Dataquest Presents

Semicondutors '96

$300 Billion Challenge:
Let the Games Begin!!

CONFERENCE TRANSCRIPT

The 22nd Annual Semiconductor Conference

October 24-25, 1996
Indian Wells, California

Prepared Especially for:

Gene Norrett
Dataquest

57

1996

Semiconductors '96

$300 Billion Challenge: Let the Games Begin!!

October 24-25, 1996
Renaissance Esmeralda Resort
Indian Wells, California

Dataquest

Agenda

Semiconductors '96

Renaissance Esmeralda Resort,
Indian Wells, California,
October 24-25, 1996

Thursday, October 24 General Session

7:00 - 8:00 am **Registration and Continental Breakfast** ______ **Crystal Ballroom**

8:00 - 8:15 am **Welcome**
Gene Norrett, *Corporate Vice President and Director, Semiconductors Worldwide, Dataquest*

8:15 - 8:30 am **President's Welcome**
Manuel A. Fernandez, *Chairman of the Board, President and CEO, Gartner Group*

8:30 - 9:15 am **The State of the Economy: Looking Ahead**
Donald H. Straszheim, *First Vice President and Chief Economist, Securities Research and Economics, Merrill Lynch*

9:15 - 10:00 am ***Keynote:* Olympic Long Jump: From $150B to $300B in Five Years**
Wilfred J. Corrigan, *Chairman and CEO, LSI Logic Corporation*

10:00 - 10:30 am ***Networking Break***

10:30 - 12:00 pm ***Panel:* Dataquest's Worldwide Outlook for Systems, Semiconductors, and Equipment**

Moderator:
Gene Norrett, *Corporate Vice President and Director, Semiconductors Worldwide, Dataquest*

Services—Electronic Systems Outlook
Gregory L. Sheppard, *Chief Analyst, Semiconductor Application Markets Program, Semiconductors Worldwide, Dataquest*

Semiconductor Device Outlook
Gary J. Grandbois, *Vice President and Chief Analyst, Semiconductors Worldwide, Dataquest*

Semiconductor Equipment and Capacity Outlook
Clark J. Fuhs, *Director and Principal Analyst, Semiconductor Equipment, Manufacturing, and Materials Program, Semiconductors Worldwide, Dataquest*

Agenda Continues

Agenda

12:00 - 1:30 pm **Will the Internet Break Microsoft?** __________________________ **Valencia Ballroom**
Luncheon and featured speaker
Roger B. McNamee, *General Partner, Integral Capital Partners*

Track 1, Thursday **The Battle for the Consumer Dollar** **Crystal ABC**

1:30 - 3:00 pm ***Panel:*** **Emerging Opportunities in the Battle for the Local Loop**

Moderator:
Gregory L. Sheppard, *Chief Analyst, Semiconductor Application Markets Program, Semiconductors Worldwide, Dataquest*

Robert L. Bailey, *President and CEO, PMC Sierra*
Russell K. Johnsen, *Vice President and General Manager, Communications Division, Analog Devices*
Stephen J. Edwards, *Assistant Vice President, Telco Broadband Networks and Applications, Northern Telecom*
Glenn H. Estes, *President, Telesis Technology Laboratory*
Lisa Pelgrim, *Industry Associate, Networking Program, Telecommunications Group, Dataquest*

3:00 - 3:30 pm ***Networking Break***

3:30 - 5:00 pm ***Panel:*** **Next-Generation Digital Consumer—Electronics for the Mass Market**

Moderator:
Dale L. Ford, *Senior Industry Analyst, Semiconductor Application Markets Program, Semiconductors Worldwide, Dataquest*

Dwain Aidala, *Vice President and Division Manager of Personal Mobile Communications Division, Mitsubishi Wireless Communications Inc.*
David Taylor, *Director of Marketing, C-Cube Microsystems*
Andrea Cuomo, *Vice President, Headquarters Marketing and Strategic Accounts, SGS-Thomson Microelectronics*
Todd Oseth, *Vice President, Business Development, Information Technologies of America Division, Sony Microelectronics*

Agenda Continues

Track 2, Thursday The Emerging Fabless Market | **Crystal DEF**

1:30 - 2:15 pm **Lan Switches, Round Two**
Prabhat K. "P.K." Dubey, Ph.D., *President and CEO, MMC Networks Inc.*

2:15 - 3:00 pm **Embedded DRAM Technology: The Next Major Technology?**
Prakash Agarwal, *President, Cofounder, and CEO, NeoMagic Corporation*

3:00 - 3:30 pm ***Networking Break***

3:30 - 4:15 pm **Ultra High Speed DRAMs: Unlocking the Potential**
Fu-Chieh Hsu, Ph.D., *Chairman and CEO, MoSys Corporation*

4:15 - 5:00 pm **How to Win the Digital Communications Race**
Hatch Graham, *Senior Vice President and General Manager, Personal Communications Group, TCSI Corporation*

5:45 - 6:00 pm ***Dinner shuttle departs***
6:30 pm ***Reception and Dinner*** ________ **La Quinta Sculpture Gardens**

Agenda Continues

Agenda

Friday, October 25 General Session		Crystal Ballroom
7:00 - 8:00 am	**Continental Breakfast**	**Crystal Ballroom**
8:00 - 8:30 am	**DSP Chips: The Hottest IC Segment in Sight** Joseph Grenier, *Vice President and Director, Semiconductor Device and Applications Program, Semiconductors Worldwide, Dataquest*	
8:30 - 9:00 am	**Java versus Intel: And the Winners Are ...** Raj Parekh, *Chief Technology Officer, Vice President and General Manager, Embedded Products Group, Sun Microelectronics*	
9:00 - 9:45 am	**The New Paradigm for Online Services** Mark Walsh, *Senior Vice President, AOL Enterprises, America Online*	
9:45 - 10:15 am	**The Technology Investment Outlook for '96 and '97** Charles A. "Chip" Morris, *Managing Director and President of T. Rowe Price Science & Technology Fund, T. Rowe Price Associates*	
10:15 - 10:45 am	***Networking Break***	
10:45 - 11:15 am	**Asia/Pacific and China: Looking Ahead to 1997** C.D. Tam, *Senior Vice President and General Manager, Asia/Pacific Semiconductor Products Group, Motorola*	
11:15 - 11:45 am	**What Do the USERS of Technology Say?** Scott Winkler, *Vice President and Research Director, Platforms and Operations Software Technology, Gartner Group*	
11:45 - 1:30 pm	***Luncheon***	**The Grand Lawn**

Agenda Continues

Agenda

Track 1, Friday — **Challenges in Multimedia and DRAMs** — **Crystal ABC**

1:30 - 3:00 pm ***Panel:*** **Mainstream Multimedia: Which Long Distance Approach Will Win?**

Moderator:
Nathan Brookwood, *Principal Analyst, Semiconductor Application Markets Group, Semiconductors Worldwide, Dataquest*

Michael A. Aymar, *Vice President and General Manager, Desktop Products Group, Intel Corporation*
Wes Patterson, *President, Chromatic Research*
Prakash Agarwal, *President, Cofounder, and CEO, NeoMagic Corporation*
Carl Stork, *Director, Windows Platform, Microsoft*
Martin Reynolds, *Vice President, PC Technology Program, Computer Systems and Peripherals Worldwide, Dataquest*

3:00 - 3:30 pm ***Networking Break***

3:30 - 5:00 pm ***Panel:*** **Sync or Swim: The DRAM Bandwidth Challenge**

Moderator:
Jim Handy, *Semiconductor Memories Worldwide Program, Semiconductors Worldwide, Dataquest*

Neal Margulis, *Senior Vice President, Research and Technology, S3 Incorporated*
Avo Kanadjian, *Vice President of Marketing for Memory Products, Samsung*
Tom Dye, *Director, Austin Design Center, Chief 3-D Graphics Architect, Cirrus Logic*
Dr. Michael Nielsen, *Director of Engineering and Chief Engineer, Digital Media Systems, Silicon Graphics*

Track 2, Friday — **Meeting the Manufacturing Challenge** — **Crystal DEF**

1:30 - 3:00 pm ***Panel:*** **Packaging Performance Breakthroughs: Is the Industry Ready?**

Moderator:
Näder Pakdaman, *Principal Industry Analyst, Semiconductor Equipment, Manufacturing, and Materials Program, Semiconductors Worldwide, Dataquest*

William Beckenbaugh, Ph.D., *Vice President and Director of Sector Technology, Advanced Interconnect Systems Labs, Motorola Semiconductor Products Sector*
Barry Lieberman, Ph.D., *Engineering Manager, Test Tools Operation, Intel*
Michael J. Cadigan, *General Manager, Packaging Group, IBM*
R.P. St. Clair, *Engineering Manager, Systems Development Engineering, Memory Test Division, Teradyne Inc.*
Igor Khandros, *President and CEO, FormFactor Inc.*

3:00 - 3:30 pm ***Networking Break***

Agenda Continues

Agenda

3:30 - 5:00 pm ***Panel:* 1997/1998 Capacity Status: What Will Be Tight and What Won't?**

Moderator:
Clark J. Fuhs, *Director and Principal Analyst, Semiconductor Equipment, Manufacturing, and Materials Program, Semiconductors Worldwide, Dataquest*

Arthur W. Zafiropoulo, *Chairman and CEO, Ultratech Stepper Inc.*
John Luke, *President, TSMC-USA*
Eberhard Klasse, *Vice President Sales and Marketing, Wacker Siltronic Corporation*
Evert Wolsheimer, *Vice President, Product Technology, Xilinx*

Be sure to stop by the Showcase Displays during the conference.

5:00 pm **Conference Adjourns**

The 22nd Annual

Semiconductor Conference

Indian Wells, California

October 24- 25, 1996

CHALLENGES IN MULTIMEDIA AND DRAMs

MEETING THE MANUFACTURING CHALLENGE

DISCLAIMER

PLEASE NOTE: The contents of this document are based on a verbatim transcription of tapes from the conference proceedings including speeches, audience participation, and panel discussions. A best effort was made to ensure the accuracy of names, acronyms, slang words, and companies prior to the printing of this document, which has been edited for readability.

Welcome

Gene Norrett

Corporate Vice President and Director,
Semiconductors Worldwide,
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. Before that 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 B.S. degree in mathematics from Temple University and an M.S. degree in applied statistics from Villanova University.

Session # 1: Welcome

Gene Norrett
Corporate Vice President and Director, Semiconductors Worldwide, Dataquest

My name is Gene Norrett, and I am Vice President and Director of the Semiconductor Group at Dataquest Gartner. On behalf of our 60 semiconductor analysts throughout the globe as well as our excellent worldwide global events staff here supporting us, I would like to welcome you to our 22nd annual conference focusing on semiconductors and the engines that drive them.

The theme of our conference today is: Semiconductors: $300 Billion Dollar Challenge: Let the Games Begin!! We chose this theme, first of all, because this is the year of the centennial anniversary of the modern Olympics, and secondly, it is also reminiscent of the challenges facing the manufacturers in the semiconductor industries, and also their similarities with the Olympic athletes that participated in the Olympic games back in August. For the champion gymnasts at the games good balance can make the difference between a gold medal or going home empty-handed.

On the other hand, the manufacturers in the explosive semiconductor industry, finding a delicate balance between risk and reward can be the difference between being a leader, or being a follower, having profits or losses, or being viewed as a valued partner to your customers, as opposed to having your hated competitor get the accolades from that same customer.

Your tasks are to avoid the many pitfalls on the road to the 300 billion dollar semiconductor market, that we are projecting for the year 2000. Your rewards are seizing the unprecedented opportunities for increasing growth, profitability, and, of course, market share.

As we start the fourth quarter, we are taking the first step on this road to the 300 billion dollar market that we are projecting in the year 2000, which we think will be the peak of the next business cycle. In our panel later this morning, you are going to hear Gary Grandbois, a Vice President in our group, discuss the factors and assumptions associated with this forecast. New media like the Internet, new markets such as wireless, networking, and digital consumer, are creating new channels for business, entertainment, education, and commerce.

Just as the Olympic athlete has several conditioning and training programs, we think that this conference will be one part – a very important part – of your training program, where you can learn how to balance the market opportunities against the severe risks in this market. This is what this world-class conference is all about.

This year's conference will examine the theme: "$300 Billion Dollar Challenge: Let the Games Begin!!" from a wide variety of perspectives and presentation styles. You will hear over the course of this conference, 14 industry executives give their perspectives about the dramatic changes in their segments of this explosive information technology industry, such as on-line and wireless services, computer and communications environments, new operating systems and applications software in regional and financial markets, and finally, in the semiconductor industry.

We will also have eight "fireside" panel discussions featuring 40 of the world's smartest guys, who will tackle the issues associated with the various industries, technologies, and markets that they all participate in. The subjects will range from the 'battle for the local loop' to new strategies for increasing performance in multimedia and dynamic RAMs, and on to meeting the manufacturing challenges over the next exciting four years.

We hope that this conference will exceed your expectations in terms of venue, quality of delivery, networking, and of course, deal-making. That is for some of you, one of the major reasons for attending this conference, and some of those deals can involve getting yourself a job. Happy negotiating.

Before I go on to the housekeeping items, I want to mention several last-minute changes in our program. First of all, Bill Beckenbaugh, Motorola Vice President for Advanced Interconnect Systems, and a member of our packaging panel, was taken sick recently, and he has been replaced by Reed Bowlby, who works in that department, and is manager of their ceramic business development. Mike Cadigan, IBM General Manager for their Packaging Group, is being replaced by Ray Bryant, who is on Mike's staff. Mike Aymar, Intel Vice President for Desktop Systems, had to cancel at the last minute, as a result of Intel's most recent stance on not having their executives participate in functions that involve forward-looking items.

Over the past 22 years of holding this conference, we have had a stock market crash in '87, the Loma Prieda earthquake in 1989, and now the worst calamity of all the calamities, Proposition 211. If you are not familiar with that, talk to one of our analysts. We are getting bombarded by a lot of calls. In fact, Wilf Corrigan, President & CIO of LSI Logic Corp. said last night, "It looks like maybe the tide is turning for people in Dataquest's business." He believes that we will end up doing all the talking about the future and the companies will not.

We are going to be asking you to write down your questions on the paper provided at the tables, or use the paper in your binder. Then during the presentations, or panels, we are going to have staff members walk around with a sign saying "Questions" and we ask you to hand your questions to the staff member. Then what we will do is hand them to the moderator of the panel or the speaker, and try to get those questions answered.

Welcome

Manuel A. Fernandez

Chairman of the Board,
President and Chief Executive Officer,
Gartner Group

Mr. Fernandez has served as 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, he was president and chief executive officer of Dataquest Incorporated, an information services company and a subsidiary of The Dun & Bradstreet Corporation. Before joining Dataquest, Mr. Fernandez was president and chief executive officer of Gavilan Computer Corporation, a laptop computer manufacturer, and Zilog Inc., a semiconductor manufacturing company.

Mr. Fernandez serves on the board of directors of Individual Inc., SACIA (The Business Council of Southwestern Connecticut), and Norwalk Community Technical College and has previously served on the boards of A.C. Nielsen Co., EMU Systems, Macmillan Inc., and ViewTech Inc.

Mr. Fernandez holds a bachelor's degree in electrical engineering from the University of Florida. He completed postgraduate work in solid state engineering at University of Florida and in business administration at the Florida Institute of Technology.

Session #2: President's Welcome

Manuel A. Fernandez
Chairman of the Board, President and CEO, Gartner Group

GENE NORRETT: I would like to start this year's conference off with a very good friend of Dataquest, none other than Manny Fernandez, President of the Board, President and Chief Executive Officer for Gartner Group. Mr. Fernandez has served as Chairman of the Board of Gartner Group since April of 1995, Chief Executive Officer since April of 1991, and President and Director since January of 1991. Before joining the Gartner Group, he was President and CEO at Dataquest, as many of you know. Before joining Dataquest, Manny was President and Chief Executive Officer for Gavilan Computer Corporation, a laptop computer company, and of course, President & CEO of Cilog. Manny has a bachelor's degree in electrical engineering from the University of Florida. He completed graduate work in solid-state engineering at the University of Florida, and in business administration at Florida Institute of Technology.

MANNY FERNANDEZ: It is truly great to be back. It was six years ago, the last time I was at a Dataquest Semiconductor Conference and as I was telling a couple of folks before coming in, it has been like I have been in a time capsule. Six years ago, memory pricing was going down, the DRAM forecast was also going down. There were also issues with Japan and the U.S.

Now it is like somebody woke me up six years later and here I am and everything looks the same. All of the friends and all of the stories are the same, so it is unbelievable. But anyway, it really is great to be here.

What I would like to do over the next 5 or 10 minutes, besides welcoming you and thanking you for being here, and for your loyal support of Gartner Group and Dataquest, is to give you a very quick picture. I will not talk much about last year, but I will make a few comments. I would like to look a little bit ahead, and give you a picture of some of the people. We in the semiconductor business have always been rather focused on that primary customer that we have had. Today I would like to give you a message from a different customer, from your customers' customers, and how they are going to drive your future and my future.

First of all, it is clear that 1996 so far has been an incredible year. Sure, the PC business continues to grow, 19% up from a year ago. Clearly Intel is sure doing incredibly well and their dominance of the microprocessor business continues.

The story of the year has obviously been the debacle of the DRAM business. Not only has this cycle experienced the largest and year-to-year swing from being up 80% in 1995 to being down 40% for this year, but it also has experienced the fastest decline from peak to valley that we have ever seen in this industry. The previous fastest decline was in 1984/85 when it took 14 months to go from peak to valley. When we contrast the previous decline to our current decline we see that it has taken only nine months to go from peak to valley. This has been a very difficult period for both manufacturers and their customers.

While the Dataquest folks did an excellent job forecasting the inflection point where this was going to happen, I think that everyone missed the severity of this particular inflection. Clearly we are paying for the excess capacity now, and that will continue in the DRAM business. I do not think that there is a foreseeable short-term change or shift. But enough about the forecast on that end because the Dataquest folks are going to do their normal tremendous job all this morning going through forecast by forecast on our view of what is happening in the industry as a whole.

I would like to shift and give you the user view. Scott Winkler is going to take a look at the Chief Information Officer and he is going to give you some of the things that they see, from their point of view.

But let me just give you a bit about the issues that exist. In the information technology world as a whole, expenditures are going to increase from about 650 billion dollars to about 1.3 trillion dollars by the year 2000 – a huge increase. It is not only that the real dollars are going to increase, but expenditure of information technology as a percent of revenues of the world corporate companies is going to increase from about 5.8% to 9.5%.

Now, it is very clear that all of a sudden the IT world, that used to be relegated to be just a staff position, is now in the middle of the center stage of the world, and all of a sudden there are huge questions that are coming from the top, from the CEO, from the Board of Directors, "How are you going to be spending this money? and how am I going to be successful? and how am I going to change the top line by the information technology revolution, not to be able just to take a hit in the bottom line by that extra 4% that is going to be spent on information technology?" Clearly, information technology is being moved from a staff role to basically a market share role.

In addition to that particular problem, the ultimate user of yours is having two other problems. Problem number one is that the life cycle of our business products converts itself into the life cycle of information technology end user products. We have seen the average life cycle of 14 to 16 years drop to the 5 to 6 year level and quickly moving into the 18 month level. At the same time, 1996 has been marked as an 80% increase on the number of new products introduced to the end user.

So think of this end user, the individual that eventually buys all of the products that you make, being bombarded from the top by this budget pressure, bombarded from the bottom with the life cycle, and the number of products available. What is happening is that there is an incredible confusion at the CIO and the CEO level that has created an incredible instability in the organization side of the business. That instability in the organization has had direct costs because of demand in the last 9 months.

We have 7,000 different enterprises that we do business with, and 90,000 clients on a worldwide basis, so we think we have a good statistical sample about anybody that buys information technology. If you ask them, "What are the three or four key issues that you are dealing with?" they will tell you that technology obsolescence is number one. They are now figuring that where software obsolescence used to be around 12 to 15 years, we now think that 70% of all products that are presently under development, or being purchased in package form, will be obsolete by the year 2000. That is an incredible impact to the user world.

The second thing they will tell you is that there is a huge movement to mobility. To give you some numbers, the average budget that is being spent by companies in the United States, for mobility or work outside the office space, is about 4% of their total budget. We are now forecasting that by the year 2000 that number will be 30%. Thirty percent of all the IT budgets will be for outside the office environment technology – huge change.

Then there is always what I would call the "killer virus," the actual things that may happen that can blow everybody away. Some one-time killer problems, and obviously the year 2000 is one of those. The impact in the technology business is going to be about 600 billion dollars of expenditures to be able to make the conversion. A common European currency conversion is going to be somewhere in the neighborhood of about 250 billion dollars. Now, what will all these things lead to? Well, actually, there is one more. And that is, I believe, total negligence in our user environment today; the infrastructure is being absolutely not looked after. Because the infrastructure that is going to be required to have the out of the office work, and the ability to put in these new products, is not being taken care of, there is going to be again the need, the requirement, to be able to invest into that whole infrastructure a significant amount of money.

So what do all these things lead to? I believe that all this leads directly into a very healthy amount of demand for products from 1998 through the year 2000. I really believe that by the time we enter the cycle of '98, '99 and the year 2000, you are going to see the biggest boom that we have ever seen, in the information technology world from the demand. We will be out of silicon again and all of those things. So the question is between now and then; I have left that to Dataquest to tell you what is going to happen over the next 12 months.

I think it is important to have one caution, one red flag, and that is with all this demand, and all the one-time killer problems that may appear, and the requirements to invest a significant amount of money that has to be invested, not all of the money is going to be available. So there are going to be incredible trade-offs being made. This is going to reflect dramatically upon you and your clients, the people that you directly sell to.

I believe there are going to be companies that will have tremendous products, but will not make it through this next four years because they are going to get caught in a trade-off problem. There is not enough money to do the year 2000 conversion, common currency, mobility, and so on. There are going to be some issues. I urge you to look at that, and make sure that the partnerships you have with your clients are ones that are based appropriately on what the end user product eventually is, and what the demand really is going to be.

In closing, I just want to thank you on behalf of the 2000 people at Gartner and Dataquest. We have dedicated ourselves to continue to bring the best in this business at a fast pace. This year alone we have added over 100 analysts to our business. We will add another 100 analysts next year to the Dataquest-Gartner family, and the serving of 90,000 clients on a worldwide basis is a critical element for us.

GENE: That was very inspiring. I am shocked at some of those statistics. I know things are changing fast in the industry, but when you reflect it on what the end user of all this technology has to concern himself or herself with, it is pretty awesome.

The State of the Economy: Looking Ahead

Donald H. Straszheim, Ph.D.

First Vice President and Chief Economist,
Securities Research and Economics,
Merrill Lynch

As chief economist and primary economic spokesman for Merrill Lynch since February 1985, Dr. Straszheim is responsible for Merrill Lynch's global economic research and analysis effort, with economists located in London, Frankfurt, Paris, Tokyo, Hong Kong, Singapore, and Sydney, as well as in New York City.

Dr. Straszheim was voted as a member of the *Institutional Investor* All America research team in each of the last nine years, he is a frequent writer and speaker on the economy and financial markets, and is a regular guest on CNN. He is also a regular contributor to *Credit Week* magazine and a member of Public Securities Association Economic Advisory Council.

Before joining Merrill Lynch, Dr. Straszheim had senior economics positions with Wharton Econometric Forecasting Associates Inc. in Philadelphia, the economic think tank associated with the Wharton School of Finance, University of Pennsylvania; Weyerhaeuser Company, the forest products firm in Tacoma, Washington; and Investors Diversified Services Corporation, the money management firm in Minneapolis, Minnesota.

Dr. Straszheim is a graduate of Purdue University. He received a B.S. degree in 1963, an M.S. degree in 1967, and a Ph.D. in 1971.

Session #3: The State of the Economy: Looking Ahead

Donald H. Straszheim
First Vice President and Chief Economist, Securities Research Economics, Merrill Lynch

GENE NORRETT: As is traditional at this conference, we like to have a noted economist give us his projection of the outlook for the economies of the world, and we are very pleased to have with us Don Straszheim. Don is First Vice President and Chief Economist, Securities Research and Economics Department at Merrill Lynch. As the chief economist and primary economist spokesman for Merrill since February of 1985, Don is responsible for their global economic research and analysis position, with the economists supporting him located around the globe.

DON STRASZHEIM: Good morning. The economy is not really as dull as people make it out to be. In fact, it is kind of an exciting place because so many things are changing and what is really driving that change is represented by all of you in this room.

First, let me describe our good looking economy because the current economic circumstance in the United States is probably as good as it has been in maybe a quarter century. Second, I want to talk about our changing economy. I want to pick out one aspect of it, to get you to think about all the things that are changing in the environment around you, and that aspect is the demographics and the aging of the baby boomers. The third thing I want to talk about is the changing global economy, and give you some highlights of what we in the research department at Merrill Lynch see going on in Europe, in developing Asia, in Latin America, and elsewhere.

Last, I want to talk about technology, but I give a lot of speeches for Merrill, and when you do that you learn one thing very quickly. When everybody else in the audience knows more about the topic than you do, you should be very careful. That is the circumstance I am in, on technology today. I will give you a glimpse of technology as a macro economist sees it and how it is affecting our day-to-day lives.

First, our good looking economy. We are, I think, somewhere near equilibrium in a broader context. We have had ups and downs in our growth rate, the second quarter growth rate was 4.7% and the markets got all uptight, "Oh, this is too strong, and now it looks like growth in the third and fourth quarters is going to be back in the 1.5 to 2% range which is a little bit below our longer run equilibrium level." But the recovery now, since March of '91, is 67 months old, and we had an eight-year long recovery with a soft landing embedded in the middle in the 1960s, and an eight-year long recovery with a soft landing embedded in the middle in the 1980s, and it looks like we are doing it again.

We are not all that far from an equilibrium. That does not mean we will not have another recession. I am sure we will – I just do not know when – and it will come about as a result of bad luck or bad policy or probably some combination thereof, and when that happens the economy will probably get too hot. Inflation will start to rise, interest rates will rise, and the Federal Reserve will say, "Wait a minute, this is a dead end street, we cannot have this." They will start to tighten monetary policy, choke it off, we will have a little garden variety recession, and then we will start the process over again.

That will happen one of these days, but it is not as far out as our headlights can see. We are not all that far from equilibrium in terms of growth. The unemployment rate is now in the low 5% range, and most economists would say that is near what we would describe as full employment. Most people who want a job have a job, and in fact, we are already beginning to see some evidence of wage inflation. That is, we are close to getting this economy a little bit overheated in terms of the labor market. But in terms of labor we are not far

Graphic materials for this speaker were not available at the time of publication.

from equilibrium.

The inflation rate has been stuck at about 3% plus or minus now for 14 years. We had 10% or 12% inflation in '81-'82, called a halt to that, brought the inflation rate down, and it has been roughly in this range now for over a decade. It goes up and down a little bit – at about 6% or so in 1989/1990, brought it back down, went back up to 4%, now it is down to maybe 2.5% or 3%, something like that. There are a lot of inflation fears but there is no inflation, so that side looks pretty good.

Corporate earnings: 1993, 1994 and 1995 we had corporate earnings for the S&P 500 on an operating basis up about 12% or 15% each year. In 1996 and '97, the earnings growth is going to be slower and we are looking for earnings to be up maybe 3% or 4% or 5%, but the estimates from the Wall Street analysts – if you add those all up – they come up with numbers that are still 10% or 12% or 15% percent. There is probably no security analyst in the room right now so they are not here to defend themselves but they are going to be wrong. Earnings are not going to be that rapid, and you know what happens to stocks when earnings come out that are disappointing. The stocks just get hammered and it looks like that kind of environment. It is going to be a very good investor who can avoid all of these down-drafts that are going to come in different places in equities because of the too high earnings.

In any case, corporate America is dramatically more competitive and you can see that by the good earnings performance, and interest rates – both short term and long term interest rates – are not all that far from what one would regard as equilibrium in a longer term context. If you take the nominal inflation rate, and pull out a normal inflationary premium, and you look at the so-called real rate, they are not far from equilibrium.

So that looks pretty good as well and I mentioned our competitiveness – we are dramatically more competitive than we were a decade ago. I never thought I would see American cars, or any car produced in America, again being a reasonable buy. After a lot of cost cutting and hard work we have made a lot of progress on getting our budget deficit under control. Again, I suspect most people in this room, including me, think we ought to do more, and I think we will do more, and we could probably argue again about who gets the credit for that.

We have made more progress than any of our competitor nations, and the trade deficit is not nearly the issue that it was a few years ago. In a whole host of areas our economy looks fairly near equilibrium, and our view at Merrill Lynch is that '96 is going to be a decent year in terms of the economy. Growth will be in the 2% to 2.5% range and the same with '97, and inflation is not about to rise a great deal. It is up a bit but not really a lot.

The concern regarding inflation, is wages, and the reason is that wages and salaries compensation is between two-thirds and three-fourths of national income in every economy in the world, rich or poor, including ours. Wage rates are now rising about a percentage point faster than they were in '92 and '93, and in the long run, wages compensation and inflation go hand in hand. It is important to keep watching the wage numbers and employment statistics, for guides as to how the economy is going to do in the next six, twelve, eighteen months, and in terms of what interest rate response we might see from the Federal Reserve. Our view right now is that the Federal Reserve will leave short term interest rates unchanged at the next few policy meetings. The next key meeting is November 13th and there is another one December 15th and long term rates, the long bond yield, the thirty-year treasury bond, what we call the long bonds, are yielding now about 685 or so.

If the inflation rate stays under control and the economy grows only weakly in the next few quarters, there is a chance that bonds will rally a bit further and we will see yields back down to maybe 6.5%. Bonds are going to essentially be in a trading range until there is an event in the economy, and the event will either be the economy gets a lot stronger, in which case rates will rise a great deal, or it gets a lot weaker, or it just simply gets a lot weaker first, and in either case you would see bond yields respond in the normal expected fashion.

This is not a bad looking economy, it is a good looking economy, and has been for some time. Let me make one other comment before I talk about some of these changes on the demographics. About Washington – the President is a parade leader, I do not think he is a parade organizer. The parade in the last decade has been a fiscal conservatism parade, and the parade for the next decade is going to be a fiscal conservatism parade because the population is aging and becoming more conservative in many ways, and that extends to fiscal policy. If that is where the country is going, you can be sure that the White House, and the House, and the Senate, are going to follow in the same direction. So we are not as concerned about this issue as some people are.

Now let me talk about our changing economy, the demographics. The baby boom from 1946 to '64 was so important because it was preceded and succeeded by a baby bust in the depression years and World War II, in the Vietnam War, and the women's movement. So we have this big hump in the population that is getting one year older every year. Marry that one piece of empirical information with one other observation we all make – that young people behave differently than middle-aged people, and middle-aged people behave differently than older people.

This is not economics, this is arithmetic. If you have changing behavior that is predictable at different ages, and a changing share of people of different ages, you are going to have changes in the overall economy, and for the last 30 years you had all these baby boomers trying to accumulate consumer durables – a great consumer boom – and they were buying their TV sets and their VCRs and their sofas and their tennis rackets – it is over. Now they are a little older, now they are trying to accumulate their mutual funds – I am not singling out mutual funds, I am talking about the shift from spending towards savings. This phenomenon is going to last about a generation because the baby boom lasted about a generation, and it is going to be a big deal. It raises all kinds of questions. You wonder about the retail capacity we have in this country. Who is going to use all these little stores in the strip malls? There are a lot of companies geared up for a future that may never arrive.

So let me just mention a couple of specifics here. Think of the number of forty-year-olds in this country. From 1975 to '85 the number of forty-year-olds grew by a million a year, and then '85 to '95, a million a year, so if you were selling tennis rackets, "What a great tennis racket salesmen I am, I am selling all these tennis rackets." Well, the next twenty years, your market is not going to grow by a million a year, it is not even going to be flat, it is going to be down. So if you are geared up for twenty more years of boom like that, you are Chapter 11.

Now conversely, take the sixty-year-olds, and these are the same people, by the way, that were forty. The sixty-year-old market has been stagnant for twenty years, "What a crappy market this is, I cannot sell anything." So now you are geared down for twenty more years of stagnation and your market is just going to blossom, it is going to grow by a million a year for twenty straight years. So if you are geared down, you are going to leave all the growth on the table for the guy across the street who is geared up, who sees it coming.

This is an extraordinarily powerful phenomenon that is going to last, and you can see this through all the demographics. Think about the products that you sell, and the markets you are in. You want to be selling downstream, swimming downstream. It is a lot easier to play downwind, and that is the way it is with all these markets. So the changing demographics are extraordinarily important.

Now let me talk about the changing economies around the world, and I want to start this by talking about inflation. We live in increasingly globalized markets, so inflation and economic activity, is not just a domestic phenomenon anymore. It is increasingly a global phenomenon now. There has been extraordinary improvement in the inflation rate around the world in the last few years. In '97 the inflation rate is going to be lower worldwide, according to our forecasts, than at any time in the last quarter century.

Take Europe as an example. The highest inflation rate we can see in Europe in '97 is about 7% for Greece. Every other European country in '97 is going to have inflation that is under 4% – one or two or three percent rates. In the ten major economies in developing Asia that we follow, eight of them are going to have lower inflation rates in '97 than in '95. The so-called transition economies – Russia and the formerly centrally-planned economies – inflation is down dramatically the last few years and going to go down more, we think. There is great progress on inflation in Latin American; Japan's inflation rate is right around zero. We do not have an inflation problem.

Let me list eight factors that are causing this. First of all, central bankers around the world are increasingly committed to inflation control. They have concluded that just letting the inflation rate slip up a half a percentage point, or a percentage point, every year does not really buy you any more jobs, and any more output. Second is globalization. We are all looking for the lowest cost sources of raw materials, and intermediate products, and labor, and so forth, and the companies that have the best growth prospects in general are those that are looking at this global market. We are very rapidly getting to the point where you no longer think of General Motors as the American car company, you think of it as General Motors, the car company, and it is not Toyota the Japanese car company, it is Toyota, the car company, and that globalization process is helping to hold down inflation because it is increasing competitiveness.

Technology is a great plus for inflation control around the world. Fiscal restraint – we have made progress. Other countries have not made as much progress as we have but they have made a lot as well, and there is a pretty good commitment, even in the developing economies. I just spent last week in India and they are talking about fiscal restraint as well as the emerged economies. Cost-cutting and restructuring is linked back to globalization, linked back to technology, we have seen it, it is everywhere. Finally, the cost-cutting and restructuring that we have been doing, really since the end of the '81-'82 recession, is beginning to take hold in Europe, and slowly in Japan as well, and even in many of the emerging markets.

