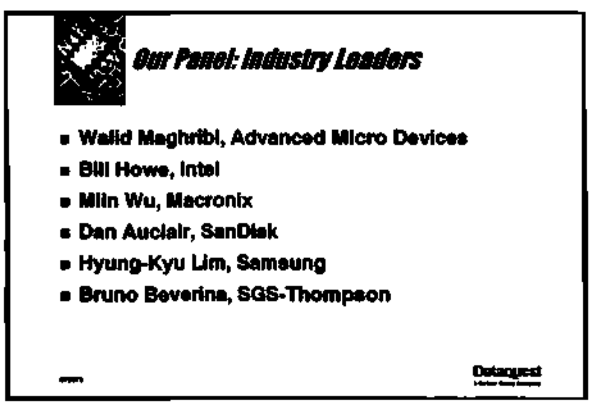

There are some standards battles that are being waged right now. The foremost one is between CompactFlash, Miniature Card and SSFDC in terms of what type of miniature flashcard or compact flashcard, small form factor flashcard should be used for digital film. Another question is are the DRAM manufacturers going to end up owning the flash market? They're very good at making memory already. Maybe they're going to end up owning the flash market also. Finally, when is the flash market going to hit $10 billion?

We have an incredible group of people here. These are all the leaders in the flash memory industry. If one of these airplanes that are buzzing around up here hits the building, we're going to set flash memory back five years. You have the distilled essence of flash memory up here on stage with me.

We have Bill Howe in the second seat —we'll come back to the first seat in a second—from Intel. Intel is the world's largest semiconductor manufacturer. They're the largest manufacturer flash memory by our count. Bill has been the general manager of the flash memory division since 1994. He's been with Intel since 1979 which is 18 years Miin Wu is going to be representing Asia/Pacific for us today, is the founder and President of Macronix. Primarily up to this point, Macronix has been making ROMs and EPROMs but now is into flash. Macronix is the largest nonvolatile memory maker in Asia/Pacific.

Mr. Dan Auclair is Senior Vice President of Operations and Technology for SanDisk. Dan has over 20 years of experience in the mass storage industry and SanDisk has a mass storage focus on flash, that's very appropriate. SanDisk is probably the world's largest manufacturer of mass storage type of flash.

H-K Lim from Samsung has been with Samsung since 1976 and really has been bringing up memory at Samsung. Samsung is the world's largest memory manufacturer. They're the world's largest DRAM manufacturer and of course, they'd like to repeat this feat with flash memory.

Mr. Bruno Beverina from SGS-Thompson is the Vice President of the memory operation for SGS. He's the general manager of the flash division also. He's been in the semiconductor industry for 37 years so he's been through a couple of cycles. He's turned the EPROM operation into a profitable manufacturing operation and also be the world's largest EPROM manufacturer. He's quite proud of his success with that. Now, of course, they're bringing on a bunch of new flash fabs. They are a force to be reckoned with.

Finally we have Wally Maghribi. AMD has a very important position in life. What they do is they keep Intel honest. They also lead in a number of different areas. He's going to be our first speaker today. I think what we're going to have to do is whatever the order that's on this computer we're going to have to do. Bill, you're on.

Bill Howe: First of all, I'd like to say thanks for inviting me here today for a panel session on the flash memory market and where it's going and what Intel's perspective is. I'd like to start off this session today by giving Intel's view of where the flash memory market is and where it's going. First of all, I'd like to see a little bit of history. Up to this point in time, flash memory has essentially been an EPROM replacement used for many code stored applications, sort of classic, non-volatile, merry applications where your flash memory was programmed a couple of times.

This has been an interesting market, clearly the market which has been driving the initial growth of flash memory. Typical examples of that are things like digital cellular phones which are reaching very high volumes today and growing even faster. However, that market is relatively limited as we look in the future. To really trigger the growth of the flash memory business, we really have to break into this data market, the use of storing data in your flash memory. That's really a much larger market and that's mainly what DRAM is used for today in your PC in your memory system.

Flash memory people including Intel have always been searching for this huge market of data storage with flash memory and I think many of you will remember, there were some early predictions that flash memory was going to replace hard disk drives, even eventually DRAM. Now, we've toned down our expectations from that from those original days and now we're looking for various applications that will use and have an interesting usage model with flash memory but nothing at the size of the disk drive market or the DRAM market itself.

You can see on this chart various handheld applications such as digital cameras and voice records, etcetera—PC companions that will use flash memory. I think everyone's waiting saying when's this really going to happen. We all remember the Newton and it really was not a market driver for flash memory or for anything else for that matter. We look at these devices and say when are these things really going to happen and when are we really going to have this digital data storage? I contend it's happening now.

Devices like the Pilot have become very popular and some of these other devices now are really starting to take off. The real key reason is compared to the days of the Newton versus today is the Internet. Now that you have Internet, you have PC connectivity with your handheld device. I look at all these handheld devices not as replacements for your personal computer but really for adjuncts or companions or complimenters for your PC system. Now that we have internet connectivity and in many cases wireless internet connectivity, you can get your digital data back to your PC environment. That's very important. That makes these devices much more useful. As standalone devices, they're interesting but when you hook them up to your PC, they get very, very powerful. That is why I believe that now is the time that digital data will happen.

The form that's going to occur is a combination of using your flash memory for both code storage and for data storage. That is not a separate device to store your digital data while you already have a flash memory for code storage in your handheld device but one device that can do both. That's what we call this new emerging code plus data market.

It just happens that Intel has introduced a new flash memory called Stratoflash using multilevel cell technology which gives you the advantage of a high density, low cost flash memory. We recently introduced a 64 megabit device at a low cost with all of the classic flash features that customers are used to. That is the fast random read speeds of your NOR or your ETOX memory with the high density and low cost that you can get using other memory technology such as NAND technology. We really think is a breakthrough memory technology and this is exactly the thing that will stimulate a new surge in growth in the flash memory market. What we want to do with Stratoflash memory is two things. The first thing we want to do is replace all the single bit high density flash memory. That's what we want to do. The second thing is we want to create this new market for high density flash memory in this code plus data or combination of code plus data space.

An example of the first one is something as your voice recorder where you can record up to, for example, four hours of voice storage of a voice file on your voice recorder using a Stratoflash in a miniature card. As a standalone device, voice recorders have been around a long time. What's the big deal? Well, the big deal is when you take your miniature card out and you put it back in your PC platform and now with, for example, applications such as the IBM Viavoice, you can have simultaneous speech to text translation and it is a useful input device for your personal computer. Before, because you didn't have enough memory because the cost per bit was too prohibitively expensively or you didn't have enough power in your PC platform.

Now with the new Intel platform with the Pentium processor, with MMX technology of course, no Intel executive can possibly have a presentation without saying Pentium processor. You can get this simultaneous real-time voice to text capability which is very powerful. That's an example of replacing today's flash memory applications.

In the future, in this emerging code plus data market, we see Stratoflash providing the read speeds that customers are accustomed to for the code storage market with the low cost per bit, with the high density that they want for the data storage in one device. With Internet connectivity and the fact now that you've got object oriented languages such as Java, you really can't tell what's code and what's data and you need both, we think Stratoflash is the flash memory for the future.

Moderator: Thank you, Bill. Next one is Walid.

Walid Maghribi: First of all, I would like to make a little correction for you, Bruce. AMD does not only keep Intel honest, in a lot of cases, eat their lunch especially in flash. I would like to today to present AMD view of the flash market.

Flash Market - 2000

- Today's applications continue to drive growth
  - Mass Storage grows to 20% of TAM
- Different applications require unique device capabilities
  - Multiple products & architectures
- The base of applications expand beyond 2000

Dataquest Semiconductors '97, October 22-23, 1997

As Bruce said, the flash market is continuing to grow rapidly albeit at less rate than what's predicted in the past. For AMD, forecast 1997 will close at around $2.8 billion, also approaching about $5 billion at the end of the century and exceeding $5 billion in the year 2001. The growth of the unit however is much larger and the switch to higher density is a driving, significant, much larger growth rate in bits. We're forecasting about 55% to 60% compound annual growth bit from 1997 through the year 2000.

What's important to note—and I was about to jump in and correct some of the statements that Bill has made—is that, let's face it, today's market, this driving flash, is code storage. About 95% of the market of flash is code storage. Less than somewhere about 5% to 7% is mass storage market. Here in some cases, the flash did replace EPROM but as you know, EPROM has never grown to $2.8 billion market. That growth is true flash market because the flash, as we all know, is enabling technology and EPROM is still $700 billion, $800 billion market. Flash may replace a little bit of EPROM but

Flash Market Issues

- The Flash Oligopoly continues
  - Ranking changes
- "Enough" capacity is in place through the year 2000
  - Dynamics in various market segments may change supply/demand equation
- Emergence of one single power supply standard Technology
  - NOR-MLC is only acceptable for audio or image storage

Dataquest Semiconductors '97, October 22-23, 1997 definitely was not created to replace EPROM 100%. What's driving the growth today is code storage. What will continue to drive the growth at least through the year 2001 will continue to be the same market that exists today. It is true mass storage that representing under 10% today, will grow, will represent a bigger percentage somewhere in the neighborhood of 20% by the year 2000, 2001.

What is important about our flash market—and that's different from a DRAM market—is that in the case of a DRAM, the industry is shifting basically one density whether it's 16 megabit or 64 megabit. What is in flash, it's the various market segment that determines what kind of a product you ship. Some market segments require one megabit, other market segments require 16 megabit. Some market segments require very high speed, other require very low power and so on and so forth. That's what makes a flash market totally different than a DRAM market. It's requires multiple architectures. The flash market is a summation of several markets. One of them is mass storage.

AMD/Fujitsu Flash Leadership

- **Continue increasing market share**
  - **33% in 1996 and gaining**
  - **Bit shipments quadruple through 2000**
- **Customer/application driven technology and products**
- **Flash is a strategic and significant business for AMD**

Dataquest Semiconductors '97, October 22-23, 1997

If you look at the flash market, because of the price—and it has dropped over 50% in the past one year—the base of application is expanding for sure. However, the big percentage continues to be in the year 2000 whatever it is today, mainly cellular, networking and PC and peripheral.

If you look at the players today, they really are the same players of four years ago. The AMD, Fujitsu camp and the Intel Sharp camp, that continues to ship an excessive 75% of the market. However, if you look within those four players, the ranking has changed where Intel, in 1991, represented in excess of 90% of the overall flash memory.

Today, they are definitely below 40% of the market. The people that have gained this decrease in market for Intel, of course, has been AMD and Fujitsu. We believe that the market between AMD and Intel, assuming Intel fab in Israel continues to be a flash fab, we believe that there is enough capacity in the world to supply the demands through the year 2000. However, dynamics in the market may change if, for example, the cellular telephone requirement was to increase from 8 megabit to 16 megabit. That will add about 10% to the overall capacity requirement. The single power supply is definitely standards and technology. All newcomers are duplicated in a single power supply. Over 50% of the product that's being shipped today is shipped with single power supply. As a comment on a MLC, MLC is good. If you could provide a product that has exactly the same attributes and you could store more than one bit per cell, we like MLC. We believe it's great because it reduce the cost of the bit.

However, for the application that Intel is talking about, mass storage, we believe NOR-MLC will not win because it does not fit the market and every market requires different attributes. The data storage requires high reliability. Portable market requires low voltage and low power. Even video and data storage requires small sectors. Very few applications require 64 megabit. You don't need to have a recording in your home that has 64 megabit because you could store a lot of talk time in 4 megabit only. We believe even though it can be used in some applications to meet the mass storage markets, there are today an existing technology that will deliver as much of a cost advantage without sacrificing any of the attributes or any of the requirements of the market.

For AMD and Fujitsu, as you know, we are partners. We obviously, because of our product attribute, have grown our market share from zero in 1991 to 33% in 1996. According to WSTS, we will be in excess of 40% during this quarter. Our bit shipment continues to increase significantly. ~~We are~~ forecasting by the year 2000, our bit shipment will quadruple from 1997. Our success is mainly because we drive our technology and our product to satisfy our customers. We don't tell the customer what he needs. We ask the customer what he needs and we build the product to fit the customer requirements.

It's a very important point that flash is very strategically important for AMD. It represents about 30% of our revenue. It represents a huge percentage of our profit thus, AMD is here to stay to service our customers.

Moderator: Thank you, Walid. Mr. Miin Wu from Macronix will give us the Asia /Pacific perspective on flash memory.

Miin Wu: Thank you, Bruce. In this lineup I have a hard time to claim for Macronix. I think I can say this. We are the youngest company here. Because we are young, we have no burden so we can develop interest in technology. Today, I'm trying to show our compatibilities in term of technology and direction but before I go in there, maybe let's examine the general philosophy on how the memory is going to be developed and sell.

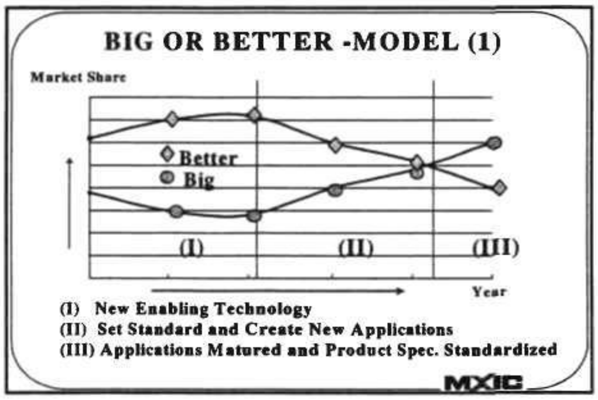

In this chart, basically, you're going to see we're trying to define what is big and better. We all know that when a new company has better technology they occupy the market. That's why Intel in the beginning has a very large sharemarket but gradually when the standard is coming, then on occasion become a driving force, then the market will come down and a new player will come in. Until finally, the standards are set and also the application becomes mature and the large company like current DRAM makers will probably come in and try to take over the market. So we will see how this relationship we can compare.