The collapse of communism. There are maybe three billion people around the world who used to work in economies that were command and controlled economies. Now they are making this transition more or less rapidly to market systems, and that is helping the inflation rate in a lot of places. There is excess capacity all over and not much pricing power in many areas.

Last is privatization. In many economies around the world, not just the centrally planned economies, but even in the mixed economies and market economies, there were operations that were public-sector enterprises now becoming private-sector operations. So there is great inflation progress. Do not worry about a resurgence in inflation any time in the near future because again these are global markets, and you cannot go it alone. If you try to go it alone, you try to pump up your economy, and the cost of it is higher inflation, the market will just leave you. Increasingly, countries around the world understand that we are all in this together, and it is an extraordinarily powerful and positive thing.

Now let me say something about a few of the economies around the world and the regions. Japan: '92, '93, '94 and '95, the four worst years back to back in Japan in perhaps a century; growth between zero and 1%. '96 and '97 had better growth, maybe 3% or thereabouts, but Japan's heyday is over. Japan is an economy of the past, not an economy of the future, and the reason is that the characteristics that Japan had, that commended it and brought it this wonderful growth record in the '50s and the '60s, increasingly now characterize many of Japan's neighbors. Japan is now a relatively high-cost place in the world to produce rather than a low-cost place to produce.

The markets that look best of all are Japan's neighbors: Korea, Taiwan, Hong Kong, Singapore, the four original tigers, and now of course it is spreading elsewhere. India has a great looking economy. I am sure many of you have operations in India now, if you do not you are probably looking at it, and if you are not, you should. They made major economic reforms in 1991 and after the election upset earlier this year, there are questions about whether the reform process is going to continue or not, with a coalition government.

I was a little nervous when I went over there. I am confident now that the reform process is going to continue. There are 250 million people in the middle class in India – that is an economy entering the body of emerging nations around the world and we like that economy. Their population will be larger than China's by about the year 2012, or thereabouts. We like China. It has been a great growth story, it is going to continue to be a great growth story. There are all kinds of questions, intellectual property rights, and all the rest, but they are gradually making this transformation from a command and controlled economy, to an economy that has at least more market principals involved in it.

Do not worry about Hong Kong in '97. On July 1, 1997, control of Hong Kong reverts back to the mainland. If on July 1, 1998 (one year after the transition), the business people and the investors of the world conclude that Hong Kong was a great place while the Brits ran it but now that the mainland has taken it over they have turned it sour, what will happen? There will be no more capital flow to Hong Kong, and there will be no more capital flow to China. It will go to Indonesia, or India, or Turkey, or Brazil. They need foreign capital, they need foreign technology, and they need foreign talent to achieve their longer run economic growth objectives. So we are not fearful of Hong Kong's transition at all. It is going to be different, it is going to change from being THE gateway to being A gateway, but it is not a place that you should be worried about.

Latin America is doing much better after having a couple of bad periods. Mexico's growth rate this year is going to be 4 or 5% (probably next year as well), after being down 5% in 1995. Argentina and Brazil are looking pretty good, especially Brazil. Chile, along with New Zealand (I realize that is not in Latin America), are the best little success stories around the globe. They are little stories, but they were both basket-case economies in the early '80s and made some dramatic economic, fiscal, social, and structural reforms, and now they have all the economic statistics going their way.

Last is Europe. We think Europe is going to be the slowest growing economy in the world for the rest of the century and probably beyond. It is a high-cost place to do business. Their managements are more resistant to change than managements here. Their governments are more resistant to change than here. Their labor unions are more powerful than here, and their societies are a little bit more rigid and ossified. They are now making progress, and there is this rush to mastery, this monetary union and great progress really is being made, and great political momentum for it, and it is going to happen. But that does not change the fact that it remains a high-cost place to do business. So we see growth there maybe at the two percent range for the rest of the century – not as high as here, not as high as Japan, and half or less of what it is going to be in Latin America, and a third or less of what it is going to be in the economies in developing Asia.

Let me now make a few comments about technology, and I will relate this to some things that you see around you, without being a technology expert. The reason you are seeing an enormous banking industry consolidation is the computers and staff of one bank can handle the transactions of two banks and it is nowhere near over. Fifteen thousand banks in the country a decade ago, and now 10,000, headed to some very low numbers.

The insurance industry has not yet really started to consolidate but technology is the key to that one as well. The technology intensity in insurance is similar to the technology intensity (if I can use that term), in commercial banking.

The edge cities, like Naperville, Illinois, are threatening the downtowns of the world, because technology allows people to do business in different places, enormous impact of technology on economic dispersion. Places like Telluride and Naples, Florida, Pinehurst, North Carolina, and Aspen; people are locating facilities here because they do not have to locate in the places that they used to have to locate.

Technology is flattening organizations. You see this, you are driving it. In the old days you had this pyramid, and the guy in the middle would get this information fed up to him from these people below, and they would massage it, then they would feed it up the line, and they would keep doing that. Now there is an infinite web of information that is available everywhere in the company if you want to make it available, so what do you need all these people for? This flattening is only beginning. The phenomenon of the rich getting richer in this country – what is really driving that is skills. People with the requisite job skills are having their wages bid up and those without are falling off the bottom. This information technology revolution is just in its infancy; this phenomenon is going to continue. It is going to last for a long time and probably be even more aggravated.

Two last points – companies that fall behind in these information technologies that you are driving, are going to be Chapter 11. A retail firm that does not have point-of-sale data capture, has no chance of surviving. A manufacturing firm that does not have goods, or just-in-time inventory procedures, has no chance. A distribution wholesale firm that does not know exactly where everything is at all times, has no chance in this world, and similarly with individuals. Individuals who do not keep up with these technologies have no chance, because the competition is going to keep up with these technologies, and increase productivity so rapidly that people who do not keep up will be unemployed.

The single characteristic that is going to be more important than any other in distinguishing the winning companies from the losing companies is, the winning companies are those that understand the power of these new technologies. Companies that can figure out how to adapt these new technologies to their own particular corporate purposes, and proliferate the technologies through the organization, cut their costs and beat up on the competition. Understand, adapt and proliferate, and if you do not do that, you are going to get left behind.

Q: The people in the audience would like to know what your forecast is for information technologies spending, as a percent of GDP trends, if not the absolute level?

A: Information technology spending really depends on how you define it, and everybody probably has different definitions, so a number is not going to be very useful but I will give you a couple of simple little anecdotal statistics. It is going to continue to take an ever greater share. In 1970, the share of capital spending that was composed of the computers and information processing was about seven percent of total capital spending. Last year under the same definition, it was about 45%. Seven percent to 45% in the last 25 years and that is a very narrow definition of high-tech in capital spending, because a lot of low-tech capital spending has a lot of high-tech content. So this increase is much greater than that, and I see no reason to think that is going to slow any time in the foreseeable future.

One other point about the capital spending and technology is 25 years ago, most all of technologies was really capital spending. Increasingly now, a greater share is consumer purchases. If more of technology spending is a consumer purchase, it means the technology cycles will be smaller in the future than they have been in the past.

Q: Could you comment on what a company that is now doing business in Japan should do, given your outlook for the lower growth rate in Japan?

A: They should just do the best they can. It is not going to be the worst economy. The Japanese still have a great labor force, good job skills, good work ethic, they have the tools and the only problem is really a cost problem. It has become a relatively high-cost place to do business. In the '50s and '60s, up to the first oil shock in October of '73, we got used to 8% and 10% growth rates in Japan, year after year, and that is over. Since then the growth rate has been about 4%, 4.5%, and the last four years have been terrible. They are doing better now, they have their discount rate at a half a percent, the lowest of any central bank since the 1930s, trying to lift the economy. It will work, but the growth rate is going to be more like

3%, not the big numbers in the past. So I guess there's a non-answer of 'do your best.'

Q: On this Japanese question, given the fact that Japan has been slower to adopt technology in their major manufacturing industries, do you see a major uptake in the usage of the technology tools that are pervasive now in the United States and Europe?

A: We lead the world in proliferating these various technologies, and Japan is way behind. There is part of an explanation of what has happened to Japan, and it is also a great opportunity. I see Japan going after these technologies in a much more aggressive fashion than they have in the past, but it again will take a long time.

Q: Concerning Asia Pacific – what is going on with the current export slowdown. Can you discuss that for a few minutes and give us your outlook of when that slowdown may end?

A: Many of Asia's markets have not been very good. We just talked about Japan and that has become a fairly important market for the rest of developing Asia. Europe has not been very good either, so those two key markets have been weak, and their export opportunities have not been what they were before. There is also an inroad into some of these Asian markets that we are making, so all three of those factors – the first two are probably the key ones – markets are weak, exports do not do as well.

Keynote:
Olympic Long Jump: From $150B to $300B in Five Years

Wilfred J. Corrigan

Chairman and Chief Executive Officer,
LSI Logic Corporation

Mr. Corrigan is chairman and chief executive officer of LSI Logic Corporation.

Before founding LSI Logic in 1981, he was president, chairman, and chief executive officer of Fairchild Camera and Instrument Corporation in Mountain View, California. He joined Fairchild in August 1966 and held a series of management positions before becoming president and CEO in July 1974. The position of chairman was added in May 1977. Before that, he was director of Transistor Operations at Motorola Inc.'s Semiconductor Products Division in Phoenix, Arizona.

Mr. Corrigan is a director and postchairman of the Semiconductor Industry Association. He is a member of the Board of Directors of Silicon Power Corporation and LucasArts Entertainment Company.

Mr. Corrigan graduated from the Imperial College of Science, London, England, with a B.Sc. degree in chemical engineering. In addition, he holds an honorary degree, doctors of laws, from the University of Calgary, Calgary, Canada.

LSI Logic Corporation (NYSE:LSI), The System on a Chip Company, is a leading supplier of custom high-performance semiconductors, with operations worldwide. The company enables customers to build complete systems on a single chip with its CoreWare® design program, thereby increasing performance, lowering system costs, and accelerating time to market. LSI Logic develops application-optimized products in partnership with trendsetting customers, and operates leading-edge, high-volume manufacturing facilities to produce submicron chips. The company maintains a high level of quality, as demonstrated by its ISO 9000 certifications. LSI Logic is headquartered at 1551 McCarthy Blvd., Milpitas, California 95035, (408) 433-8000, http://www.lsilogic.com.

Session #4: Olympic Long Jump: From $150B to $300B in Five Years

Wilfred J. Corrigan
Chairman and CEO, LSI Logic Corporation

GENE NORRETT: Our next speaker is certainly no stranger to any semiconductor conference worldwide. Wilf Corrigan, Chairman and Chief Executive Officer of LSI Logic, has happily agreed to give us his outlook on the next four years and what it is going to take to get this industry to the three hundred billion dollar market size that we are projecting.

Before founding LSI Logic in 1981, he was President and Chairman and Chief Executive Officer of Fairchild Hammer & Instrument, and, of course, in Mountain View, I think most people know that. He held a series of management positions there at Fairchild and reached the level of CEO and President. Before that he was Director of Transistor Operations at Motorola Semiconductor Products Division in Phoenix, Arizona.

Mr. Corrigan is Director and Past Chairman of the Semiconductor Industry Association. He is a member of the Board of Directors of Silicon Power Corporation and Lucas Arts Entertainment. Wilf graduated from the Imperial College of Science in London, England, with a Bachelor of Science degree in Chemical Engineering. In addition, he holds an Honorary Doctors of Law degree from the University of Calgary.

WILF CORRIGAN: I was asked to talk about the Olympic long jump that we seem to have ahead of us, and I think the Olympic long jump is a good metaphor for what the semiconductor industry will do in the next five years. One year ago, 300 billion dollars in five years looked easy. Several years of 30% growth created an illusion that perhaps the semiconductor down cycles were behind us. One down year later, over-pessimism has set in.

Now let's face it, this is a cyclical industry. It always has been, it probably always will be. But it is cyclical with a high secular growth rate – something in excess of 15% plus. In the long view, the inventory correction of '96 will be a minor blip on the long-term growth of our industry.

I first became aware of the cyclicality of this industry in '61, but if somebody would have used the word 'cyclical' I would have thought they were talking about a bicycle. The technologies were coming on so quickly, and yields would bounce from 1% to 30% to 40% in very short timeframes. But in the '60s, we had almost an annual circle, and companies went in and out of this business so fast that it was just incredible. By the late '60s, we settled into something like a four to five year cycle. What is important is the amplitude of that cycle.

Now today, the amplitude is much less, and companies are better able to handle it. They are organized differently. The global nature of the business, just like investing in stocks globally, tends to cushion any individual country's cycle, in the same way the global nature of the business cushions the overall cycle, because the countries do not operate in sync at this time. I am talking about end demand now. This current slowdown has been mainly a U.S. phenomenon. As the U.S. has slowed down in its demand for semiconductors, it has had effects as far away as Japan and Taiwan. But in reality, the end demand, with our Japanese customers, our European customers, really has not changed much because the U.S. cycles in its management of inventory. The Japanese always manage their inventory very, very tight, so Japan does not have that much of an inventory cycle because they do believe in right-on-time. I do not think the Europeans have gotten around to it yet, so they do not have an inventory cycle either. It is really a U.S. phenomenon.

Over the last 20 years, we have seen what I would call the "delamination" of the business. We now

Graphic materials for this speaker were not available at the time of publication.

actually do have an equipment industry. Twenty years ago, we had an almost non-existent equipment industry. Many of the big companies manufactured the equipment themselves. We do not do that today. So this is what I call the delamination of the industry – that we have an equipment industry, a wafer fabrication industry, we now have foundries, assembly and test sub-contractors. We have fabless companies taking advantage of this, and an increasing software content; patent law is much more rigorous and applied today so that there are some relatively firm walls between companies and markets. It is much easier to maintain a proprietary position today than it was 20 years ago.

All of this acts as a cushion and it makes the amplitude of the cycle less. To return to the present, I would say we are at the beginning of the end of the slowdown. This has been an unusual one because it has been an excess capacity driven slowdown, and there has not been that big a change in the growth of the end markets. That is different than if the end markets collapse, like we saw in '74 or we saw in the '85 timeframe. I am very confident that within the next six months (and it might be quite a bit sooner than that), this industry will be back on its march to 300 billion dollars at the turn of the century.

That is a doubling in five years. That is pretty hefty secular growth. Few major industries will achieve that. To some extent, one of the influences that is reducing inflation around the world, certainly in the industrialized world, is the price of semiconductors. As those prices drop, the prices of PCs drop, the prices of mainframes drop, the prices of anything that is electronic drops, far more rapidly than anything else. In fact, many of the other factors seem to just increase in price. Even with the currency devaluing over time, generally the absolute prices of semiconductors have dropped phenomenally compared to virtually any period a year ago, five years ago, ten years ago, and this will continue at the same dizzying rate.

How can I be so confident in this relatively gloomy environment? A little more than two decades ago the worldwide semiconductor industry shrank from 5.2 billion in '74 to 4.9 billion the following year. Now, that does not sound so bad, 5.2 to 4.9, but you have always got to look at these things by quarter, and in late '73 we were going through one of these phenomenal growth periods and suddenly in the middle of '74, it abruptly stopped.

That is not what we have seen this time around, and I would say '85 was probably halfway between. After that very catastrophic 1974, and to put that in perspective, I was CEO of Fairchild at the time, we started the year with 30,000 people and finished the year with 15,000 people. That was a fairly traumatic change.

However, the very same semiconductor industry had doubled in size by 1979, so there is a pattern emerging here, you get this pause, then you double. You get another pause, then you double. A pause is not a heavy price to pay for a double. Global revenues plummeted from 26 billion in '84 to 21 billion the following year, and again, it was the same phenomenon. That did not seem that bad, even though it is pretty significant, but compared to the peak rates in the peak quarters in the previous year, to the low quarter in the next year, it was very dramatic.

Many small public companies were absorbed by bigger ones and just disappeared. In the ASIC business, going into that period, we figured that we had about 200 competitors. When we came out of that period in '86, the number of competitors had dropped to about half. Nobody ever figured out where they went to.

But again, we had this massive correction in the mid-80s, and by 1990 we were back to more or less another double. Then there was a significant pause in the 1990-91 period, and then a double again, and really the industry bounced back with a vengeance. 1993, '94, '95 were just phenomenal years. Now the hypergrowth in '93, '94, '95, with the collateral growth in earnings and market capitalizations, would have pulled in capital from the moon, and it did.

Too rapid a capacity build caught up with demand, end-markets have actually continued strong. The inventory correction that started, really began in Q3 of 95. It was not immediately obvious because people were carrying big backlogs, but it probably started with the computer industry. The computer industry has been around a long time, and they have lots of people that are very familiar with these cycles, so they tend to jump on it very early.

The networking industry, which again is an important industry to all of us today, had been growing so rapidly, that things like inventory management and MIS systems, were not as high a priority as the computer industry, so it took another quarter to two quarters before the networking industry said, "Hey, our inventories are not turning over, lead times have come down, we can start cutting work in process inventories, and inventories before the production line."

It has been a rolling adjustment. I believe the computer industry is through this period now. Some of the networking industries already have their inventories in line, but in another quarter, the adjustment will be done.

This is a lot different to a demand slowdown. Most of the inventory slack is already taken up. We are almost ready to resume growth. This story of a down cycle followed by doubling of the industry in the following five years, is about to be repeated. The bad news is that '96 is a down year. The good news is that there will not be another one until the next century. If you are in the investment business, that is a fairly easy thing to understand. Some skeptics will look at a $300 billion number at five years hence with somewhat of a jaundiced eye and ask, "Where will this demand come from and who will buy the chips?"

The real demand is still there, and it really has not slowed down, and in many cases the new products are going to accelerate that rate of demand. Prices have dropped precipitously in the memory market. That probably corrects for a suspended learning curve on pricing over the last few years. If you take the historical memory price learning curve, even with the 75% drop in the past year, you are not too far off what you would have anticipated from the learning curve. I believe prices they are going to pop up a little, but in reality it has been a suspension of belief and those margins, perhaps, in memory, got unrealistic.

However, the net effect of these lower prices today falls through very quickly into the end-market and is accelerating the end-use of memory. Once people go to larger memory sizes, that is a one-way irreversible process. They are not going to turn around in six months time, if the memory prices go up just a little bit, and say, "You know, I am going to buy a computer with half as much memory as the last one." It is not going to happen. You have built a habit with these people, and that is one of the good things about lowering the prices.

Another good example is platform video games – very sensitive to the retail price point. These 32 bit and the 64 bit video games were introduced in Japan at around 3,000-4,000 yen. These are products that initially could be sold in the $300-$400 range. If those prices come down to $250, to $200, and are now being discounted, soon they are going to be at the $150 range. As you look out to '98 they are probably going to be approaching $100, and so on.

What enables this to happen, is partially the normal learning curve and also the availability of large amounts of cheap memory. When you have a lot of memory in a box and suddenly that memory is 75% or 80% cheaper, this is a dramatic change. It means you can reveal a lower retail price point, and at that point your volume explodes. Each time you take these $50 moves, volume explodes, and ultimately that both increases the demand for semiconductors, and increases the semiconductor content in the hardware. This is happening in '96 and '97, rather than '98.

That is a phenomenon that is broadly happening right now, not just in those products but in a lot of other products. That will pay off very soon as those prices get into the psyche of the consumer. I agree with the earlier speaker, that suddenly the consumer is much, much more important in our world today. Clearly a cellular phone is a consumer product. I am not sure we talked about telephones as consumer products ten years ago, and obviously the more consumers you have, who can afford to pay the price, the bigger your market is going to be. Believe me, we are creating a lot of consumers around the world at this point in time.

Lower prices are fueling the customer demand, particularly consumer demand. For the financial industry, which tends to be very focused on the U.S., consumer products have been a bad word on Wall Street for a long time, ever since the U.S. companies got out of the consumer products business in the '70s – in fact it began in the '60s – and moved onto other things. However, the next generation of consumer products make heavy use of U.S. component technology and that is a big change. The new consumer business is a good business, it is a profitable business, and it is a sustainable business, where you can differentiate the products, you can differentiate your technology.

Let me give you some examples from the past. The mechanical calculator was made into an antique virtually overnight by the introduction of single chip processors in the early 1970s. We had quite a good business at Fairchild at that time, making custom chips for these rather specialty digital calculators. All of that custom business, which was actually the beginning of the ASIC business, dried up overnight and we had to shut down the whole operation. It was not until 10 years later that the ASIC industry started to take off again. There were lots of sudden casualties, some of them quieter than others, but the mechanical calculator business just went away overnight. When a single chip solution arrives, all of the economics in the industry change, the volumes change, and the previous players usually disappear. Sometimes they are able to make the transition, but generally they are replaced by the next generation of players.

Another example is in '91 you could buy a Compaq 386, 20 megahertz laptop, for just over $4,000. Today a Compaq Pentium processor, 166 megahertz, 16 meg of memory, can be bought for about $3,000. That is why the number of PCs has grown from 700,000 in the early '80s, to approaching 200 million in '96. By the year 2000, the same computer will cost less than $1,000, and the performance will rival mainframe computers.

The product drivers that will fuel the growth of the global semiconductor industry are generally being produced by trend-setting customers in key vertical markets, but the main drivers are communications, consumer and computer. The three "C"s.

Lets take the consumer electronics market for example. Most U.S. investment themes, if they involve personal computers, workstations, storage, networking, have their roots in the United States because the key suppliers are in the United States. The key end-markets are in the United States. But the consumer electronics market for the last 20 years has been driven outside of the U.S. and are much more important in Japan and Europe. Even though the U.S. has been a big participant as a buyer of these products, very little of the manufacture and design has been done in the U.S. – and that is changing. The new generation will carry a lot of American, and a lot of Japanese semiconductors.

So what are these products? DVD: Digital Versatile Disk it was called, now Digital Video Disk. DVD is a very important product. It will replace the VCR, it will replace the CD in the PC. I just came back from Japan, and I got the perspective of all the major Japanese companies on what the numbers would be by the year 2000. Many of them a little different than what Dataquest is estimating, and significantly different to each other. They all had different ideas in what the quantities would be. But the range is mostly between 30 and 40 million units will be shipped in the year 2000. But almost to a man there is a belief that that number could be a lot higher and certainly in early 2002, 2003, we could be looking at an annual sales perhaps double that, as the applications proliferate.

This is a market that is certainly in the early going, is dominated by a small handful of Japanese companies, perhaps soon to be followed by Korean companies and Taiwanese. But this will replace the video cassette recorder. It will replace the laser disk, and it will supplant all the CD-ROM drives within the next few years. This gives you a 15 to 20 times increase in the capacity of a CD-ROM, and even though they will be priced higher to begin with, if you look out three, four, five years, I think the price is going to be very comparable with the existing CD.

It is estimated that DVD will achieve the same U.S. household penetration in five years that it took the VCR 11 years to get to. It took the audio disk about 7 years to begin to exceed tape. That will probably happen in four to five years with the same phenomenon in video. So the VCRs will be replaced much quicker than the old tape players for audio. This is going to be a major business. Fifty million VCRs are sold a year, and this is simply a replacement market for a market that already is pretty much saturated. So this is going to be a big, big market.

This also enables products like the "James Bond" automobile map. It takes about six disks to give you a complete street map of Japan. With DVD you get a complete street map of the whole U.S. on a single disk. So you never have to change the disk, you just put the disk in and drive your car around, and wherever you are, you know where you are, to within four feet. These products will probably sell in the $1,000 to $2,000 range, by that timeframe. They will be as pervasive in automobiles as cellular phones are today.

Another example is digital video cameras – the next generation of hand held video camera. The Japanese consumer is willing to pay quite high initial prices. For something like this, $3500 to $4000, which is definitely a non-starter in the U.S. But that can be done for a year or so and those prices will steadily come down. Once you start reaching U.S. consumer and European consumer price points, once you start to get below $1000, you start to see some real volume. That is going to happen, probably within 12 months, with DVC, and then, as time rolls on, you get down to the $500 price point, again volume starts to explode. These have much higher resolution than analog cameras. You can freeze-frame images, download them into computers, you can show them on your TV set. It is a fairly radical change in the way you think of a camera, certainly a video camera.

Then we have digital still cameras. You can get many, many more shots in your camera, play it back on your PC, or play it back on your TV, or print out the pictures, or you can make up your own album. Market research indicates that about 48 million cameras will be shipped in 1999 – these are serious numbers.

Let's talk about Internet browsers. Nobody was thinking about Internet browsers two years ago. Everybody accepts that today. This is going to be a big business in itself. We do not really know exactly how many Americans go online – it depends on which survey you get – but we know the numbers are at least 35 million, and it may be as high as 48 million, and this thing has not even started yet. This is just people that have PCs. One new product that has hit the market is Web TV, which allows you to use your TV screen as an Internet access device and it works just fine. There are going to be a lot more of those products. There must be a dozen companies in Silicon Valley, all coming up with what look like rather similar products, and they will be moved through the usual consumer channels. That growth will skyrocket, and, of course, once that infrastructure is enhanced by these little boxes, all of the services that now start to flow into those boxes start to become an even bigger industry.

If I look at the five billion dollar a year semiconductor content in networking, which is experiencing unprecedented growth in the last couple of years, this is going to be driven even faster by the existence of the Internet, and the existence of these boxes at the many different access points.

A 0.5 micron will always beat 0.8 micron, and 0.35 micron on cost will beat 0.5 micron, and 0.18 micron will beat 0.35 micron on cost, no matter what you put into these fabs. In reality, the capital investment in the semiconductor industry has not really changed much in the last 15 years. All that has changed is the granularity. If you measure the capital investment in terms of the percentage of the semiconductor companies sales that they put into capital, which is probably around 20%, it has been at that for some time. All that has happened is that you now have a single process, CMOS, whereas ten, twenty years ago there was still a wide variety of different processes. The granularity of the investment has changed, so instead of building five 200 million dollar fabs, you build one 1 billion dollar fab.

A major trend is "system on a chip," one thing that my company focuses on today, is that within a one-centimeter chip with three, four, five layers of metal, you can place millions of transistors, which means you are almost forced to suck the whole system onto one chip. Once you do that, you have now reduced the cost of the system to a photographic process. That radically changes the cost of the end system. If you can keep photoreducing these things, you can take these systems, and instead of getting 200 on a wafer, you get 500, and instead of getting 500, you get 1,000, but fundamentally you are still taking the same photograph, and that is the driver that we are seeing in this industry.

The semiconductor content in the electronics products is increasing rapidly. Just four years ago. semiconductors made up 12% of the $585 billion global electronics business. This year, the semiconductor content is expected to approach 20% of the world's $850 billion electronics market. By the year 2000 we are projecting semiconductors will represent 28% of all electronics products. That is a long way from single digits, 20 years ago.

The most important sales driver for this $300 billion long jump, is the consumer. Things are moving towards being consumer driven, which means the more people you have, which is one variable, and how much money those people make is another variable, and both things are moving in a positive direction. Only 4% of the world's population resides within the U.S. The three primary markets for the things that we make historically have been bought by less than 10% of the world's population. Suddenly we are starting to see the rest of the world – significant pieces of the rest of the world. There now is a 200-million-person middle class in India, for example, and the beginnings of a middle class in China needing all of these things that the middle class needs, is a big shift in global demand.

Approximately one out of every two chips manufactured by U.S. companies are sold overseas. This has been a fact through the '90s with the percentage continuing to increase. It is over 50% at this point in time, and it is going to be quite a bit more in the next five years.

Let's look at satellite TV. 60% of U.S. homes are wired for cable, and maybe 80% of them have reasonable cable access if they want it. That is very unusual. That is not true in Europe, and it is not true in the rest of the world. The emergence of the satellite technology means that the medium of choice to get this programming in the future is going to be by satellite. Again, a huge opportunity for satellite set top boxes. India, with a population of more than 900 million, has just 200,000 satellite TV subscribers.

In China, there is a population of more than 1.2 billion, there are only 250,000 TV satellite subscribers. These markets represent huge opportunities. The majority of the people in the world today have never made a phone call. All we have to do is have a few of them, on a percentage basis, doing that and the number of telephones in place start to explode. We have seen an explosion in phones just by selling the second, third, fourth and fifth to Americans and Europeans and Japanese. Just imagine if we can just sell one to 10%, 20%, 30% of the Asians.

To do these things we have to have free and open access to global markets. I have spent a lot of the last 10 years involved with the SIA on this subject. The Japanese, and more recently the Koreans, are starting to hone in on the same issues that we are, as they see much of the market being outside of their own countries. Suddenly it is very important to them that there are not barriers to sell into Asia. It is very important that there is copyright protection. So suddenly our issues are starting to become the same issues, and a lot of the discussions today are much more constructive than they were 10 years ago. We have to continue to push open trade. The governments are recognizing that this is a major issue. What we are going to see in the next several years, is even more drop in tariff barriers, starting with Europe, and then Korea, and then we will really start to have a global market.

We may soon be the economic bell-whether for the nation. If we look at the pursuit of the market that we represent, we are rapidly approaching 1%

of the gross world product, which is a big percentage. It is something that you can now start to track, once you get to that 1% range.

Electronic instrumentation, computers, and the rest of the electronics industry, now represents 3.5% of the GDP of the United States. If you look at the total semiconductor industry, 1% of world GDP in the year 2000 is very dramatic. The key drivers are certainly the customers, particularly the consumer customers, futuristic products, advanced technologies, system on a chip integration, rising chip content, global markets.

This year will be looked at as just a pothole in the road, on the path to a $300 billion industry. We will not reach $300 billion by bemoaning our fate. I would like to close on a quotation from President Teddy Roosevelt, who might have been talking about the semiconductor industry when he said, "Far better to dare mighty things, to win glorious triumphs even though checkered by failure, than to take rank with those poor spirits who neither enjoy much, nor suffer much, because they live in the great twilight that knows not victory nor defeat." Thank you.

Q: This is really something that impacts a lot of system users, systems manufacturers and users in this audience. You talked about the system on a chip – what are the positives and the negatives, at that level of integration, as far as the manufacturers of systems, and the users of systems?

A: As these systems get very complex, if you can get everything on one chip, it is infinitely more testable, infinitely more manufacturable than if you are going to take even four or five or six chips. You can test it in the wafer form, be sure the whole system is going to work, and then putting the thing together is very simple at the system level. All the interconnect problems get sucked back into the chip itself and fundamentally those interconnects become just a photographic problem. So it really is a major reduction in the cost, and that is why you are seeing these products that would seem extremely complex a few years ago, being sold at these very low price points. Intel has reduced the computer to a single chip, even though the memory is separate, and so on, but the actual intellectual property of that PC is all defined by that single chip, and that is going to be true in most products in the future.

Panel:
Dataquest's Worldwide Outlook for Systems, Semiconductors, and Equipment

Moderator:

Gene Norrett

Corporate Vice President and Director,
Semiconductors Worldwide,
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. Before that 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 B.S. degree in mathematics from Temple University and an M.S. degree in applied statistics from Villanova University.

Panel:
Dataquest's Worldwide Outlook for Systems, Semiconductors, and Equipment

Gregory L. Sheppard

Chief Analyst,
Semiconductor Application Markets Program,
Semiconductors Worldwide,
Dataquest

Mr. Sheppard is 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 (MCSAM), and the Communications 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.

Mr. Sheppard is a specialist on the end use or application of semiconductors. His scope of analysis includes both economic and technical trends regarding the semiconductor opportunities in electronic equipment.

As an analyst Mr. Sheppard's current focus includes LAN, WAN, and modem markets; telephony and public infrastructure markets; and set-top boxes.

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. Before that, he worked in engineering management at GTE Corporation specializing in communications systems design and decision aid systems.

Mr. Sheppard received a B.S.E.E./C.S. degree from the University of Colorado and an M.S. degree in system management from the University of Southern California.

Panel:
Dataquest's Worldwide Outlook for Systems, Semiconductors, and Equipment

Gary J. Grandbois

Vice President and Chief Analyst,
Semiconductors Worldwide,
Dataquest

Mr. Grandbois is vice president and chief analyst for the Semiconductor Industry and Manufacturing (SIM) program of Dataquest's Semiconductors group. Besides covering the broad semiconductor market, Mr. Grandbois monitors the discrete device, optoelectronic, analog, and mixed-signal IC markets. He also has responsibility for the following programs:

Semiconductor Equipment, Manufacturing, and Materials program provides the look at wafer fab, fab equipment, materials, capital spending, and manufacturing issues.

Semiconductor Contract Manufacturing program covers the issues surrounding the wafer subcontract (foundry) business and the trend toward the outsourcing of semiconductor manufacturing, both for fabless companies and fully integrated suppliers.

Semiconductors Worldwide program provides a broad view of the device market with in-depth looks at issues that affect all devices. Mixed-signal ICs, power, RF, and packaging are a few of the issues covered. Companies and their competitive positions are also covered.

Semiconductor Supply and Pricing program monitors the pricing and availability of key semiconductor devices. This service also reviews technical and industry issues that may impact device pricing and supply.

Mr. Grandbois has held a broad spectrum of application engineering and marketing positions in the semiconductor industry. Prior to joining Dataquest, Mr. Grandbois was vice president of marketing and sales at Teledyne Semiconductor. Numerous technical articles authored by Mr. Grandbois have appeared in trade publications. They include some of the earliest use of microprocessors in analog/digital applications.

Mr. Grandbois holds a B.S. degree and an M.S. degree in electrical engineering from San Jose State University.

Panel: Dataquest's Worldwide Outlook for Systems, Semiconductors, and Equipment

Clark J. Fuhs

Director and Principal Analyst,
Semiconductor Equipment, Manufacturing, and Materials Program,
Semiconductors Worldwide,
Dataquest

Mr. Fuhs is director and principal analyst for Dataquest's Semiconductor Equipment, Manufacturing, and Materials (SEMM) program in the Semiconductors 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, 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 B.S. degree in chemical engineering from Purdue University in West Lafayette, Indiana, and received an M.B.A. from the University of California at Berkeley.

Session # 5: Dataquest's Worldwide Outlook for Systems, Semiconductors, and Equipment

Moderator:
Gene Norrett
Vice President & Director, Worldwide Semiconductor Group, Dataquest

Panelists:
Greg Sheppard
Chief Analyst, Semiconductor Application Group, Dataquest
Gary Grandbois
Vice President & Director, Semiconductor Industry & Manufacturing Group, Dataquest
Clark Fuhs
Principal Analyst, Semiconductor Equipment Materials & Manufacturing Service, Dataquest

GENE NORRETT: In the semiconductor group, we look at three different segments of the food chain. The first segment that drives the chip industry is the end user, and the issues for the manufacturers, the vendors, of those systems. What are their issues? What is the semiconductor content? What is the technology trends of those boxes? That, in turn, will drive the semiconductor manufacturers in what they need to be manufacturing to meet the needs of these end systems. Last, we follow the equipment and materials industry. Clark Fuhs is going to be talking about the outlook for that particular industry.

Our first speaker is Greg Sheppard. Greg has been with us at Dataquest for almost 10 years. He has worked in the industry for Fairchild and GT&E and he is a MSEE from University of Colorado.

GREG SHEPPARD: Today I would like to paint a picture for you for what we believe the main environment is going to be, particularly with focus on next year, but also a long term view through the end of the decade. We feel that demand for semiconductors, in terms of gates and bits of memory, transistors, and units of things, is going to remain very healthy. I think the main question in everyone's mind is what price are we going to fetch for those particular functions.

Here is our worldwide electronic equipment production forecast. This is the key way that we measure the demand environment for semiconductors. After growing at a double digit rate over the past two years, 1994 and 1995, the market for production of electronic equipment has slowed to 6.2% in this year (we are projecting), reaching a value of $840 billion – not a small market by any means. The reasons for the decline are two-fold. 50% has to do with the way the dollar and the yen get valued. The dollar has appreciated against the yen, and since we have a great deal of electronic production based in Japan, that helped pull down the overall number. About half of the decline is an actual softening of the production of electronic systems. As a rule of thumb, semiconductor markets grow at roughly two to two and a half times what the electronic production shows, except in times of upswing and downturn, and that is where inventory building and pricing factors contribute.

We expect the demand environment for chips – as modeled, or based, on production – to stay in the

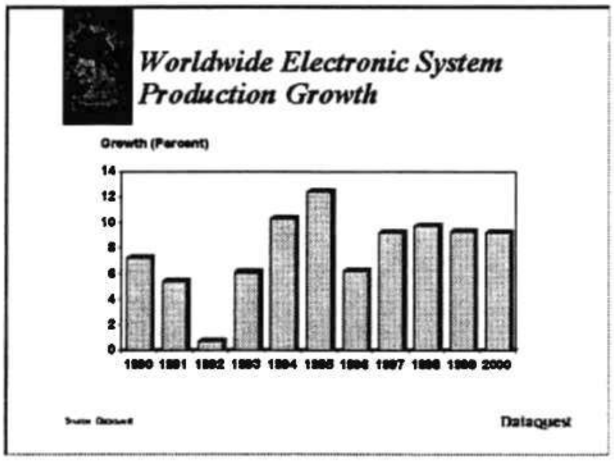

high rates, relative to the previous part of the decade, but certainly over the 8% range. This sets up an environment where the chip market should grow in the 18-22% range through the end of the decade.

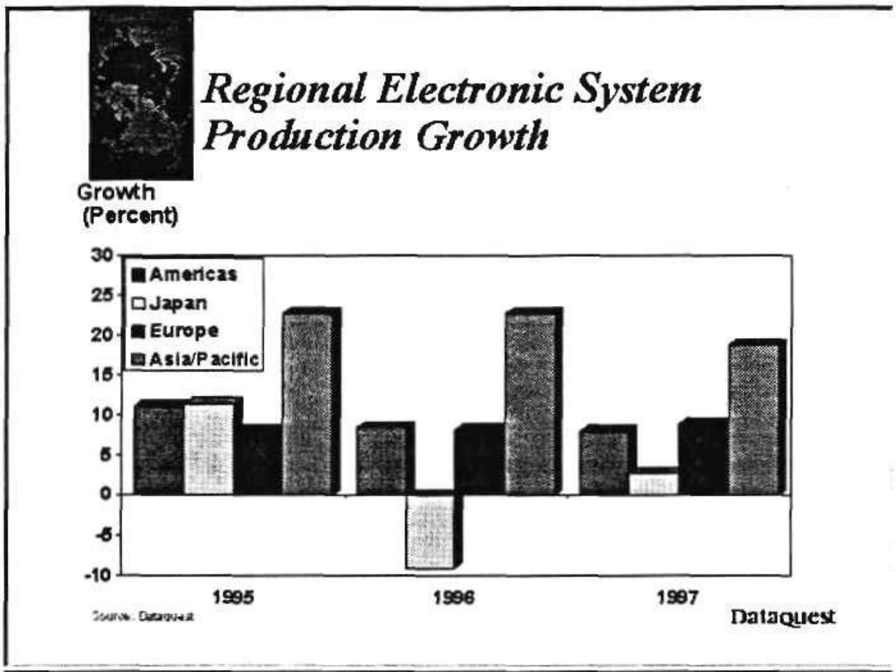

We measure roughly 100 or so different types of equipment at Dataquest. We maintain very detailed models on unit projection by region, semiconductor content, and in general, how this is driving and changing the demand patterns in the chip business. We segment the market into these basic six high level areas: data processing, which includes computer and peripheral equipment; communications, which is both LAN/WAN, public systems, and mobile communications, (that is where we count, for example, cellar telephone technology); consumer, which is consumer entertainment systems; industrial, which is basically systems that make other systems, or measure performance of other systems, we also include medical electronics in this category; military aerospace, which is both military or defense electronic systems as well as civil aerospace, as well as commercial space, i.e. satellite systems; and automotive, which is automotive and truck electronics.

We have shown an overall decline in all categories in 1996. The consumer segment was the hardest hit as there has been a slowdown in the shipment of consumer systems globally. Some of the economies around the world have been weak that absorbed those systems. In general we have all of the categories forecasted to rebound in the coming year.

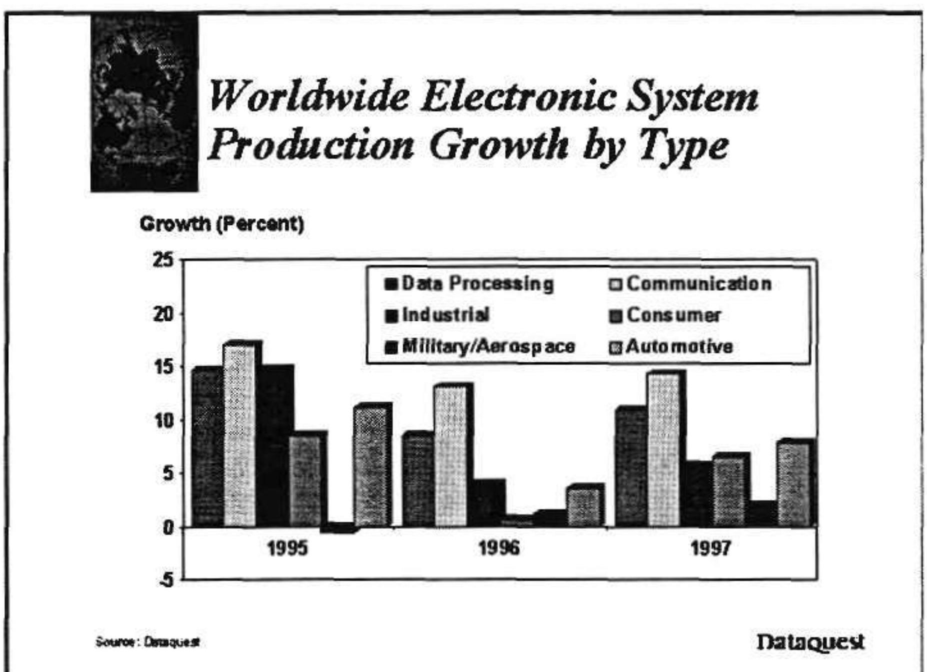

From a regional perspective, if you take a look at 1996 and notice the Japan bar being negative, this is the appreciation factor that has hit the Japan electronic equipment production. But the fact of the matter is that we had slowing in both the U.S., or Americas, market, as we describe it here, and that is one clarification I should make – we do measure all of the Americas in our data now, so we

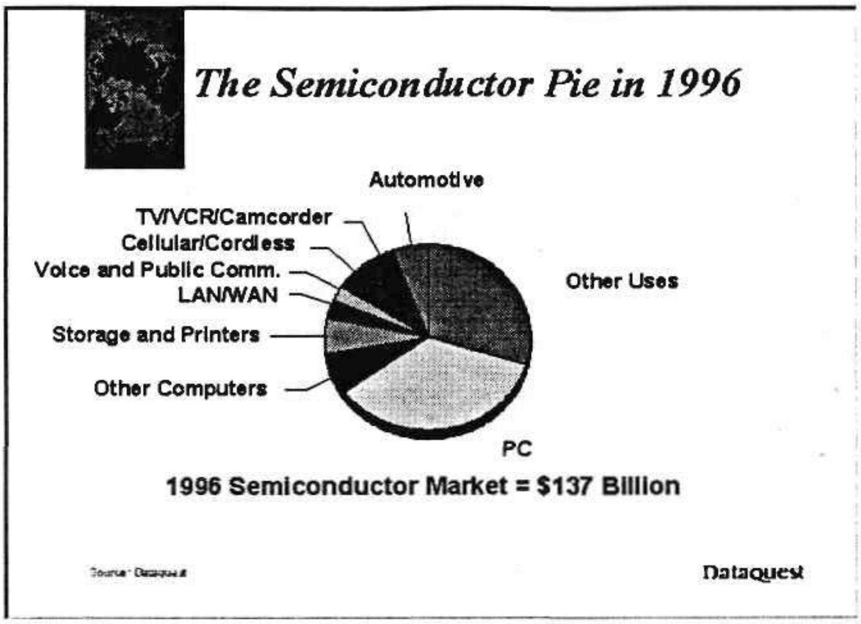

have broadened our measurement of that area. Likewise, Europe includes Middle East and Africa. Europe remained pretty much flat year to year, but Asia-Pacific, which is normally the big double digit growth engine, has slowed somewhat.

For next year we are predicting that all of the regions will see some growth, except for Asia-Pacific, where we have it, in fact, declining in growth rate, although it is still positive growth rate, but slowing somewhat. At some point, this is basically the law of large numbers that kick in, and where Asia is starting to be in the base of numbers is getting so large that in fact it had to start slowing in the growth basis.

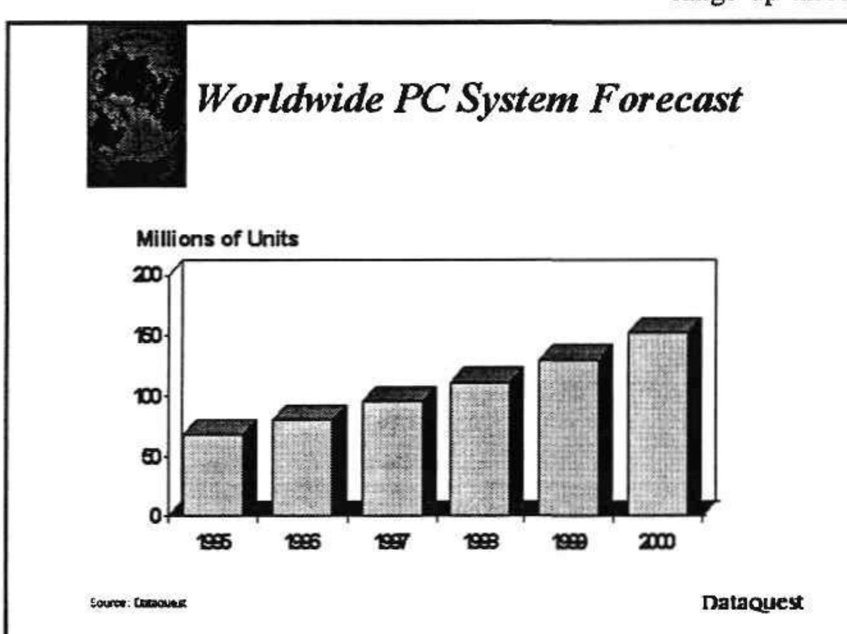

If we take the semiconductor market, and now I am going to direct my comments more at semiconductor demand, more so than electronic equipment production, so take a little bit more directly into the chip demand that we expect from different segments. But if you take the pie chart of the industry which is going to be, we figure, $137 billion this year (1996), and segment it up into the systems that are sucking up these chips, we see that the PC market accounts for roughly a third. This is actually chips that are going inside the PC box itself. That is excluding the storage systems, monitors, peripherals, and so forth, which we count in the storage and printer section.

The second largest demand for chips is other computer systems. It is still quite a sizable market: workstations, server systems, traditional mid-range up through supercomputer-type systems, is still quite a big demander of chips. Likewise LAN/ WAN voice and public communications systems is another good slice. Cellular and cordless handsets. The infrastructure actually is not included in this particular number. On through to automotive. It is important to note that for other uses, which is about a third of the market here, there is literally tens of thousands of applications out there which are literally impossible to measure.

I think that Joe Grenier, in past years, has given you some nice overviews of these weird and unusual applications, but the fact of the matter is that the chips are ending up in tennis shoes and all kinds of unusual things never thought about before.

Let's focus in on the PC area. This was a third of the chip market as I showed earlier. This is a forecast from Dataquest's PC service, and we note that – and this is the forecast of worldwide PC shipments on a unit basis – we note that in 1995 the market had grown 26% on a unit basis. We are projecting that this will be in the 20% range for 1996. This group is gathering data quarterly, and they feel that we are pretty much on track for that number. Obviously, the big variable is Q4 as to whether the systems will sell or they will stay on the shelves, but that is one of the variables, or risks, certainly, to this forecast.

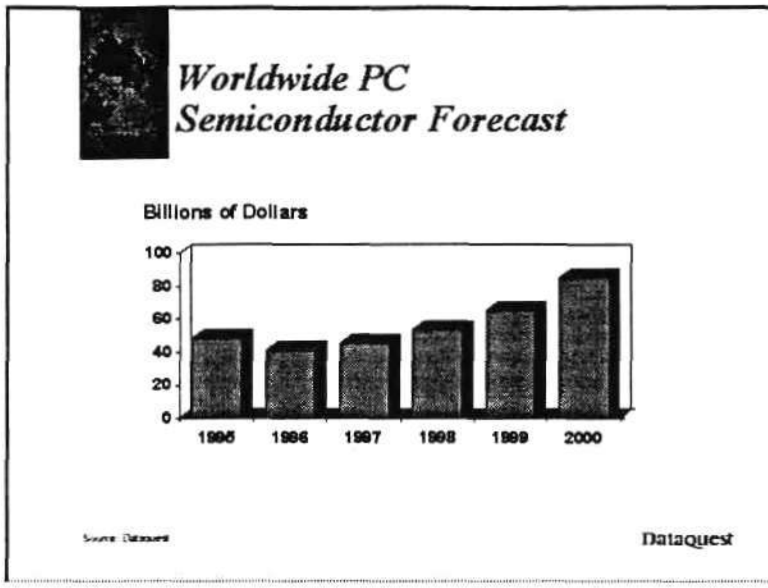

Looking at 1997, they are expecting the market for PCs to grow at 18%, so a further slowing and then kind of hovering in that range, out through the end of the decade, as the unit shipments of PCs approaches 150 million units annually.

If we look at the demand of chips and PCs which are comprised of DRAM microprocessors, graphics, multimedia chips, sound communication chips that are imbedded in the PC, we note that the overall market for chips consumed by PCs will decline on the order of 10% this year. This is driven mainly by the decline of the memory, particularly the DRAM market. This also shows that the elasticity of the DRAM market, i.e. more bits shipping because they are cheaper, was not enough to overtake the decline in pricing, and so thus the market takes a dip.

We expect this market to recover and start to grow at a more normal pace, relative to that unit demand forecast through the year 2000, and drive and stimulate the recovery of the overall semiconductor market.

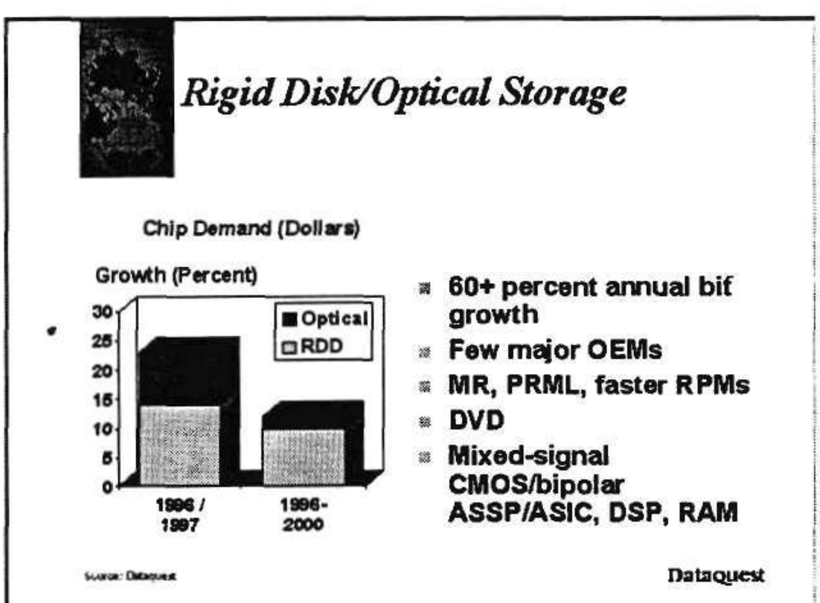

Let's take a look at a few other markets in more focused detail. This first one is a look at the rigid disk and optical storage area. What this stacked bar chart is showing is that in the upper part of the stacked bar is what the optical drive market is growing at, and the next lower bar is what the rigid disk drive market is growing at. Bottom line is that optical storage chip demand will grow at almost 25% next year, and then rigid disk drive demand will be somewhere in the 15% range.

We are expecting these to both cool off over the long range as new product is introduced and prices come down. This is an extremely cost competitive area for chip companies to play in. There are companies literally designed in and out every nine months on the basis of features, and what they are really beat up on is cost. So the demand on this tends to be somewhat suppressed based on that. But it is still a fantastic unit market, it fills fabs. Overall, if you look at some of the things going on in the rigid disk area, we have the densities of bits you can put onto any disk platter growing at 60% a year.

It is controlled by a few major big OEMs on the order of half a dozen or so. There are some new technology opportunities in magneto resistive, or MR, partial response maximum likelihood, PRMLs, I have noted here. We are spinning these drives faster and faster so we can get the data off faster. These are all contributing to new types of chips that will be developed. Also, I think there is the DVD market, which can be used both as an optical storage mechanism for PCs and as a consumer player. We are right in the ground floor, if you will, of companies developing chips to go into these to drive the costs up. The types of technologies that will be needed from the chip community I have noted here below. ASSP stands for Application Specific Standard Product. ASIC is Application Specific IC.

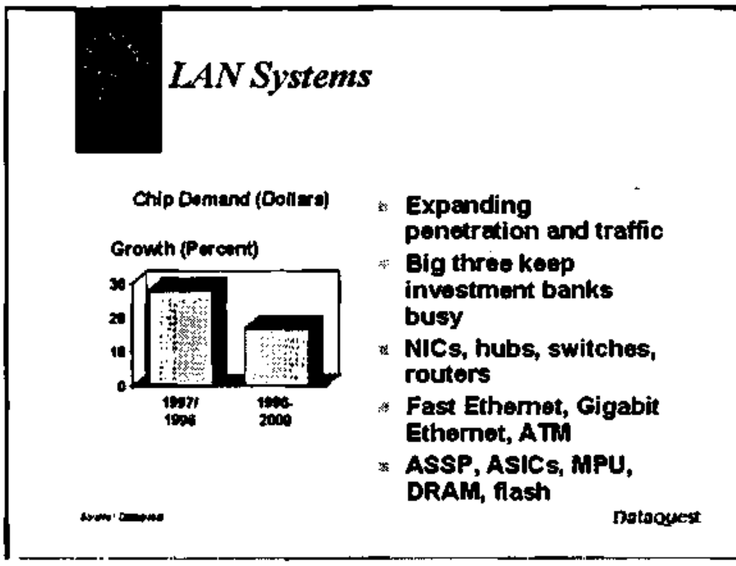

Now I will shift over into the communications space. Local area networks systems continues to be a very healthy market for chips. It is going to grow in the 25+% range for next year, and continue to click along in the 15+% range thereafter. What is driving the local area network market? Well, expanding desktop connectivity, and also upgrading of existing systems to higher performance, higher speed. There is three or four big companies that are the giants in this area: Cisco, 3Com, Bay Networks and Cabletron.

Key equipment, a NICS, or network interface cards, these are the cards that go inside the PC to connect it to the network. Hubs, and switches, and routers, are the guts of the network, if you will. These are all very rich in semiconductors and all very hidden from us, in our daily lives. Something in the order of $500 of equipment is installed behind the scenes, behind each PC if you will, to get it connected to the network, to the LAN or WAN. This is running with 20 to 30% chip content.

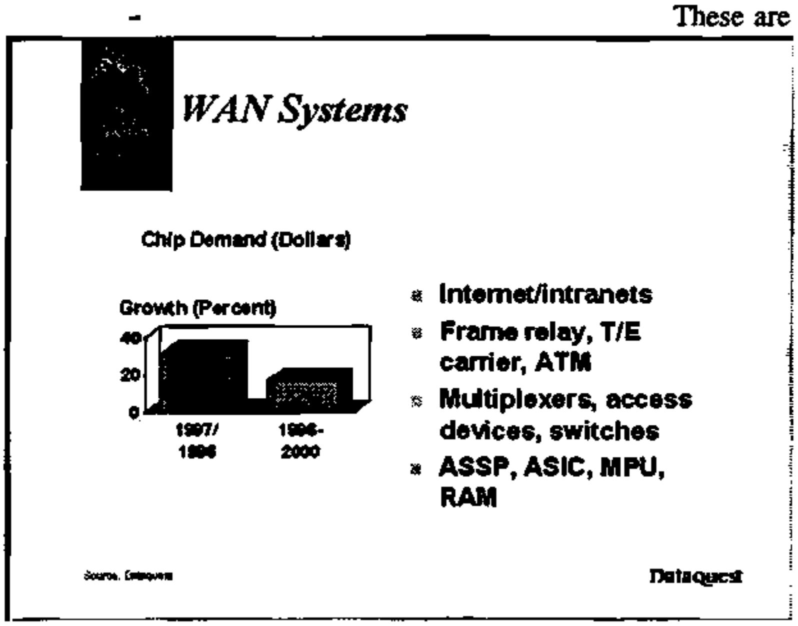

Other happenings here are the movement towards 100 megabit Ethernet technology known as fast Ethernet, and now the one beyond that known as gigabit Ethernet is entering the scene. Then the asynchronous transfer mode market, ATM, as the next generation technology, and I have noted some of the impact areas for

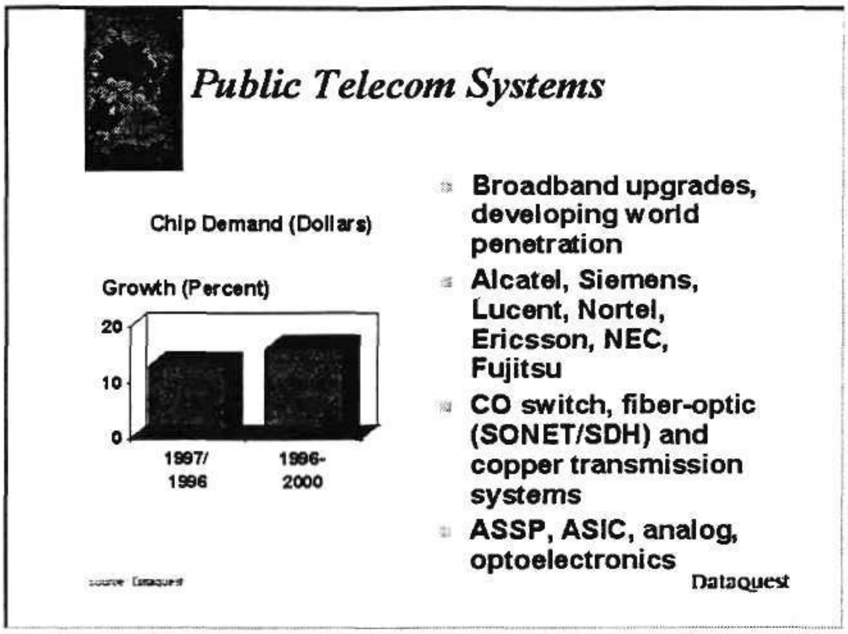

standard chips sets, or ASSPs, ASICs, and so forth.

Over in the WAN side (Wide Area Networking), is where the Internet is making its play basically. We are seeing services being rolled out that support Internet, and then equipment being put into place to support those services. The most common Internet hardware out there now is the modem. We think this market will continue to clip along from a chip perspective in the 20% range or so. In addition to this, there is some new equipment being introduced – frame relay. It turns out that a good percentage of the world's websites are hooked up via frame relay.

So we are seeing silicon opportunities targeted at switches and access devices in this area. T&E carrier are the digital transmission lines, that are in the bowels of the public telephone company, but also being delivered out to businesses, and asynchronous transfer mode is also a technology used here. So we are seeing both a technology and a set of services that are being rolled out.

Over in the public telecom side, this actually has a different profile. We see it actually accelerating toward the end of the decade in terms of chip demand. This is due to the fact that we have numerous projects underway around the globe – some are just basically getting phone line access out to the average home. There are something like 500 million, or so, access lines worldwide currently, and that is projected to grow to over a billion lines by the year 2000. This is addressing the market in India and China, putting in phone access for the first time. Chip demand is drawn from that as well. Likewise we are going to see in the developed countries of the world, upgrading to broadband capability to be able to handle fast Internet access and video multimedia services. This market is also controlled by a small set of big companies. Key equipment involved are switches, CO (central office), there are fiber optic transmission systems involved, and also a great deal of optoelectronics as well, particularly related to the fiber electronics.

Mobile communications: this is cellular, cordless, and we count cordless technology in this as well, as well as pagers. As these systems go digital, that market can grow exceeding 40% next year, and

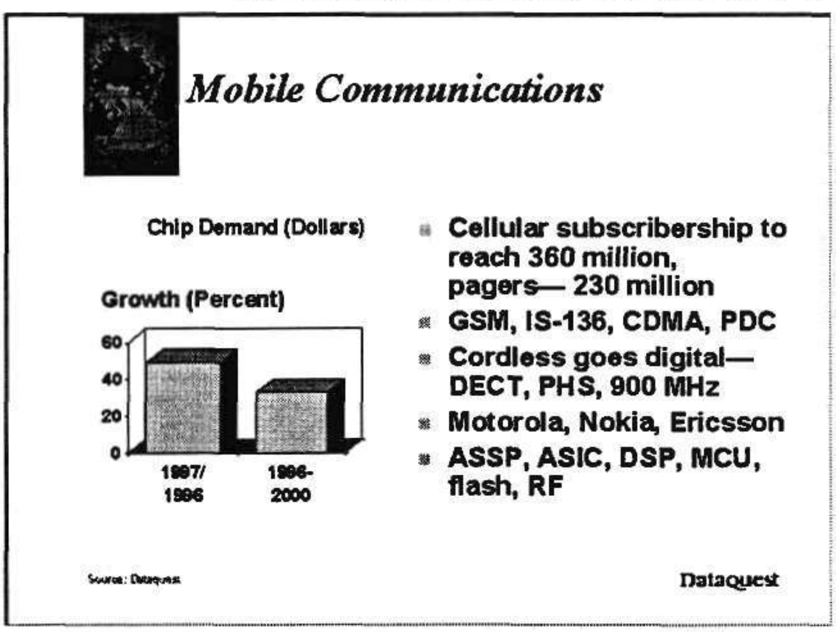

then starting to slow somewhat into the 35% range out through the end of the decade. This is just the digital part of it. The overall market is seeing some inventory problems, if you throw in the analog cellular technology as well, and there are some suppliers that are doing well, and some are not so that has to be kept in mind.

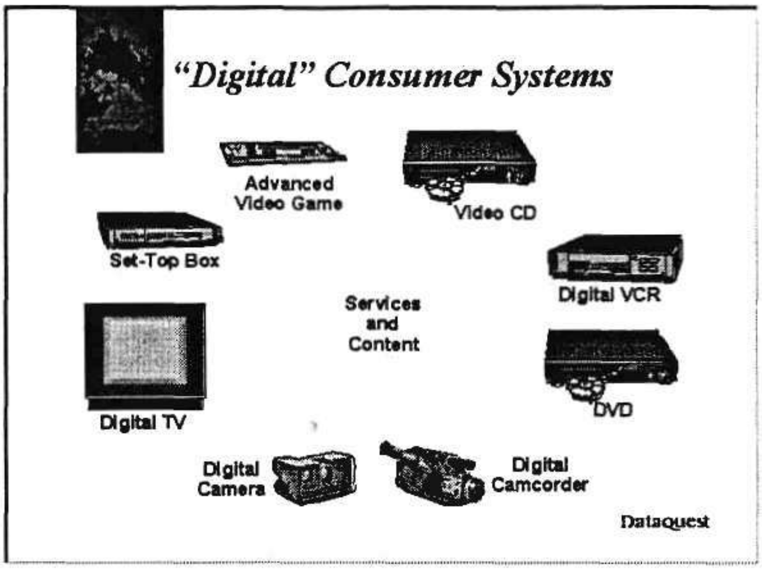

This forecast is predicated on the fact that we will see cellular subscribership tripling between now and the end of the decade. Our wireless, our personal communications group is forecasting that we will see upwards of 360 million subscribers, by the year 2000, or during the year 2000, we will have 360 millions subscribers, likewise, 230 million pager users. I have also noted here the key digital cellular standards, GSM, is going to be the majority owner of digital standard technology around the world that we can tell, and likewise, cordless technology in the digital form, is just now getting rolled out.

The consumer electronic giants are turning to digitally enhanced, next generation systems as a way to regain profitability, to get sales back up, particularly into their current markets. So we have seen a variety of new systems come out, although these are all dependent on the availability of services, particularly in the set top box case and the availability of content, for example, in the DVD case. Those are all key issues that we will be tracking. One of the two hot items for this year is the digital still camera, as that is rolling out in earnest. We are looking at some cameras inside our shop right now, trying to figure out what the chip technology is, and where that will be going. DVD, which is being introduced this year, is probably going to slide somewhat into next year as there are still a lot of issues surrounding that.

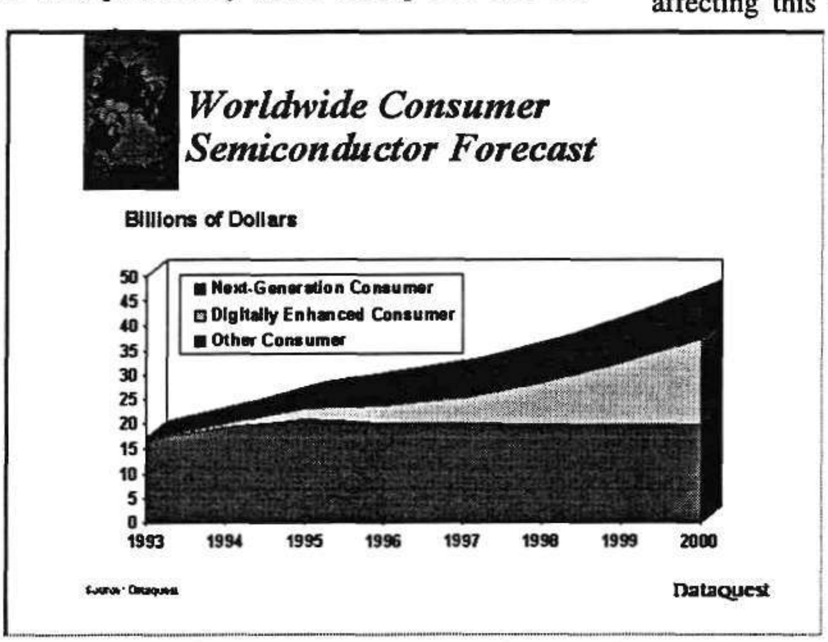

If we look at the semiconductor demand for consumer electronics, we note that in 1996 we are going to take a dip because of pricing declines in the analog and in discrete areas particularly, it is affecting this market as well. But the real true growth is in digitally enhanced, next generation systems, which include DVD, and set top boxes, and so forth. That market is actually growing at 3.5 times the rate of the overall top line.

Just a quick word on other usage areas. Automotive: after taking a bit of a pause this year we are expecting vehicle production and electronic content penetration to accelerate starting next year. In 1998 we will start to see navigation systems become regular options. BMW, I think, is going to be the first one who will have it as an option you can get right from the factory. We think that the navigation systems are really going to be the next big driver of electronic content within cars, as airbags and anti-lock braking systems have been in the most recent past.

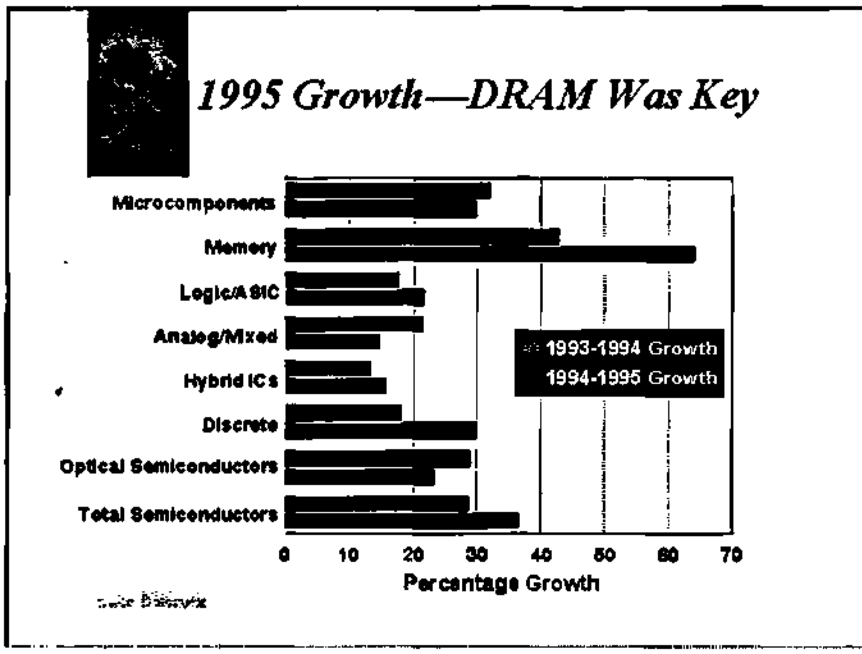

Military aerospace: basically we are in the throes of all the world's aviation plates being updated, and the upswing in satellite building supports all the new direct broadcast satellite services. In the appliance world, the movement towards electronically controlled appliances will be an important factor. The last point is the network computer. We think it is an opportunity for semiconductor suppliers, although we are not so convinced as to how big an opportunity it will be, until the market issues sort themselves out.