This chart, basically, we'll show is the braver company certainly is trying to increase their leads and trying to join the market share. What can they do? They create the new applications just like Bill and Wally said. They will create the new applications and the new fields. By doing that

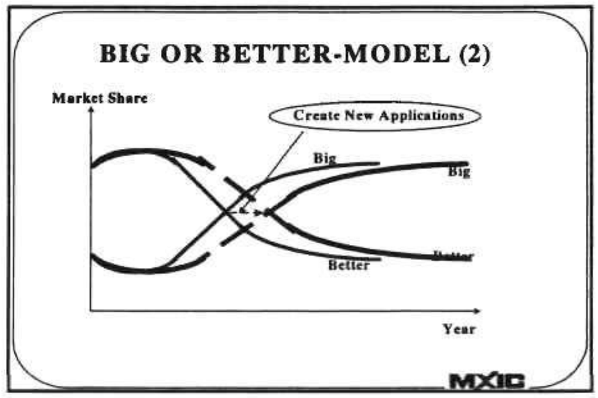

actually, you just have to create new market segment and then create the technology for that.

Certainly, if we look at the bigger company, what they have to do is move the competition and the next page basically is trying to show you how they're going to do that. They have to create the leverage. The big company has many other areas they can leverage. For example, they are not only selling this flash. They can use DRAM or for example, the logic

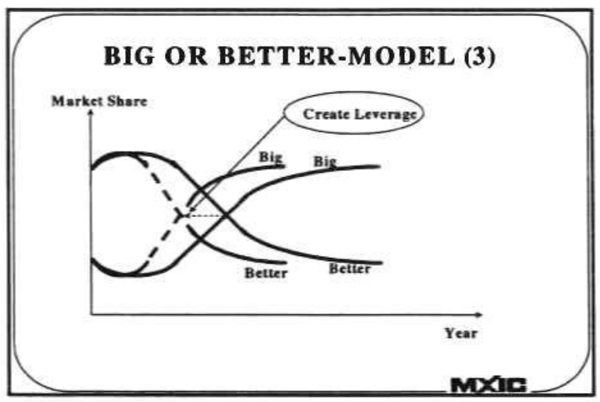

device could do that. At the same, they can really use packaging and different capabilities to set up new standards to try to move that. Right now, Macronix, what we're trying to do is really to follow this pattern. We're trying to create the new technology and also working with the people setting up new applications to preserver opposition.

Very obviously, from this chart, we can see they are trying to move into the beta flash right now. That's what we're going to do. In the next graphic we'll show you we're in opposition. Intel is making a big splash on their Stratoflash technology. Very obviously, with MLC, they can reduce their sale by

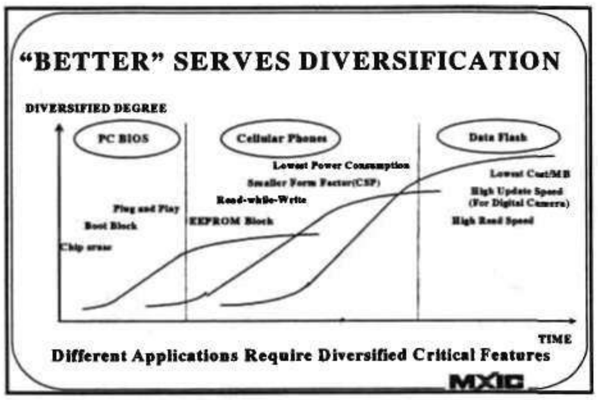

half effectively. Macronix, we develop a very interesting technology we call the PAC-AND. This is called the Pair Array Contactless. That means we're really actually combining the best of the Nengate and Norgate together. We believe that is really the future technology. Not only can we serve the code and the data storage but also we can make a chip very small and we can also make it fast.

If we look at our technology, actually without even making the push on the technology, we can move from the pack end to what we call the asymmetrical virtual ground, Engate. In fact, we already have the chip size like Intel has. It will consider the app is a minimum feature size. At the same time, I think with the new technology like that, you will be making the chip smaller, you put the pressure on the manufacturing on how to control the prosper meter so you can have a better yield. Our approach is we actually can introduce MLC concept and then we

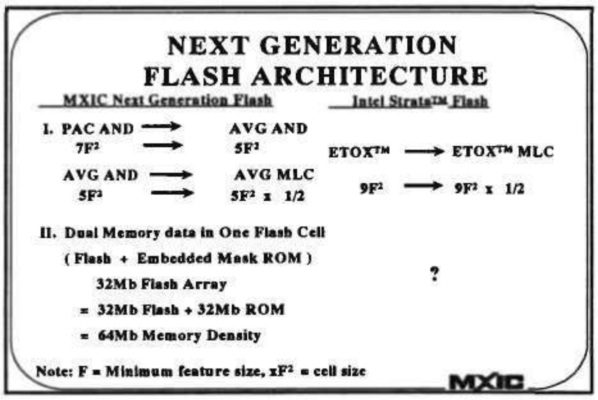

make chips very small. If you look compared to that data, you can see we can actually make about 2.5 f square compared to 4.5 f square like Intel approach.

At the same time, we actually have another approach. We believe in most applications, you probably don't need the auto flash there because a large coat has been fixed. All you have to do is add in some of the variable into the new device. We actually designed an interesting cell where we can have fewer memory data in one transistor as well. Similar to Strata, instead I have 100% flash devices. We have half of the flashes and half of the embedded mass ROM because on a lot of occasions, you don't need to change. By doing that, we don't have the push technology. At the same time, effectively we gain our capacitor rating. In that case, with older technology, we can prolong our life and at the same time, we can solve a lot of application problems for the system users.

We believe Macronix is better. We have better technology compared to NAND and NOR. Consider

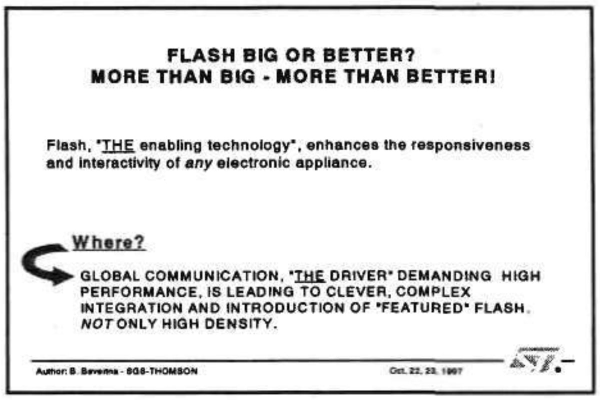

we're only eight years old and we're already the world's six largest memory supplier in the world. Part of the reason is we believe we have this proprietary and fairly compatible array architecture. This architecture will offer smaller dye size, better

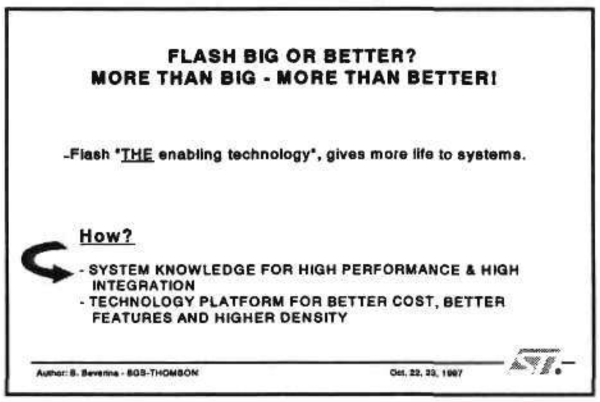

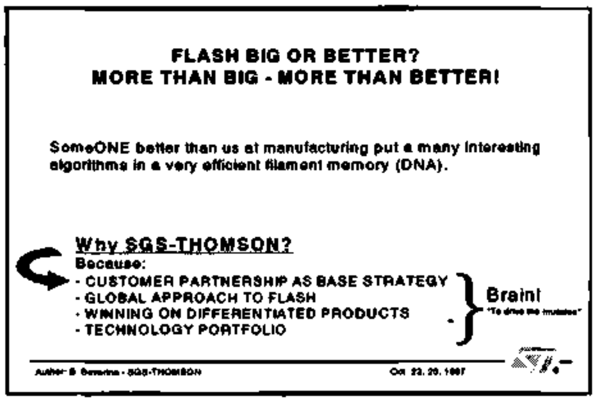

performance, easy to manufacture—that means it's a better yield. At the same time, we have a tremendous IP portfolio. We have 550 patents and the 16 patents already granted. We believe in the next couple years, we will have a lot more new technology, get into IP protections. Also we believe we are one of the very few total solution supplier in the memory area. We offer flash, we offer multiple time programming EPROMs, we offer the ask ROM.

What that means is we're really able to become a player and each player of the migration path. That means we are offering pin-to-pin compatible products between the different types of the devices. We also provide the product to product compatibility. By doing that, we can learn a lot of things from the simpler device. At the same time, we can offer an option for the customers.

We also have ATA controller and also microcontroller processor DSP with the embedded flash. By doing that, we actually can offer a lot more application-driven type of the business later on and we understand one thing very clearly. In order to grow the company, we really have to work with the system coupling and then try to do the application. Recently, we signed a deal with Philips. Basically, we are providing flash technology by working with them to put into controller. I think this is one of the examples we can grow the company. Another example, we work with Nintendo in a very interesting area where we provide the flash card for different type of business. Those are the examples of how we work with the key customers.

We believe, in the long-term, not only are we trying to address the standalone flash devices, most important, we'd like to get into a little bit more system level integration. Thank you.

Moderator: Thank you, Miin. Bruno, you win the lottery.

Bruno Beverina: It's a lot of very long volatile memory of our industry. I still feel young, like a flash. It's a continuing new product. It's really what you need to grow a system, a segment that is so exciting by the one of the flash. We've been hearing a couple of different approaches. I'll probably present another angle of the same market.

At first is it a question, big or better? Big or better what—the prop or the player? You understand it's a very ambiguous interesting question. Let me start with the problem. Since their first appearance, the beginning of the 90s—the flash—we all knew that the flash would be an enabling technology and not just another number memory. It's been said already by Wally, it's been said by Bill to some extent. We know it. Those which plays since the beginning, we know.

Big or better flash, what does it mean? It would be easy to reply. As it is written here, they are much more in depth, yes. Both—big and better and much lower. What I think is that beyond the classical categories where you take the performance, they're talking of the flash you're saying got low energy, higher speed or you think high density and the density can be at the chip level or at the board level. Let's not forget the discussion of yesterday evening on the packaging. It's quite interesting because we have to see the density as an overall situation. The categories of the low cost—Wally just said that the market this year, it's unfortunate to be down. You would expect it from us the market would have been up, 30%. Instead it's more or less flat because we add the price down at least 30%. The cost is down,

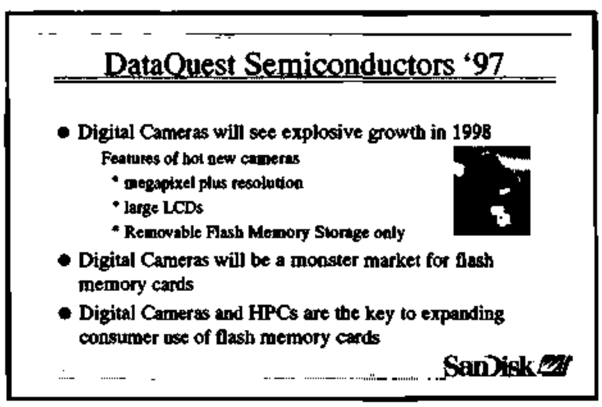

should be lower. We should find a way of reducing the cost.

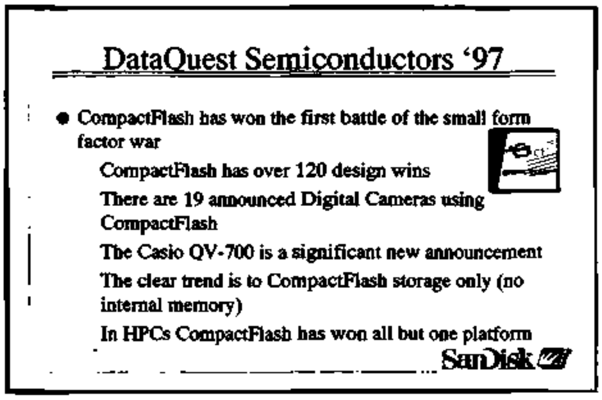

Beyond this categories, I see a clear path to a widespread integration to enhance the responsiveness and the interactivity of the overall electronic appliance. Because of this history of the semiconductor memories since when at the beginning of the '90s has been introduced, all the people tend to think like a commodity. Then there is a surprise. Like in the Dataquest report, yeah, it's a strange flash, it's not a commodity. In the beginning, everyone would have expected this. Usually, when you have a new product, you have a plethora of solutions. Even today, after probably 10 years of introduction, a lot of solution of flash to exist on the market. Several of them winning. The very reason is said that flash is not a commodity and in my eyes,

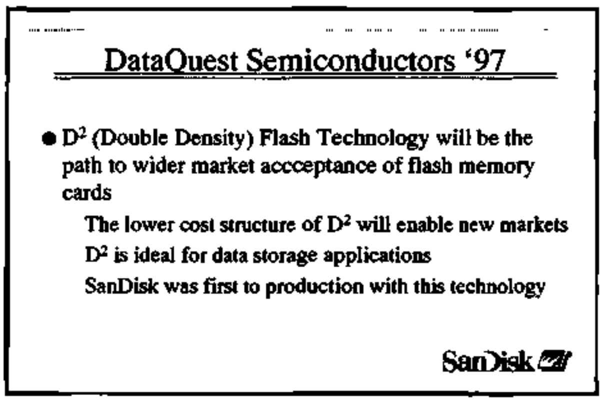

will never be a commodity. It's flexibility. The application dynamics and the enormous power of the software-driven approach will more and more identify flash in my eyes as an application-driven standard product for which both the classical paradigm of the commodity—low cost and substitution—coexist with the imperative or an application standard product.