In summary, we think next year, from a demand perspective, is going to be good. The number of units of systems that are pulling the gates, the bits, and transistors, is going to be in good shape. The PC and mobile communications areas will continue to be drivers of growth; high speed LAN and WAN systems, public telecom systems. Digital entertainment is growing at 3.5 times the overall consumer electronics area, and will be important, and then the resurgence of automotive over the next couple of years as well.

GENE: Gary Grandbois, who is Vice President and Director of our Semiconductor Industry and Manufacturing Group, and is a major contributor to the device forecast. His personal expertise is in the analog and mixed signal area. Gary has worked for many industry companies including Precision Monolithics and Siliconics, and has an MSEE from San Jose State.

GARY: I am going to talk about semiconductor devices. I really want to talk about where we have been in 1995, this relapse we are seeing in 1996, and when the rebound will occur. It will occur very soon.

1995 was an exceptional year. It really capped three years of very strong growth. We saw 31% growth in 1993, 29% in 1994, and 37% in 1995. That is a three-year compound annual growth of 32%, a doubling of revenues in three years – pretty substantial growth. And since we got to $151 billion markets in that time, it seems that it would be a cake walk to get to $300 billion in five years. But we have hit the pothole in the growth to the $300 billion, and we are not likely to see this kind of growth in the near future.

In 1995 we saw tremendous growth, led by DRAM essentially, because the normal price per bit did not decline. It was stagnated for a while, it held up due to capacity limitations, and so it was an unnatural market. That unnatural market was also driven by the very strong growth in PC units – 70% growth in Japan in 1995 really led the growth for the worldwide growth of 25-26%. That has slowed only slightly into 1996. Other factors have caused the decline that we are seeing now.

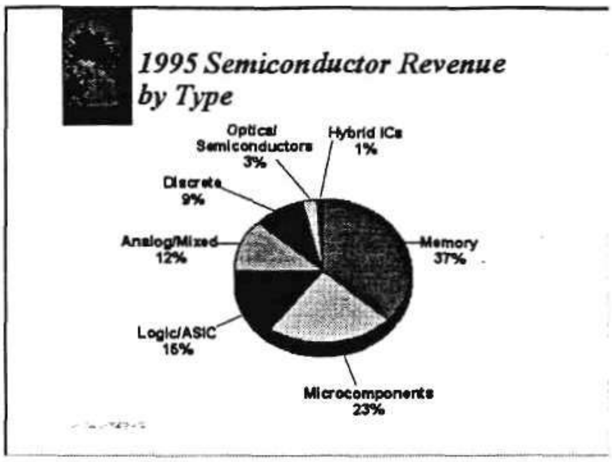

I would like to talk about the semiconductor mix in 1995. You see it here where, memories and micro component ICs, the PC-driven products of the past decade, now take almost 60% of the semiconductor revenue. Ten years ago, in 1985, memories and microcomponent revenues together accounted for only 29% of the market, and it has doubled its share of the markets in just a decade, pushing all other products into what seems like an other category. In fact, the products we are talking about are also very strong growers, but not to the degree that we have seen in memory and microcomponents. So how many multiples of the revenue in 1985 did we see in 1995? It was more than twelve-fold for microcomponents and memory ICs. That is a 29% compound annual growth rate.

When you consider that the PC market only grew at a 15% compound growth rate over that time period, that is phenomenal, substantial growth, and of course, an increase in semiconductor content in personal computers at that time. Will it continue into the future? Probably not at that rate, because we have seen a growth and emergence of the PC market that probably cannot continue in that manner. Over the decade, semiconductor revenue grew by 20% compounded.

In the past few years, DRAM was the king as the normal price decline was suspended, and microcomponents have settled into a more natural growth rate, rather than what was seen in the past. In the past three years, the market has really been spurred by the DRAM growth.

What is happening in 1996? In fact, the PC market continues strong, 19-20% unit growth. Even with the inventory problems we had earlier this year, we are still seeing a growth in the bit consumption. It is growing again, about 67-68%, over 1995. So the market from that standpoint is still looking very strong. What has changed is that we have gotten back on the normal price per bit decline curve and done it very rapidly.

If you look at this chart, the left side shows a by month price per megabit for DRAM and you can see how consistently it has maintained, resulting in about a $25/megabyte price in 1995 and a very rapid decline since. Because of that constant price per bit in 1995, we saw a continual growth, month by month, throughout 1995 in revenues, peaking at the end of the last quarter of the year. With a rapid decline, of course, no amount of bit increase in 1996 is going to be able to hold revenues up.

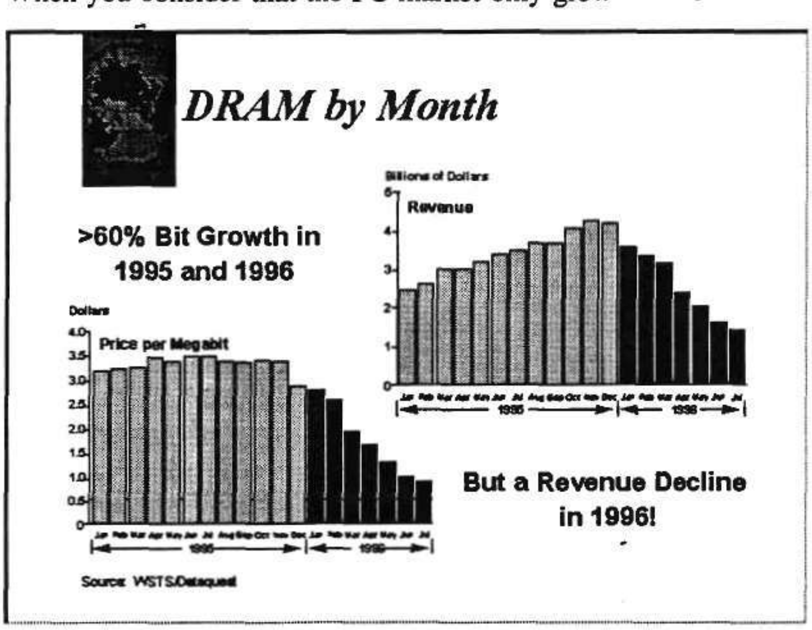

With this are a lot of expectations. We expected Windows 95 to drive the demand for memory. Windows 95 is doing well, but Windows 3.1 is still the operating system of choice, so it took off a little more slowly. We expected very strong PC shipments in the fourth quarter that were not realized. The market was already slowing down, as mentioned earlier. That led to the excess inventories in early 1996, inventories that we had to work off. Coupled with the inventories dumped into the marketplace, we had a 4 to 16 megabit transition and that freed up a great deal of capacity. Because the area per bit in a 16 meg device is much less than for a 4 meg. That added to the capacity that was already coming on line, and we got into a capacity glut in 1996. To add a little insult to injury some of the markets have slowed down, the PC market probably the least, but some of the other markets have come along and impacted other products, if not DRAM.

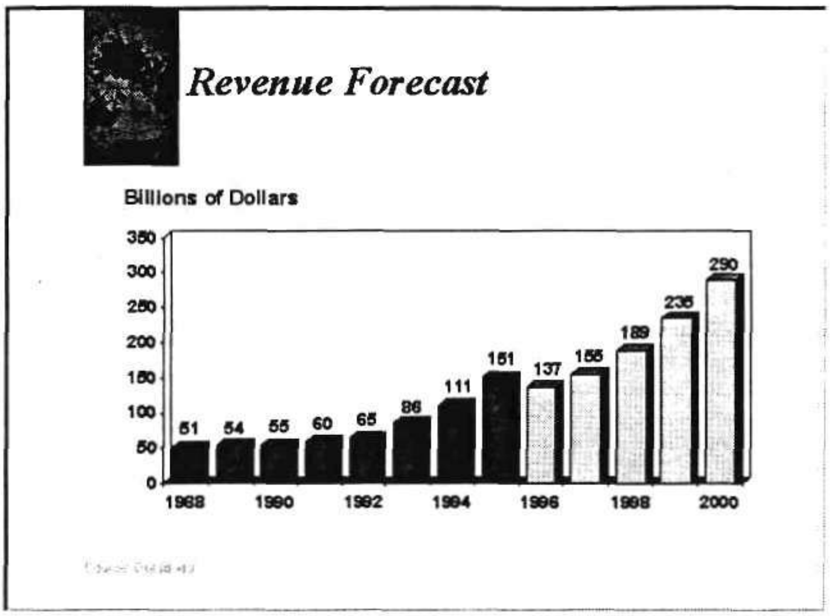

How severe is the downturn? This slide shows 4 different periods of MOS memory downturn. It plots two years by month. You see that 1996 is as bad as 1985, it is a very severe turndown, very rapid. In fact, earlier this year, and at the end of last year, Dataquest was expecting a softer downturn, one more like 1989. So our forecast was a little high, to say the least, the market turned down far faster and far deeper than expected. The white line that continues is essentially our forecast for the next few months of the year.

So this is how the revenue plays out by year for the next five years. 1996 will see a 9.4% revenue decline, $137 billion. Then we will see a 13% rise in 1997, and from then on, actually, we get back into a very strong growth phase, where we will see better than 20% growth in the final three years of this forecast period.

How did our forecast change in the last six months? We had come down substantially in our April forecast, but we were not expecting the deep downturn that I just showed you a few slides ago. It has been much steeper; we were expecting a flat revenue growth. We are expecting this year a 46% revenue decline for DRAM. That translates to a change in our forecast of about 11%. The slowing of markets, especially the industrial and consumer market, has additionally impacted the non DRAM products more than we expected. That is another 6% that has changed our forecast. So that is the 16% swing in six months in our forecast.

This is how the forecast plays out, by quarter, over the next couple years. We are just starting the rebound. On this chart, we expect to see about

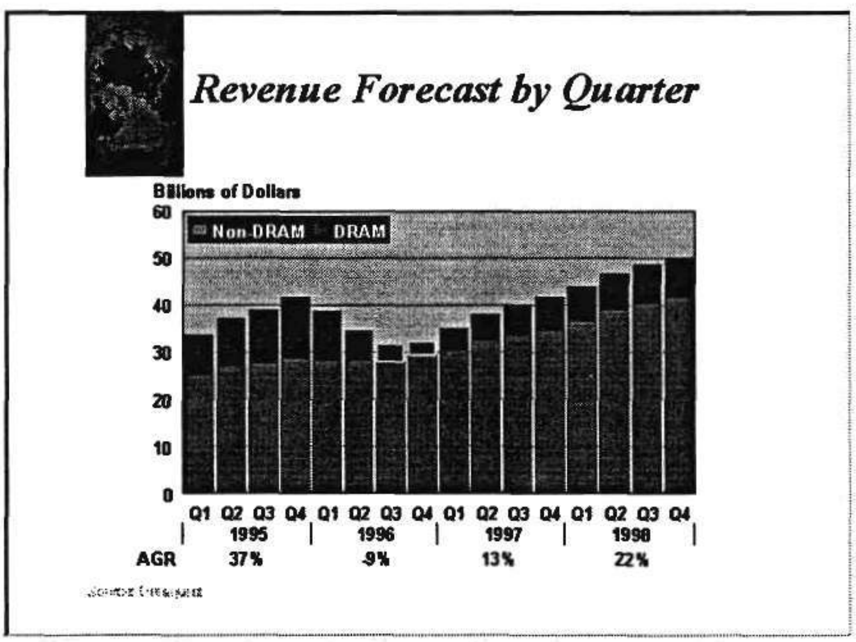

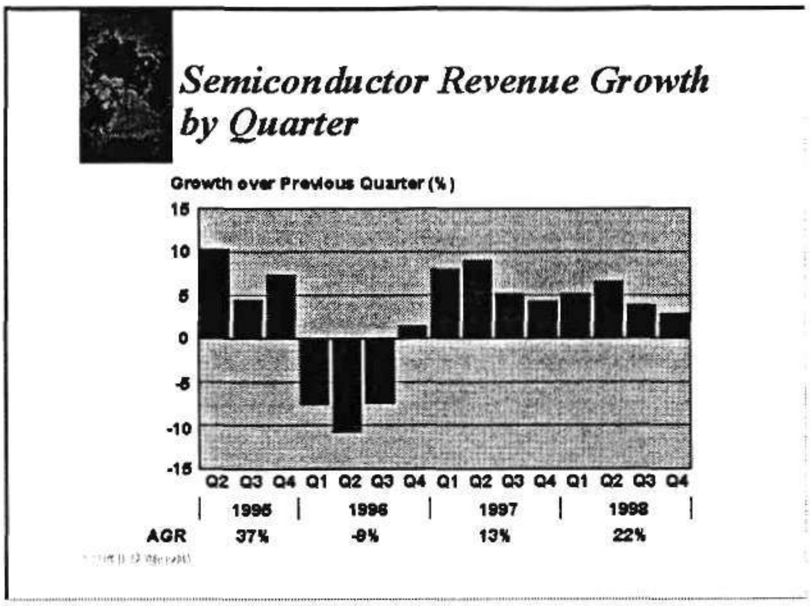

a one percent increase in revenues in the fourth quarter of 1996. If you look at the bottom portion of the bar, you can see that non DRAM products actually have been continuing to grow – that is a 5% growth in 1996. Certainly it is a slowdown over what has been done in the past, but it is a continuing growth that will rebound also, as the markets strength in the coming years. So what we are seeing is largely a DRAM driven downturn, and the size of it is substantial. We are at the bottom, we are on the way up, and we are very optimistic about this forecast.

This is how it looks on a quarter by quarter growth comparison. We are looking at about 1% in the fourth quarter of 1996, and then about an 8% growth in the first quarter of 1997.

By product: there is a very large downturn in memory. The bright spot here is the microcomponent category, especially driven by microprocessors and microperipherals, two products doing very well on the 20% PC unit growth in 1996. Those products are more impacted by the declines in the consumer market and industrial market. The microcontroller portion of microcomponents is also impacted by that, which pulls the microcomponent IC growth down to what might be 18-19% through the 14% level. We are expecting all the products to rebound into double digit growth in 1997.

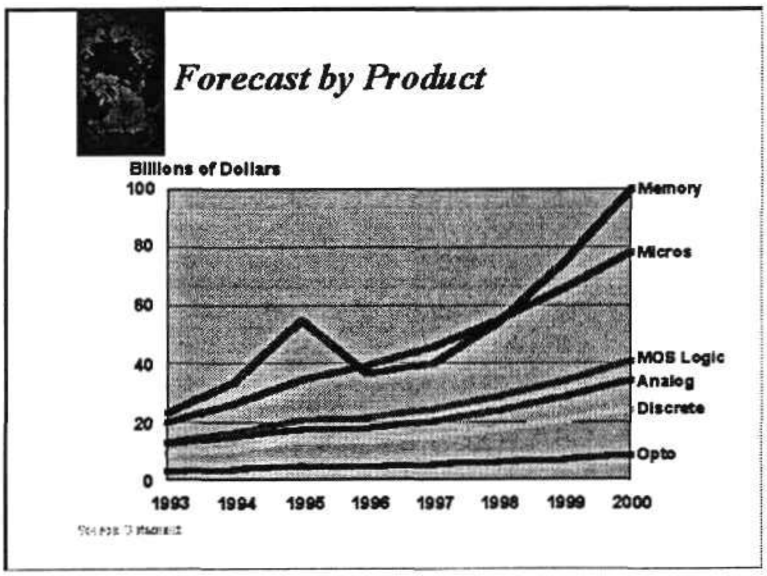

These bars compare the five-year growth forecast versus what was the five-year historical growth. Excluding memory and microcomponent ICs – pretty similar growth. Microcomponents is settling into a very decent 18-20% growth over the forecast period. The market can support that kind of continued growth for that category. Memory looks like it is down quite a bit, and it is, because of this pothole of 1996. But if instead of looking at the 1995 to 2000 compound growth, we look at from where we are standing now, from 1996 to 2000, actually our forecast is a 29% compound growth over four years. We think memory will start to come back. Certainly we are going to have a capacity glut for the next couple of years, but in 1998 we start to see very strong revenue growth, and memory coming back, strongly contributing to semiconductor revenue and growth.

Looking at it by region, all regions share the misery in 1996 – Japan more so because of the change in exchange rates. There is simply fewer dollars per yen, so that accelerates the current downturn in Japan. We see very good growth in 1997 by region as most of the regions show double digit growth.

A rebound is not a single region rebound, but the market will return in all products in all regions. What is different about this forecast is the change in Asia-Pacific growth. Asia-Pacific growth drops virtually in half over the forecast period. There are a number of reasons for that. The market is maturing, it is becoming a more significant part of the total market, and can not continue that kind of growth. Another factor is a lot of the Asia-Pacific growth was due to Japan manufacturing moving into Asia-Pacific sites. With a weaker yen, that manufacturing movement, that migration, is slowing down. We do not expect to see as much migration to Asia-Pacific. The market itself is getting larger, and becoming a more significant portion of the total market. By the year 2000, the Asia-Pacific market will be one-quarter of the total of the worldwide market. So it cannot continue this kind of 35-34% compound growth when it gets to be that size.

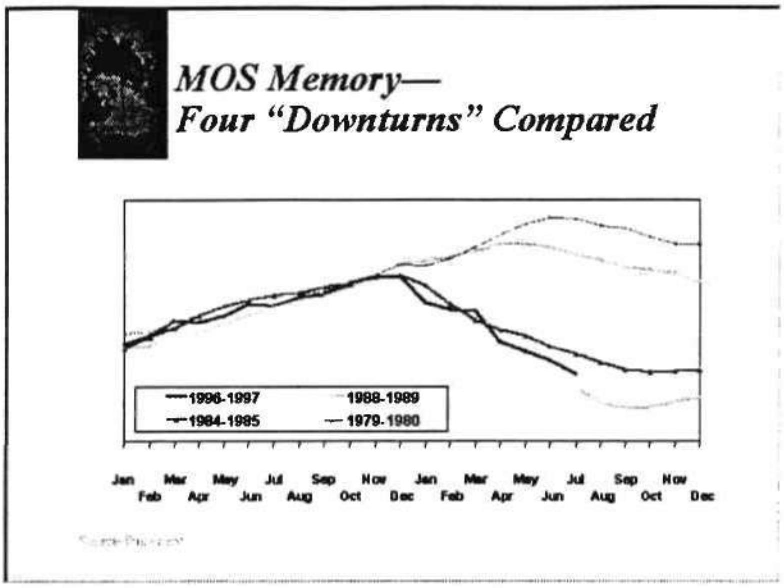

Semiconductor Forecast 1996: What Changed?

	Component (%)	Growth or Change
April 1996 Forecast		*+7.5*
	DRAM Prices	-10.9
	Yen/Dollar Rate	-0
	Non-DRA Products	-6.0
October 1996 Forecast		*-9.4*

In the past, we have forecast that Asia-Pacific will grow and surpass Japan within our forecast period. This forecast differs from that we see a little slower growth in Asia-Pacific than previously forecast and a consistent market share for both North America and Europe, 33% and about 18% respectively. If you take Japan and Asia-Pacific as a unit, they remain about 48-49% of the total market over the forecast period.

In summary, we have seen another DRAM cycle. We had hoped that we were not going to see another one, and that we were not going to see one as deep as we had seen in 1985. Nonetheless, it occurred, and we are seeing a very deep dive in 1996. The good news is, we are on the way up. We will see a little bit of revenue growth in DRAM in 1997, non DRAM products have con-

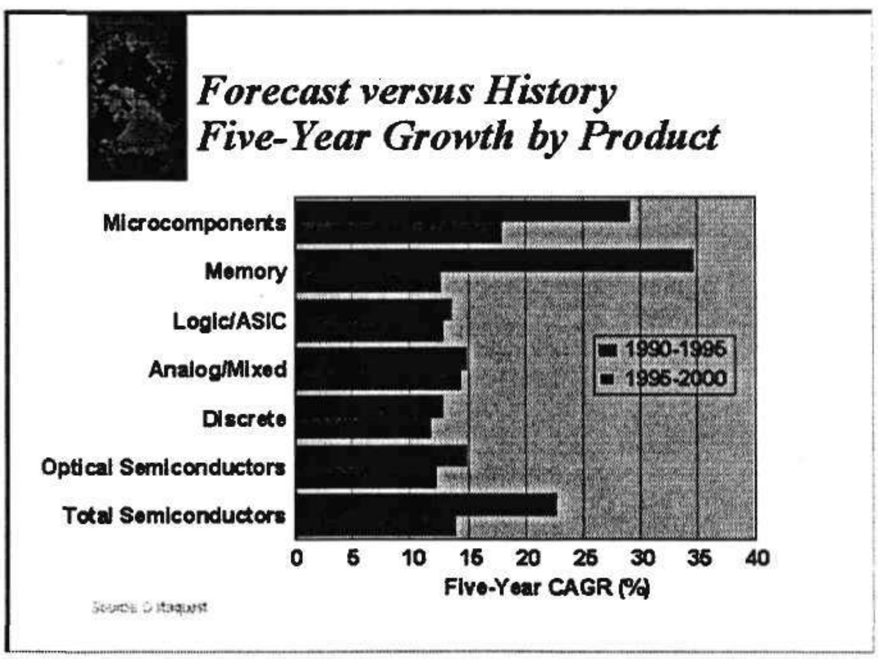

Regional Consumption as a Percentage of Total

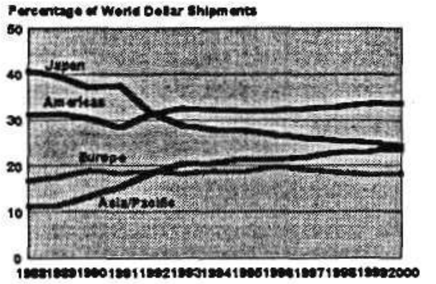

Regional Revenue Forecast

	1996 Billions of Dollars	1996 Percent	1997 Billions of Dollars	1997 Percent
Americas	44.2	-8.6	50.5	14.2
Japan	36.6	-13.3	40.5	10.7
Asia/Pacific	29.1	-10.3	34.0	17.0
Europe	27.1	-4.3	29.6	9.2
Worldwide	137.0	-9.4	154.6	12.9

tinued to grow strongly, and recovery is starting right now. This transition of memory and micro-component ICs taking more and more market share is probably going to slow down. Some of the products delineated, the advanced consumer etc., will demand more discrete, more analog, so we see a rebound in some of those markets. We are looking at the convergence of computers, and consumer entertainment, and communications. The market is changing but the market is looking good. We are looking at $290 billion in the year 2000. A little below our expectations in the past but still very optimistic.

GENE: The next speaker is Clark Fuhs. He is Director of our Semiconductor Equipment Materials and Manufacturing, as well as our semiconductor contract manufacturing service.

CLARK: For those of you who are not familiar with our service, we tend to cover in detail all of the equipment segments associated with front end manufacturing – basically anything having to do with the fab. Turning silicon into money, and translating money into silicon. We do have a lot of detailed forecasts available.

Supply side economics has often been called voodoo economics, and occasionally people look at me like I am talking voodoo, but we will try to bring these issues down to some simple concepts for you. The capital equipment industry is currently in a downturn, driven by, of course, overcapacity.

I thought it would be interesting to show our forecast conclusion slide, from a year ago at our mid-year seminar, unedited. The question that we raised was: Are we ready to call the silicon cycle dead? Our answer was: Yes, yes and no. It had to do with the three areas that have typically been associated with spending downturns in the past. The electronic equipment market, which is normally caused by a worldwide economic recession, we did not see at that time, and we still do not see that.

Dataquest Conclusions from SEMICON/West 1995 (July)

- **Worldwide economic recession not on horizon**
- **PC unit shipments should continue to grow at a 16% CAGR**
- **Silicon cycle is not quite dead**
  - **Narrowing scope to DRAM business cycle**
  - **Magnitude of cycles should diminish over time**
- **Watch for clear signs of transition to 16Mb**
  - **1x16 configuration reaching 60 to 65% yields**
  - **DRAM price index falls**

A PC unit slowdown was not expected, based on the drivers that we were seeing. However, the silicon cycle was not quite dead because we saw that it was going to be narrowing to the DRAM cycle. Something that we had mentioned a year ago is that the magnitude of the cycle should diminish over time. I know that is the conventional wisdom, but I am beginning to question that wisdom. The things that drive the cyclicality in the business happen to be the DRAM, and that is concentrating in our industry, it is not diminishing over time. In the early 1980s, DRAM was roughly in the high single digits, as far as percent of the market. Today it is in the mid-20% range. Semiconductor consumption within the PC was under 15% just five years ago. It is about a third of the semiconductor market today. That is the cyclicality driver in our business and I do not think it will diminish in severity, at least in the next cycle, beyond this one.

Last year, we were going to watch for clear signs of the transition to the 16 meg DRAM, and 60-65% yields. The question that I have gotten lately is: How can this industry go from undersupply to oversupply so quickly? There are a couple of reasons for this. The first one is supply driven. We were interested in the timing of the transition, going from 4 to 16 meg DRAM because that creates silicon efficiency. The key is bits per square inch. If I have a 4 meg DRAM that has an average die size of 50 square millimeters, and I have a 16 meg DRAM that has an average die size of 80, I can increase the bits per unit four times and only consume 60% more silicon. So the bits per square inch goes up by a factor of three just by doing conversion from 4 to 16.

When the industry is stuck at a product unit density (such as a 4 meg DRAM), bit demand, which averages about 60% a year, tends to drive silicon very directly. Typically this industry is in undersupply. When we have a typical 60% bit demand, silicon demand increases by roughly 35% into the DRAM sector. This creates the need for more equipment and a capacity-driven spending cycle. Capital spending cycle slows in the bottom half of the chart, during times of bit oversupply and the industry migrates to the next generation of DRAM. We are in the lower right hand corner of this circle today. We just crossed the time when the industry is taking that step function jump upward in bit capacity. During these times where you have an average bit capacity of 60%, silicon demand is flat to down. We did some modeling, and starting in the third quarter of this year, for three quarters, wafer start requirements will decline going into the DRAM sector, with a bit demand of about 70% of the market for growth. The complicating factor of the increased expectations is that we have a double barrel supply and demand-driven slowdown. This has increased the severity, hopefully not the length, of the downturn. We think that this downturn is going to be shorter that the average two-year slowdown, and the reason for that is the strong end use markets.

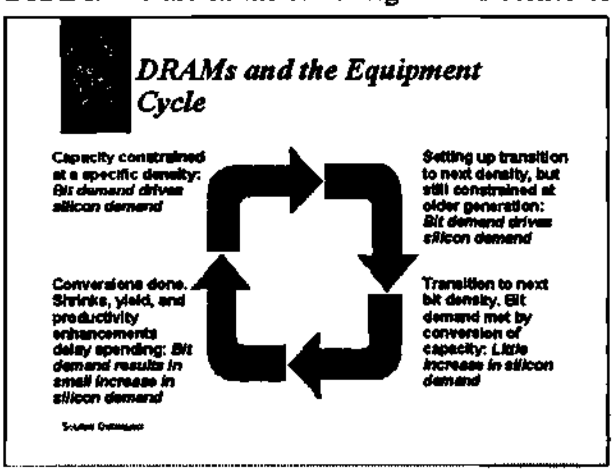

The past series of slides basically explains the difference in our forecast today versus a year ago. The two forecasts basically have the same profile, as we built in a supply-driven DRAM slowdown in '96 into '97. However the better than expected 1995 capital spending picture is leading to a more difficult 1997. A 1.5 to 2 year pause in the equipment industry in 1997 and 1998 is still in our outlook. The next major growth in spending will start in '98, into '99, driven by capacity.

1996 is a year which turns from growth to decline, but will net out with about 17% growth this year. The only reasons 1996 remains a double digit growth year, are that there were strong backlogs coming out of 1995, and a few companies such as IBM, TI, Siemens, and some of the dedicated foundries, have continued to grow investments in the first half of this year. Korean companies have also grown, but in moderation.

A Look at the Previous Wafer Fab Equipment Forecasts: Little Change

	July 1995 (Percent)	July 1996 (Percent)
1995	52	77
1996	21	17
1997	-5	-16
1998	2	6

- 16Mb DRAM transition occurring
- DRAM pricing has fallen
- A better 1995 means a more difficult 1997
- Only a modest recovery in 1998 seen

The Japanese companies have already cut back, but they are still investing in shells. Shells are where a fab is built, you build a building and you put in a minimum equipment complement in order to test the process, setting up for the next generation. So the number of fabs will remain inflated, as there are more shells being built right now, but the equipment will not go into those fabs at nearly the rate that was anticipated originally. DRAM companies everywhere are currently cutting back the budgets. We do not think that 1996 spending levels can be maintained. In particular, we expect in 1997 that investment into the DRAM area for equipment will be cut as much as 30%. Asia-Pacific and Japan will be hit the hardest, as these are the DRAM production centers. We do expect spending on advanced logic capacity to be maintained and begin growing toward the back half of next year.

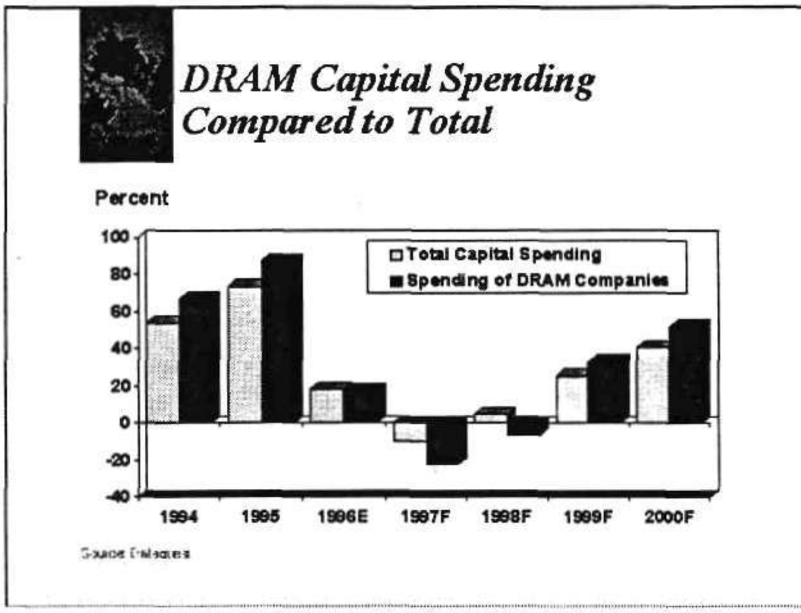

In general, companies will continue to invest in 0.25 micron technology and equipment, in order to gain experience for the shrinks that are about to come later in the decade. So we have just gone from a capacity-driven upturn to a technology-driven spending pattern.

Another question is: Are there too many fabs being built? When we analyze the amount of silicon (in terms of square inches), that is required in order to meet the near $300 billion market in the semiconductor forecast, there should be roughly about 35 new fabs per year being built. Currently by today's announcements there are too many fabs coming on line in the next two years. We would expect several more delays, or cancellations of plans, to emerge into '97. Many of the fabs starting in '96 and '97 will remain shells, with a minimum equipment complement. These fabs will not likely be fully equipped until 1999, meaning the number of new fabs starting out later in the decade will actually decline a little bit, but equipment spending will increase.

What our analysis also says is that there needs to be more new fab announcements to occur before the end of the decade. We have the distinction of having the only forecast which has a plus 40% or more growth year in the time horizon. We are a DRAM driven industry and therefore the swings in spending for DRAM, the capacity will be more volatile. In this chart we highlight the expected changes of DRAM companies shown against overall capital spending growth. The bars on the left represent the growth rate in capital spending for '94 into the year 2000. The bars on the right are the growth rate of spending for the DRAM companies specifically (TM, Micron, the three Korean companies, the top six Japanese companies and some new Taiwanese players). You will notice that we expect this group to lag the spending recovery, meaning DRAM is viewed to be the last segment to be coming out of oversupply.

Our forecast for front end wafer equipment profile follows the profile noted on this slide. The recovery pattern starting in mid-'97, and progressing through '99, will depend on how today's overcapacity will trickle through the industry. After the wafer fab equipment market contracts in 1997, it stabilizes in 1998 to about the '95 levels. We do see resumed strong growth, as the need for the capacity will again become the driver, leading to an over $40 billion market in the year 2001. This profile is essentially the same as the semiconductor forecast, but lags by one year.

In order to quantify and characterize the dynamics of the near term market, we have established a top line quarterly shipment forecast for wafer fab equipment. We expect a decline to accelerate into 1997 and should perhaps bottom in the first half of the year.

The next slide shows the sequential quarterly growth rate, which is derived by taking the current quarter revenue and comparing it to the previous quarter. For example, the first sequential decline quarter is Q3 of 1996, at -3%. This means that the revenue in Q3 was 3% below the revenue for Q2. These are also seasonally adjusted.

Our current forecast shows the first sequential growth quarter will be the fourth quarter of '97. The quarterly outlook has produced the annual forecast where we show 1997 declining by nearly 16% in 1998 as a single digit growth year. But how will the industry recover? What are the anticipated dynamics of how this recovery is going to take place?

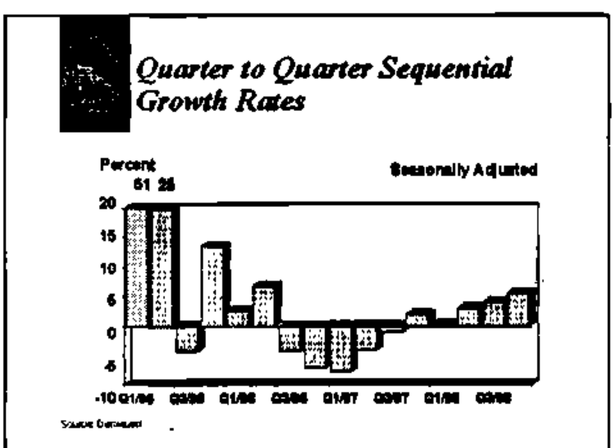

The industry is currently experiencing a DRAM oversupply coupled with a product transition. In order to determine how capital spending may recover, it is important to understand how this excess capacity will migrate, or trickle, into other areas. In general, capacity can be split up into three major areas in the semiconductor manufacturing infrastructure. The first is DRAM, the second is advanced logic, and the third is power discrete. Capacity is very fungible, or very changeable, going from product to product, so it is really impossible to do a supply/demand issue on E-proms or Flash, or SRAMs, or ASICs, or any other specific advanced logic area, because a stepper does not really care what kind of chip is on the wafer, all they care about is what the wafer is, and the wafers that are going through.