The strength of the technology joined to each strategic role in the system is symbolically joined with the sector which drives its development. The question is what is the segment driving the technology. Global communication is driving, today, this technology. Global communication—it means to have mission of the voice of the daytime, of the image, of the picture. The flash designer should learn more and more how to communicate the low energy with the speed, with the density in an environment of increased value-added at lower and lower cost.

The cellular phone— it has already been said—is today the biggest eater of flash and cellular phone, while they're striving for the low energy—unfortunately striving also for the lower and lower price, fortunately for us —but the cellular phone is demanding the first integration of world memory fashion. The flash to emulate this EPROM, the flash to emulate the standard ROM.

In the near future, we'll see the flash in a total integration that will substitute the logic or will integrate the logic, the baseband and maybe in the future, even the energy management. The PDC-PCS networking will also move out from the classical under the values of the old PC or the today PC needing big super performance, direct access flash.

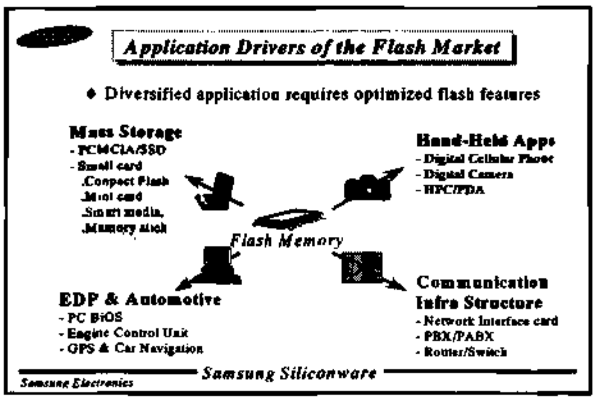

I disagree a little bit with Walid and I tend to agree more with Bill. I clearly see in this environment of the total of what they call the mobile communication—a big space for a big expansion. Very much likely, in three or four years from now, we'll see this as a new driver of the flash technology.

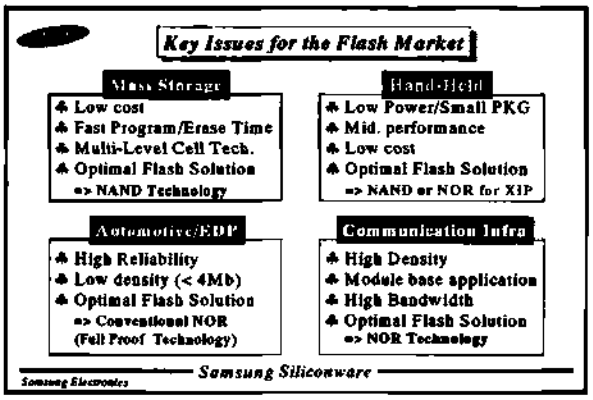

Let's take even another segment—automotive that everyone knows has been the first to use extensive with the flash at the beginning of the '90s. Even automotive is moving from the classical power train in body only to incorporate GPS, to incorporate radar, to incorporate cellular phone and the radio, to incorporate—computers, maybe tomorrow. Now, all this and even more even the consumer environment is changing.

All these do require special flash and again, I insist to make a point—flash is not a commodity. We all have a different point. Finally, we see that there is no other flash. Flash is the memory that allows what they call the time to system for the customer. It's the

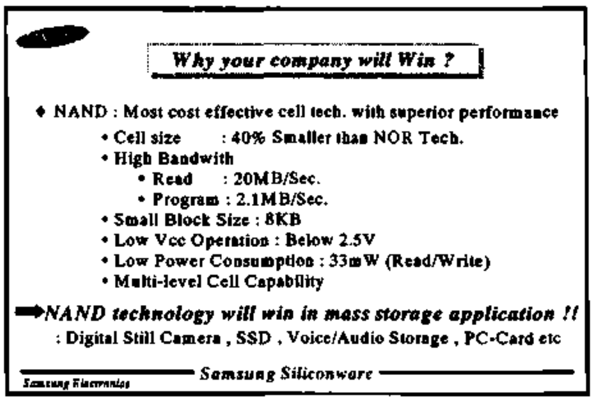

memory that allows the generation of a dream in the semiconductor since life the expert system. If we have to think in 10 years from now what will be the driver or flash, today, the global communication is what drives the flash market. A few years from now, what I call a mobile computing which is driving the market. Later on, the expert system is really what probably will drive the future.

Flash memory is fundamental to give more life to a system. It allows them to remember what they are and what they should do. It gives this system the life because the system code of the behavior can be changed by the system itself. Like life, it allows the system the propagation and the production. I see the future of the flash in the field of the future of the flash in a world on application oriented standard. Then how the flash designer should learn even more to work in symbiosis with the system designer. This either to convince the system designer to use the good flash we are making today or to take and being convinced by the system designer of a new product that has to be designed for them. At this point, it will come out as a new standard product or a new super integration or a new super device what they call flash in the system. In the end, what we have to win is the system user competitive advantage. Beyond the time to system including also the system total figure of cost and performance.

Beside this knowledge of the system, another important point is the technology platform. If we look at today's flash available and we fundamentally ignore as Walid said that today is the main driver, it's the main part of the market. If we look at the product today existing, today's architecture can be scalable down to four gigabit or down to .12 micron technology. Pay attention. Below the .25 micron, the cost could be a surprise. We love to solve the issue of the voltage—the voltage scalability, the array efficiency, the test efficiency, the huge cost of the future factories, then the availability of a single technology platform.

Eventually, a model of platform could be the key to solve the issue. The architecture of the whole process will be playing an important role. If we are able to develop a module or a process by which the high voltage transistor and the cell can be as a model up-front to the low voltage transistor, you come out with an architecturally process that is perfectly compatible with the high performance logic. In that case, you gain a lot. You have to gain because you are able to fill your fab. At that moment, practically, you don't really care how big the market is and how big the cost is because you have the possibility either via super integration or via the standard devices to really grow the market.

Obviously, the technological design solution for low power, minimum access the time, shortening the time to design should converge with the need of integration and availability to allow a reasonable cost notwithstanding the multibillion dollar cost of the next technology and fab. For standalone high density flash, I believe the multilevel approach will play a significant role but besides that, I clearly see other technological approach—multi-layer. I see the incorporation of new design technique that are not the wide array of the flash input. At the end, we're not the wrong interface.

I said at the end, you have a couple of issues—the knowledge of the system and the technology platform. This can be the issue but can also be the reason for success. This has been the reason for our success in the past. Now, the fact that you are able to develop for a competitive event of your customer, you're able to follow a virtual one extended enterprise concept. The fact that you have been very strong in a differential product, the fact that you are able to take product, innovation, commitment is something that could drive you in a possibility to win.

System ownership, global approach to flash—this is what we believe and the total commitment in driving the system. We have been quite strong in the past. Technology portfolio is one of our strengths together with the partnership. At the end, if I take back the ambiguity, would the bigger or better win? I see that the best will win. Obviously, we'll be big. Thank you.

Dan Auclair: SanDisk is a company that is exclusively focused on the data storage market for flash memory. We're a young company that has pioneered a number of new markets and new technology. One of the key areas that is emerging as a monster market is that of digital cameras. Some of the features of the hot new cameras are the megapixel resolution, large LCDs so you can have instantaneous feedback.

Examples of this is the new Kodak DC210 and the previous generation Kodak, the DC120. The 210 actually is over a megapixel and the 120 is a little below a megapixel. This year, in '97, they'll be about 600,000 to 800,000 digital cameras shipped with a substantial portion of that having only internal memory, no removable memory. In 1998, we think that will reverse and a majority of cameras will have only removable compact flash as storage and the hot cameras will be the high-end cameras—the ones with megapixel+ resolution. Along with HPCs, we see digital cameras as the key to expanding the consumer market for flash memory cards.

Our view is the CompactFlash has won the first battle for the small form factor cards. Today, we've got in excess of 120 design wins for CompactFlash. That's just not SanDisk but including other manufacturers of CompactFlash. They're actually now, since I made this slide, over 20 digital cameras that use CompactFlash. One of the most recent is the Casio QV-700. Up until the 700, Casio had only internal memory and with the QV-700, it now has CompactFlash as the only storage for the image. The very clear trend is to use CompactFlash as the only removable memory. In the HPC area, CompactFlash has won every design win except one. I would predict that that one will convert to CompactFlash in the not too distant future.

Another new technology that we absolutely agree with Intel on is that double density is going to change the vision that we have of flash memory in the future. There's no question it will drive down cost for data storage. Double density, or as Intel calls it MLC, is ideal for data storage applications with the right architecture. San Disk was first with this technology and we intend to remain the leader in double density technology in the data storage area. We can coexist nicely with Intel with their MLC although I do not believe that any chip that is designed primarily for code storage will be successful in the data storage application. That's a certainly a view that we differ quite strongly with on Intel and so far, I think our track record has demonstrated. This year, we've had very strong growth compared to some of the other flash markets that have had their relatively flat quarter revenue. Thank you.

Moderator: Last but not least, the world's largest memory manufacturer and the world's largest DRAM manufacturer, Mr. H.K. Lim.

Hyung-Kyu Lim: I think since you're saying "world's largest memory manufacturer," it might give the wrong impression. I didn't come here to sell all the market. I'm pretty sure we are one of the better technology companies. I think the present memory market will be led by those four application areas like the first mass storage market like PCMCIA. The second is handheld applications presented by Digital Cellular himself and conventional, traditional EPROM market like EDP and automotive and communication infrastructures. Those applications require different device features.

For the mass storage market, the low cost is the most important factor. Also this application typically requires like a millisecond program/erase endurance. The current app is the leading technology for this application. Cellular handset—lower VCC capability is very important. For the automotive, it's the reliability issue. The communication infrastructure requires high density and high speed flash memory.

For the last seven years, Samsung was concentrating on that mass storage. We expect we can remain as a major supplier in that area because we have a flash memory technology which we believe is the first technology for that application. The cell size is smaller than NOR by 40% and it has high bandwidth in read and write operation. It's also capable of achieving lower VCC. Also, surprisingly, it has the MLC capability, multi-level cell capability as we demonstrated in this year's ISSCC with 120 megabit chip.

Another point is we can utilize all the events of the DRAM process technology and the manufacturing facilities for those mass storage flash. We expect early next year, we can achieve $2 per megabyte with our 64 megabit chip. In late 1999, we can break $1 per megabit cost.

Who will win the market? They're trying to give us some point here. It doesn't matter each company's capabilities but one thing I'd like to bring to this session is the flash memory is still a memory. Although they have many unique features, it's still a memory. A lot of think without the memories, especially the process technology, we don't need any additional new process development for the flash. We just use NAND process technology. Also we can share 100% the facilities to DRAM. There are other areas which we can show the core competence. This is an advantage for the companies having other memories. To win the market, you should assess by your customers. Big is not enough all the time, you should be better all the time. Thank you.

Moderator: The floor is now open for questions.

Q: One question I would start off with for Dan Auclair of SanDisk is that all the other gentlemen on stage here sell integrated circuits. You don't. Does SanDisk plan to sell flash chips in a discreet matter in the future?

A: Let me correct that, Bruce. We actually do sell flash chips today and are active in the market and will continue to be active in the flash chip market. Today, it's a chip set, a controller plus of flash memory chip but obviously in the future, that could change.

Q: There's one question that keeps coming up here which is flash capacity that was touched on by all the panel members. Maybe if each of you could give a quick answer to do you think that currently and in the future, the flash market is undersupplied, oversupplied or maybe just right? Walid?

AMD: Today, there's a balance between the supply and demand. We are obviously adding capacity in anticipation of the demand. If the market continues to grow as forecasted, based on what we know today, the equilibrium would be maintained. However, you've got to look that there's a dynamic at each of the market segment and a sudden change in the requirement of a particular market could put a significant effect on the overall capacity. We are bringing in a brand new fab that's coming online by the end of this quarter but what I said in case of the cellular telephone which represents 25% of the overall market, it's using mainly 8 megabit flash right now. If all of a sudden they were to change to 16 megabit or 32 megabit, it could put a significant additional demand. It could change the balance all of a sudden.

Intel: There's an oversupply today of flash memory. That's driven down the prices in cyclical memory business fashion. When that happens as you all know, with the new price point, you generate a new set of applications kick in, the new set of design comes in and we expect the next wave of designs, for example, for the Stratoflash memory to kick in around the middle of next year and then we expect to be back in more of a balanced situation, perhaps even constraint. From a supply side, Intel's strategy is just continuing to add capacity chunk by chunk. The other thing I'd like to say though is with today's pricing for flash memory, it is more clear that you have to use leading edge technology and you have to move to the latest geometries. For example, moving from below .5, .4 to .25 next year as a requirement for success in this business and it's no longer a business of just taking some old factories and converting them to flash memory which was the case a few years ago. The whole dynamics of industry has changed and now it's a much more high stakes game. Major investments are required to stay competitive.

Macronix: I agree with what Bill's saying. That's why we build the age with the fab for the next generation flash technology. This new fab actually has a 40,000 wafers output capacity right now will only fill up 15,000. We believe today, we own that small portion of flash market but in the next few years, we will have tremendous potential to grow not only from our own technology improvement but also by working with the customer. We should be able to gain the market not trying to reduce our price or compete with the people on this floor here.

Sandisk: I agree that today there is adequate capacity. Although in our particular segment, data storage flash, demand's been very strong and we've been a little bit of allocation at the high capacity CF20 and 24 meg have been extremely strong. That's why SanDisk has invested heavily in capacity. We recently made a minority investment in a UMC joint venture. Right from the beginning will be .25 compatible or capable. I have to agree with Bill that you have to be on the leading edge to be competitive in this business.

SGS: I believe that all of us, we agree more or less. Similar situation. Big capacity today, my view is not really sure but is not really over. At least the flash. I believe it is more or less anyway. I'll give another explanation of the price war, what they call DRAM syndrome. I do agree with Bill, anyway, that we better go as big as we can to the .25 micron otherwise that's a problem. That poses a problem. .25 micron is not the cost, the capital investment or the .4. Then we'll have to sink a little bit. My view is that the investment that is now being done by the major player should keep the capacity reasonably aligned through the year 2000 unless something strange happens. Suddenly, as Walid said, if we change to 60 meg or whatever or if Intel changes the scenario.

Q: Got a couple of questions on Intel. You have a high profile here. There's one good question. They say why would Intel build a $1 billion fab to make flash devices to sell for $5 when it could use that capacity to sell $300 microprocessors?

Intel: That's a good question. Intel is looking for new areas to grow. Currently, we are the world's largest semiconductor company. We have a lot of money and we're looking for new business to grow. Intel is in a much different situation today than we were back in the famous 1984, 1985 situation where we had to make the tradeoff between microprocessors and DRAMs. We can basically afford to do both and we're committed to our flash memory business. We believe we've got an excellent technology strategy and we think we can make money in this business.

Q: Actually a similar question for AMD about everyone saying the business is flat. There's a question also on AMD's commitment to flash memory if it in fact flat even though you addressed it in your presentation.

AMD: Everyone's business may be flat, not AMD's business. AMD's business, from Q3, we grew our unit shipment 10%. We grew our dollar shipment from Q2 to Q3. Our fab is co-owned by us and Fujitsu is totally sold out. That's obvious because we are getting marketshare and Intel is losing marketshare. When you lose in marketshare, of course, you have excess capacity laying around.

Q: One other question here—Mr. Lim from Samsung, there's a question on gross margins of flash compared to DRAM. Do you expect the profits to be about the same between DRAM and flash?

Samsung: Depends on the situation. Two years it was great for DRAM. Now DRAM price is very down and the current leader of flash is better but because of that development investment, DRAM revenue should be higher. That's what we expect.

Q: There's another question for Bill on Stratoflash and MLC. I suppose this might apply to some different people too. Do you expect to apply MLC to lower densities or only to higher densities?

Intel: You get the most leverage from applying it to the highest density. We will continue to manufacture the mid-densities with the single bit per cell but eventually, the emphasis, we believe will be on the higher densities. That's where you really get the bang for the buck. Basically, we have a 32 meg, a 64 meg and that'll be sort of upwards from there.

Q: One of the questions that's been asked before and I would ask this to all the panel members is that what is the most important criteria for success in the flash market from both a business and a technology standpoint? Walid?

Intel: You have to customer driven. You have to know what are you designing a product for? What market are you serving because every market segment requires different architecture and has different requirements and may require totally different technology. If you think that you could approach the flash market like a DRAM and design a product where one density could meet 50% or 60% of the overall market and then the only thing that you have to have is the lowest cost in order to win, you are mistaken. Flash is several markets. Each one requires different architecture, different product. You have to be attuned otherwise you will miss the boat. If you're always a follower and copy whatever product is available, you're going to be totally behind and never catch up. You have to be very close to the customer and design every single product to meet the market segment that you're serving.

Macronix: You've got to do those and you've got to do more than that. From my perspective, you've got to grow the new markets. You've got to create the new applications that can take advantage of flash memory. That's the utility in flash memory. Intel has traditionally done that. Essentially, we created the flash memory market. It's our challenge to keep it growing. We're not satisfied with your $6 billion number. We don't like it. We want it to be bigger than that. We want it to be $8 billion or $9 billion and we're committed to go make that happen.

Sandisk: We believe the application is important. Keep a close relationship with the system customer is important and we believe the cost structure is very important. In order to achieve that, we have to create a technology which we can address different application a little more flexibility. I believe with that kind of technology, we'll win this market.

A: To be successful in this market, you've really got to focus and you've got to understand the target market you're going after. I'm repeating basically but saying it differently. In our case, we really have focused on data storage and looked at the specific applications and we what they need—not just from a device viewpoint but from an overall system viewpoint. The results have been extremely encouraging. You can't always just look at the individual component. You've got to look at the broader picture and understand what the problems are that the customers are trying to solve and make sure you've got the solution for them.

Samsung: In the area of mass storage flash, seems to me that market development is more important now. Actually, when we started to develop another type of flash, we were thinking about kind of small disk replacing hard disk. Today, we're selling most of our flash to those digital still camera and some audio storage. Quite different application compared to the original target. Currently, that mass storage area is really the market limited area. We need to promote efforts to develop a new market for mass storage and flash memory.

SGS: Three words—innovation, partnership and commitment. The partnership and the customer approach on top of the three.

Q: On the mass storage, I have a question for SanDisk and Samsung. On the mass storage, do you think that flash will ever get as cheap as disk? It's a hard disk drive because it's magnetic disk drive.

Sandisk: If you look at 10 gigabtye disk drives and compare it flash, you're not going to equate the same cost per megabit or megabyte. What we do envision though, today an entry price disk drive is $150 to $200 for a two gigabtye disk drive. If you don't need two gigabytes, if you need 40 megabytes or 100 megabytes, what we're starting to find are people much more willing to buy flash memory than they are our disk drives. There are many applications that will use 100 megabyte flash memory card instead of a two or three or four gigabyte hard disk drive. In that sense, yes, we'll take away some of the hard disk drive business but not directly, certainly not on a desktop but in mobile platforms.

Q: Would it ever be as cheap as disk?

A: No, it will never be as cheap as disk.

Intel: I've got to tell this. Just before the seminar I was up in the hotel room doing my email on my PC and my disk drive crashed. It was the most irritating thing in the world. I was sitting there right in the middle of this big, long email and the thing crashed. All I want is a simple light PC that I can write, receive my email messages and I think there's finally a market there for that. I'll pay a lot of money for that thing. I'm tired of lugging around that heavy PC and having that disk crash. It is just too irritating.

Samsung: I would agree with Dan and like that will be the major market for mass storage but we should have more market in digital consumer area—storing more audio/video data. There are always new applications coming from some unexpected areas. That's something we really want to have.

Q: Do you think that flash will ever be as cheap as magnetic disk drive?

Samsung: No. Maybe in low density area, like 100 megabyte type but gigabtye type, not in the near future.

Q: I have a question from a concerned buyer of flash memory in the here and now. It's to both Intel and AMD. They say that they can't use the same part. They can't either use an Intel part and put an AMD part in—they can't have the socket and use both parts. Do you foresee standardization in the flash memory industry in the foreseeable future? Walid?

AMD: We always ask Intel to design products compatible with AMD. The reason we are getting marketshare is because our customers like our product solution. We went from zero marketshare to an excess of 40% marketshare where Intel went from 90% to under 40% marketshare. The customer is saying they like the AMD technology, period. Now, for the same price, I could design my product over Intel at any time at lower power, at higher reliability. We're the only one in the world to guarantee 100,000 cycles with capability to do 1 million cycle—the only one in the world. Higher reliability, lower power, lower voltage, single power supply. Who needs more than one power supply to run his flash?

Q: Your customers never come to you and say please make yours compatible with Intel?

AMD: Absolutely not. As a matter of fact, they're going to the competition and asking them why can't you design a product compatible with AMD.

Q: They would ask you the same thing—why can't you make your product—

A: They never have.

Intel: I think in the early days, this was an issue. There was a lot of talk about a lack of standardization in the flash memory market was thwarting its growth. It was going to be another bubble memory market as a result. We've proven today it's a $3-ish billion market and will continue to grow. There's a need for different devices and that's what the customers want and the collective industry is providing those different solutions.

Moderator: Thank you, gentlemen. Do we have any more statements? Ladies and gentlemen, thank you very much for attending the Dataquest Semiconductor Conference for 1997. We appreciate your being here all these days. Thank you very much. Bye.

X

Index Table

Index Table

Sellside Analysts' Forecasts and Other Fairy Tales

Larry Bowman
Founder
Bowman Capital Management

In March 1997, Bowman Capital Management became the new General Partner of the Spinnaker Technology Fund and the Spinnaker Offshore Fund.

Mr. Bowman began his career with Texas Instruments in Dallas, Texas. In 1981 he joined Apple Computer spending the next four years managing production engineering and new product startups. After graduating in 1987 from Stanford Graduate School of Business, Mr. Bowman joined Fidelity Investments and for six years covered numerous industries as a high-technology analyst. His responsibilities included management of two technology mutual funds, Select Computer and Select Technology. In 1990, these two funds were both ranked Lipper Number 1 in their respective categories. From 1991 through late 1993, Mr. Bowman managed the newly formed Fidelity Emerging Growth fund and during his tenure the fund grew to over $750 million assets, while becoming Fidelity's Number 1 performing fund and history's fastest growing mutual fund.

Mr. Bowman joined Tiger Management, the $7 billion hedge fund in New York City as a managing director responsible for high-technology investments in 1993. His portfolio consisted of up to $1 billion of domestic and foreign technology holdings.

In January 1995, Mr. Bowman founded SoundView Asset Management where he launched the Spinnaker Technology Fund. In March 1997, Bowman Capital Management became the new General Partner of the Spinnaker Technology Fund and the Spinnaker Offshore Fund.

Mr. Bowman graduated from Lehigh University with a bachelor's degree in engineering in 1980.

Panel Discussion—Semiconductor Interconnect: The Market Differentiator

Panelist:
John Novitsky
Vice President of Marketing
MicroModule Systems

Mr. Novitsky is MicroModule Systems' vice president of Marketing.

Mr. Novitsky worked at Intel Corporation for more than 11 years on the 386, 486, and Pentium processor products in a variety of architectural, engineering, marketing, and management roles. He began his career at Intel in 1982.

Mr. Novitsky received a bachelor of science degree in computer science from Michigan State University. He is a member of the IEEE, ACM, and SIGARCH, and he is also a member of the editorial review board of *Microprocessor Report*.

Panel Discussion—Semiconductor Interconnect: The Market Differentiator

Panelist:
Steven Anderson

Senior Vice President
Product Marketing
Amkor Corporation

Mr. Anderson is senior vice president of Amkor's Corporate Product Marketing. He has been with Amkor since 1988 and has 20 years of industry experience.

Before joining Amkor, he held various positions with Texas Instruments and three start-up companies involved in semiconductor and electronics.

Mr. Anderson holds a bachelor of science degree in electrical engineering and an M.B.A. degree.

Panel Discussion—Internet Appliances: Are You Ready for the Cyber-Consumer?

Panelist:
David Limp

Vice President of Consumer Marketing
Network Computer Inc.

Mr. Limp is vice president of consumer marketing at Network Computer (NCI). He is responsible for marketing for the NCTV products to the consumer retail, cable, and satellite marketplaces. Additionally, Mr. Limp manages all consumer content partnerships for NCI. Before merging with NCI, Mr. Limp was the director of marketing for Navio Communications.

Mr. Limp spent nine years with Apple Computer Inc. Most recently, he was director of Apple Computer's North and South American PowerBook division. In that position, he oversaw the daily operations of the business, manufacturing, and channel, guiding Apple's efforts to renew the PowerBook line of products. Mr. Limp also held numerous roles in product marketing and sales during his tenure at Apple.

Mr. Limp holds a bachelor of science degree in computer science and mathematics from Vanderbilt University and a master of science degree in management from Stanford University School of Business.

Panel Discussion—
Will Big or Better Win the Flash Market?

Moderator:
Bruce Bonner

Principal Analyst
Semiconductor Memories Worldwide Program
Semiconductors Group

Mr. Bonner is a principal analyst for Dataquest's Semiconductor Memories Worldwide program. He is responsible for the forecasting and analysis of memory products and markets, including both volatile (DRAM, SRAM) and nonvolatile (flash, EEPROM, EPROM, ROM) types.

Before joining Dataquest, he was a marketing manager for Flash Memories at Intel Corporation in Folsom, California. Before Intel, he held various marketing and management positions in the rotating mass storage industry with Western Digital, Mitsubishi Electronics America, and Applied Magnetics.

Mr. Bonner received a bachelor of science degree in electronics engineering, from Cal Poly in San Luis Obispo, California.

President's Welcome

Manny Fernandez

Chairman of the Board, President, Chief Executive Officer, and Director
Gartner Group

Mr. Fernandez has been chairman of the board of Gartner Group since April 1995, chief executive officer since April 1991, and president and director since January 1991. Before joining Gartner Group, Mr. Fernandez was president and chief executive officer of Dataquest, an information services company that was acquired by Gartner Group in 1995. Before joining Dataquest, Mr. Fernandez was president and chief executive officer of Gavilan Computer Corporation, a laptop computer manufacturer, and Zilog Incorporated, a semiconductor manufacturing company.