There are two general blocks of capacity now available from today's conditions. These two blocks are being redirected into other semiconductor product areas today. Based on how this capacity is migrating, the first area of spending recovery will be in the advanced logic area, as early as mid-'97. So in '97, metal etch will perform better than polysilicon etch. The microcontroller, analog, mixed signal, and telecom chip capacity will mean extra recovery, but not until the end of '97 and into '98. The DRAM segment is expected to be the last, probably later in 1998.

The first block of capacity that has become available is comprised of old format DRAM fabs that cannot run 16 meg chips. These are limited to two-level metal, and are at 0.6 to 0.8 micron technology. Microcontrollers, telecom chips, consumers chips, mixed signal, analog ICs are using this

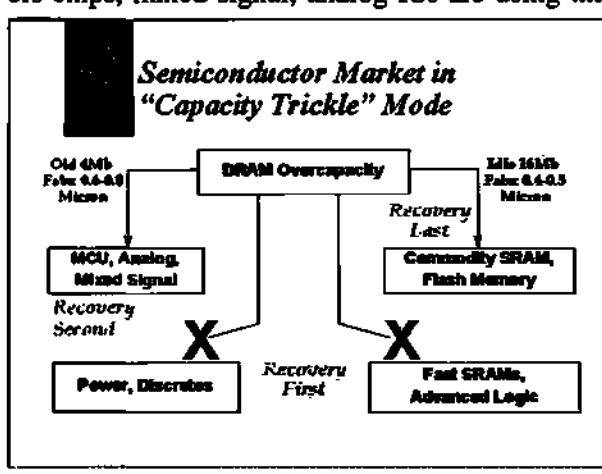

capacity. Right now it is likely that most of these fabs in Japan, and some in Korea, will migrate to these or have been migrating to these product areas.

Power and discrete chips have very specialized processes not found in old DRAM fabs. So these segments are relatively isolated. It is not impossible to convert a DRAM fab to run power discrete, but an investment into some additional process flows is needed, and a learning curve in order to convert these fabs. We expect this capacity to be relatively isolated, but we expect capital spending patterns to be closely tied to demand in these product segments.

The second block of capacity is comprised of idle, or underutilized, advanced 16 meg DRAM capacity, which is limited to two-level metal, but is at 0.4 to 0.5 microns. Because these fabs generally lack the process sequences of self-aligned silicide, which is required for advanced logic, and three levels of metal or more, they cannot effectively, immediately, be redirected into advanced logic or fast SRAM. Therefore, they are limited to commodity SRAM, Flash, and other non-volatile memory, or a limited span of logic products that do not require three levels of metal. We expect most of these plants to actually remain somewhat idle, and we have gotten confirmation that wafer starts into DRAM factories today are actually declining.

We believe the first areas of spending recovery will be in the advanced logic area, since this will be one of the more isolated places, and equipment companies positioned for these markets will have more moderate slowdowns. The companies that are positioned for the DRAM factories are going to have more severe slowdowns. Microcontroller, analog and mixed signal, and telecom chip capacity will be next to recover, and Motorola, among others, will be a key company as far as a spending pattern recovery here. The DRAM segment, the root cause of the problem, is expected to resume robust spending at the very last.

In summary, 1996 is a transition year for wafer fab equipment, industry contraction. The industry should bottom and start recovery as early as mid-1997. But the downturn is severe and likely to be shorter than normal, because the PC unit growth rates are very healthy, and very strong. We do expect a resumption of capacity spending to come back in the 1998 timeframe.

Q: Since forecasting the specific level of the market in 2000 is quite a difficult task, what is the envelope, or the upside or downside, of your particular forecasts, and some of the issues affecting that upside or downside?

GREG: In my area, we are talking a trillion dollar market by that timeframe so the variability is going to be fairly small. Some of the key fundamental economic assumptions are the developing economies of the world demanding more electronics, and the ongoing upgrade of information technology in the rest of the world. Plus or minus 10% from the demand perspective is reasonable.

CLARK: DRAM represents a factor that could take the forecast down 10% by the year 2000. What is important is what is going to happen to DRAM pricing, and how will it stabilize? Will we see another period where the price essentially stagnates and remains constant as the demand catches up with capacity? We are expecting DRAM pricing to not decline, but if it were to decline at a faster rate than has been anticipated it could take the forecast down substantially.

GARY: With regard to the equipment manufacturers it really depends upon how the DRAM market plays out. Not so much the level of DRAM, or the demand of DRAM, it is the timing of the next product migration. If the product migration to the 64 meg happens sooner rather than later, there is more downside than upside. If the lengths of the cycles are actually getting longer, then there is more upside than downside. If you lengthen the cycles for the product migration, you are going to flatten the curve for cost per bit decline, which means more capital spending per dollar revenue. That is what will present some upside. The reason we are in a slowdown today is in part, due to the 4 to 16 meg transition. If going to 64 meg happens quicker, there will be another block of silicon efficiency, and capital spending will be subdued. That is good for profitability of the DRAM manufacturers, and it is not good for the equipment industries. It is in the best interest of both those to push out the 64 meg transition.

Q: What potential new technologies could change your forecast up or down over the next four to five years? Do you see any dislocations coming?

GREG: We are going to see some exciting possibilities coming out of displays, and how they are integrated into other electronics systems, potentially Web browsers. The other area is battery technology. It is probably 100 years old and you still cannot get your laptop computer to last more than three hours because we keep raising the power requirements on these as they put faster processors in. In general the whole issue with batteries and also energy – the green movement – is going to become the driving factor, globally, that will impact new systems that need to be designed, and likewise the types of chips that will go inside these systems.

GARY: In terms of the devices certainly there are things in the end equipment – the convergence of PCs and television and entertainment. But the semiconductor side will be a winner either way; the semiconductor content will grow along with it. The 64 meg DRAM is a transition that could impact the forecast substantially. That kind of a transition and the technology in that area could affect it. System on a chip could impact certainly the way these products are categorized so that could change very dramatically. You have to look at the market in terms of the end applications, or the systems, that define it. That can change the by-product forecast substantially and could impact the total market also. We see probably the biggest potential for change in the DRAM area.

Q: Do you see any dislocations coming as a result of forward integration trends?

GREG: I do not see any dislocations, I think those will complement. We will have to migrate from one particular product area to another, but in terms of the revenues that it is not going to be a substantial difference. Analog for instance, is an area that gets cannibalized by other product types. Certainly analog is a functional block, and various kinds of functions do not disappear, but you find them in a lot of microperipherals.

CLARK: As far as the manufacturing issues are concerned, I can see one demand-driven and two infrastructure-driven issues. One on the demand side is the system on a chip, the very nature and definition of that whole realm means that you are going to silicon a little bit more efficiently. So as that trend begins to emerge there might be some issues with the volume requirements. It should also spur demand so it is going to be a plus and minus issue. The two infrastructure issues have to do with the first being the foundry industry and the fabless foundry model. Right now about 10% of the world's capacity is being served by the foundry industry and maybe in fifteen years, go to 25 or 35%. There is concentration of capital. The fabs are getting bigger, but there is also a more efficient use of a piece of equipment. The supply is being used much more efficiently, and that is why that whole issue will win. About 15% of all capacity that is being added this year is at 0.8 micron and above. So there has been a new emerging segment of the equipment industry to specifically address the lagging edge technology markets, and this will be an interesting one to watch in the future.

Q: Wilf said he did not see much of a slowdown in the demand in 1996. At the same time you had a reduced growth rate for electronic equipment production, does this seem kind of contradictory? Could you explain how you could have both at the same time.

GREG: In 1996 we had production slowing, but half of that slowing is the exchange rate, the dollar appreciating against the yen, and the other half was due to true slowing of the growth rate of systems production. The PC market grew at 25% in the previous year on a unit basis globally for shipments of PCs. That slowed to 20% growth this year. That is 20% difference in growth rate, and put another way, it translated into planning for inventory levels that were much higher, particularly in Q4 of last year. That all had to be worked out of the system, even though everything was in a positive growth mode in the first couple of quarters. Ordering of chips that go into PCs were lowered strictly because of inventory adjustment, because of a lower, albeit very positive, growth rate, so that type of factor contributed to the year as well.

Q: Can you give us some idea what you think about long term in the compound growth rate per units?

GARY: The IC unit compound growth rate is about 9.9-10%. It tends to run 3 to 4% lower than the revenue growth.

Q: We have seen a slowdown certainly in the dollar portion of the equipment industry and the chip industry, and you talked about the capital spending still increasing by 16-17%. How does that equate with slowdowns in the other part of the chain?

CLARK: There is a lag time of 9 months associated with when we see demand fall off, and actually see spending patterns cut back. That happens for a couple of reasons. One is there is just simply too much momentum. You are building a fab, you are installing factory capacity, and if you responded in capital spending to every little ripple that happened, you would be jerking around your capital spending plans for a while. The second issue is that capital spending does tend to be somewhat emotional. Once you start a downturn there is a continuation of spending because of that. Spending cuts get delayed but they happen, and they happen at a much more severe pace than the end use markets do.

Q: In terms of the fab, is there a relationship between the back end to the front end? What do you think the growth rate may be, and what is the implication for back end vendors of equipment.

CLARK: We do not have a lot of specialty in the back end versus the front end. I have made a cursory look at some of the issues to determine whether or not front end was leading back end, or back end was leading front end, or whether or not there was any correlation at all. The biggest correlation I can come up with, in terms of the test market, the test market is driven by design starts. In terms of test equipment, once you start yielding a part, the test equipment would tend to slightly lead the front end equipment. In the downturn in '95-'96 that indeed did happen. There was a lead. The front end market is much bigger than the test market, so to presume that there is a strong correlation is really not the right thing to do. If anything, test leads, the test is more relevant to design starts. The mask business is booming right now. Equipment companies that are servicing the mask business are booming right now. The mask business is driven by design starts, so as the applications proliferate, and the number of applications and the chips get more complicated, that end of the market will just continue to have fairly long term growth. The test equipment market should grow at relatively the same rates to the front end equipment market. Capital spending, as a percent, should remain relatively the same. There are more complicated issues emerging in the test area that make me believe the ASPs might go up faster than the front end.

Luncheon and featured speaker
Will the Internet Break Microsoft?

Roger McNamee

General Partner,
Integral Capital Partners

Mr. McNamee is a general partner of Integral Capital Partners, a family of private investment partnerships that invest in the securities of expansion stage private and growth stage public companies in the information and life sciences industries.

Before the formation of Integral in 1991, Mr. McNamee worked for nine years at T. Rowe Price Associates in a variety of research and portfolio management positions before the founding of Integral. From September 1988 through August 1991, he was portfolio manager of the top-ranked T. Rowe Price Science & Technology Fund, a position that included responsibility for the firm's investment strategy in technology. From October 1987 through August 1991, Mr. McNamee was comanager of the T. Rowe Price New Horizons Fund, a diversified emerging growth fund with assets of $1.3 billion. As a member of the T. Rowe Price equity research department, he covered such industries as computers, software, aerospace, telecommunications, and office supplies.

Mr. McNamee is a frequent speaker at industry and investor conferences and is a member of the Marketing Committee of the Nasdaq Stock Market.

Mr. McNamee has a B.A. degree from Yale University and an M.B.A. from the Amos Tuck School of Business Administration at Dartmouth College. Mr. McNamee is a Chartered Financial Analyst.

Session #6: Luncheon Presentation - Will the Internet Break Microsoft?

Roger B. McNamee
General Partner, Integral Capital Partners

TRANSCRIPTION FOR THIS SESSION IS UNAVAILABLE.

Panel:
Emerging Opportunities in the Battle for the Local Loop

Moderator:
Gregory L. Sheppard

Chief Analyst
Semiconductor Application Markets Program,
Semiconductors Worldwide,
Dataquest

Mr. Sheppard is 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 (MCSAM), and the Communications 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.

Mr. Sheppard is a specialist on the end use or application of semiconductors. His scope of analysis includes both economic and technical trends regarding the semiconductor opportunities in electronic equipment.

As an analyst Mr. Sheppard's current focus includes LAN, WAN, and modem markets; telephony and public infrastructure markets; and set-top boxes.

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. Before that, he worked in engineering management at GTE Corporation specializing in communications systems design and decision aid systems.

Mr. Sheppard received a B.S.E.E./C.S. degree from the University of Colorado and an M.S. degree in system management from the University of Southern California.

Panel:
Emerging Opportunities in the Battle for the Local Loop

Robert L. "Bob" Bailey

President and Chief Executive Officer,
PMC Sierra

Mr. Bailey joined PMC Sierra of Vancouver, British Columbia, as president and CEO in November 1993. PMC Sierra is a leading maker of broadband networking integrated circuits. In September 1994 he helped facilitate the merger of PMC Sierra, then a privately funded company, with Sierra Semiconductor as a wholly owned subsidiary. PMC Sierra has gone from $8 million of revenue in 1993, when Bailey joined the PMC Sierra team, to approximately $40 million in 1995. He is also presently a director of Teltone Corporation in Bothell, Washington.

In 1979 after graduation, Mr. Bailey began his career with Texas Instruments in Dallas, Texas, where he worked for 10 years. Most of his tenure at Texas Instruments was spent within the Semiconductor group in a marketing responsibility where he launched several product lines, which accounted for over $2 billion in revenue. In August 1989, he accepted a job in AT&T's Microelectronics Strategic Business Unit as vice president of Application Specific Integrated Circuits. This position was one of several that was created to assemble a turnaround team to help transform AT&T's Metal Oxide Semiconductor business unit into a merchant market success from an unprofitable captive supplier. Mr. Bailey's division went from an unprofitable $100 million of revenue in 1989 to over $500 million in 1993 and was solidly profitable when he departed. Additional product lines were started within his division such as the Hard Disc Drive ICs, the "ORCA" FPGA, and the Fast Ethernet LAN IC businesses. AT&T-ME's business is now a renowned merchant market semiconductor company with the ASIC business accounting for over 50 percent of AT&T-ME's revenue and retaining the No. 1 market share position in the competitive Standard Cell ASIC market.

Mr. Bailey received a B.S. degree in electrical engineering from the University of Bridgeport, Connecticut, in May 1979 and a master's degree in business administration from the University of Dallas, Irving, Texas, in 1983.

Panel:
Emerging Opportunities in the Battle for the Local Loop

Russell K. Johnsen

Vice President and General Manager,
Communications Division,
Analog Devices

Mr. Johnsen is the vice president and general manager of Analog Devices' Communications Division, headquartered in Wilmington, Massachusetts. He is chartered with managing the youngest and fastest-growing division within Analog Devices. The Communications Division was established in 1994 to address opportunities in digital wireless communications and in broadband telecommunications markets.

Before joining Analog Devices in 1993, Mr. Johnsen was the general manager of National Semiconductor's Wide Area Networks Division. He has also held key management roles in wireless communications, microprocessors, and in Southeast Asian manufacturing operations.

Mr. Johnsen holds an S.B. degree from MIT.

Panel:
Emerging Opportunities in the Battle for the Local Loop

Stephen J. Edwards

Assistant Vice President,
Telco Broadband Networks and Applications,
Northern Telecom

Mr. Edwards is assistant vice president responsible for the sales and marketing of Nortel's Telco Broadband Networks and Applications products and solutions to North American telephone companies as well as global responsibility for the business management and development of broadband access solutions.

Mr. Edwards worked for British Telecom in access network planning until 1977 when he emigrated to Canada to join Bell Northern Research. There he was involved in the early deployment of fiber-optic technology into Bell Canada's network. In 1982 he joined Northern Telecom (Nortel) and has worked in a variety of transmission, switching, and access product management, marketing, and sales roles in Canada, the United Kingdom, Holland, and the United States.

Mr. Edwards was born in England and educated at Chatham House Grammar School and the University of Surrey, where he graduated in 1973 with an honors degree in electronic engineering with economics.

Panel: Emerging Opportunities in the Battle for the Local Loop

Glenn H. Estes

President,
Telesis Technology Laboratory

Mr. Estes was appointed president of Telesis Technology Laboratory in April 1994. Telesis Technology Laboratory (TTL) is a subsidiary that provides advanced technology support to all units of Pacific Telesis group and pursues exploratory work on strategically important emerging technologies.

Mr. Estes is also a vice president of technology in the Corporate Strategy and Development organization at Pacific Telesis group. Before joining the company, he held various positions in technology development at Bell Laboratories and Bell Communications Research.

Mr. Estes led the national and international standardization of the Synchronons Transfer Mode protocol. He is past chairman of the Technical Committee, vice president of technical affairs of the ATM Forum, a worldwide, cross-industry consortium. In addition, he is also a past member of the board of directors of the IEEE Communications Society, an international telecommunications technical society. He is also the U.S. technical representative to the UN-sponsored Telecom Forum technical conferences.

Mr. Estes was awarded a bachelor's degree in electrical engineering with highest honors from New Mexico State University in 1972, a master's degree in electrical engineering from Polytechnic Institute of Brooklyn in 1973, and completed a computer science curriculum at Bell Laboratories in 1980.

Panel:
Emerging Opportunities in the Battle for the Local Loop

Lisa Pelgrim

Industry Associate, Networking Program,
Telecommunications Group,
Dataquest

Ms. Pelgrim is an industry analyst within the Networking program of the Telecommunications group at Dataquest. She researches areas of wide area data networks including modems, cellular WAN equipment, and digital network equipment. Ms. Pelgrim is also involved in conducting primary research, consulting, and customized research products. She is recognized as an expert in the modem market and has been cited in various publications including *Communications Week, Computer World, Infoworld, Forbes, LAN Times, Network World, Newsweek International, PC Week,* and *PC World.*

Before joining Dataquest, Ms. Pelgrim was a member of the marketing organization at Time Electronics, a nationwide semiconductor and electronic component distributor. Her marketing efforts focused on telecommunications and electronic companies located within the Silicon Valley.

Ms. Pelgrim holds both an undergraduate degree in finance and an M.B.A. with an emphasis in market research from Santa Clara University.

Session #7: Emerging Opportunities in the Battle for the Local Loop

Moderator:
Gregory L. Sheppard
Chief Analyst, Semiconductor Application Markets Program, Semiconductors Worldwide, Dataquest

Panelists:
Lisa Pelgrim
Industry Associate, Networking Program, Telecommunications Group, Dataquest
Glenn H. Estes
President, Telesis Technology Laboratory
Stephen J. Edwards
Assistant Vice President, Telco Broadband Networks and Applications, Northern Telecom
Russell K. Johnsen
Vice President & General Manager, Communications Division, Analog Devices
Robert L. Bailey
President and CEO, PMC Sierra

GREG SHEPPARD: You are going to hear about one of the hottest marketing opportunities in all of information technology and the impact on the chip business is tremendous. Over the past ten years the IT industry has taken great strides to produces fantastically powerful PCs and workstations. We have seen the capacity – the ability to process information increase at least 100 times or so, over the last ten years. What we have in place on the desktop is something equivalent to what the original Cray-1 computer was back in the late 70s. At the same time, local area network systems have done a pretty good job – they have increased on the order of fifty to a hundred times as well. However once you deal with the public network, that has not kept up quite as fast. We have a gap forming in terms of processing power at the desktop versus the ability of the public infrastructure to handle that.

Let me take a few minutes to go through some terminology here. There are a lot of acronyms in this area. First off, we have POTS, which stands for plain old telephone service, which is the original analog twisted pair of lines coming to your home. This is also what analog modem technology takes advantage of. Recently we have seen announced initiatives to take this technology up to 56 kilobits for essentially a doubling of what is currently available. ISDN is a digital type technology capable of up to 128 kilobits per second. ADSL stands for asymmetrical digital subscriber loop and this is a technology that enable these twisted pair of lines going into the home to handle much higher data rates, particularly the data rates coming into the home, and some variance of this up to as high as 50 megabits per second potentially. This is one of the technologies that is part of what we call XDSL. There is another technology known as HDSL, which also is being used as a technology in what we call the T-carrier marketplace. Cable modem technology has to do with attaching a radio frequency-based type modem, to a coaxial cable line – in this case, the one that the cable TV service provider has in place. HFC and SDV are variations on a theme: HFC, hybrid fiber coax technology involves blending both fiber optics and coax. SDV, switched digital video, is also a similar technology although it involves dedicating bandwidth and switching it directly to each home. ATM, asynchronous transfer mode, is a set of technologies that provide a connection orientated capability from point A to point B virtually through a network scaleable. It is also a set of services that can be provisioned and there is a lot of variation, different types of services. SONET is the fiber optic technology that is being deployed primarily in the U.S. The international variant of it is known as synchronous digital hierarchy, or SDH, but the fiber optic technology going in there is SONET. MMDS, and LMDS are essentially wireless ways to deliver high bandwidth to the home. At Dataquest, we look at these markets and the impact on chip demand. In upgrades alone, we see about a $4 billion chip market by the year 2000 based on the shipments of equipment implementing these technologies and services.

Our first speaker is Lisa Pelgrim. She is an industry analyst in our Dataquest Telecom group.

LISA PELGRIM: In 1995, the Internet met the consumer world and it had a profound effect on the modem market. The modem market grew that year by 53% and it shipped over 18 million modems to North America so it was quite an exciting year for a lot of the modem vendors.

Now we are starting to see some competition in that area of modems, and analog modems are starting to get ready to compete against ISDN TAs cable modems and XDSL modems. XDSL includes ADSL and other types of DSL technologies in there. Over the past couple of years we have seen significant speed improvements for analog modems, and those speed improvements are going to be quite drastic in the future in the way consumers are going to be able to get more speed to access the Internet.

This shows a bit of the timeline and how those speeds are presented to the consumer market. This chart outlines what that speed is going to buy you.

The main factors driving the modem industry are really coming out of the consumer and the SOHO market. The consumer market is absolutely huge – there are 94 million households out there and not all of them have PCs and not all of them have modems so there is quite a large installed base. The Internet has been a phenomenal success, or driver, for this market and the other big driver right now, which touches more into the corporate world is the telecommuter market. Working at home, (which Dataquest defines as someone who is working at home three days a week or more,) is growing very quickly and we expect it to grow by 20% until 1998. That is going to put a lot of demand for different types of devices for people to have access. The telecommuters are different from the Internet users in that they have more symmetric needs. The Internet users oftentimes are not sending very much information, but they have a need to download a lot of information and that is where some of the new technologies are going to be very beneficial. Both cable and ADSL right now are asymmetric and that fits this model perfectly. As far as telecommuters, a lot of them have much more symmetric needs in downloading a whole bunch of files, making a few changes and sending the entire thing back. That would change what their demand is. So what does that mean? That means you have a great forecast for all these different products.

Even though analog is starting to see competition, it is still going to be king through the end of this century, and you can see that toward the end of the forecast period, it does start to flatten out as the other products start to increase. The reason that analog is maintaining such a large share of this market is that it can be used almost everywhere. If you are a mobile user, you cannot depend on having ISDN, ADSL, or cable. When you are traveling around, you are going to have to fall back on analog. One of the other key factors in this is that the PC companies are including analog

Bandwidth and Performance

Bandwidth Object	Analog 28.8 Kbps	Analog/ISDN 56 Kbps	ADSL/Cable 1.5 Mbps	ADSL/Cable 6 Mbps
Two Text Pages 64Kb	2.2 seconds	1.1 seconds	0.04 seconds	0.01 seconds
Simple Image 2Mb	1 minute	30 seconds	1.3 seconds	0.4 seconds
Complex Image 16Mb	8 minutes	4 minutes	11 seconds	3.2 seconds
Full-Length Movie 4.2Gb	41 hours	21 hours	8 hours	14 minutes

Modem Technologies Forecast Unit Shipments to North America

Modem	1995	1996	1997	1998	1999	2000
Analog	18,829	25,056	34,654	43,078	46,550	46,376
ISDN TA	130	220	344	481	636	805
Cable	13	25	80	210	640	900
xDSL		36	210	980	3,200	7,400
Total	18,972	25,337	35,188	44,749	60,026	65,480

Increases in Modem Speeds

- Analog modems
  - 14.4 Kbps in 1992
  - 28.8 Kbps in 1994
  - 36.6 Kbps in 1996
  - 56.0 Kbps in 1997
- ISDN
  - 56 Kbps to 128 Kbps in 1995
- ADSL
  - 1.5 Mbps (downstream)/176 Kbps (upstream) in 1997
  - 6.0 Mbps (downstream)/640 Kbps (upstream) in 1998
- Cable
  - 10 to 27 Mbps Shared (downstream)/various upstreams 1997

modems in the majority of the PCs that are now being shipped. That is a direct reflection of the demand that is coming from the consumers, they want to get on the Internet. That connectivity is very important. ISDN is very strong in Europe, very strong in Japan and although we are predicting growth in that market, we do not see it hitting the big volumes that some of the other products will eventually hit. One of the things that are starting to put some pressure on ISDN are the new 56K modems that have recently been announced by all the key players in the industry: Rockwell, Motorola, and North Robotics in conjunction with TI.

The 56K is not significantly different from a single B channel at 64K transmission. Consumers are addicted to speed; they want more speed but they are looking for something less expensive. Across the country, it costs an average of $50 a month to get Internet access, or have ISDS service. In addition to that, fees for an ISP and per use charges tends to be too expensive for mass market.

Cable is another product that is starting to be rolled out. There are a couple of pockets around the country where commercial service is actually available. The price on that right now ranges from about $35 to $40 a month, and a user can surf all day long if they want to and not have any per use charges. It also includes the cost of the equipment. XDSL is something that is a little bit farther off but we do expect to see some commercial climate beginning in 1997. One thing to note on this forecast is that XDSL numbers include both the end user modem and the central office modem.

xDSL Technologies

Technology	Speed	Mode	Distance Limitations (24-gauge wire)
ADSL	1.5 to 9 Mbps 176 to 640 Kbps	Downstream Upstream	18,000 feet (12,000 feet for fastest speeds)
HDSL	1.54 Mbps (T1) 2.048 Mbps (E1)	Duplex Duplex	15,000 feet
SDSL	1.54 Mbps (T1) 2.048 Mbps (E1)	Duplex Duplex	12,000 feet
VDSL	13 to 52 Mbps 1.5 to 2.3 Mbps	Downstream Upstream	1,000 to 4,000 feet depending on speed

There is some confusion in the market today with XDSL, ADSL and all the other acronyms. This is just a quick way to look at it and some of the differentiators between them. In our forecast, we are not including HDSL which is important to remember when looking at these numbers. We have another forecast for that which is separate. HDSL is being used more as a T1 replacement service; it is more of a business market. It is it not something that is going to consumers.

The cable model is different than the rest of the other technologies in the fact that it is a shared technology. It works similar to an Ethernet at work where if a lot of users are using it at the same time, speeds are going to decrease. The cable companies claim that they could put in additional equipment and divide the group in half so that half the number of people were now sharing data, or sharing the speed. Right now only 10% of the cable infrastructure is upgraded to handle the two way data.

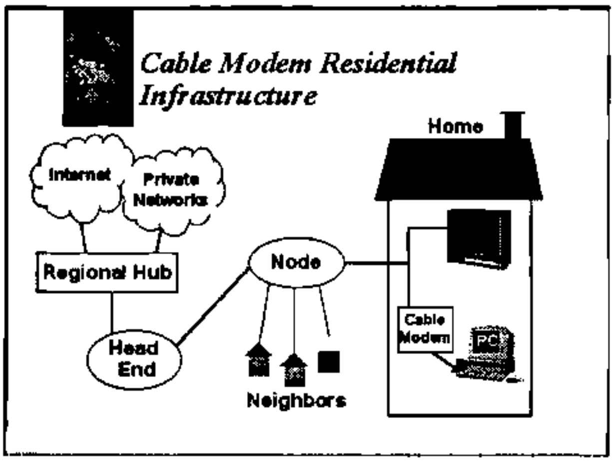

Cable is facing a lot of challenges and that is one of the reasons we do not see it growing as quickly as ADSL. One of the key challenges for the cable companies are they do not have experience in managing a network. They have been set up on a broadcast network. Their users can simply tune in based on whatever channel they want to watch. They are not measuring how often people are using it and a lot of times they do not know that lines are down unless somebody calls them up. In doing Internet access that is something they are going to have to learn about and that is going to be part of a learning curve. It is not simple technology, it is very complex and it is going to be quite an obstacle for them.

One of the other challenges that they have is competition from the direct broadcast people so they are very motivated to do this. Direct broadcast is taking away a lot of the cable market share and a lot of their subscribers are leaving the cable world and subscribing through satellite. In order to continue to have a business plan, they have to have a new product that they can continue to attract consumers with.

GREG: Our next speaker is Glenn Estes, with Telesis Technology Laboratory. Glenn was appointed President of Telesis Technology Labs in April of 1994. TTL, as it is known, is a subsidiary that provides advanced technology support to all units of Pacific Telesis' Group and pursues exploratory work on strategically important emerging technologies. Glenn is also an ABP of technology to the corporate strategy and development organization of Pacific Telesis' group. Prior to joining the company, he held various positions in technology development at Bell Labs and Bell Corp.

GLENN ESTES: I am going to focus on some of the realities of the telephone network. When Greg went over the list of acronyms, with maybe one exception, Pacific Telesis' group is using all of those technologies; we are either in trial or in service and that includes cable modems. We are not involved in switch digital video at the moment. There is a tremendous challenge underway, we are here mostly talking about technology today, but even more important for people to be aware of, at least here in the United States, is telecom deregulation legislation which was recently passed. Rulings from the FCC passed on August 8th, and how those rules get interpreted and the resulting willingness or unwillingness to invest by people in the infrastructure business will probably have a great deal of impact on your businesses. Pacific Telesis' group is a holding company – our largest and best known subsidiary is Pacific Bell which is a local exchange carrier in almost all of California. The telecommunications industry, with its deregulation, is going through tremendous change and when we talk about the impact of new technologies, we have to take this in context of everything else that is going on, including in the public policy arena.

In this figure, XDSL is the direction that Pacific Telesis' group and a number of other large local exchange carriers in the United States are headed toward at the moment. We have a number of motivations to invest in this technology. We do recognize the market need for some of the technology that we are willing to deploy. One of the issues we face today is the very long call holding times that data users need. Holding times have increased, maybe 20% a year on average over the last couple of years, and as there is greater penetration of on-line services, we are going to see those trends continue. Some of this really does

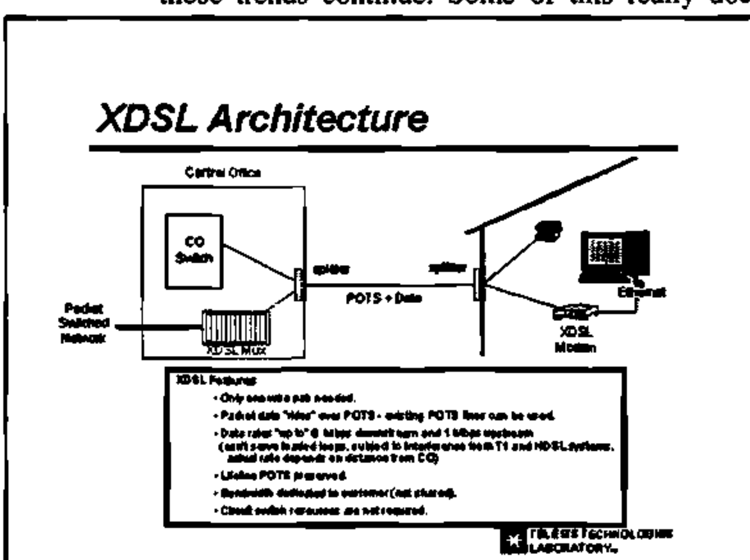

relate to regulatory and public policy issues. In California, for instance, there is unlimited free local calling, within twelve miles of the central office. That allows an on-line service provider to go out and place modem banks in selected wire centers. By covering less than 100 wire centers, you can cover almost all the population in California. These people are located in such a place that they would have unlimited local calling – no charge, no revenue. That makes for a strange business model when you consider that we are tying-up our central office switches with effectively free calls. Suffice it to say that we need some regulatory relief on that. We are also working on terms of what technology solutions could be brought to bear.

The architecture would allow us to use the existing wire pairs to carry data as well as analog voice. We use FDM to put data in the spectrum above the POTS (the plain old telephone service, which is, say, 3300 hertz). At the central office we have a splitter which splits off that spectrum, and sends it to a channel bank (shown as an XDSL multiplexer). Then only the voice calls would go through the central office switch and tie up that critical resource.

There is a great deal of concentration done in today's voice network based on the old Erlang model of three minute holding time per calls. There are a lot of challenges though, and that involves working through the outside plant. In the context here, the outside plant is the connection between the central office, a wire center, and the individual subscriber.

By going to some fairly old data, I have shown some statistics based on the old Bell System, the subscriber loop distribution. Not many loops are less than 18 Kft. If your modem technology could reach out to 18,000 feet and operate reliably, then you are going to cover over eighty percent of the subscribers in the United States. In different regions of the world, there are different loop distributions, but in general they have more wire centers, shorter subscriber loops, and therefore are easier to reach in terms of the technology you need for the modems.

Twisted pairs have distributed capacitance between the pairs of wires as they reach further out, and to keep the impedance constant inductors can be installed. They are distributed at set intervals throughout the outside plant. That loading of loops, is only done beyond 18,000 feet. Our basic analog phone requires a DC current to drive the microphone and the handset, and you get 20 some milli-amps, so we have a resistance design limit of about 1300 ohms in our longest loops. Depending on the gauge of the wire, that is going to set a length limitation also. There are many different technologies being proposed for XDSL. The technologies we are intending to deploy initially are asymmetric digital subscriber line. There is a U.S. standard which ETSI in Europe also standardized on a modulation scheme called discrete multi-tone and that is the direction we are headed as a company.

We also have hybrid fiber coax. Pacific Telesis' group has launched an initiative to build out a consumer broadband network throughout the state

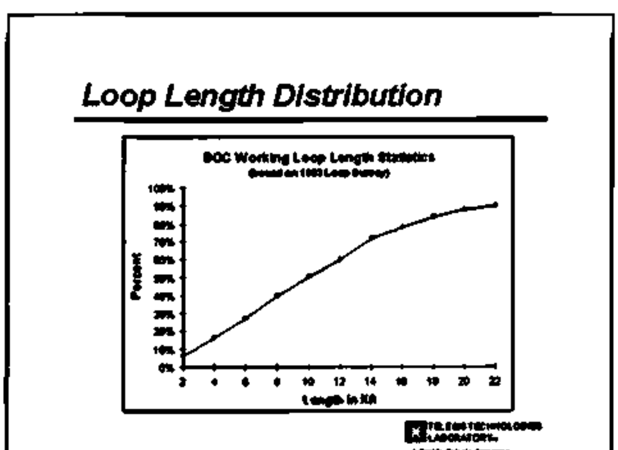

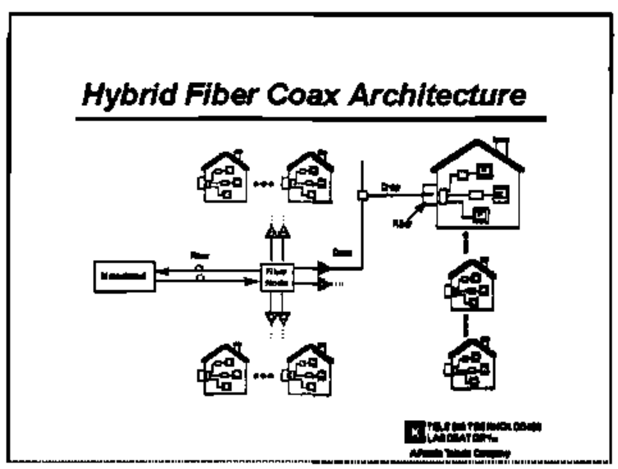

of California. We have wired past about 180,000 homes now but we have not turned all those customers up. We are in the process of putting both telephony and video services on that facility. We have two beachhead areas, one in San Diego, and one in the south San Francisco Bay area, particularly around the San Jose area. In San Jose, we have cable modem trials underway, and almost a hundred users, on their second generation cable modem.

Technology Comparison

ISDN	ADSL	CABLE MODEMS
•Already Available to 95% of Pacific Bell's 15 Million Subscribers •Works Over Existing Copper Network •Standards in Place •Over 500 CPE Manufacturers With Equipment Widely Available For As Low As $200 •Extremely Flexible, Supports Data, Multiple Voice Lines, Video, Digital Audio and Point of Sale (Card Swipe)	•Not Available Yet, Trials Only •Works Over Existing Copper Network •Limited Manufacturers,Equipment Not Readily Available, CPE Questions Open •Supports High Speed Data And One Analog Voice	•Not Widely Available, Only Trials or Severely Limited Introduction • Full 2-Way Requires New Coax Network or Expensive Upgrades •Lacks Standards, Limited Manufacturers, Equipment Not Readily Available, CPE Questions Open •Supports High Speed Data

TELESIS TECHNOLOGIES LABORATORY
A Pacific Telesis Company

The architecture shown here is one typically used by MSOs or cable operators, and also by companies such as ourselves. From the central office out to the neighborhood node (called the fiber node), we use coaxial cables to go down through the local streets and connect to the premises.

In our hybrid fiber coax, we are trying to integrate telephony onto that facility from the very start. Ours is not a plan to merely put in HFC, offer analog video, then later continue to use our twisted pair. Our architecture was to use voice from the beginning, so we constructed perhaps the best hybrid fiber coax networks that can be built today economically.

It is possible for ingress to occur there and while impulse noise will not be a terrible problem to overcome, there are noise sources which can affect many of the upstream protocols used for a two-way service. Two way cable modems using a POTS return or a telephone return with high speed data downstream is probably a more pragmatic way to go.

In our particular plant, we have done some things which we feel are giving us some advantage. It is still a challenge for us, operating cable modem technology. ISDN is happening and we do have a very active data market in California. We were forecasting over 50,000, new ISDN lines this year and we are running ahead of rate by about 30%, so I would project that if we can get the installs done, we will end up the year installing more than 50,000 lines of ISDN just in our market.

We have ISDN capability throughout California. We have trials today where we have tens of customers connected, running at 64 kilobits upstream and one and a half megabits downstream. Those installs have gone very well, but we do not have the infrastructure backhaul support in place to scale that into roll-out service right away. The practicality of the cable modem issue is that it takes a lot of capital investment to make the infrastructure capable of supporting that in a robust manner.

I anticipate the dollars will flow into that sort of an investment.

GREG: Our next speaker is Steve Edwards with Northern Telecom. He is currently Assistant Vice President with responsibility for sales and marketing for Nortel's Telco Broadband Networks and Application products and solutions to the North American telephone companies. He has responsibility for business management and development of broadband access solutions.

STEVE: The local loop has been remarkably stable for well over a hundred years – it is still a twisted pair into the home, or into the business.

The bandwidth of that circuit has remained unchanged at less than 4 kilohertz of delivery. Clearly the bandwidth requirements on that access into homes, and small business are driven by the requirements for high speed data.

What is driving this demand for high speed data? It is the growth of home PCs and there is tremendous growth in North America in both small business and home offices. There is a significant business opportunity for cheap connectivity for small businesses as well. Second, the massive acceptance of the Internet, explosion of usage of services such as America OnLine, and the increased desire to work at home. This is driven both by us as individuals and our companies who want to save on facility expense, and governments that want to save on transport, infrastructure and pollution costs. There are three drivers in that area and they are creating frustration for users with slow modems, downloading files, and tied-up telephone lines. An independent estimate suggested that to actually upgrade the switch infrastructure in North America to support the longer holding times, would cost the local exchange companies $6 billion in equipment purchases. Chief financial officers do not like that business case, that is capital expenditure with no incremental revenue. There is a need to migrate the power users to high speed services.

There are in homes and small offices in excess of 20 million modems in use today in North America. My current ISDN projection would suggest, that possibly there would be 500-600,000 lines in use by the end of this year. This is somewhere in the order of 45 million data access users, and

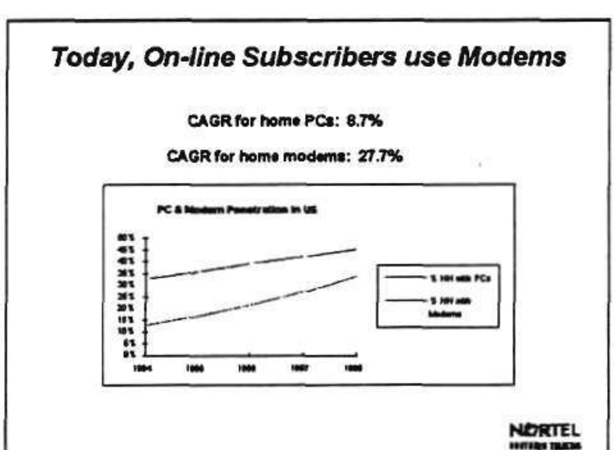

about 35 million of them will be analog modem based. Our forecast suggests that ISDN may be at the 3 million level by then, and high speed access with ADSL might be around 5.5 million, with modem actually being stronger than indicated at around 2 million users.

NorTel's high speed network solutions cover both coax and copper-based infrastructures. They incorporate intelligent networking which is a key part of the overall service delivery and they are based on the notion of providing end to end ATM connectivity, which is very important to scale to mass market proportions.

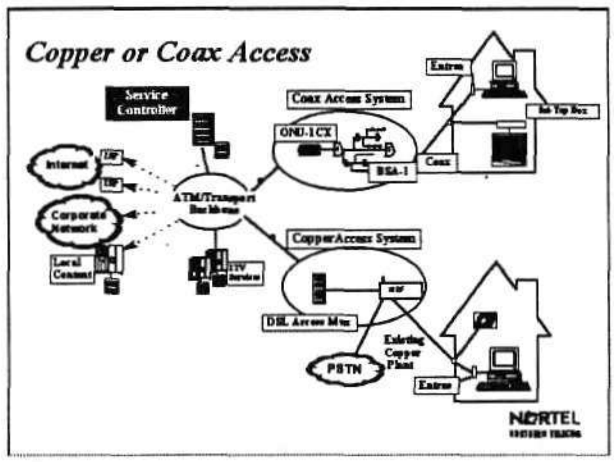

An end to end service perspective is required for the service providers, whether they be companies like competitive local exchange carriers, inter exchange carriers, or the cable companies. They will need to have a system perspective. It is not sufficient to just deliver high bandwidth, you have to be able to deliver services that work.

Our approach here is what we call a four layer approach. It consists of the access network itself, both coax and copper access networks, using ADSL technologies. We are providing ATM connectivity from the home or small office through an ATM modem back into the network, through the SONET networks into the ATM (the backbone) environment and then for connection to content there is a service intelligence layer. Then the application layer which is things like World Wide Web browsing, or work at home applications etcetera. Focusing on the coax axis we have been proposing, and have deployed, a system that delivers 20 megabits of symmetric service to anywhere from 50 to 250 homes in a shared bus environment. This is a symmetric approach, and we have been very cognizant of the issue that was talked about in telecommuting that one person's download is another person's upload into the network. We have tested this system under high bit rate transfer of 9 megabit files, 14 PCs doing that simultaneously. We found performance that was 40 times in excess of what you could achieve on a 28.8 modem. This solution set is in play today in New Brunswick Telephone, in eastern Canada. Consumers using the service for fast Internet access are very ecstatic about the solutions being delivered to them. We have many good comments about its use there.

Moving onto to the copper access, we are working to deliver a system that matches very closely against the solution that Glenn showed in his network. ADSL technologies are ideal connectivity solutions for small companies that are buying T-1 and HDSL solutions. If the tariff of Internet service rates is $30 to $50 a month, it is going to be easy migration when paying a $1000 today. This local loop is available to other providers, and what we are seeing is the competitive access providers using it to gain access to service. Service control is very important, for the scalability of the service delivery networks and is providing both session and service control for the services being delivered to either the home, or the end user.

NorTel's Solutions are based on delivering ATM to the home, for residential or small business use. We do support high speed access, both downstream as well as upstream in coax based implementations. These solution sets are ready for interactive TV, when that comes back. The notion of the service controller does provide a new way, to a new connectivity. Networks are proposed to provide ease of network management and offer a high degree of end to end security. They addressed the fast Internet access requirements to ease the switch congestion and then creating the revenue opportunities, not providing more opportunities for free local service.

GREG: Our next speaker is Russ Johnsen, who is with Analog Devices. He is Vice President and General Manager of their Communications Division, based in Wilmington, Massachusetts. He is chartered with managing the youngest and fastest growing division with Analog devices. The communication division was established in 1994 to address opportunities in digital wireless communications as well as broadband telecommunications markets. His prior job was with National Semiconductor where he was General Manager of their Wide Area Networks Division, and he has held various roles, working in communications, microprocessors and out in South East Asia, with manufacturing as well.

RUSS JOHNSEN: We have had in the last few years a lot of things that are really stacking up to provide a tremendous amount of opportunity for us in our industry. First, we have heard about a pent up demand, a tremendous amount of people that are focusing on Internet access. So we have market demand out there.

Second, new transmission technologies, whether it be DMT, CAP, QAM, QPSK, or DWMT, various transmission technologies are coming down the pipe and lining up right now to provide the technology to conquer these problems with Internet access and also adding, enabling, new carriers in the local loop as well.

The government is deregulating the local loop, deregulating the telecommunications infrastructure all around the world, and that is putting new pressures on the phone companies. At the same time, it is breeding a lot of competition that is going to help enable a lot of the markets that we are talking about. Governments are working in concert with industry and markets, to be able to

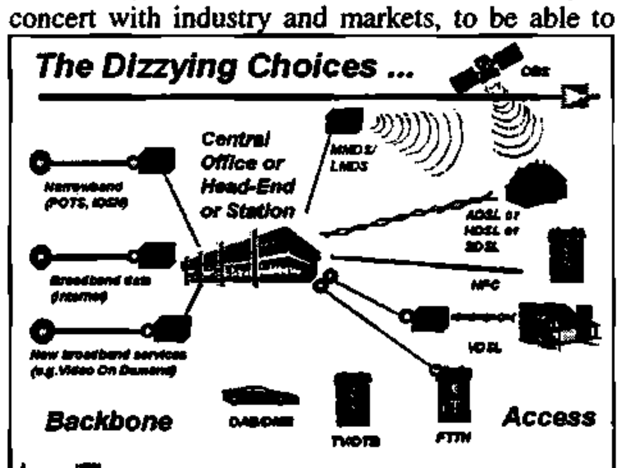

provide a lot of new opportunities.

There are a tremendous amount of choices: ADSL; HDSL; SDSL, which is symmetric digital subscriber loop hybrid fiber coax; VDSL, sometimes called switch digital video coming down the pipe which will get network access up in excess of 51 megabit. Consumers want POTS, they want POTS access, they are going to want broadband data, they are going to want other new broadband services. When we started talking about broadband into the home, we were talking about video on demand, and now it is Web access.

There are a whole host of different delivery techniques. To start with, there are about 600 million lines of copper twisted pair that are already out there, on loops that already exist in the infrastructure. They are being installed around the world to the rate that we should see 900 million lines by the end of the decade. Copper is here to stay, it is not going to go away. Ultimately companies such as PacTel and others are going to be putting in much stronger fiber infrastructure in addition to twisted pair. Technology such as ADSL can take advantage of the existing infrastructure and at the same time relieve the overly congested central office. As fiber is deployed around the world, companies can begin putting in VDSL, even higher speeds into certain selected areas of businesses and consumers. Some of the satellite TV companies are talking about Internet access over DBS. In Europe, a new technology is being deployed called digital audio broadcast. Initially it is being set up for CD quality sound but it is a 1.7 megabit per second wireline pipe that goes into mobile units. A lot of these solutions are going to co-exist.

As a semiconductor manufacturer, I feel as if we are the tail wagging the dog. Which standards are going to triumph and how are we going to get at those? There is a very strong argument for coexistence. It is not the technology that will answer the question on how it will be deployed. It will vary from country to country and from country to county, it will get down to the local public utilities commissions in terms of how much money companies are going to be permitted to spend. We in the semiconductor industry will have to be prepared to answer a lot of these. Population demographics are going to be important too. Companies are going to have different levels of technology for different businesses and consumers. Timing is going to shift in terms of when all these things are going to roll out. ADSL and VDSL are going to come out a little later, there will be interim technologies but a lot of these things are going to co-exist well into the next decade as well.

What are the applications that are going to be out there? Is it going to be purely Internet access, is it going to be symmetric requirements or is it going to be asymmetric requirements? Work at home is going to require a higher level of symmetry and the will be able to choose as time goes forward, is once again, leading credence to co-existence. It is a question of what makes sense where, within the local loop, or within the infrastructure.

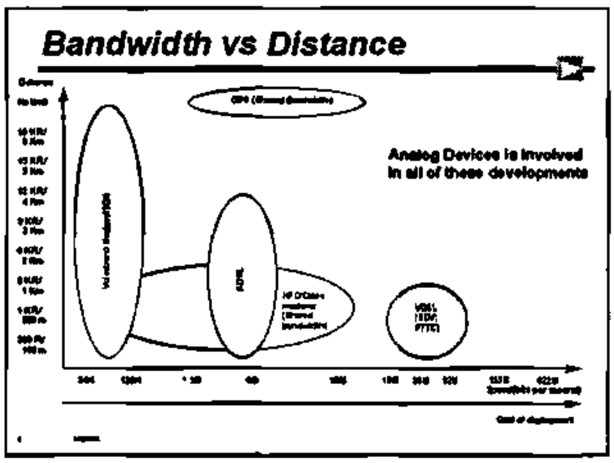

It has to do with reach of the loops, what technologies can reach out to those 80% of the loops that are 18,000 feet or less, and how much does infrastructure have to increase and improve your as time goes forward? Voice band modems are already there, but even at 56 kilobit, it still takes a long time to download full motion video, which is going to be demanded by each and every one of us. Other applications are going to be needed. ADSL can service in the 6 or 8 to 10 megabit per second range on the downstream side, rate adaptive ADSL can become symmetric in the 3 or 4 megabit per second, depending upon loop lengths. VDSL is also very real and in areas that have fiber within 1000, 2000, even 3000 feet, of a business or residence, VDSL is going to be a very real application that will start deploying in the '98 to '99 timeframe and well into the next decade. VDSL is an asymmetric solution, very likely, although there are ways it can be made more symmetrical. It is going to have 26 to 51 megabit per second into the home, and six or eight megabits per second, depending upon what upstream technology is used going back the other direction. That is a lot of bandwidth that will also co-exist for a long, long time.

Analog Devices is working on many of these applications right now. Whether it be ADSL, or VDSL, or hybrid fiber coax, there is going to be an evolution of this technology, there is going to be a lot of problems associated with deploying them. That is why the phone companies are putting together very close trials. Phone companies are going to have to learn a lot about their own infrastructure in order to fully deploy it. Deployment is going to happen; it is going to go a little slower, with more care than we thought, but there is going to be $50 to $100 of semiconductor content in each modem. It is a tremendous opportunity for our industry.

Semiconductor companies are having to take a step more than in the past. You have to do, not only the systems on a chip, but you have to add on to it the systems architecture, the software that has to bundle around it, in order to create an entire solution. That is how you are going to win, and we have an opportunity to do that.

GREG: Our last speaker today is Bob Bailey. Bob is President of PMC Sierra which is a subsidary of Sierra Semiconductor. He is based in Vancouver, British Columbia, where he has been since November of 1993. PMC Sierra is a leading maker of broadband networking integrated circuits and with Sierra moving out of the modem business, this is now certainly akey thrust area for them. In 1993, PMC Sierra had revenues of $8 million and this has climbed to the $40 million range as of 1995. Before working at PMC Sierra, Bob was VP of the ASIC Strategic Business Unit, at AT&T Microelectronics who are now Lucid Technologies.

BOB BAILEY: Well, I got in very late last night, it was after midnight. This is my first visit to Palm Springs. When I got into the rental car and I noticed they had one of those global positioning things. It gave me different directions than the Avis guy did to get here. So I decided I was going to follow what the machine was telling me to do, and this robotic voice would say, "Turn right up1/10 of a mile." I accidentally made an honest mistake, went the wrong way, and it was telling me, "You're going the wrong way," and I was very impressed. So I turned around and I was going the right way, and as I was getting closer to the hotel, it told me to turn. I thought this was odd as it looked like a residental street, but it gave me the name of the street and again, it was very impressive. I was driving through and I said, "This cannot be right!" Right, but you know it is run by a computer, it has to be right. So anyway, I go to this dead end and there was a wrought iron fence there and in the distance I can see the hotel! And then the machine said, "You have arrived." If I was a crowI had arrived. So I turn around and I am back tracking to get around and the machine kept saying, "You're going the wrong way, you're going the wrong way, you're going..." so I am arguing with the machine.

This global positioning system with this satellite infrastructure, was put in place by the Department of Defence in the Cold War build up. They have to do something with these things, and there are a whole bunch of commercial applications. Here is an infrastructure with insufficient intelligence attached to it, which is exactly the opposite of what we have in the Internet, which is an overabundance of intelligence with insufficient infrastructure.

We have to crêate this infrastructure for the digital age, and the opportunity is really unprecedented. There is a tidal wave of digital traffic building up and the information infrastructure, just cannot handle it. One study said in the next twenty-five years there will be over a trillion dollars in capital investment in the world to handle this infrastructure. Digital age is transforming our global economy where the product that is represented in streams of bits or bytes is transferred over a wire or piece of fiber, and will have a profound effect on all of our lives.

There is on-line ubiquity, we all hear about it how many people are getting on-line. It is getting to the point where even today, you cannot really run your business without being on-line. This is creating seemingly unlimited choice not just in products, but in information, in political candidates. Everybody's got a Website. Maybe you have a business that is very non-technical, and you do not want to know anything about a Website. If you type the basic text of what you want, and put it on a floppy in the mail, you will have a Website. The equipment is not even on your premises, so this is really pervading everything.

A big opportunity is the merging of networks – the cable networks, the regional Bell operating companies, the carriers, and this new entity called an ISP, an Internet Service Provider, that five years ago did not even exist. One networking executive told me that the cable companies are losing market share so quickly right now, to this direct broadcast phenomenon that if the current trend continues, within only two or three years their basic business model will not work. With the merging of networks there is going to be merging of equipment and chip people like that because our integration capabilities come into play. There is also massive restructuring of traditional distribution channels as a result of that, cutting out the middle man. This is direct mail in a turbo mode, which creates an incredible leap in market efficiencies. Everything is faster, better, cheaper, there is no room for inefficiencies. The government is trying to figure out how are they going to do the typical things they do to government that do not work anymore. Communist China has one approach: make the Internet illegal. That really is not going to work, but it is going to make things very exciting.

One of the mega trends that are going on, on-line ubiquity, is going to become a necessity almost like a social security number, companies will increasingly leverage the Web with sound clips and videos to market their products on the Web. Everyone tries to do it, but it is not very feasible with the speed of the modem technology today. There are tens of millions of LAN connected PCs installed every year and that traffic has to be managed over the wide area network. There is U.S. telecommunications deregulations, and there is another thing going on, especially in Europe, and that is, corporate welfare, as it is called, is now unaffordable, which means the big vertically integrated telecommunications industries in Europe can no longer afford to be isolated from competition because they would rather spend their money on healthcare, and retirement pensions, and things that they are having a hard time supporting right now. That is going to open up a lot of competition. Vast repositories of information are coming on-line, and are going to make the whole thing come together.

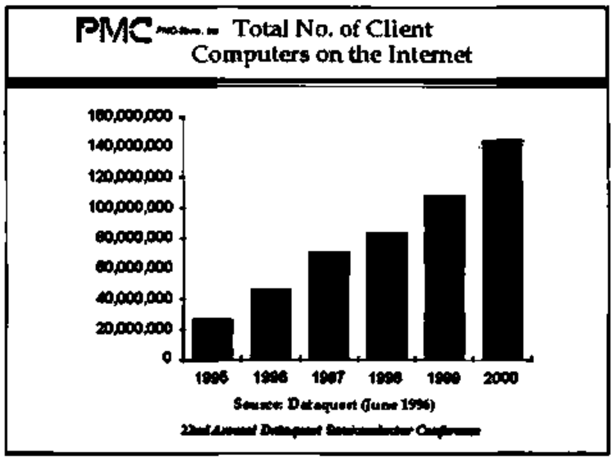

This is the chart shows that the total number of client computers on the Internet will increase 5 times from 1995 to the year 2000. The most prevalent modem technology used in 1995 was the 14.4. In the year 2000, some of the other forecasts show that something on the order of a T-1 or a full speed ADSL modem will be prevalent. That is 1.5 meg, about a 10 times increase in bandwidth at each node, so that is a 50 times increase of traffic, and that probably does not include all the corporate networks. Video and sound is incredibly more complex, I think it could be at least a million times the traffic in just 5 years and that is a big number.

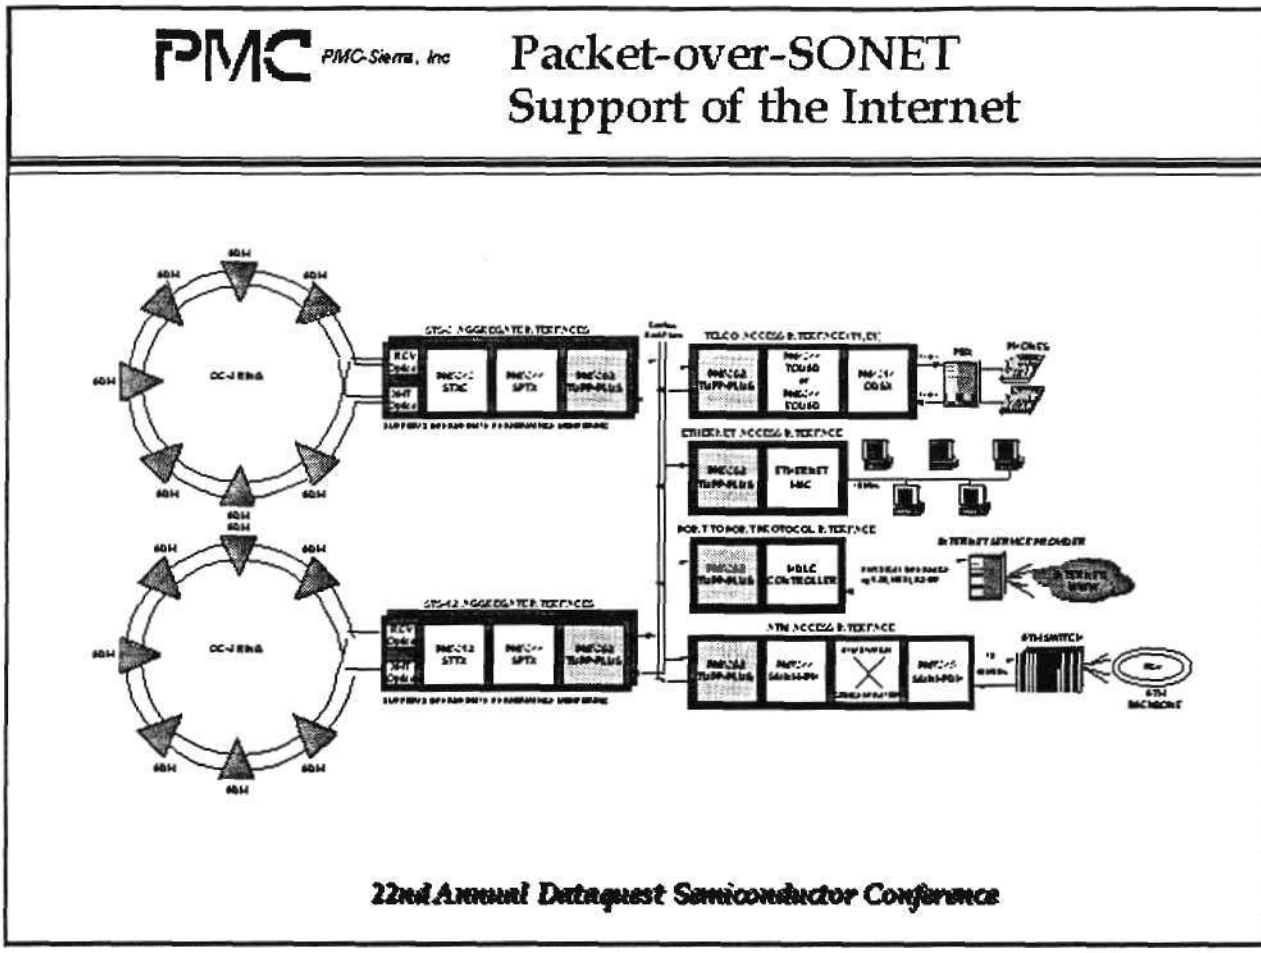

So, what kind of technologies are being used to address this? There is IP, which is internet protocol, switching or tag switching, from Cisco frame relay from companies such as Northern Telecom, Newbridge Networks, Stratacom, Cascade and others. ATM in the wide area network, the ISP carrier modem, DSP modem banks, and packets over SONET, which looks like an exciting application.

Frame relay is incredibly popular and is growing at over 100% a year right now. One of the reasons is the cost of a frame relay switch has gone from about $3,000 to about a $100 a switch port in only about three years. The service for frame relays can be up to T-1 rates, 1.5 megabit per second rates, for only a $1000 a month.

PMC has connected all its branch offices with frame relay and behind them, in the backbone, aggregating that traffic, are ATM switches – ATM has already won. The old TDM, or time division multiplex switches are being transitioned, or overlayed, with ATM switches, and in the future they will be all ATM switches and everyone agrees on this, this is a consensus. The corporate world is connecting all their factories and all their places with frame relay, as opposed to lease lines, which were much more expensive

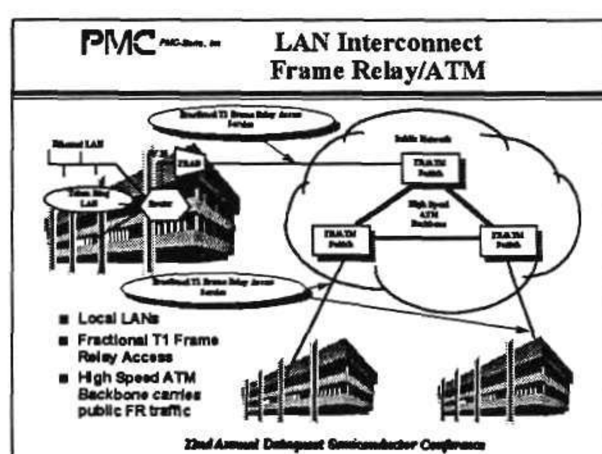

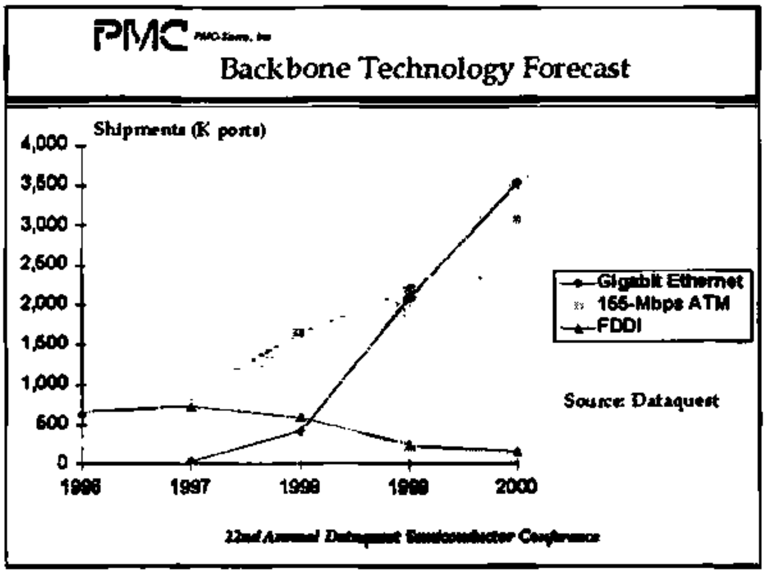

and inefficient.

Dataquest came out with this report on gigabit Ethernet, and said that gigabit was going to be a $2.9 billion equipment market by the year 2000. It shows ATM in the LAN, and in 1997, ATM will surpass FTDI in terms of switch ports shipped. FTDI had been pretty much, THE backbone technology. In 1999, it says that gigabit will then surpass ATM in terms of port shipments and in the year 2000, it will be slightly more. But if you look more at the investments that the Ciscos, Cabletrons, 3Coms and others are making, they are betting on both. ATM is growing quite nicely, even though there is this gigabit product being developed as well, and is going to probably have first products out next year.

I had heard that one of the big four, I do not want to give you their names, but in their first year, they shipped $100 million of ATM network equipment, in their first year. So, the answer is both, both win.

If you look at all this modem traffic we talked about, and it goes into the public network, it is aggregated quite often on T-1 links, into the ISP and then aggregated through the Web, across sometimes IP or through IP routers, sometimes Cisco routers. There is big opportunities here as well in silicon, making all these interfaces happen. The new phenomenon is packets over SONET. The ISPs are putting SONET equipment into interface to leverage the SONET infrastructure and the carriers. You can put voice over that, Internet data over the World Wide Web, IP can be over that, as well as ATM cells can be over SONET, so it is a universal transmission technology.

Those semiconductor companies that best address the needs of those networking OEMs, with enabling technologies, will be the primary beneficiaries of the digital age and as we said, a trillion dollars of capital will be invested in the next 25 years to manage this tidal wave of digital traffic. Semiconductors will provide the technologies to enable this digtial age.

GREG: Of the technologies talked about to provide these high speed services to the home, which one will be the most prevalent and why?

RUSS: I can tell you what we are investing in and obviously that gives you some clues on what we think are going to be important markets. We are investing very heavily in inside analog devices and discrete multi-tone based ADSL. And that was what was referred to earlier, it is an ANSI standard. We think it is going to be deployed and why it will win as much as anything, is exactly what we have talked about, this congestion inside the central office on the voice band switches, as well as because Internet Web requirements are coming down very quickly. We are investing very heavily there, and in certain areas video cell fiber up to within 3,000 feet of a home or a business is going to start to be deployed although it is going to be well into the next decade before it is really ubiquitous.

LISA: One of the things that I really like about the ADSL is that it separates the DSO channels, so you still have your voice call going over, and you are able to use your data and telephone over the same line which is not the case right now with the cable modems. When people are on their modems, the average time spent is over an hour, and that is incredibly long compared to three or four minutes for a voice call. That is one of the reasons I think the telephone companies are very motivated to move forward with ADSL.

STEVE: I agree, but I would like to put some cautions on the table here. The technology has to prove it can deliver end to end service capability. There are some issues in home wiring that are yet to be addressed by the industry. Clearly we have not talked about interoperability. Although it is a standard, it is only a standard at the transport layer. We need to work through many probability issues so this industry can rise to the challenge of putting modems into circuit cities and consumers can buy them off the shelf, phone up their friendly service provider, and get service. There are some challenges that need to be addressed, but if ADSL does prove to be a technology that can deliver the goods on data delivery, and can migrate into some form of interactivity at the video level, solutions that require digging up infrastructure of North America will be hard put to displace it. That is the experience that you certainly had in San Jose, where there was a lot of push back from communities on digging up the streets to put these technologies in. The ones that can already use the infrastructure in place, or right over the ether, are the technologies that will win into the next century.

BOB: We will never be able to deploy anything without your companies. Our economics are such that we do things in scale, we have to have forward pricing, and you do not get forward pricing until you get the chips. We are starting to see people building, they have the piece parts and they can build them in the systems. Many of these ADSL systems we are talking about actually could be providing ATM in-points, and the issue has been that there are not enough ATM in-points out there to continue some of the investments that corporations have been making. When you start counting in these access loops as ATM in-points the numbers start getting very intriguing.