He is presently serving on the board of directors of Brunswick Corporation, Getty Communications, SACIA (The Business Council of Southwestern Connecticut), and Norwalk Community Technical College and has previously served on the boards of Individual Inc., A.C. Nielsen Co., EMU Systems, Macmillan Inc., and ViewTech Inc.

Mr. Fernandez holds a bachelor's degree in electrical engineering from University of Florida and completed postgraduate work in solid-state engineering at University of Florida and in business administration at the Florida Institute of Technology.

Panel Discussion—
Will Big or Better Win the Flash Market?

Panelist:
Walid Maghribi

Group Vice President
Memory Group
AMD

Mr. Maghribi is group vice president of the Memory group at AMD.

Mr. Maghribi joined AMD as a product line manager in 1986 and was promoted to product line director in 1989. He was later named vice president and general manager in 1991 and assumed his current position in April 1997.

Before joining AMD, Mr. Maghribi was director of Operations at Seeq Technology for four years and a senior engineering manager at National Semiconductor for seven years. He also worked as an engineer at Intel Corporation.

Mr. Maghribi has a bachelor's degree in electrical engineering from San Jose State University and a master's degree in computer science from Santa Clara University.

Panel Discussion—
Will Big or Better Win the Flash Market?

Panelist:
Dan Auclair

Senior Vice President
Operations and Technology
SanDisk

Mr. Auclair is senior vice president of Operations and Technology at SanDisk and has over 20 years of experience in the mass storage industry.

Before joining SanDisk, Mr. Auclair was vice president of engineering at Anamartic, a company that utilizes wafer scale technology to build DRAM mass storage systems. He also was vice president and general manager of the OMTI division of Scientific Micro Systems, a leading supplier of disk controllers and disk controller chips to the disk drive industry.

Mr. Auclair has a bachelor of science degree in engineering physics from the University of Maine and a master of science degree in computer science from the University of Santa Clara.

Panel Discussion—Semiconductor Interconnect: The Market Differentiator

Panelist:
Tom DiStefano

Vice President of Marketing and Founder
Tessera

Dr. DiStefano is vice president of Marketing and founder of Tessera, a company devoted to advanced system interconnection products. At Tessera, he has developed technology addressing key issues in module level integration for small, high-performance systems. With the development team at Tessera, he has developed products chip size packaging for high-performance substrates and for systems cooling.

Before Tessera, Dr. DiStefano was a senior manager in Manufacturing Research with responsibility for manufacturing systems for packaging and semiconductors. He established and headed the Measurement Science and Technology Department, providing manufacturing technology to 22 IBM sites during his seven years in the department. He joined the IBM Research Division in 1970, where he led research projects in device interface physics, optical storage, and test technology. He has received IBM awards for his work on NLC Test Technology and on Photoemission Imaging of MOS Devices.

Dr. DiStefano has authored or coauthored more than 18 patents and 46 technical journal articles. He received a Ph.D. in applied physics from Stanford University in 1970, and a master of science degree in electrical engineering in 1965. At Stanford, he held an NSF Graduate Fellowship in electrical engineering. He received a bachelor of science degree in electrical engineering with highest honors in 1964 from Lehigh University.

Panel Discussion—Internet Appliances: Are You Ready for the Cyber-Consumer?

Panelist:
Raj Parekh

Vice President and General Manager
Volume Products Group
Sun Microelectronics

Mr. Parekh brings more than 20 years of engineering and management experience to Sun Microelectronics. He oversees the division's research programs, international R&D, architectural innovation, and Internet-enabling components to support Java and other network technologies. As vice president and general manager of the division's Volume Products group, Mr. Parekh oversees the proliferation of 32-bit SPARC and JAVA processors into the emerging embedded network marketplace.

Most recently, he was vice president of engineering and chief technology officer for Sun Microsystems Computer Corporation, were he was responsible for the computer system strategy, technology, and international design and development and before that was vice president of the company's Advanced Workstations division. Before joining Sun, Mr. Parekh spent 10 years in various engineering and general management positions at Silicon Graphics Inc.

Mr. Parekh holds a master of science degree in electrical engineering from the Polytechnic Institute of New York, and a B.E./B.S. in electrical engineering from the L.D. College of Engineering, India. Mr. Parekh also holds U.S. patents for bias control circuit for a substrate bias generator, EPROM reliability test circuit, and a segmented channel field-effect transistor.

Panel Discussion—
Will Big or Better Win the Flash Market?

Panelist:
Hyung-Kyu Lim
General Manager
Memory Division
Samsung Electronics Co.

Since 1976, Dr. Lim has been with the semiconductor business, Samsung Electronics Co. in Kiheung, Korea, and is currently responsible for all memory business including product development and research as general manager of Samsung's Memory Division. From 1978 to 1981, he was engaged in the development of bipolar linear integrated circuits and CMOS watch chips. After finishing his Ph.D. study, he worked mainly in the area of high-density MOS memory development. Starting from a 64Kb EEPROM design in 1984, he led various memory device research and development projects that include 256Kb EEPROM, 16Mb mask ROM, 1Mb high-speed static RAM, and 1/3-inch CCD image sensor.

Dr. Lim has authored or coauthored over 20 technical journal and conference papers and holds 23 patents. He is a member of the IEEE Electron Device Society.

Dr. Lim received a bachelor of science degree from the Seoul National University, a master of science degree from the Korea Advanced Institute Science and Technology, and a Ph.D. from the University of Florida, Gainesville, all in electrical engineering, in 1976, 1978, and 1984, respectively.

Panel Discussion—Internet Appliances: Are You Ready for the Cyber-Consumer?

Panelist:
Farid Dibachi

Chairman of the Board
Executive Vice President of Development
Diba

Dr. Dibachi brings more than 15 years of industry experience to Diba, most recently as chief executive officer and founder of Wavetron Microsystems. Before Wavetron, Dr. Dibachi was director of engineering and operations at Analogic Corporation where he led product development and oversaw the day-to-day sales, marketing, and manufacturing functions.

Dr. Dibachi began his career at Hewlett-Packard where he spent seven years in a variety of technical and management positions.

Dr. Dibachi holds graduate degrees in mechanical and electrical engineering from Stanford University and Cornell University, respectively.

Panel Discussion—
Will Big or Better Win the Flash Market?

Panelist:
Bill Howe

Vice President and General Manager
Memory Components Division
Intel Corporation

Mr. Howe is the vice president and general manager of the Memory Components division at Intel Corporation headquartered in Folsom, California. He joined the Memory Components division in 1994. Previously, he was president of Intel Japan with responsibility for building the country's business for Intel microprocessor and communications products. Before that, he held sales and marketing management positions for Intel's European operations.

Mr. Howe graduated from the University of Waterloo with a bachelor's degree in mathematics, and he holds a master's degree in business administration from Harvard.

Panel Discussion—Semiconductor Interconnect: The Market Differentiator

Panelist:
Bill Howe

Vice President and General Manager
Memory Components Division
Intel Corporation

Mr. Howe is the vice president and general manager of the Memory Components division at Intel Corporation headquartered in Folsom, California. He joined the Memory Components division in 1994. Previously, he was president of Intel Japan with responsibility for building the country's business for Intel microprocessor and communications products. Before that, he held sales and marketing management positions for Intel's European operations.

Mr. Howe graduated from the University of Waterloo with a bachelor's degree in mathematics, and he holds a master's degree in business administration from Harvard.

Panel Discussion—
Will Big or Better Win the Flash Market?

Panelist:
Bruno Beverina

Vice President, Memory Product Group
General Manager
FLASH Memory Division
SGS-Thomson Microelectronics

Mr. Beverina is vice president of the Memory Product group and general manager of the FLASH Memory division of
SGS-Thomson Microelectronics. Before his current position he was the general manager of the EPROM division.

Mr. Beverina began his career in the semiconductor industry in 1961 as a process engineer and occupied various functions in R&D, engineering, and marketing.

Mr. Beverina has a degree in electronic engineering and an M.B.A.

Cable TV: Delivering Digital Services to the Home

Henry Nicholas
Chief Executive Officer and President
Broadcom Corporation

Mr. Nicholas cofounded Broadcom Corporation with Dr. Henry Samueli. As chief executive officer and president, Mr. Nicholas is responsible for the strategic direction of the company and the day-to-day operations. His experience spans over 15 years in VLSI design technology and broadband data communications.

Before founding Broadcom, Mr. Nicholas was director of Microelectronics at PairGain Technologies. While at PairGain, he directed the development of the world's first high bit-rate digital subscriber line (HDSL) transceiver. Before working at PairGain, Mr. Nicholas held senior positions at TRW and was responsible for leading the development efforts for waferscale and ASIC programs.

Mr. Nicholas is very involved with the business community and was awarded 1996 Regional Entrepreneur of the Year for Electronics.

Mr. Nicholas holds a bachelof science degree and a master of science degree in electrical engineering from UCLA.

Intellectual Property: The Enabler of System-Level Integration

Dr. Walden C. Rhines
President and CEO
Mentor Graphics

Dr. Rhines is president and chief executive officer of Mentor Graphics, a leader in worldwide electronic design automation with revenue of $448 million in 1996.

Before joining Mentor Graphics, Dr. Rhines was executive vice president in charge of Texas Instruments' Semiconductor group with responsibility for over $5 billion of revenue and more than 30,000 people. He joined TI in 1972 and held a variety of technical and business management positions, primarily in the Semiconductor group, but also in the Consumer Products division, Central Research Laboratories, and Data Systems group. From 1985 to 1987, Dr. Rhines was president of the Data Systems group. During his career at TI, he was responsible for development of products including TI's first speech synthesis devices (used in "Speak & Spell") and the TMS 320 family of digital signal processors. He managed TI's microprocessor and application specific integrated circuit (ASIC) businesses from less than $10 million in annual revenue to nearly $2 billion.

Dr. Rhines served as chairman of the Semiconductor Technical Advisory Committee of the Department of Commerce, as an executive committee member of the board of directors of the Corporation for Open Systems, as a board member of the Computer and Business Equipment Manufacturers Association (CBEMA), and as a board member of Sematech. He is currently Chairman of the Electronic Design Automation Companies (EDAC) and a board member for the Oregon Independent College Foundation and of Lewis and Clark College. He also serves on the Boards of Cirrus Logic and Triquint Semiconductor.

Dr. Rhines holds a bachelor of science degree in metallurgical engineering from the University of Michigan, a master of science degree and Ph.D. in materials science and engineering from Stanford University, and a master of business administration from Southern Methodist University.

Big Opportunities in Wireless—Are RFICs Up to the Challenge?

David A. Norbury
President and CEO
RF Micro Devices Inc.

Mr. Norbury is president and chief executive officer of RF Micro Devices located in Greensboro, North Carolina. RF Micro Devices went public in June 1997 and is traded on Nasdaq under the symbol RFMD.

Mr. Norbury has more than 20 years of design engineering and general management experience in the RF and microwave field. He has worked for Watkins-Johnson and Frequency West and spent more than 12 years at Avantek where he was the division vice president in charge of several RF/microwave component product lines and Avantek's microwave subassemblies group. Before joining RF Micro Devices in 1992, Mr. Norbury was president and CEO for a multichip module start-up company in Santa Clara, California.

Mr. Norbury received a bachelor of science degree in electrical engineering from the University of Michigan, a master of science degree in electrical engineering from Stanford, and an M.B.A. from Santa Clara University.

Organizational Opportunism—Utilizing Change as a Competitive Advantage

John McCartney

President and Chief Operating Officer
U.S. Robotics

Mr. McCartney is currently U.S. Robotics president and chief operating officer. He joined the company in 1984 as vice president of finance and chief financial officer. He became an executive vice president in 1988 and was named executive vice president of international operations in 1990. His accomplishments include growing U.S. Robotics' international operations from $6.6 million in 1990 to $176 million in 1995.

With the closing of the 3Com-U.S. Robotics merger, Mr. McCartney will become president of the Client Access Products Business Unit for 3Com Client Access Products. He will also work closely with the Interface Products group, spearheaded by Doug Spreng, who will join his senior executive team.

Mr. McCartney holds an M.B.A. in operations and finance from the Wharton School and a bachelor of arts degree in philosophy from Davidson College.

Panel Discussion—
Emerging Technologies in Equipment and Processes

Panelist:
Eiji Takeda
Department Manager, System LSI Development Office
Semiconductor and IC Division
Central Research Laboratory
Hitachi Ltd.

Dr. Takeda was the department manager in the ULSI Research Department of the Central Research Laboratory at Hitachi until February 1996 and is now department manager of the System LSI Development Office of the Semiconductor and IC Divisions of Hitachi in Japan. Since 1975 he has been working for the Central Research Laboratory on VLSI devices and process physics and technologies. He has been managing VLSI memories (DRAMs, SRAMs, and Flash nonvolatile memories), BiCMOS, and advanced submicron MOS device-process groups. After working on the marketing of DVD (digital video disc), he is now working on such systems as LSI including the embedded DRAMs aiming at the "system on a chip," and its applications will be e-D graphic engines and PDA.

Since 1979, Dr. Takeda has been working on VLSI device physics and process technologies and memory applications including 1Mb to 1Gb DRAMs. Due to activities on hot-carrier effects, he received the 1994 IEEE Cledo Brunetti Award and since then has moved on to the system solution business.

Dr. Takeda was a visiting research associate at Cambridge University in the United Kingdom from September 1983 to September 1984. He has published and presented more than 140 international technical papers. He was also the program chairman of the 1994-1995 Symposium on VLSI Technology and he is a Fellow of the IEEE.

Dr. Takeda received a bachelor of science degree, a master of science degree, and a Ph.D. in applied physics from the University of Tokyo, Tokyo, Japan, in 1972, 1975, and 1987, respectively.