There is a great deal of business opportunity a few years down the road for figuring out a home gateway device and the internal network issue for inside the residence. But to try and actually answer the question, we are going to leverage our investments, we have that money sunk in the ground. The U.S. local exchange carriers have hundreds of billions of dollars basically tied up in this plan. That is our lifeblood and that is what we have to work with, so we are going to continue investing in that to the extent that public policy and the stock markets allow. I do not want to discount the difficulties that we foresee in rolling it out. We have to make the ADSL service simpler than ISDN and that includes the installation procedures, etc. We also have to get to the point where we can get the low price customer premises equipment and that requires interoperability. Our company has joined together with three other regional Bell companies in going out for joint R&D for ADSL technology in trying to get forward pricing, and set a direction. I hope other companies will step up and we can get the trend going. Right now, I think we are succeeding in that so far. We do see the ADSL as happening.

BOB: The risk of the ADSL is, there is such a threshold and an upgrade to get there and by the time we get there, where it is pervasive, where it is easy to get, where if somebody picks the phone up and within at least 30 days they can get it into their home, that 1.5 meg asymmetrical will not be enough. That is the risk, and I am not talking about cable modems here, I am talking about another technology altogether where somehow you get Ethernet and voice on one line, and then a splitter once you get in the home.

RUSS: By the way the ANSI standard is 6 megabyte, not 1.5 ADSL.

BOB: But the number of years it will take to get there, you will have to shoot for something that is moving and have to lead the target. It is going to grow and there is infrastructure that will be put in place but I do not know if it is going to be the panacea, or if it is going to pervade the whole Internet world.

LISA: The forecast that we have for ADSL or XDSL, in the year 2000, only represents 2% of the PSTN lines in the U.S., so while our forecast looks very aggressive, it is still a small percentage of what is really out there. Another interesting number that came out of Bell-Cor was that 34% of all Internet traffic worldwide originates in California. You can see the demand is a little skewed. And 25% of Internet traffic originates in the Bay Area so when we are looking at deployment guess where it is going to be. We can really see where certain pockets are going to move ahead earlier.

GREG: There are several issues in ADSL that still have to be worked out. What are the most important hurdles?

GLENN: There are physics issues on the outside plant about getting it to work. To us it is scaling this out to where you can get the installs done. We have a lot of systems issues, support to customer care platforms which actually gives us an advantage. The inside wiring issues we will initially deal with that by doing direct wiring to the room in the home. We know how to do that quite well, we do that every day. For us it is a business case issues that gets back to willingness to pay and the penetration rates. There will be a certain level of penetration for a $30 a month service that you will not get if it is a $150 a month service. That has to do with this circular argument that goes around, that need the chips and we need the scaling to be able to do this in volume. The issues around ADSL, is we might have to go slower than faster to begin with. It is more important for our company to be able to reach more customers than it is to be able to go lickety-split fast. If you go to an ISP and say, "I am going to bring you a firehose and 1.5 megabit from all your subscribers," then ISP is going to go upside-down and the Internet itself would have tremendous problems. We are not the only game in town, literally, that has to get lined up in this process to succeed. I see it as an incremental step by step. We are looking for technologies that would eventually let us get to the place where we can do stream video, etc., and that is where these long term bets are. I really hope the DMT work you are doing is as promising as we think it is.

GREG: What are your expectations out of the semiconductor industry, the supplier? How can they best help you achieve your objectives in this area?

STEVE: Clearly the expectation of our customers, is that we will provide forward pricing in many of these technologies so that we can stimulate the economy of scale of these devices. What we are looking for is a rear view that there is uniformity in the standards that are being proposed so the industry as a whole can take the steps that we need to take in integrating, fairly massively, devices into single chips so that the price points come down. I think that really is a challenge for all for us, settling on that standard. Glenn's company along with the three other RBOCs have taken a good step in that direction to, getting everyone's minds focused on. What the real problem is, and that is hitting a certain price point and doing it as quickly as possible. The only way to do that is to work in conjunction with your industry, so that we can drive the cost of material that is in those devices down to a point where it makes sense for us to build them. That is the fundamental challenge: understanding what the end game is, bringing investments in, the integration of them, and getting on with the job.

GREG: What do you think it takes for a semiconductor supplier to be successful in this market space moving forward?

RUSS: Just integrating hardware together is not cutting it anymore. Probably a third or a half of our investment in DMT is in silicons, systems and software. Integrating the hardware systems and software, optimizing the partitioning between the software and hardware, between analog and digital has to be critically important. Unfortunately, these are long-term investments and by the time we anticipate seeing large-scale returns, we will be in our third or fourth generation of circuits. In 1998, we will be in our third or fourth generation already and that is how we are going to get the cost up and we recognize that we have to do that.

BOB: To be really successful as a semiconductor company in the communications and networking segments, it is very different from the segments such as the PC segment. In the past, the customers would tell the semiconductor companies exactly what they want and then we would develop that precise thing the way they wanted. With communications and networking there are more standards, so the customers expect products that adhere to those standards and are already completely de-bugged, and will enable them to make their end product as a result of it.

Panel:
Next-Generation Digital Consumer—Electronics for the Mass Market

Moderator:
Dale L. Ford

Senior Industry Analyst,
Semiconductor Application Markets Program,
Semiconductors Worldwide,
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 B.S.E.E. degree in electrical engineering from Brigham Young University.

Panel:
Next-Generation Digital Consumer—Electronics for the Mass Market

Dwain Aidala

Vice President and Division Manager,
Personal Mobile Communications Division,
Mitsubishi Wireless Communications Inc.

Mr. Aidala is vice president and division manager of the Personal Mobile Communications Division of Mitsubishi Wireless Communications Inc., which develops and markets wireless voice/data telecommunications products and advanced mobile communications applications. He is responsible for the tactical marketing, product management, and sales of these products, such as the Mobile Access phone, as well as strategic planning for the development of new products and services.

From January 1985, Mr. Aidala was employed by Mitsubishi Electronics America Inc.'s (MELA) Electronic Device Group. As vice president of marketing until 1992, he was responsible for semiconductor products ranging from DRAMs and other memory devices to laser diodes and gallium arsenide devices for telecommunications. In 1994, Mr. Aidala began and headed the MELA North American Multimedia Business Center, a business incubation group focusing on wireline and wireless multimedia business. Before joining MELA, the Los Altos, California, resident worked for nine years at NEC Electronics America, where he held product management positions in microprocessors and microcomputers. From 1970 to 1976, Mr. Aidala was a member of Arthur D. Little Inc.'s professional staff, consulting in the area of new products and systems development.

Mr. Aidala earned a bachelor's degree in electrical engineering and a master's degree in business administration from Northeastern University.

Panel:
Next-Generation Digital Consumer—Electronics for the Mass Market

David Taylor

Director of Marketing,
C-Cube Microsystems

Mr. Taylor joined C-Cube in 1992 as director of marketing for Broadcast Products. While at C-Cube, he has spearheaded the marketing activities for the company's MPEG-1 and MPEG-2 encoders as well as its set-top decoders. His responsibilities include business and product management, market development, and product definition for the broadcast, video teleconferencing, and desktop video marketplaces.

C-Cube lead the industry in defining and implementing the MPEG and JPEG standards, and is today the leading supplier of VLSI solutions for digital video applications in the consumer, communications, and computer markets. C-Cube has worked closely with many international broadcast and consumer electronics companies to develop MPEG-based products for videoCD, direct-broadcast satellite (DBS), digital cable, and the emerging market for digital video disc.

Mr. Taylor's career has been spent in the semiconductor industry focusing on both business and technical issues related to digital signal processing. Previous employers include Cypress Semiconductor, Zoran, and Advanced Micro Devices.

Panel:
Next-Generation Digital Consumer—Electronics for the Mass Market

Andrea Cuomo

Vice President, Headquarters Marketing and Strategic Accounts, SGS-Thomson Microelectronics

Mr. Cuomo became SGS-Thomson's headquarters' marketing and strategic accounts director within the Headquarters and European region in 1994. In this position, he has responsibility for corporate strategic marketing, which identifies future semiconductor trends and develops advanced systems, as well as combining the diversified know-how and wide-range technologies of the company, and also holds responsibility over sales to the company's strategic accounts.

Mr. Cuomo joined SGS in 1983 as a system testing engineer, and from 1985 to 1988 he held various positions up to marketing manager in the automotive, computer, and industrial product segments within the Monolithic Microsystems division of SGS-Thomson. Consequently, during this period, he accumulated a wide experience in the majority of semiconductor application segments.

In 1989, when the division became the Dedicated Products group, changing its focus and identity to better suit the needs of the market and of its customers, Mr. Cuomo was placed in charge of strategy and market development with responsibility over strategic marketing and strategic planning. Since then, a special mission he covered was that of developing strategic alliances with top customers.

Mr. Cuomo was born in Milano, Italy, in 1954, and studied at the Milano "Polotecnico" in nuclear sciences, with a special focus on analog electronics.

Panel: Next-Generation Digital Consumer—Electronics for the Mass Market

Todd Oseth

Vice President, Business Development,
Information Technologies of America Division,
Sony Microelectronics

Mr. Oseth has over 16 years of experience in the computer and semiconductor industry. He is currently vice president of Business Development for Sony's Information Technologies of America Division, located in San Jose, California. In this capacity, he is responsible for product planning, technology acquisition, and business modeling for Sony to maintain competitive advantages in the consumer computer marketplace.

Before joining Sony, Mr. Oseth was president of Enhanced Memory Systems, a subsidiary of Ramtron International, where he developed the first high-speed specialty DRAMs and computers that utilized the memory technology and established an international sales and marketing team. In addition, he worked as vice president of Business Development for Ramtron International where he was a business unit manager for ferroelectric nonvolatile memory, and assisted in raising capital for the company.

Mr. Oseth was a founder of two different start-up companies that focused on software simulations of large software applications and computing systems. The first start-up specialized in the hardware super computers required to execute the software. He started his career with Honeywell's System and Research Center in Minneapolis, Minnesota, where he held various engineering positions and became Honeywell's VHSIC Insertion program manager where he managed numerous IC and system product developments.

Mr. Oseth has a bachelor of science degree in electrical engineering and computer science from the University of Minnesota and a master's degree in business administration from the University of St. Thomas. He has published over 10 technical papers and has taught at the University of Minnesota and Johns Hopkins Applied Physics Labs.

Session # 8: Next Generation Digital Consumer Electronics for the Mass Market

Moderator:
Dale L. Ford
Senior Industry Analyst, Semiconductor Application Markets Program, Semiconductors Worldwide, Dataquest

Panelists:
Dwain Aidala
Vice President and Division Manager of Personal Mobile Communications Division, Mitsubishi Wireless Communications Inc.
David Taylor
Director of Marketing, C-Cube Microsystems
Andrea Cuomo
Vice President, Headquarters Marketing and Strategic Accounts, SGS-Thomson Microelectronics
Todd Oseth
Vice President of Business Development, Information Technologies of America Division, Sony Microelectronics

DALE FORD: On this slide, I have tried to represent the flow of new digital content coming into our homes. We can come down here and on the right hand side here, you will see what I would term as the digital wave coming into the home laden with content. We have made significant advances. On the upper end we have what is termed as fixed content, and in this fixed content we have moved from tape to optical. In between, we have gone through various other delivery forms, such as diskettes, and even prior to tape we had vinyl. We have also had, coming into the home, wireless delivery. We have gone from terrestrial-based systems to satellite systems and back again to a more advanced terrestrial systems. In the cable world, we have cable, the Telco companies, and even the utilities competing to deliver digital content into the home.

With that digital wave coming into the home, we also have the digital surfboard: the television, the video games, the set top boxes, and video CD players. in the future DVD players, and Internet access devices, providing the platform where consumers will be able to view this digital content ranging from movies and videos, to audio sound entertainment, video games, and the Internet access appliances.

In the consumer electronics world we have gone through a period of what might be considered a calm, even depressing, period of relatively flat growth in many mature markets, but if we look at

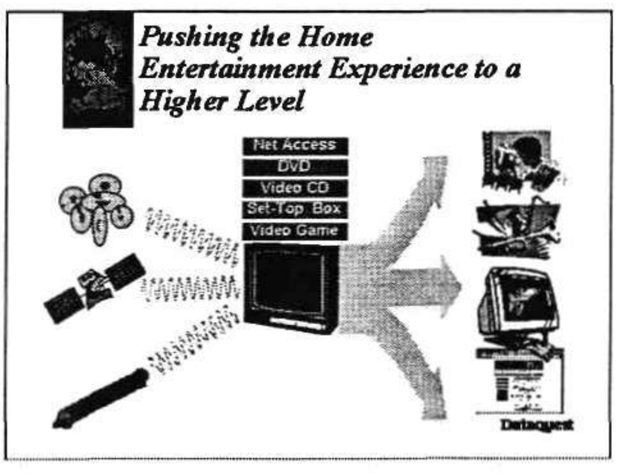

this new equipment as a digital surfboard, the surf is definitely up now, and we have new surfers ready to catch this wave, and ride it. Chip companies and content providers that are ready to go with it will reap great benefits.

That really is one of the hidden stories behind all of this – the chip. Many consumers are familiar now with the term "digital," and many consumers now equate the term "digital" to "good," or "better," or "best," but those of us in this room know that behind that "digital" is a semiconductor, and it is the semiconductor and the advances in the semiconductor that have enabled all of this progress. We have the opportunity, going forward, to reap the benefits that come from competing in this semiconductor market.

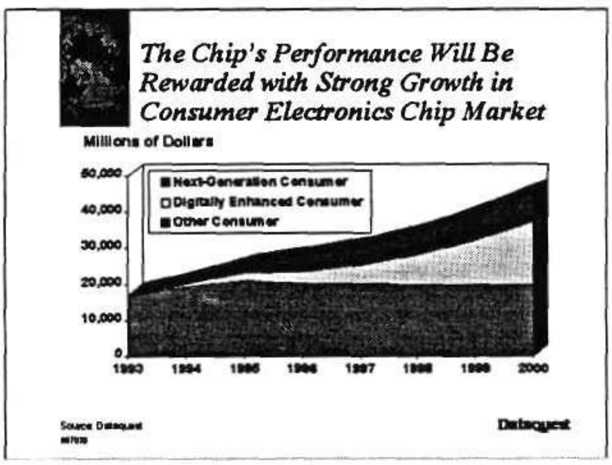

This is an updated slide reflecting the most recent Dataquest forecasts for chips consumed in consumer electronics equipment, and here we reflect the decline that took place in the market during 1996. Much of the decline that took place in 1996 came as a result of the drop in memory prices. Whereas in the PC we had bit flexibility, the ability to add increased memory content into the PC, and buffer some of the blow from the drop in DRAM prices, that same flexibility in adding additional memory content does not exist in most of the consumer electronics equipment. Of the \$1.9 billion drop from 1995 to 1996, \$1.7 billion of that came as a result of the decreased memory prices. That came in the midst of this calm, almost mildly depressed, consumer electronics market, that remained essentially flat from 1995 to '96, growing only a billion dollars, from \$170 billion to \$171 billion. So it was very difficult to keep the growth path of chips going into this market, with the flat equipment demand, and the decline in the memory prices.

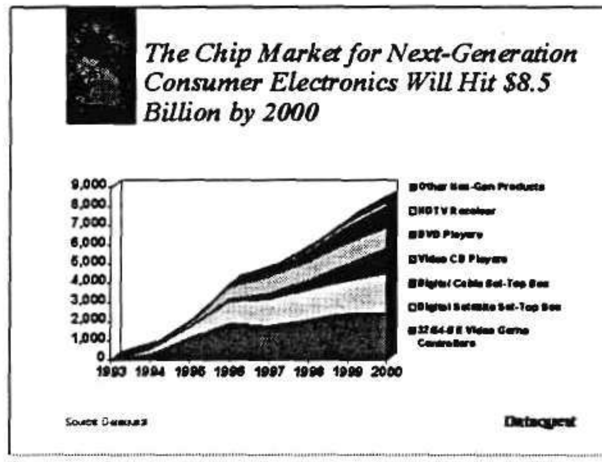

The decline in memory prices came at a very opportune time for this category of next generation consumer electronics. This next generation requires increased memory content and the high DRAM prices were creating a barrier to acceptance of these products at an acceptable price into the home of the consumer. In a way, the decline in memory prices was good medicine for this section of the industry. It was needed medicine that will enable the industry to grow, based on the growth of the next generation, and the digitally enhanced products, as we move toward the year 2000. Had memory prices remained where they were, it would be difficult to keep a forecast as optimistic as the one that we have today, where we have the solid growth in place of over almost 12% compound annual growth, from 1995 to the year 2000.

I will explain briefly the categories here. We have next generation consumer electronics, which consists of this new generation of equipment from direct broadcast satellites, digital cable, DVD players, video CD players, 32 and 64 bit video games, etc. There is also an important category called digitally-enhanced consumer electronics. These are products we are already familiar with, most of us have many of these products in our home, but now there are generations coming out that take advantage of the new digital age. So we now have camcorders that have motion compensation, and have digital zoom; VCRs that will come with video guide capability built in; televisions that have closed caption capability; and in the future, even Vchip in the U.S. All of these products really represent the growth opportunity for retailers selling these products, and the growth opportunity for semiconductor suppliers, supplying chips into this arena. The traditional consumer electronics market represents an essentially flat market going forward, with the decline in chip prices and the way these markets are shaping up.

To provide a little more detail on this category that I have called next generation consumer electronics – in this category, we have all the products that I have mentioned previously. We more than doubled the revenue for semiconductors from 1995 to 1996, going from $2 billion to $4.2 billion in products. Much of that jump came from the direct broadcast satellite, and the 32 and 64 bit video game generation that was introduced with Sega and Sony leading the way, and now Nintendo coming into the market. Moving forward, set top boxes will represent an important market opportunity by the year 2000 – 41% of this $8.3 billion pie. They will be followed closely by the 32 and 64 bit video game controllers at 21%. Video CD players especially in the Asia-Pacific region, in Japan, will capture 13%. DVD players, we expect to ramp quickly, with 11%, and then other products, such as HD TV, digital cameras, digital camcorders, digital VCRs, DVD audio, will represent the rest of our forecast.

Dwain Aidala, joins us from Mitsubishi Wireless Communications, where he is a Vice President and Division Manager of the Personal Mobile Communications Division of Mitsubishi Wireless Communications, which develops and markets wireless voice data telecommunications products and advanced mobile communications applications.

DWAIN AIDALA: I have spent 17 years of my career on your side of the fence, marketing and managing of semiconductor products, but for the last three years I have moved over to run a North American multimedia business incubation organization for Mitsubishi Electric, and today I am going to be talking to you as a customer, not as one of you, the suppliers of semiconductors.

We talk about consumer electronics and its movement to digital. There are two main changes that are ongoing at this point in time. One is the movement from analog to digital, and the other is the movement from stand-alone to network appliances. Just as the PC computer marketplace continues to become more network oriented, that trend is starting in the consumer electronics marketplace, so I have titled my small, short presentation as "The North American Multimedia Dilemma," because that movement from analog to digital, and then more importantly, from stand-alone to network, has very important ramifications on what we need as an end customer, and what you need to supply as a semiconductor supplier.

That movement from network digital could be one of the reasons for the flatness in the consumer electronics marketplace. If you look at the delivery of interactive TV into the home, with the different standards of DBS, switch digital video, HFC, and how those standards come into play, to some extent, come from who the players are. Is it a telco? Is it a cable company? Or is it a new industry such as what has happened in direct broadcast satellite? As well MMDS is coming into play. In those four areas of networked consumer electronics there are four major infrastructure investments that are not only driven by semiconductor content that is out there, but, more importantly, by people defining what the end consumer will pay for, and what will finance those infrastructures.

Additionally, continuing the confusion factor in the North American marketplace is the battle that will be waged between the TV as the gateway and PC as the gateway; and the hybrid device that gets developed – an Internet TV or is it a PC TV? From the Internet side, in that network, do you start providing an Internet phone capability, or do you move into the wireless side?

On the wireless infrastructure there are other buzzwords: TDMA, CDMA, GSM, PAK, FLEX, and many others. As you get into broadcast TV: Is it HD TV? Is it digital NTSC? All of these provide very significant challenges and obstacles to the growth of our business.

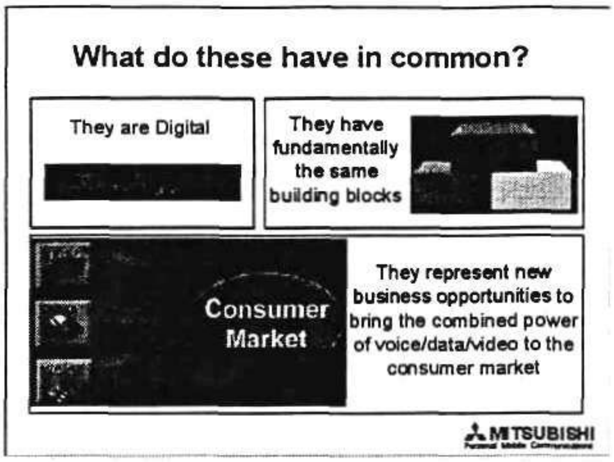

What do they all have in common? They all are digital. They all have fundamentally the same building blocks, and they all represent brand new business opportunities for integrating voice, data, and/or video for the consumer electronic marketplace. It is a huge marketplace from end equipment, and therefore, from the semiconductor side, but they definitely have obstacles.

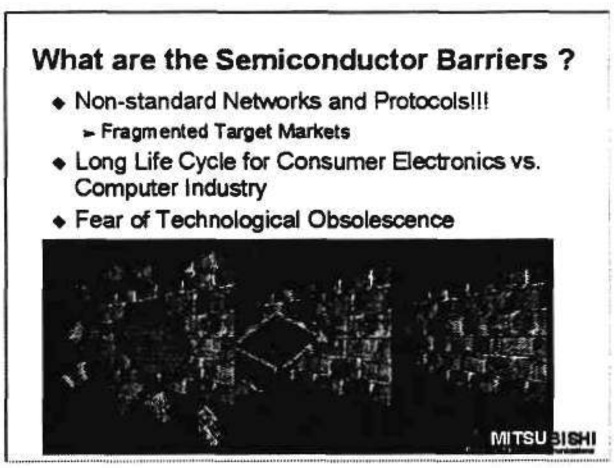

The barrier to semiconductors is the non-standard standards of networks and protocols. This entails several things: one is a very fragmented marketplace, and it also leads to a high confusion factor in the end customer. The life cycle for consumer electronics is significantly different than what has historically been in the computer industry, what has been primarily targeted at the corporate marketplace. When we sell a large screen TV, we do not expect to see that customer for somewhere between five to seven years. Even audio equipment, that cycle is typically three to five years – much different than the two-year, going to twice per year, type of product life cycles that we see in the semiconductor industry. Those things have to be taken into account, and our business planning, as we develop and build new equipment and even more important, business models for new equipment. Additionally, in the consumer marketplace, because of the rapid changes in the semiconductor side, there is a fear of technical obsolescence in the consumer and in fact, that also is a factor in retarding their purchase decisions. All of these are factors that we, as equipment suppliers, and you, as a component supplier, need to take into effect in your planning and in your product development.

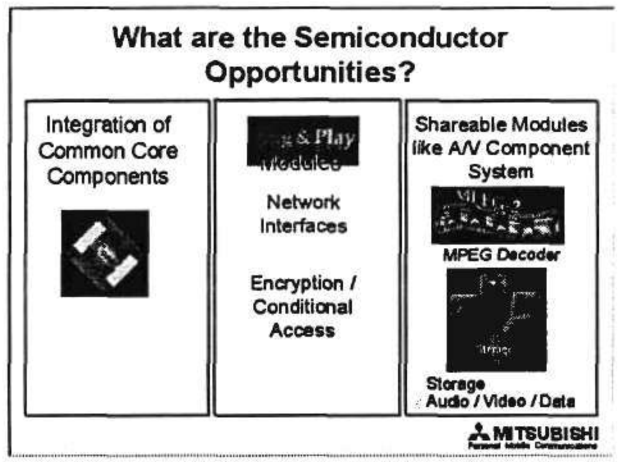

On the positive side there are definitely opportunities. What we would be looking for from semiconductor suppliers, including our own in-house capabilities, is integration of core components. The system on a chip, is definitely an opportunity for the common parts. We have the ability to move out of the multiple devices, into continuous integration for the right set of components, and part of that is choosing the right set. Are you able to, in the set top box area for example, build a digital entertainment terminal, or a DET, with the ability to put network interface modules on the outside? There are opportunities in the integration of core components, and there are also opportunities of building plug and play modules in the consumer electronics area; again, carrying over some of the benefits of what we have learned from the PC and computer industry.

One of those opportunities, is network interfaces. Another area which changes from component to component from the equipment side, is the encryption and conditional access capabilities. In

North America we have multiple standards. Other geographic regions such as Europe, have been working toward a more common standard. Finally, the ability to take key technologies, such as storage, and move those technologies into a very modular approach at the end equipment side, such as the AV component side, whereas things like MPEG decode capability probably should be built inside the box. The message that I am trying to present here for you, as semiconductor suppliers, and for diskussion on the panel, is: How do you approach this diversity of standards in a way that is cost-effective, but still leads to a rapid growth in the marketplace?

DALE: Our next speaker is David Taylor, from C-Cube Microsystems. David joined C-Cube in 1992 as a Director of Marketing for broadcast products. While at C-Cube he has spearheaded the marketing activities for the company's MPEG-1 and MPEG-2 encoders as well as its set top decoders. His responsibilities include business and product management, market development, and product definition for the broadcast, video teleconferencing, and desktop video marketplace.

DAVID TAYLOR: While the end result of our products is typically semiconductors, we focus on some standards for video compression to provide products that compress and decompress video, usually motion videos, sometimes still-frames, sometimes video teleconferencing as well. We address markets like the broadcast industry and consumer markets like video CD, and the DVD market as well. We are very focused on those areas, and in 1995 we had the fortune of getting awarded an Emmy from the Academy of Television Arts and Sciences for the work that we did on MPEG coding.

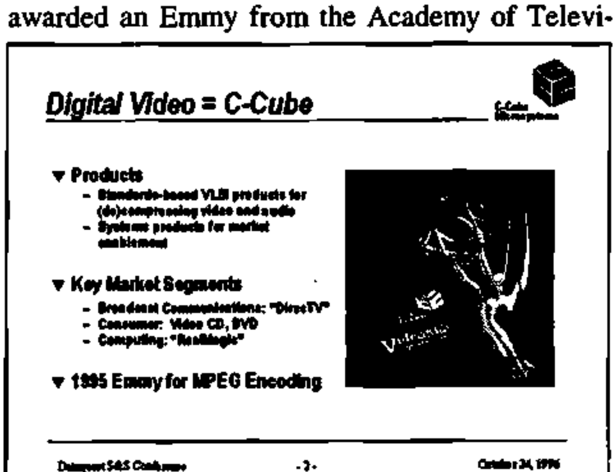

Before I talk about MPEG, I wanted to bring you back in time and show how some consumer products ultimately became adopted by the market and some of the timeframes that it took to achieve adoption.

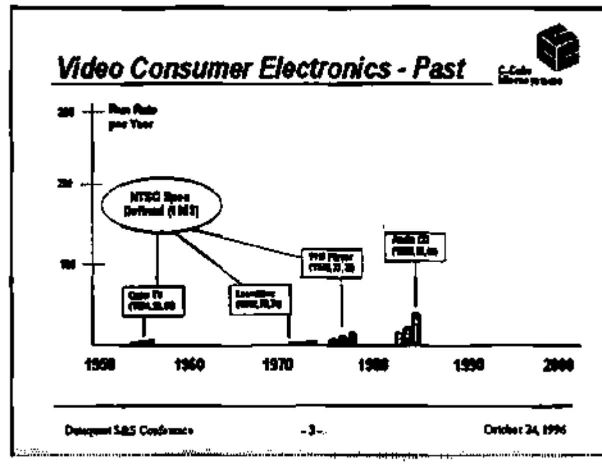

We have to go back over 40 years to see the adoption of the North American color television standard known as NTSC. It took some fifteen years from that point until color televisions became even somewhat mainstream and it was into the '70s before there was dramatic penetration. We saw laserdisk products introduced in the early '70s, we saw VHS players introduced in the late '70s, all took a significant amount of time, and most have a market saturation at this point except for laser disks. Audio CD was faster when it was introduced in the early '80s but it still took three years or so before we saw a million units ship. A slightly different market, it was not necessarily an international standard; it was certainly a de facto standard. But it takes time for products like this to achieve mainstream success in the consumer marketplace.

In 1992 the International Standards Organization ratified a standard called MPEG-1. It took a few years before the first products were introduced using this technology. In fact, the first products that really have become successful in a fairly narrow market – that being China, a large market, but a fairly narrow market. In 1995 there were some two million plus video CD players shipped into China. While 1996 is not yet over, we are expecting that perhaps something near seven million players will be shipped into China this year, based on this video CD technology, using MPEG-1. It turns out that these players are extensions of the concept of an audio CD player, except that they have video on these disks as well

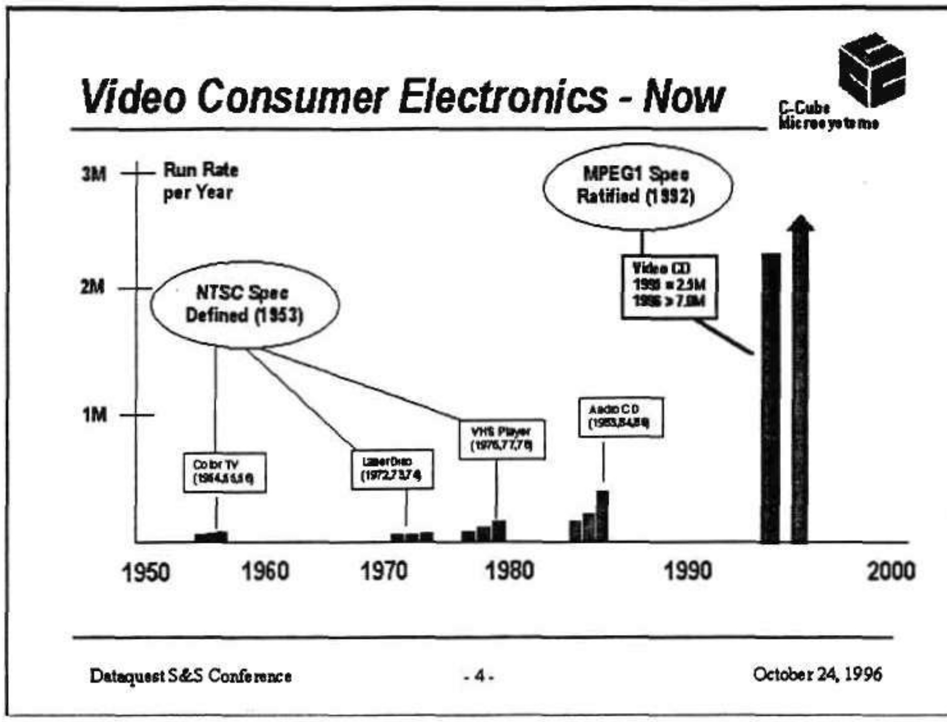

The timeframe for adoption and the final numbers, become very large, very quickly with some of the new technology. People understand video today, people understand disks today, so some of the technological barriers, or even psychological barriers, are reduced.

In 1994 ISO ratified what is known as the MPEG-2 standard. MPEG-2 addresses a slightly different market. MPEG-1 addressed a low bit rate, CD-ROM style market, 1.5 megabit per second, which is basically a CD-ROM data rate. MPEG-2 is designed to address markets that can handle higher bandwidths, say 2.5 to 3 megabits and on up, and it was intended for broadcast-style applications. There were 140 companies that ratified the MPEG-2 standard with zero abstentions and zero nays. It was probably one of the most successful standard activities ever. In fact, the same year that MPEG-2 was ratified, DirecTV began shipping boxes, with the DSS system. It was reported at the time, to be the fastest consumer electronic product, to a million units of at least this caliber of product. There are some two million units perhaps for DirecTV in this year, and this year

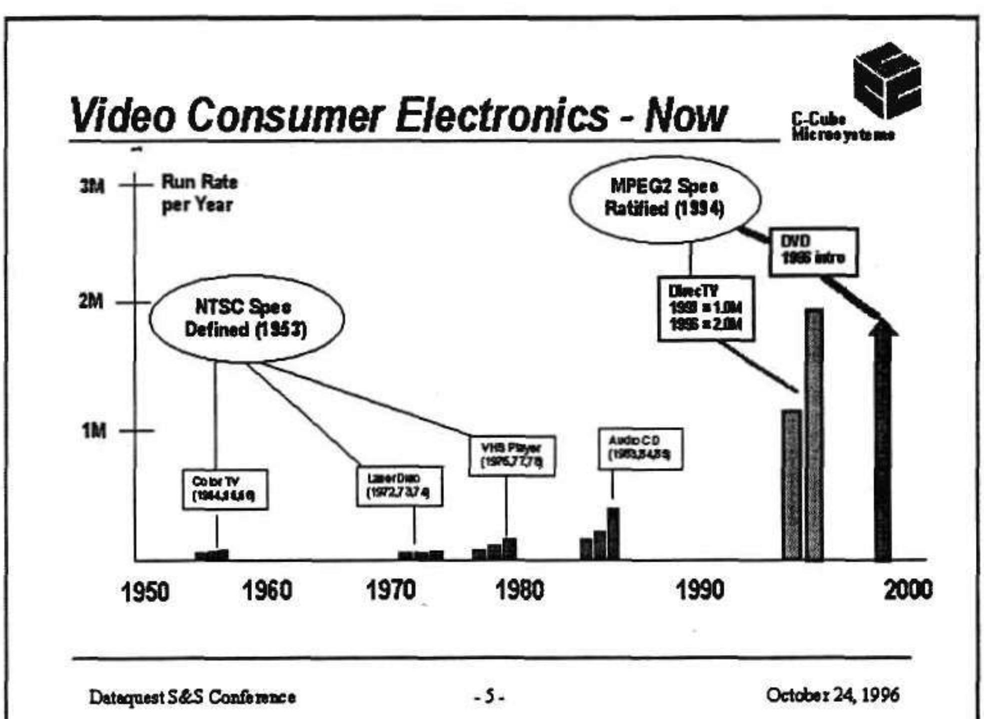

perhaps four to five million set top boxes within the same mold as DirecTV. It is clearly an international success story, not just a U.S. success story.