Panel Discussion— Emerging Technologies in Equipment and Processes

Panelist:
Dale R. Harbison

Vice President, Semiconductor Group
Manager, Manufacturing Science and Technology Center
Texas Instruments Incorporated

Dr. Harbison has been with Texas Instruments for 20 years. He is presently responsible for TI's Manufacturing Science and Technology Center. In this role he manages all of TI's equipment and supplier programs, factory and equipment systems integration programs, manufacturing operations and methods programs, contamination free manufacturing activity, and design/start-up of new wafer fabs.

In addition to this assignment, Dr. Harbison was responsible for the design and start-up of the TwinStar Semiconductor TI-Hitachi Joint Venture wafer fab in Richardson, DP1/DMOS 5 wafer fab in Dallas, TECH Semiconductor fab in Singapore, and the Avezzano wafer fab in Avezzano, Italy. During this time he was also heavily involved in both the TI-Acer and KTI start-ups. He has lived in Korea, Japan, Italy, and Singapore.

Dr. Harbison received a bachelor of science degree, a master of science degree, and a Ph.D. in electrical engineering from the University of Texas in Austin, Texas.

DSP—Be All That You Can Be

John Scarisbrick

Senior Vice President, Semiconductor Group
Worldwide Manager, Application Specific Products
Texas Instruments

Mr. Scarisbrick is senior vice president of Texas Instruments Semiconductor Group and worldwide manager of the Application Specific Products business. He is responsible for such product families as digital signal processors, microcontrollers, ASIC products, microprocessors, and networking.

Mr. Scarisbrick was named to this position in August 1996. Before this assignment he was president of TI Europe, responsible for all of TI's business operations in the European region, including its business units, manufacturing services and sales, and marketing offices in 16 countries.

Mr. Scarisbrick joined the company in 1976 as a field engineer in the United Kingdom and has since held a variety of positions in TI's Semiconductor Group. From 1984 through 1989 he was responsible for TI's Digital Signal Processing activities, based in Houston, Texas. Between 1989 and 1994, as vice president of Texas Instruments Europe, he was responsible for TI's European Linear business, based in Bedford, England. From February 1994 through February 1995, Mr. Scarisbrick managed TI's worldwide Computer Component organization, responsible for TI's SPARC and x86 microprocessor and networking business, based in Dallas, Texas.

Before joining Texas Instruments, Mr. Scarisbrick worked as a design engineer in Marconoi Space and Defense Systems, Frimley, England, and Rank Radio International, London, England.

IC Manufacturing 2000-2010: A Decade of Momentous Change

C.M. (Mark) Melliar-Smith

President and Chief Operating Officer
SEMATECH

Mr. Melliar-Smith was named president and chief operating officer of SEMATECH in November 1996. He joined the semiconductor manufacturing research consortium as an assignee from SEMATECH Member Company Lucent Technologies (formerly AT&T Microelectronics).

Before joining SEMATECH,, Mr. Melliar-Smith served as executive director of Integrated Circuits Division in AT&T Bell Laboratories and Chief Technical Officer for Lucent Technologies. In his 25-year career with AT&T, he has worked in a wide variety of assignments including fundamental research, electronic and photonic device development, and manufacturing and business unit management.

In 1970, Mr. Melliar-Smith joined AT&T Bell Laboratories in Murray Hill, New Jersey, as a member of the technical staff in the research area. His work involved materials research for advanced silicon integrated circuits with particular emphasis on reduced design rules and increased scales of integration. He was named head of the Semiconductor Laser Department in 1982 and transferred to the Reading Works, where he managed the final development of the semiconductor lasers used in the optical fiber communications system, and the rapid scale-up of laser production.

Mr. Melliar-Smith became the director of Lightwave Device Laboratory in 1983, overseeing the development of all optoelectronic devices used for lightwave communications systems. In 1984 he transferred to AT&T Technology Systems. As director of Engineering, he managed the overall development, manufacturing, and factory engineering at the Kansas City Works, including silicon IC manufacturing lines that produced the 256K and one-megabit computer memory chips. Mr. Melliar-Smith returned to Bell Laboratories in 1987 and was named executive director of the Electronic and Photonic Devices Division. In 1989 he assumed the position of vice president of Lightwave Strategic Business Unit with AT&T Microelectronics. His assignment involved acting as the chief operations officer of the Lightwave SBU. His responsibility for the unit covered manufacturing, research and development, marketing, strategic planning, and financial results.

Mr. Melliar-Smith has served on the SEMATECH Board of Directors since 1990 and serves as chairman of the SRC (Semiconductor Research Corporation) Board of Directors.

Mr. Melliar-Smith was born in England in 1945. He earned a bachelor of science degree and Ph.D. in chemistry from Southampton University in 1967 and 1970, respectively. In 1986, he obtained an M.B.A. from Rockhurst College, Kansas City, Missouri.

Panel Discussion—Semiconductor Interconnect: The Market Differentiator

Panelist:
Ed Fulcher

Director
Package Development
LSI Logic

Mr. Fulcher is the director of Package Development at LSI Logic. He was previously director of Package Development, Assembly Engineering, Assembly Production, and Memory Systems Development at UNISYS and manager of Microprocessor Systems Design at RCA.

Mr. Fulcher is a member of the Sematech National Technology Roadmap for Semiconductors representing packaging, and of the Semiconductor Research Corporation (SRC) for university graduate level packaging research.

Mr. Fulcher has authored and presented papers on MCM's computer system wiring rules and analysis, and designing integrated circuits to reduce switching noise. He has received several patents.

Mr. Fulcher did his undergraduate work in electrical engineering at Princeton University and the University of Maryland. His graduate work was in electrical engineering at the University of Florida and in business at Stanford University.

Panel Discussion—Internet Appliances: Are You Ready for the Cyber-Consumer?

Panelist:
Kevin Fielding
StrongARM Product Line Manager
Digital Semiconductor

Mr. Fielding is manager of Digital Semiconductor's StrongARM product family. The StrongARM product line is focused on delivering best-in-class microprocessors for embedded consumer applications such as PDAs and other smart handheld devices, Internet appliances, 3-D games, and interactive video systems. The group has design and marketing operations in Palo Alto, California, Austin, Texas, and Hudson, Massachusetts, in the United States, and in Reading, in the United Kingdom.

Mr. Fielding joined Digital Equipment Corporation in 1989, as an integrated circuit (IC) designer, and later became the product marketing manager responsible for Digital's Alpha microprocessor chips. He previously worked in the area of IC design as a principal design engineer for Philips BV in Europe, and as a research scientist at the National Microelectronics Research center (NMRC) in Cork, Ireland.

Mr. Fielding holds a bachelor of science degree and a master of science degree in electrical engineering from the National University of Ireland (Cork, Ireland), and an M.B.A. from Northeastern University, Massachusetts.

Panel Discussion—
Emerging Technologies in Equipment and Processes

Panelist:
Inseok S. Hwang

Senior Vice President
Semiconductor R&D Division
Hyundai Electronics Industries Co. Ltd.

Dr. Hwang has been with Hyundai Electronics in Korea since 1989, and is presently a senior vice president heading up the semiconductor R&D division, where he is responsible for research and development of semiconductor products and technologies.

Dr. Hwang has broad academic and industrial experiences. Before going to the United States in 1976, he had worked on TV circuits at Taihan Electric Wire Co., Ltd. In Korea. After earning his advanced degrees in the United States, he was an assistant professor of electrical engineering at the University of Minnesota, Minneapolis, where he researched diagnosable computer systems as well as taught logic design and microprocessors.

In 1984, he moved to AT&T Bell Laboratories. During the years there with the Signal Processor Laboratory at Whippany, New Jersey, and the VLSI Design Laboratory, Allentown, Pennsylvania, he was involved in the development of the EMSP and GSPA, both large-scale modular signal processors based on the data flow concept, and in the design of the WE 32200 microprocessor and the DSP16A digital signal processor. He was also engaged in the VLSI implementation of electronic switching subsystems for broadband ISDN.

Dr. Hwang received a bachelor of science degree in electronics engineering from Seoul National University, Korea, in 1972, and master of science and Ph.D. degrees in electrical engineering from the University of Wisconsin, Madison, in 1978 and 1982, respectively.

Panel Discussion—Semiconductor Interconnect: The Market Differentiator

Moderator:
Mohan Warrior

Director
Strategic Final Manufacturing,
Sector Manufacturing, SPS
Motorola

Mr. Warrior is the director of the Strategic Final Manufacturing, Sector Manufacturing, and SPS division at Motorola. His principal assignments include managing an advanced bipolar wafer fab, directing the advanced interconnect effort in SPS to drive bump technologies, and developing new multilevel structures. Mr. Warrior has authored several publications in these areas.

Mr. Warrior's areas of expertise are thin film technologies, MLM structures, especially for front-end/back-end integration in semiconductor manufacturing and advanced materials.

Mr. Warrior has a master of science degree in chemical engineering with 13 years of diversified experience in the semiconductor industry.

Technology and the Global Economies

Donald H. Straszheim

President
Milken Institute

Dr. Straszheim has recently joined the Milken Institute as president.

Dr. Straszheim previously served as chief economist of Merrill Lynch & Co. since 1985. He is a frequent writer and speaker on the economy and financial markets and has been a regular guest on CNN and CNBC. As primary economic spokesman for one of the world's largest securities firms and the architect of its global economic viewpoint, Dr. Straszheim's expertise is in the transnational economic issues that have increasingly become the focus of the Institute. His interests are globalization, education, capital markets and their interrelationship, as well as the crucial ways technology is changing the way we live and work. From 1981 to 1985, Dr. Straszheim was at Wharton Econometrics with responsibility for U.S. operations. Before this he was the chief economist of Weyerhauser Corporation from 1979 to 1981. From 1972 to 1979 he worked in investor services at a money management firm, and from 1970 to 1972 he was the chief economist at Fluor Corporation.

Dr. Straszheim received a master's degree in 1967 and a Ph.D. in 1971 from Purdue University, where he wrote his doctoral thesis on education's role in the economy.

Panel Discussion—
Emerging Technologies in Equipment and Processes

Panelist:
John E. Kelly III

Vice President
Strategy, Technology and Operations
Microelectronics Division
IBM

Dr. Kelly is the vice president of strategy, technology, and operations for the IBM Microelectronics Division (MD). Dr. Kelly joined IBM in 1980 as a senior associate engineer at IBM's Poughkeepsie, New York, location. He held a variety of management and technical positions in development and manufacturing and in 1990 was named director of the Semiconductor Research and Development Center (SRDC) with worldwide responsibility for the research and development of MD's semiconductor process technology. He was named vice president of business process re-engineering in 1994 and in 1995 was named the vice president of systems, technology, and science. It was in this position that Dr. Kelly established the Austin Research Lab for advanced microprocessors. He was named to his current position in 1996.

Dr. Kelly is a senior member of the Institute of Electrical and Electronics Engineers and is on the board of directors, Center for Integrated Electronics, at Rensselaer Polytechnic Institute (RPI). In addition, he is a member of the operating committee of Dominion Semiconductor, IBM's joint venture with Toshiba, the Dean's Engineering Council at Union College, and the foundation board of directors at St. Francis Hospital.

Dr. Kelly holds a bachelor of science degree in physics from Union College, a master of science degree in physics from RPI, and a Ph.D. in materials engineering from RPI.

Panel Discussion—
Emerging Technologies in Equipment and Processes

Panelist:
Peter Chang
President, U.S.C.
United Semiconductor Corporation

Dr. Chang joined the United Semiconductor Corporation in 1996 as the president of U.S.C. (an affiliated company of UMC) and specializes in foundry business. Dr. Chang's expertise is on process module development and process integration.

Dr. Chang was born in China and received his basic education in Taiwan. He came to the United States in 1969 and after furthering his education worked for various companies such as Hewlett-Packard, Zilog, Seeq Technology, and Paradigm Technology in the I.C. processing field. He has been involved in several wafer fab start-up situations.

In 1989, he returned to Taiwan to work for UMC. He began as a technical consultant and then, in 1992, was appointed to director responsible for the operation of Fab II. In 1994, he was promoted to vice president of the Business Group I in charge of Fab I, Fab II, and Q.S.

Dr. Chang received a bachelor of science degree in electrical engineering from Chenk Kung University in Tainan and a master of science degree and a Ph.D. in electrical engineering from the University of Texas at Austin.

Conference Opens

Gene Norrett

Corporate Vice President and Director
Semiconductors Group
Dataquest

Mr. Norrett is corporate vice president and director of Dataquest's Semiconductors group and is responsible for all worldwide semiconductor research, including Asia/Pacific-, Europe-, and Japan-based semiconductor research. Before this, he was director of marketing, responsible for the worldwide marketing strategies. Previously he was general manager for all North American technology services. Mr. Norrett was also the founder of Dataquest's Japanese Semiconductor Industry Service.

Before joining Dataquest, Mr. Norrett spent 14 years with Motorola's semiconductor product sector, serving in various marketing and management positions. Mr. Norrett was also a founder of the World Semiconductor Trade Statistics Program and was chairman of the Board of Directors of the Statistics Committee. He speaks frequently at Client Industry and Trade Association conferences. In 1987 he was voted by the *San Jose Mercury News* as one of Silicon Valley's top 100 influential people.

Mr. Norrett's education includes a bachelor of science degree in mathematics from Temple University and a master of science degree in applied statistics from Villanova University.

Panel Discussion—Tracking the "Food Chain": Dataquest's Worldwide Outlook for Key Technologies

Moderator:
Gene Norrett

Corporate Vice President and Director
Semiconductors Group
Dataquest

Mr. Norrett is corporate vice president and director of Dataquest's Semiconductors group and is responsible for all worldwide semiconductor research, including Asia/Pacific-, Europe-, and Japan-based semiconductor research. Before this, he was director of marketing, responsible for the worldwide marketing strategies. Previously he was general manager for all North American technology services. Mr. Norrett was also the founder of Dataquest's Japanese Semiconductor Industry Service.

Before joining Dataquest, Mr. Norrett spent 14 years with Motorola's semiconductor product sector, serving in various marketing and management positions. Mr. Norrett was also a founder of the World Semiconductor Trade Statistics Program and was chairman of the Board of Directors of the Statistics Committee. He speaks frequently at Client Industry and Trade Association conferences. In 1987 he was voted by the *San Jose Mercury News* as one of Silicon Valley's top 100 influential people.

Mr. Norrett's education includes a bachelor of science degree in mathematics from Temple University and a master of science degree in applied statistics from Villanova University.

Panel Discussion—Tracking the "Food Chain": Dataquest's Worldwide Outlook for Key Technologies

Panelist:
Gregory L. Sheppard

Chief Analyst
Semiconductor Application Markets Program
Semiconductors Group
Dataquest

Mr. Sheppard is chief analyst and manages Dataquest's Semiconductor Application Markets program responsible for coordinating worldwide semiconductor applications and system-specific devices market research for Dataquest. He oversees the research for the Semiconductor Application Markets (SAM) Worldwide, PC Semiconductors and Application Markets (PSAM), Consumer Multimedia Semiconductors and Application Markets (MSAM), and the Communication Semiconductors and Application Markets (CSAM) programs. He has also participated in various customer-directed research projects concerning company positioning and emerging semiconductor markets for application-specific products.

Before joining Dataquest, Mr. Sheppard was Worldwide Business Analysis Manager at Fairchild Semiconductor Corporation. In that position he coordinated the worldwide product and market plan that drove investment decisions. He has also been a participant in the World Semiconductor Trade Statistics (WSTS) organization and the American Electronics Association. Previously, he worked in engineering management at GTE Corporation specializing in communications systems design and decision aid systems.

Mr. Sheppard received a bachelor of science degree in electrical engineering and computer science from the University of Colorado and a master of science degree in system management from the University of Southern California.

Panel Discussion—Bandwidth: Who, What, When, and Where?

Moderator:
Gregory L. Sheppard

Chief Analyst
Semiconductor Application Markets Program
Semiconductors Group
Dataquest

Mr. Sheppard is chief analyst and manages Dataquest's Semiconductor Application Markets program responsible for coordinating worldwide semiconductor applications and system-specific devices market research for Dataquest. He oversees the research for the Semiconductor Application Markets (SAM) Worldwide, PC Semiconductors and Application Markets (PSAM), Consumer Multimedia Semiconductors and Application Markets (MSAM), and the Communication Semiconductors and Application Markets (CSAM) programs. He has also participated in various customer-directed research projects concerning company positioning and emerging semiconductor markets for application-specific products.

Before joining Dataquest, Mr. Sheppard was Worldwide Business Analysis Manager at Fairchild Semiconductor Corporation. In that position he coordinated the worldwide product and market plan that drove investment decisions. He has also been a participant in the World Semiconductor Trade Statistics (WSTS) organization and the American Electronics Association. Previously, he worked in engineering management at GTE Corporation specializing in communications systems design and decision aid systems.

Mr. Sheppard received a bachelor of science degree in electrical engineering and computer science from the University of Colorado and a master of science degree in system management from the University of Southern California.

Panel Discussion—Tracking the "Food Chain": Dataquest's Worldwide Outlook for Key Technologies

Panelist:
Joseph Grenier

Vice President and Director
Semiconductor Device Group
Dataquest

Mr. Grenier is vice president and director of Dataquest's Semiconductor Device group. He is responsible for managing the semiconductor device research for Dataquest's Memory, Microcomponents, ASIC/SLI, and Semiconductor Worldwide programs.

Before joining Dataquest, Mr. Grenier was marketing manager at GCA Corporation for the reactive ion etch program. He was also International Marketing Manager at GCA and was responsible for the overseas marketing of wafer-processing equipment. Previously, he worked as a product manager at Varian Associates/ Instrument Division, as a systems engineer at the USAF Satellite Test Center, and as a test engineer at General Motors' Noise Vibration Laboratory.

Mr. Grenier received a bachelor of science degree in electrical engineering from the University of Detroit and an M.B.A. from the University of Santa Clara.

Panel Discussion—
Major Device Trends and Forecasts: Looking Ahead

Moderator:
Joseph Grenier

Vice President and Director
Semiconductor Device Group
Dataquest

Mr. Grenier is vice president and director of Dataquest's Semiconductor Device group. He is responsible for managing the semiconductor device research for Dataquest's Memory, Microcomponents, ASIC/SLI, and Semiconductor Worldwide programs.

Before joining Dataquest, Mr. Grenier was marketing manager at GCA Corporation for the reactive ion etch program. He was also International Marketing Manager at GCA and was responsible for the overseas marketing of wafer-processing equipment. Previously, he worked as a product manager at Varian Associates/ Instrument Division, as a systems engineer at the USAF Satellite Test Center, and as a test engineer at General Motors' Noise Vibration Laboratory.

Mr. Grenier received a bachelor of science degree in electrical engineering from the University of Detroit and an M.B.A. from the University of Santa Clara.

Panel Discussion—Tracking the "Food Chain": Dataquest's Worldwide Outlook for Key Technologies

Panelist:
Clark J. Fuhs

Director and Principal Analyst
Semiconductor Equipment, Manufacturing, and Materials Program
Semiconductor Contract Manufacturing Services Program
Semiconductors Group, Dataquest

Mr. Fuhs is director and principal analyst for Dataquest's Semiconductor Manufacturing group, which includes the Semiconductor Equipment, Manufacturing, and Materials (SEMM) program and the Semiconductor Contract Manufacturing program, which covers the foundry industry in the Semiconductor Manufacturing group. He is responsible for research and analysis of semiconductor materials and trends in IC manufacturing techniques along with forecasting capital spending and the wafer fab equipment market. He is also responsible for directing worldwide research activities in semiconductor manufacturing including foundry, fab capacity, epitaxial silicon, and silicon supply and demand.

Before joining Dataquest, Mr. Fuhs was strategic marketing manager for Genus Inc., a manufacturer of advanced chemical vapor deposition (CVD) and high-energy ion implantation equipment. During his 10 years at Genus, he held positions of product manager, several responsibilities in product marketing, and process engineer in the metal CVD group. In his most recent position, Mr. Fuhs was responsible for correlating process techniques with demand for equipment and materials. He has been involved with the Modular Equipment Standards Committee of SEMI, a trade organization, as chairman of a task force, authoring a standard. His experience also includes Chevron Oil, where he was a process engineer in the Richmond, California, refinery responsible for the hydrogen manufacturing plant.

Mr. Fuhs earned a bachelor of science degree in chemical engineering from Purdue University in West Lafayette, Indiana, and received an M.B.A. from the University of California at Berkeley.

Panel Discussion—
Emerging Technologies in Equipment and Processes

Moderator:
Clark J. Fuhs

Director and Principal Analyst
Semiconductor Equipment, Manufacturing, and Materials Program
Semiconductor Contract Manufacturing Services Program
Semiconductors Group, Dataquest

Mr. Fuhs is director and principal analyst for Dataquest's Semiconductor Manufacturing group, which includes the Semiconductor Equipment, Manufacturing, and Materials (SEMM) program and the Semiconductor Contract Manufacturing program, which covers the foundry industry in the Semiconductor Manufacturing group. He is responsible for research and analysis of semiconductor materials and trends in IC manufacturing techniques along with forecasting capital spending and the wafer fab equipment market. He is also responsible for directing worldwide research activities in semiconductor manufacturing including foundry, fab capacity, epitaxial silicon, and silicon supply and demand.

Before joining Dataquest, Mr. Fuhs was strategic marketing manager for Genus Inc., a manufacturer of advanced chemical vapor deposition (CVD) and high-energy ion implantation equipment. During his 10 years at Genus, he held positions of product manager, several responsibilities in product marketing, and process engineer in the metal CVD group. In his most recent position, Mr. Fuhs was responsible for correlating process techniques with demand for equipment and materials. He has been involved with the Modular Equipment Standards Committee of SEMI, a trade organization, as chairman of a task force, authoring a standard. His experience also includes Chevron Oil, where he was a process engineer in the Richmond, California, refinery responsible for the hydrogen manufacturing plant.

Mr. Fuhs earned a bachelor of science degree in chemical engineering from Purdue University in West Lafayette, Indiana, and received an M.B.A. from the University of California at Berkeley.

Surviving and Thriving in the Chipless Business

Gary Smith

Director and Principal Analyst
Electronic Design Automation Program
Online, Multimedia, and Software Group
Dataquest

Mr. Smith is a director and principal analyst for the Electronic Design Automation program in Dataquest's Online, Multimedia, and Software group. He is responsible for all research, publications, and client projects relating to the electronic design automation marketplace and is involved in research and consulting projects in the emerging methodologies in RT Level and ES Level design.

When Mr. Smith came to Dataquest in January 1994, he already had 20 years of experience in electronic design. Starting in the semiconductor industry, he was involved in some of the first attempts at customer-designed ICs. During the 1980s, he specialized in the ASIC end of the semiconductor business. While at LSI Logic, Mr. Smith became involved in the development of the RT Level design methodology, later leaving the company to become a consultant in design methodology. Mr. Smith is a current member of the EDA Industry Council and serves as a member of the System Level Design Language working group.

Mr. Smith earned his bachelor of science degree in engineering from the United States Naval Academy at Annapolis, Maryland.

Panel Discussion—Bandwidth: Who, What, When, and Where?

Panelist:
Stephen Diamond

Vice President
Worldwide Telecommunications Group
Dataquest

Mr. Diamond is vice president of the Worldwide Telecommunications group and is responsible for managing Dataquest's global research in LANs, WANs, public networking, voice communications, and personal communications. He has more than 15 years of worldwide management experience in the networking and communications industry.

Mr. Diamond joined Dataquest from AirWave Networks Inc., a wireless networking start-up, where he was vice president of Marketing and Sales. Before joining AirWave, Mr. Diamond held various senior management positions, including vice president of Marketing and Service for Retix and vice president of Corporate Marketing for Ungermann-Bass (UB), where he was responsible for launching the company's industry-leading ATM and LAN switching products. He also brings international experience to Dataquest, having lived in Brussels, Belgium, while director of European Marketing for UB.

Mr. Diamond has been a frequent speaker at industry conferences and has been a consultant to companies including IBM, Digital, Apple Computer, and NEC America. Mr. Diamond is a graduate of Boston College and holds advanced degrees in engineering and management from Tufts University and Northeastern University. He has also completed the AEA/Stamford Executive Institute for the Management of High Technology Companies.

Panel Discussion—Bandwidth: Who, What, When, and Where?

Panelist:
Bobbi Murphy

Chief Analyst, Remote Access Program
Telecommunications Group
Dataquest

Ms. Murphy is chief analyst for the Remote Access program of Dataquest's Telecommunications group and has more than 15 years of industry experience in data communications. Ms. Murphy's marketing and product management experience includes a diverse set of technologies including ATM, routing, hubs, SNMP, wireless, and SNA protocols. She is responsible for the new Remote LAN and Internet Access program launched in August 1996.

Ms. Murphy joins Dataquest from RadioLAN, a wireless LAN start-up, where she was vice president of marketing. Before joining RadioLAN, Ms. Murphy was vice president of marketing for Hughes LAN Systems, a division of GM/Hughes Electronics. Other experience includes 3Com Corporation where Ms. Murphy was director of marketing for its SYSTEMSGroup, which comprises routers, bridges, and communications servers.

Ms. Murphy has an M.B.A. from Texas Tech University and an undergraduate degree from Carnegie Mellon.

Panel Discussion—Bandwidth: Who, What, When, and Where?

Panelist:
Brett Azuma

Director and Principal Analyst
Public Network Equipment and Services Program
Telecommunications Group
Dataquest

Mr. Azuma is director and principal analyst for Dataquest's Public Network Equipment and Services program for North America, focusing on the public network service market with a special emphasis on high-speed data services for both the consumer and business market.

Before joining Dataquest, Mr. Azuma's most recent assignment was with Pacific Bell Communications as head of Data Services and New Product Development. In his careers he has held a variety of leadership positions in the marketing, product development, and engineering departments of Pacific Telesis. This included leading Pacific Bell's Advanced Intelligent Network efforts including service and infrastructure and product development; leading the creation of Pacific Bell Communications' New Product Development process; leading the deployment of Pacific Bell's ISDN Service Center; testing and certifying new telecommunications services; leading a team of senior engineers responsible for testing and certification of new switching hardware and software; leading the introduction of Pacific Bell's ATM services; codeveloping the foundation for Pacific Bell's CalREN Initiative focused at providing research grants to foster the development of high-speed data applications; negotiating the technical terms and conditions for Pacific Bell's digital switch replacement agreement; and leading the development of Pacific Bell Communications Data Strategy.

Mr. Azuma received a bachelor of science degree in electrical and computer engineering from the University of California at Davis.

Panel Discussion—
Major Device Trends and Forecasts: Looking Ahead

Panelist:
Nathan Brookwood

Principal Analyst
Personal Computer Semiconductors and Applications Program
Semiconductors Group
Dataquest

Mr. Brookwood serves as principal analyst for the Personal Computer Semiconductors and Applications Worldwide program in Dataquest's Semiconduuctor group, and also supports research in Dataquest's Embedded Microcomponents Worldwide program. He directs research on microprocessors used in computational and embedded applications, including personal computers, workstations, and servers, with special emphasis on those based on X86 and popular RISC architectures.

Mr. Brookwood joined Dataquest from Micronics Computers, where he helped market small Pentium-based computers with 130 MIPS of processing power. Earlier, he worked for Intergraph Corporation and planned the strategy for their 10.0-MIPS CLIPPER microprocessors. Mr. Brookwood has also been employed at Convergent Technologies, where he directed the marketing of a line of proprietary X86-based workstations, Prime Computer, a then fast-growing (and now extinct) vendor of 1.0-MIPS minicomputers, and Digital Equipment Corporation where he developed operating systems for powerful 0.1 MIPS minicomputers.

Mr. Brookwood received a bachelor of science degree from the Massachusetts Institute of Technology, where he had access to an advanced 0.01 MIPS IBM mainframe, and attended the Program for Management Development at the Harvard Business School.

Panel Discussion—
Major Device Trends and Forecasts: Looking Ahead

Panelist:
Bryan Lewis
Director and Principal Analyst
ASICs Worldwide Program
Semiconductors Group
Dataquest

Mr. Lewis joined Dataquest in 1985 and is the director and principal analyst of Dataquest's ASICs Worldwide program in the Semiconductors group. His focus is on analysis of system-level integration (SLI), cell-based ICs, gate arrays, and PLDs. He has responsibility for tracking and evaluating market movements, forecasting markets, and tracking technology trends. Mr. Lewis founded Dataquest's ASICs program and is responsible for instituting annual supplier and regional end-user surveys. He has spoken at numerous conferences, traveled extensively in Asia, Europe, and North America, and given consultation to a wide variety of clients.

Before joining Dataquest, Mr. Lewis was a research manager for a semiconductor market research company. His responsibilities included establishing and managing a primary research group as well as performing research on GaAs ASICs.

Mr. Lewis received a bachelor of science degree in marketing from the University of Oregon.

Panel Discussion—
Major Device Trends and Forecasts: Looking Ahead

Panelist:
Jordan Selburn
Principal Analyst
ASICs Worldwide Program
Semiconductors Group
Dataquest

Mr. Selburn is a principal analyst for Dataquest's ASICs Worldwide program. He is responsible for the evaluation and analysis of the system-level integration (SLI) market and
cell-based and gate array ASICs, including tracking and forecasting both technology and market trends.

Before joining Dataquest, Mr. Selburn was the marketing manager for ASIC products at LSI Logic, with responsibility for the portfolio of core products. Also at LSI Logic, he was the product manager for a number of ASIC products ranging from 0.6-micron to 0.25-micron, including both high-performance and consumer technologies. Before LSI Logic, he held technology management and applications positions at Cadence Design Systems, as well as microwave design and applications positions with EEsof, Harris, and Watkins-Johnson.

Mr. Selburn received a bachelor of science degree in electrical engineering from the University of Michigan and a master of business administration degree from Santa Clara University. He is currently enrolled in Stanford University's graduate program in engineering-economic systems and operations research.

Panel Discussion—
Major Device Trends and Forecasts: Looking Ahead

Panelist:
Jim Handy

Director and Principal Analyst
Semiconductor Memories Worldwide Program
Semiconductors Group
Dataquest

Mr. Handy is director and principal analyst for Dataquest's Semiconductor Memories Worldwide program. He is responsible for the forecasting and analysis of memory products and markets.

Before joining Dataquest, Mr. Handy was strategic marketing manager for static RAMs at Integrated Device Technology (IDT). Before IDT, he was product marketing manager of memory and microcomputer-based products at Intel Corporation, National Semiconductor Corporation, and Siemens Corporation and has a rigorous design background. Mr. Handy is the author of *"The Cache Memory Book"* (Academic Press, 1993) and his other work has been widely published in the trade press including *Electronic Design*, *Computer Design*, *EDN*, and *Byte.* He has spoken internationally at universities and numerous trade shows including Wescon, Electro, WinHEC, Northcon, Southcon, and the Personal Computer Design Conference. Mr. Handy is reported to be the semiconductor industry's most often quoted analyst and is frequently quoted in the electronics trade press. Mr. Handy is also a patent holder in the field of static RAMs.

Mr. Handy earned an M.B.A. at the University of Phoenix and holds a bachelor of science degree in electrical engineering from Georgia Tech.

Panel Discussion—Internet Appliances: Are You Ready for the Cyber-Consumer?

Moderator:
Dale L. Ford

Senior Industry Analyst
Semiconductor Application Markets Group
Dataquest

Mr. Ford is responsible for conducting market research and analysis for the Semiconductor Application Markets group at Dataquest. He is a specialist on the end use or application of semiconductors with the scope of analysis including both economic and technical trends regarding the semiconductor content of electronic equipment. His work also includes contributions on client-specific consulting projects.

Mr. Ford is the program manager for the Consumer Multimedia Semiconductors and Applications program and also has primary responsibility for Dataquest research in wireless communications and mobile computing semiconductor applications. In addition, he contributes to the general semiconductor applications research.

Before his current role, Mr. Ford completed major consulting projects in the telecommunications, mobile computing, and multimedia industries for Dataquest. His work included the development of forecasting models to project the development of new technologies and the growth of emerging markets. He also led the launch of Dataquest's successful Teardown program where in-depth analysis is performed on electronic equipment including PCs, workstations, cellular phones, set-top boxes, and video games.

Before joining Dataquest, Mr. Ford was employed by Sun Microsystems in its product marketing organization where he created and implemented marketing plans and joint development agreements with third-party vendors. Earlier, he was a design engineer working with real-time image processing technologies and computer-aided-engineering systems for Evans & Sutherland, a producer of graphics workstations and high-performance flight simulators.

Mr. Ford has an M.B.A. in strategic management from The Wharton School, University of Pennsylvania, and a bachelor of science degree in electrical engineering from Brigham Young University.

Special Report—Inside the Mind of the End User: What Semiconductor Users Really Care About

Jonathan L. Yarmis

Vice President and Manager
Special Projects
Gartner Group

Mr. Yarmis is vice president and manager of special projects at Gartner.

Mr. Yarmis has been with Gartner Group since 1987, most recently serving as service director of the PC service. He joined Gartner Group from General Instrument, where he was a personal computer analyst and corporate MIS; responsible for all aspects of corporate personal computer implementation including policy, support, hardware, and software evaluation and training. Before that, Mr. Yarmis was at Touche Ross & Co., where he was director of microcomputer services for the Financial Services Center. Mr. Yarmis is a regular speaker and panelist at industry conferences and trade shows, as well as a frequent contributor to major trade publications.

Mr. Yarmis has a bachelor's degree in economics from Hamilton College.

Special Speaker Panel—Bringing It All Together

Moderator:
Jonathan L. Yarmis

Vice President and Manager
Special Projects
Gartner Group

Mr. Yarmis is vice president and manager of special projects at Gartner.

Mr. Yarmis has been with Gartner Group since 1987, most recently serving as service director of the PC service. He joined Gartner Group from General Instrument, where he was a personal computer analyst and corporate MIS; responsible for all aspects of corporate personal computer implementation including policy, support, hardware, and software evaluation and training. Before that, Mr. Yarmis was at Touche Ross & Co., where he was director of microcomputer services for the Financial Services Center. Mr. Yarmis is a regular speaker and panelist at industry conferences and trade shows, as well as a frequent contributor to major trade publications.

Mr. Yarmis has a bachelor's degree in economics from Hamilton College.

2010: A Vision of the Future

Gordon Bell

Senior Researcher
Telepresence
Microsoft Corporation

Mr. Bell is a senior researcher in the Telepresence group at Microsoft. He is also a computer industry consultant-at-large working on computer-related projects, especially parallel processing, and is a director and partner of the Bell-Mason Group that provides systems for venture development.

Mr. Bell spent 23 years at Digital Equipment Corporation as vice president of Research and Development where he was responsible for Digital's products. He was the architect of various mini- and time-sharing computers and led the development of DEC's VAX and the VAX Computing Environment. He has been involved in, or responsible for, the design of many products at Digital, Encore, Ardent, and other companies. During 1966 through 1972 he was Professor of Computer Science and Electrical Engineering at Carnegie-Mellon University. He was the first Assistant Director of the National Science Foundation's Computing Directorate from 1986 through 1987. He led the National Research Network panel that became the NII/GII and was an author of the original High Performance Computer and Communications Initiative. Mr. Bell has authored books such as High Tech Ventures: The Guide to Entrepreneurial Success, and papers about computer structures and start-up companies. He is on the boards of Ambit, Adaptive Solutions, Cirrus Logic, DES, Fakespace, Microsoft, University Video Communications, Sun Microsystems, and CSC's Vanguard Group. He is a director of the Bell-Mason Group supplying expert systems for venture development to start-ups, investors, governments, and entrepreneurial ventures. He is a founder and Overseer of The Computer Museum, Boston.

Mr. Bell is a member of various professional organizations including the American Academy of Arts and Sciences (Fellow), American Association for the Advancement of Science (Fellow), ACM (Fellow), IEEE (Fellow and Computer Pioneer), and the National Academy of Engineering. His awards include the IEEE Von Neumann Medal, the AEA Inventor Award for the greatest economic contribution to the New England region, and The 1991 National Medal of Technology. He also received the 1995 MCI Communications Information Technology Leadership Award for Innovation.

Mr. Bell was born in Kirksville, Missouri, andbegan his career as an electrician at Bell Electric until he entered Massachusetts Institute of Technology where he obtained his S.B. (1956) and S.M. (1957) degrees. In 1958 he was a Fulbright Scholar to Australia.

Mr. Bell resides in Los Altos, California, and Boston, Massachusetts.

Panel Discussion—Internet Appliances: Are You Ready for the Cyber-Consumer?

Panelist:
Randall T. Littleson

Vice President of Marketing
Spyglass Inc.

Mr. Littleson was promoted to vice president of marketing in October 1996. He joined Spyglass in July 1996 as director of product marketing. Mr. Littleson is helping to build Spyglass' leadership position in making consumer devices, office equipment, and industrial machinery work with the Web.

Mr. Littleson joined Spyglass from Seagate Software, where he was director of marketing for the company's Storage Management Group. He was responsible for all strategic marketing initiatives for the newly formed $85 million division, including product marketing, channel marketing, and marketing research.

Earlier, Mr. Littleson spent six years with Palindrome Corp., makers of storage management software for high-end networks. He joined Palindrome as the company's first product manager and later served as both director and executive director of product managemen, with complete responsibility for the company's product strategy. Shortly after Seagate Software purchased Palindrome in August 1994, Mr. Littleson was promoted to vice president of marketing and product management. In this position, he successfully managed the launch of five major new releases over eight months and completely repositioned the entire product line.

Before Palindrome, Mr. Littleson worked as a systems engineer with Novell and as a systems analyst and representative with Unisys Corp., gaining valuable experience in product design, sales, and marketing.

Mr. Littleson earned a bachelor's degree in computer sciences and communications from the University of Michigan and a master's degree in business administration (M.B.A.) from the Keller Graduate School of Management.

Panel Discussion—
Will Big or Better Win the Flash Market?

Miin Wu

Founder and President
Macronix International Co. Ltd.

Mr. Wu is the founder and president of Macronix, focusing on nonvolatile memory such as ROM, EPROM, and flash and systems-on-chip integration. The company has developed many technologies such as modems, networking, DSP applications, audio, and video. Macronix maintains its position in the global market with more than 50 percent of its sales coming from the Japanese territory. Mr. Wu founded the company in November 1984 and was the vice president working on process development and process transfer to major Japanese, Korean, and Taiwanese companies. He also had sales responsibilities in the Far East, including Japan, Korea, Taiwan, and Hong Kong.

Before founding Macronix, Mr. Wu was process development manager at VLSI Technology; process development engineer and program manager at Intel; process development engineer and section manager at Rockwell International; and process development engineer at Siliconix.

Mr. Wu is also a director of the Taiwan Electrical and Electronics Manufacturers' Association (TEEMA); chairman of the Semiconductor Committee (TEEMA); executive director of the Association of Allied Industries in Science-Based Industrial Park; director of the Surface Mount Association; recipient of a Premier Award as a Contemporary Business Leader in 1993 from *Business Weekly*; secretary of the Stanford Alumni Association in Taiwan; consultant to the Electronics and Communications Subcommittee for the Ministry of Education; executive director of the Taiwan Semiconductor Industry Association; and director of the Electronic Devices and Materials Association.

Mr. Wu received a master of science degree in materials science and engineering from Stanford University and a master of science degree and a bachelor of science degree in electrical engineering from the National Cheng-Kung University.

Panel Discussion—Semiconductor Interconnect: The Market Differentiator

Panelist:
Rama Shukla

Director
Advance Interconnect
Intel Corporation

Dr. Shukla, an Intel employee since 1979, is the advanced interconnect technology development manager in the Components Technology Development Group at Intel, in the Santa Clara, California facility. He is currently involved in developing and bringing into manufacturing advanced interconnects for next-generation CPU products, developing flip chip technology with associated collaterals, and definition of long-term packaging development programs.

Over the last 18 years at Intel, Dr. Shukla has worked on developing and delivering VLSI silicon thin film-interconnect processes for NVM and logic silicon products as well as VLSI packaging/assembly materials and processes. Recently, he has also led the efforts on developing high-density substrates for MCM's, wafer bumping for TCP/flip chip, and related technologies.

Dr. Shukla has published numerous papers in the field of VLSI process and microelectronics packaging and developed and taught courses in these areas through various professional organizations.

Dr. Shukla has a master's degree in solid-state chemistry from I.I.T.-K, India, and a Ph.D. in materials science from the University of California, Berkeley.