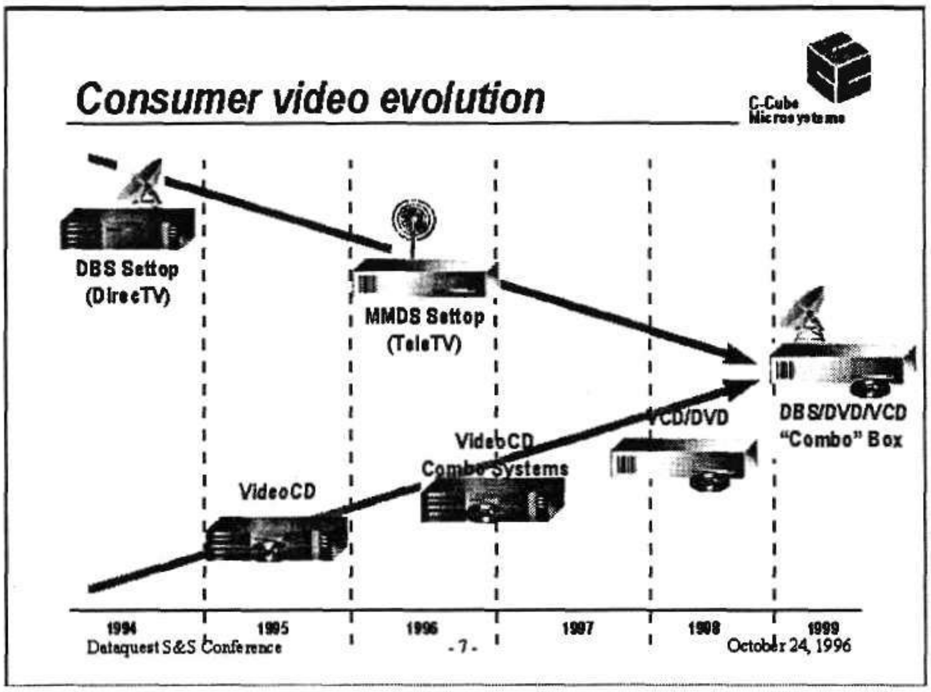

DVD has not yet been introduced. Samsung in Korea is planning to introduce boxes soon, but DVD is going to start in 1997 so that will be the opportunity we have for this new technology.

What is MPEG? It is an international standard designed to represent how you can compress video and audio information. MPEG defined a syntax so that everybody in the world, could build decoders and have them all interoperate. The decoders we build, and SGS-Thomson and LSI Logic, and others build, can all handle MPEG data and do so in a similar fashion.

Its purpose is to allow us to now use video essentially as a data type. We are very visual people still. The ability to convey information in text form and video form is compelling. We can now do this in a lot of different ways, and MPEG provides the baseline for that. It also has a significant future. MPEG is a steamroller as a foundation technology, and likely will still be here 20 years from now. NTSC has been around for 40 years, and the HDTV standard is based on MPEG-2, in fact, there is what is known as a profile, an extension to the MPEG standard, for HDTV as well. The availability of digital NTSC, and the quality it provides, may have the effect of potentially delaying any introductions of high definition as well. It is going to be around for a long time.

What I have shown here are some boxes that do exist, and boxes that likely will exist, and some that might exist, as we look out over time. We certainly see video CD players today. We certainly see DSS boxes, not just by RCA, but there are some eight or nine licensees now of the DirecTV system. The disadvantage is, there is no local content. We are going to see some convergence, from playback boxes, to cable boxes, to Internet access boxes, as we look forward as well. It is hard to predict exactly the feature set and functionality of that integration, but we do expect it.

MPEG-2 is a foundation technology on which a lot of new applications will be built It is the interactive capability, as we look forward, that has even more significance, to allow us to meld Internet and other styles of interactive applications together as well. It is a very exciting future. There is dramatic semiconductor content within all of these, as we look forward.

DALE: Our next panel member joins us from SGS-Thomson Microelectronics, Andrea Cuomo. Mr. Cuomo was called to become the Headquarters Marketing and Strategic Accounts Director in 1994, and in this position he has responsibility for corporate strategic marketing, which identifies future semiconductor trends, and develops advanced systems, as well as combining the diversified know-how and wide range of technologies of the company. Mr. Cuomo also holds responsibility over sales to the company's strategic accounts

ANDREA CUOMO: The real reason why Dale called me here is not for my job title, whatever I do in life, but just because I am an Italian head of a leading semiconductor company, and multimedia is looking more and more like Italian politics. Everybody talks about it, few people understand what it is, all of us feel it when we have to pay taxes. Sometimes we can even smell it, and I know one can digest the noise which is built around it, multimedia is communication, and content.

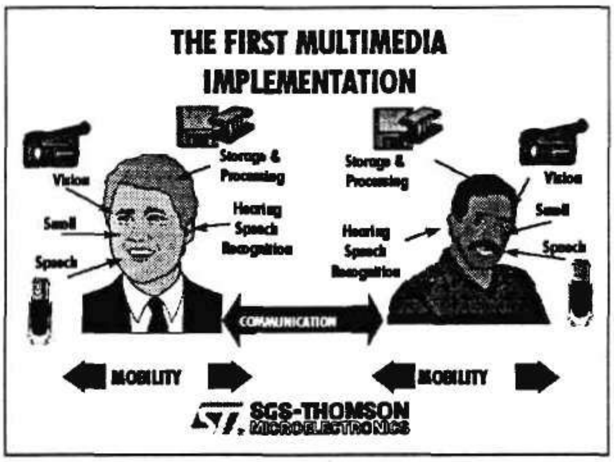

Look at this example. If both guys communicating are using different media, and one communicates only through CNN, and the other has no intention to listen to what he is talking about, there is no communication possible. In the same way, if we look at the old multimedia pictures, only Hollywood, and-the network companies, will make it happen, and it will change our lives. Just think about one aspect of the revolution happening in this world. Today, each of us, with a $3,000 PC, can produce a movie featuring, for instance, Bill Gates kissing Princess Diana, and it would look authentic. It could even threaten the throne of England. This technology will change some of the roots of our lives, as we are raised knowing that what we see, is true. Modifying the basic psychological world is going to be quite a challenge for everyone, and our role as semiconductor companies is to fuel this revolution, first for the benefit of mankind, and as a side effect we would also like to milk every single dollar out of it. Over the past 20 years, the rules for semiconductors have been the same: cheaper, smaller, faster, cooler. By building cheaper and more integrated silicon chips, we have made possible the economy of scale that will bring multimedia to mass success in set top boxes, and here is an example. In two years, we at SGS-Thomson went from eight down to two chips, cutting costs by a factor of three, and helping the worldwide leadership of our partners.

The real challenge for the semiconductor industry will be around the terminals, where most of the dollars are. Today we see DVD players, set top boxes, network set top boxes, digital VCRs, network computers, not to mention digital still cameras, digital VCRs, and the automotive car information systems. What will succeed is a good subject for a long discussion, and I am afraid it is not going to be any more productive than an argument about the sex of angels.

We are at the beginning of a new era, and different products will appear in the marketplace, and the consumer alone will, in the end, decide which one will win. Not forgetting that the consumer is not only the computer nerds, teammates I have in Silicon Valley, but also my old aunt, who lives in the country, and whose contribution to the inertia of the market goes often underscored.

What can we do in this changing environment? First, we need to focus on key technologies, such as microcontrollers, MPEG decoders, analog signal demodulators, signal re-creation algorithms, and embedded memories. Understanding eight years ago that MPEG was going to be one of the core technologies for the future, has led SGS-Thomson to a focused investment, and it paid off a worldwide leadership in MPEG chips, which is in excess of 40%, and 70% on MPEG-2.

Datacom technologies, including modems, ISDN, XDSL, ATM, wireless, and whatever is going to increase bandwidth and mobility will be the other key component for the success of multimedia, and we are back, once again, to the basic two technologies, which are communication, and display content: two markets in which we have leadership by working closely with leading partners in the world. The second pillar is to have the most advanced simulation tools, all the way from the systems, down to simulate the behavior of the chip in the system, both the simulated system and the real system.

These two elements, mastering the basic technologies, and having advanced simulation tools, the ability to bring chips quickly to the market, together with heavy close links to customers, leads us today to have systems on the chip, and to lead the fast-changing market. We shall benefit from the early phase of new products, building the profit when it is there, at the beginning of the life of product, and rushing down the manufacturing learning curve, before competition can catch up. We have to take profit from the advantage, in order to kill our own products, and build new products with more competitive devices, to continuously raise the standards of the race.

This is the name of the game: accelerate, accelerate, accelerate, to be early in the market and maintain your competitive advantage. You need to have a lean and empowered organization, with a strong, unified strategic drive, a larger local economy, to keep pace with the market, and strong and entrepreneurial leaders, able to successfully take the lead and manage, in the new race, and we, in our company, are ready. We have the technology, we have the tools, we have the people. No industry in the world has the pace, the fierce competition, and the globality of semiconductors, so I am sure that when the challenge comes, our industry shall succeed. To paraphrase John Belushi, 'when the going gets tough the tough go shopping,' and I will be glad to lead the happy crowd to the nearest mall to buy the latest multimedia product featuring the most unimaginable semiconductor technology.

DALE: Our final presenter for this session is Todd Oseth, who comes to us from Sony Microelectronics. Todd is currently the Vice President of Business Development for Sony's Information Technologies of America Division, located in San Jose, and in this capacity he is responsible for product planning, technology acquisition, and business modeling for Sony to maintain competitive advantages in the consumer computer marketplace. He is also responsible for managing the marketing and sales effort of Sony's most recent introduction of a consumer PC. Prior to Sony he was President of Enhanced Memory Systems, a subsidiary of Ramtron.

TODD OSETH: In the consumer market space, consumers require all the functionality that you can dream of, and they want it for free. We have to continue to fight that battle as we develop new and better products, every single time we go to market. Some of the new, next generation of digital consumers are people that have never experienced the linear or analog worlds. They started with computers, and will start to represent and bring together new marketplaces and new types of products. While the semiconductor has a lot to do with the next generation display technologies, we need to realize how much that will drive the future. We saw earlier that digital television and HDTV are some things that have been bogged down for years inside different political environments. There is a good chance that in 1997, digital TV will finally get ratified as a standard. With that, is going to bring a whole new

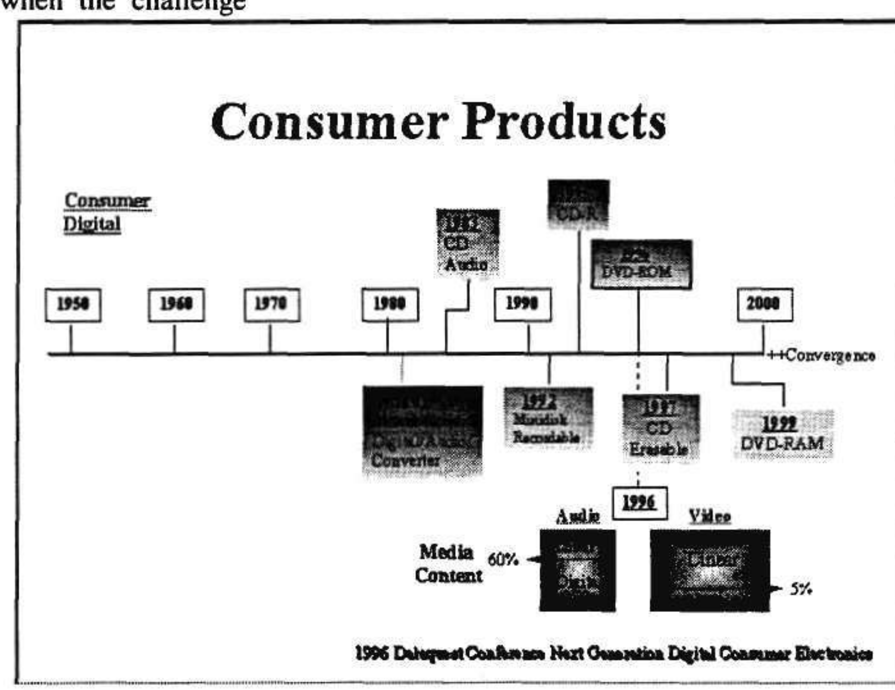

wave of new types of content, new ways of using the content, and a new experience for the users.

Let's quickly go through of a consumer products road map. We look at how long it takes to take a technology, move it forward. We see when CD audio was actually brought to market. We see when the actual CD got brought into the computer, in 1987 the response was: Every computer will have a CD-ROM. It was not until later in the 90s before CD-ROMs actually got moved into the computer space. We also start to hear a lot about new technologies. We have CDR, recordable-type technologies that use the CD formats. Everyone has been awaiting the world of DVD ROM. Remember that DVD ROM is a read-only material. Today it only does playback. The opportunity that this set in front of us is DVD RAM, which will not be available for another few years. This then allows us to give the interoperability and the capabilities of our existing systems, but on a single disk platform of 3.5 to 4 gig.

Consumer Electronic Functions

Gaming		Video		Audio		Compression/ Encryption	
Function	Performance	Function	Performance	Function	Performance	Function	Performance
•Current 3D	+40m pixels/sec >1m polygons/sec	NTSC	Low, but with Enhancement High	3D	8 MIPS	JPEG	20 MIPS
•NextGen 3D	>2.0m polygons/sec >100m pixels/sec	SVGA	4-8 Mbs	SRS	12 MIPS	MPEG I Enc	90 MIPS 185 MIPS
						MPEG II	200 MIPS
•Interactive 3D	>5.0m polygons/sec >300m pixels/sec	DTV	8-20 Mbs	AC-3	60 MIPS	MPEG IV	????

-Operations/Pixel and Fill Rate are Key Metrics

Greatest thing about SW: It Will Always Use the Cycles

1996 Dataquest Conference Next Generation Digital Consumer Electronics

One of the other major barriers that one has to deal with in the consumer space, is the media itself. How much of it today is on traditional linear type technologies and tape? In the audio world, we have at least started to produce a lot of digital formatted content. Unfortunately, in the video side, the majority of the content is still in linear-based solution. It will take a number of years for this technology to really gain the momentum, but during that time period, that is when all of these markets are going to expand.

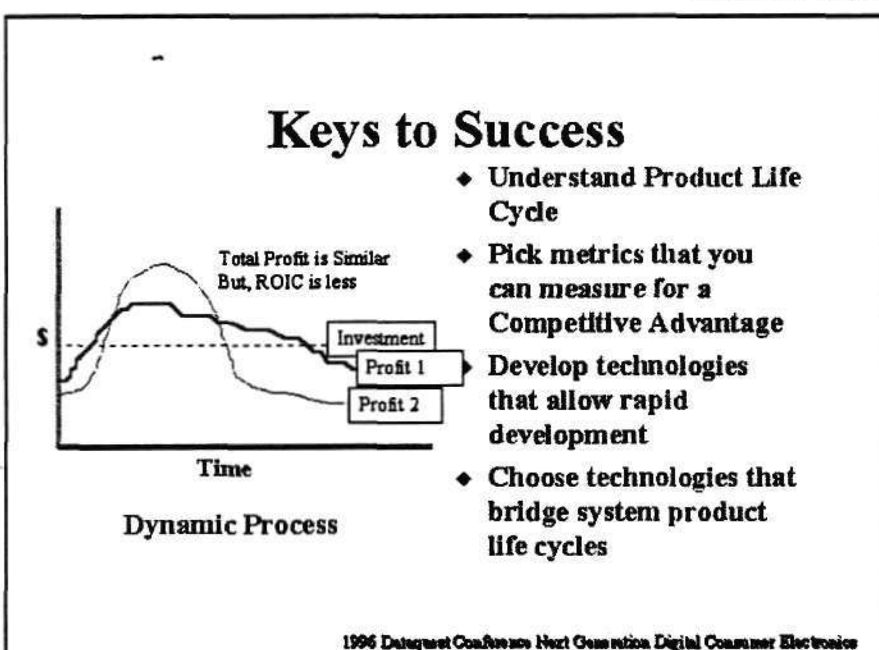

What does convergence mean in consumer products? Convergence brings out the concept of a new experience. How can you integrate between multiple different platforms and different pieces of hardware and software, and make it better for the user? That is one of the big challenges.

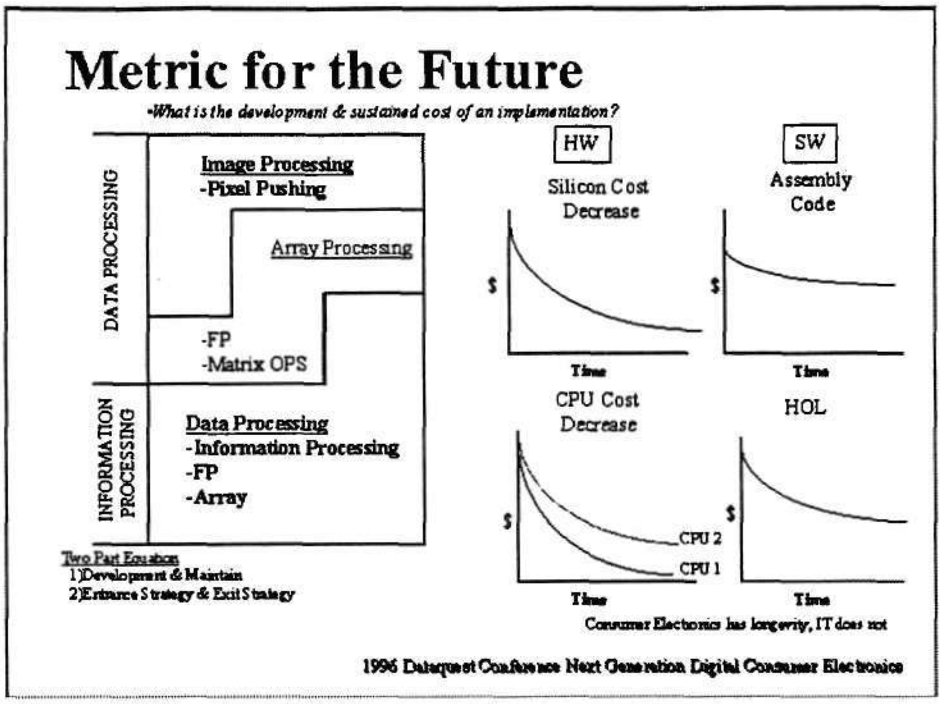

The keys to success, especially from a semiconductor side or from a systems side. People do not expect machines, especially in the consumer world, to be replaced every 18 to 24 months. They expect them to be around for five or seven years. Unfortunately, where the majority of that profit is to be made is in the first twelve months. So many of our older strategies of continually improving the costs are not necessarily going to help us in a rapidly changing consumer electronic environment. We are going to have to pick metrics that allow us to identify what is important to our consumer marketplace. We like to use "function per instruction per dollar," because at any one point in time, there are multiple different technologies that are available, that can answer and satisfy a specific positioning. We need to make sure that when we develop these technologies in the semiconductor area, we truly appreciate how rapidly they have to be integrated. We cannot wait twelve to eighteen months for a new product to be developed. The system on a chip is a good idea, as long as it can be done quickly. If it cannot, we will find ourselves using alternative technologies, other semiconductor technologies, and building systems around them. Those technology bridges, whether it be Internet or MPEG, are very important to continue to proliferate through different types of products. That is one investment area that you can see and actually get a return on investment for a much longer period of time.

We look at the functions that consumer electronics take on. In the family room, what is seen is gaming, video, audio, and a lot of the underlying technology that you see are compression and encryption techniques. In the gaming side, we today have things like Sony PlayStation, Sega Genesis, and Nintendo. The types of 3D performance that are required are staggering in order to make this experience real and reusable. The next generation of 3D technologies required for the consumer will be pushing the envelope further, upwards to two million polygons, fully filtered. We are talking about, in the next generation after that, by the year 2002, fully interactive 3D, and what does that mean? That means when you sit behind a machine, you have actually joined the game. You are not in a static position, but you are interacting and it is taking your responses and you are living inside of what that 3D experience is all about.

In the world of video today, we have a lot of NTSC, but there is a huge retro market that we have to consider. Not everyone is going to buy new televisions, but there are some five hundred million televisions in the United States available to be tapped for this new digital revolution. We need to make sure that we can bridge that path. As we move into higher and higher technologies, SVGAs, digital TVs, the overall performance requirements that are going to be set forth, especially for memory, in speed and access, as well as in density. The audio side – we have already found ourselves talking about Dolby AC3, surround-sounds, 3D. All of these take up significant amounts of cycles, can be satisfied either with individual DSP-type functions on DSP chips,

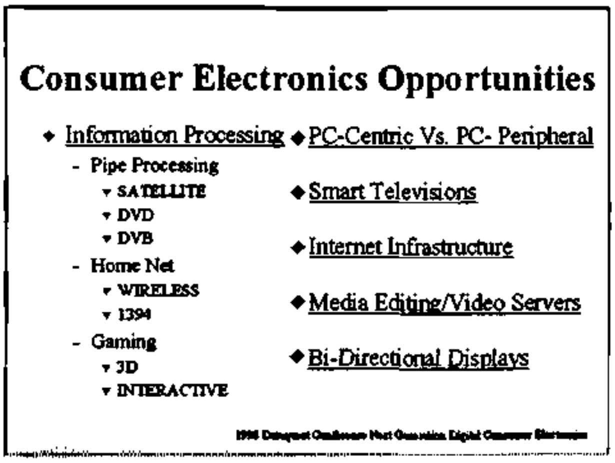

and/or hard-wired chips. We have got to be prepared to do it all.

There are actually a number of different performance metrics that we use to say what can fit in a CPU, and what should not fit in a CPU. We go back to the world of function per instruction per dollar. The compression technologies we have heard a lot about: MPEG-1, MPEG-2. We also now have to start worrying about encoding technologies as we start to get a more realistic experience with some of the games. MPEG-4, as we start to move into the next generation of broadcast. The metrics for the future – we are always continuing to identify new and better areas to evaluate. As a semiconductor producer you no longer can just worry about the cost of the chip? You have to worry about the entire system cost. What does it cost to develop and what does it cost to maintain? You look at your assembly codes, you look at your high order languages. What can be done by CPU?

In the diagram, you see kind of a broken out process of image processing, array processing, and data processing, and today we try to do too much potentially in one type of processor versus another. What we started back in the late '70s and early '80s, with DSPs, start to take on larger function in our overall scheme of things. We already know that the data processor today, through software, is able to do a lot more functions than a while ago, but we need to make sure that we divide up our architectures, and make sure that we are getting the best function per instruction dollar.

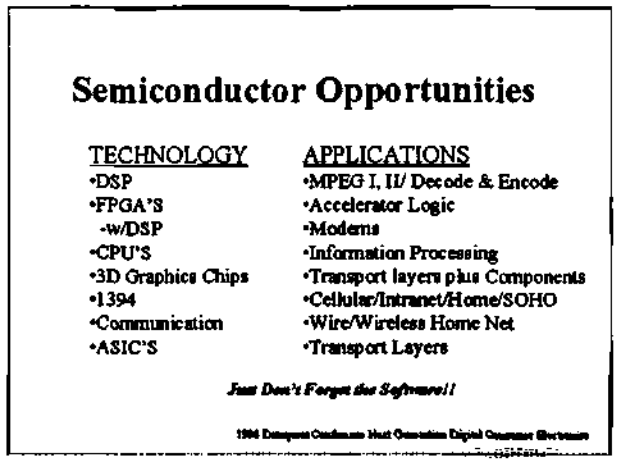

The opportunities – they are in the information pipe. These different pipes, whether they come from satellite, DVD, DVB, all require different mechanisms to actually transport mechanisms to be brought into the home. There are actually interface chips in front of every one of these big MPEG-2 decode capabilities, and those are all pure opportunities for you. We look at the home net. We have heard about wireless technologies. Most important that we get the same types of performance and speed in the world of wireless. 1394 will be one of the major changes in the home that we have seen in a long time. 1394 is a new audio and video standard that a number of companies are starting to support, and it will be the way that we network our AV equipment, in and out of PCs, as well as smart televisions, digital cameras, and digital camcorders.

We cannot forget the gaming because the entertainment side is extremely important. There is a real opportunity for all of us, in the world of PC-centered versus PC as a peripheral. Of course, the PC companies, of which we are now one, would like to see the world rotate around the PC as the center of the home, but I would sure hate to make a requirement, in order to pay for a $500 camcorder, that you have to buy a $4000 PC. It does not seem to make a lot of sense. Smart televisions are going to be available starting in '97 and '98, where it can take on some of the networking functions. The Internet infrastructure has to be fixed. From a consumer standpoint, our customers expect going to web pages at the same rate they go through channels of television. We are not even close. The overall media editing, the capabilities to take the pictures that you developed, and put them into scrapbooks, and send them around to your family. All of these technologies have to be brought together.

From our standpoint inside of Sony, we see that the DSP world is extremely important because of the functionality that it brings, and the overall performance, and the costs. FPGAs to allow us to rapidly get into the market, get through the product life cycle, through the profit life cycle, and then move into custom design if the product merits. We will continue to push the CPUs as far as we possibly can, and we will always be the ones that are pushing Intel to make even faster machines. 3D graphics, communications – these are the areas that are going to make this whole multimedia consumer marketplace come together, and we hope all of you in the semiconductor arena, are able to provide some of those functions for us.

Q: Will the market for digital video components, such as MPEG decoders, evolve as a commodity market, with many vendors providing components with little differentiation and low margins, or, as a more proprietary market with lots of product differentiation amongst vendors?.

DAVE: The market will evolve in two ways. MPEG itself was designed to be a commodity so that all semiconductor manufacturers could build products that are consistent with the standard. But MPEG itself is worth 30 points of margin. That is very important in the sense that MPEG implementation becomes a science, but the differentiation will be through adding features, that go around the components, that add value to the system. There really are going to be two paths there, but ultimately the value-added path is one of integration and functionality beyond the basic definition of what MPEG is.

ANDREA: I am in total agreement. We will see the two paths: the stand-alone MPEG, and we will see MPEG integrated with other functions as one cell of the library. There is a third path done by Microsoft, which is on the PC. MPEG would become just a routine on the main processor, so we will see different things in different applications. As usual, the application and the end market is driving the semiconductors.

Q: We have been long awaiting the launch of DVD, and have had many promises over the year regarding DVD, and as we look at the future forecasts, much of the future potential and growth will come when it becomes a true VCR replacement, which means it has recording capability. "When do you believe recordable DVD products will be available, and what has to happen in order for those products to come to the market?"

TODD: Right now the DVD that are available, and realistic to have in the home, are two different questions. It will probably be available in the next year and a half. You will see DVD RAM. To really make it a consumer product, we are looking at late 1999, probably the year 2000. It will require changes in the overall marketplace today, as we get over the political changes that we are still faced with.

ANDREA: We will see it first in the PC. The most advanced work and closest to the market are people working in the PC area to use it as a mass storage device.

DAVID: I agree with Andrea. In terms of DVD ROMs, we will see a lot in the PC space. The consumer space is where we are going to see video exploited on DVD first. Why? It is content. The DVD player is a known application. You put movies and audio on DVD, it is an easy playback mechanism. It does what CD audio did to the LP. It will be successful as a playback-only medium for the same reason that CD audio was successful, and DVD RAM will create new applications, maybe ultimately replacing just the playback application itself. Content is going to be a problem in the PC space, short-term, because movies on PCs are not interesting.

Q: As we look at MPEG-2, AC3, and, in general, the audio/video processing that is required in a DVD player and a digital set top box, when, if ever, do we get to a single chip solution for these types of products, or is it even an economically feasible solution to plan on a single chip product in these new consumer generations?

DAVID: We will absolutely get towards single-chip solutions. It is within the set top box arena, we are now integrating audio/video, transport

CPU devices, graphics, and that is the next generation. That is coming out in the next few months from a variety of semiconductor companies. In the DVD arena, we are one step behind, because the standards have been set a little bit later, but clearly all of my colleagues and competitors are also doing video and audio decoders that include MPEG audio, and AC3 audio, and DVD interface circuitry, and so forth, to the DVD ROM. If it is not in 1997, it will be the generation after, where everything will be single chip.

ANDREA: If the application is static, next year we will have the single chips. If you start adding, for example, cable modems, if you start adding graphics to the set top box, it will not be single chip. A single chip solution for the present set top box will be available sometime by the end of next year. A single chip solution for the set top box, I do not know when.

TODD: Right now one could take a high end Pentium and do a poor implementation of MPEG-2, and still do some AC3, so maybe the answer is that it actually is available today, but it is probably not a very realistic scenario in order to implement it. In order to get the kind of quality one needs, you cannot do that in the processor. Dedicated pieces of hardware will be required, and with that, probably both of them together in the next one to two years make an awful lot of sense for stand-alone consumer electronics.

DWAIN: Our expectation is that the network set top boxes will evolve into probably a three-chip system: main data processing component, graphics audio processing component, and an interface component. We would expect that in the '98 timeframe.

Q: Do you see the potential that we will have a situation where the U.S. has multiple standards, but the rest of the world comes around to one standard, DVB, much like it did on GSM?

DWAIN: We have gone in our communication infrastructure, whether it is telephony or broadcast, from a common standard that has driven the volumes into this fragmented capability where the market totally defines itself. We will suffer from that, and we do have a strong potential of DVB and GSM to be an overwhelming worldwide influence that is not as strong here in North America.

DAVID: Excellent question. In fact, not all DVB is created alike either. DVB in Japan and DVB in Europe actually are implemented slightly differently as well, and that is a problem. Certainly DirecTV, Echostar Dish, Alphastar, do not talk to one another. That is a problem. I do not know whether all of them ultimately will be able to survive. Interoperability, or lack thereof, between them is certainly a contributing factor.

Q: We have seen the great difficulty of coming to agreement on one standard to launch a major new product, as we have observed the DVD process, and we saw that process drag out over multiple years, even after there was agreement that we would come to one standard. Does the experience in trying to launch DVD in the market tell us that it will be impossible in the future to agree on common standards to keep this growth going in digital consumer electronics, or can we overcome that with launching new products in the future?

TODD: We can expect those standards to be generated, but continue to go through as much pain as we have in setting up the DVD standards. We are talking about hundreds and hundreds of billions of dollars of opportunity here. Everyone is going to have their own opinion on what is the precise mechanism in order to satisfy the requirements. I cannot expect it to be any different on the next generation as this, but we will continue to have standards set, so that we all can serve those marketplaces equally.

ANDREA: DVD has been an improvement. If we go back to camcorders, there were three different standards and one winning de facto. This time people have been discussing, have been trying to come to an agreement, and to come to one standard. It is certainly going toward an improvement. The strong force here has been Hollywood, who, I am sure, do not want to set, put three or four different products in the marketplace. That is where the standard would be driven, and where people would be forced to go to one standard.

Q: Will we ever get to an HDTV standard in the United States, or will we back off and move to more of an SDTV standard as they are pursuing in Europe?

TODD: We will get to that standard in 1997, and it will be similar to what the grand alliance had proposed. The whole broadcasting scheme has made it so that much of the content that people receive today, they believe is a right, and in order to put in a standard that does not allow for any of the history of equipment that you have, is going to upset a very large part of the marketplace. While it may have some problems, we will have a standard that is more similar to what the grand alliance has proposed, and we are hoping to see the White House back it up early in 1997.

DAVID: While we may have an HD standard in '97, the existence of a standard will not guarantee its success in the marketplace. If you ask the average consumer what the problem with television is, he does not say, "I sure wish I had a better picture." He says, "I sure wish I had better programming." Once we are fully deployed with a lot of the digital NTSC systems, at 16 by 9 resolution, and with dramatically better picture quality than today, most people are going to say, "Wow, that looks great," and the change-over to HDTV demands an economic model that I am not quite sure how it gets paid for at this point.

Q: Is it really necessary for government involvement to take place to enable this digital TV era, or will the consumer market be able to drive this forward without intervention by the government?

TODD: The consumer marketplace, with all of its market dynamics will allow, at the right price point, for all to come together. Unfortunately, it is going to be a very low price point in order to access a lot of the media that is out there today. If the government gets involved, it will not be done right, and we probably will find ourselves with the standard, not working.

Q: If the PC does not become the center of the home entertainment experience, with the set top box, the DVD, the digital camcorder, etcetera, what does become the center of that home entertainment experience

ANDREA: Why should there be a center? People often are underevaluating the inertia of the market. When I need to buy a new TV, I am buying a TV, not necessarily a new PC, so there will be a bit of everything, I think. There will be a possibility for the set top box to hook to the Internet and be able to communicate. We will be able to go directly into from a camcorder, to a TV or to a printer, and print images, and at the same time able go into a PC. We will see a bit of everything, as normally.

TODD: I would agree. What happens is that the power is shifted from these boxes, or displays, or PCs, to the home net. The network itself is where the power resides and it may upset a lot of the people that either make the television, or the PC, or maybe both. The network, making sure that no matter what you purchase allows for easy access throughout the home, really becomes the avenue that we have to seek.

DAVID: I actually tend to agree with both of the previous comments as well. When people talk about PC versus TV, you would naturally draw the conclusion, "Am I going to have a PC in my living room?" I think the answer to that, if we mean today's PC, the answer is absolutely, unequivocally, no. In a new model, where PCs look a little different, and TVs look a little different, that is where potentially the convergence of that can draw.

DWAIN: The opportunity is in a home network, without a doubt. There is also the recognition that there is, what we call, "near" content, and "far" content. There are family entertainment interactions that work in a large screen TV environment, there are individual activities, and content applications that work in a more near-term PC. 1394 was mentioned as a home network. There are opportunities for home routers to take this information and spread it across multiple display, interactive platforms.

Q: You have the involvement, with a new cellular telephone product that enables you to get content over the Internet, and as we look at these other Internet-enabled products coming into the home, we find that the most likely purchasers of those products are people who already have PCs. Perhaps we are missing that target audience that do not have a PC in their home. How do you see that we could enable this market to expand beyond the current PC-centric home?

DWAIN: TCP/IP does enable various content providers to put information out to a variety of devices and platforms. What has driven the growth in the PC market is open platforms and the ability for developers to ride applications on multiple platforms. The Internet, either on a wired form, or a wireless form, allows new devices to be developed, and move away from the graphic side of the web to where the information lies, which is in text, and move it outside of the pure PC marketplace. That is one of the efforts that we are trying to do. The network computer is another play here, of taking advantage of the power of the network. So there are opportunities to move because of price points, because of familiar form factors, to move away from the people that are knowledgeable of PC, but still give them platforms that can take advantage of the information that is on the network, whether it is the public Internet or private intranet type of networks.

DAVID: DVD will have a consumer player bent to it where it will see success, but the technology is supportable, and will also exist in PCs. Content created for consumer players can play on PCs, and vice versa, and we are starting to see a great cross-breeding from a fundamental technology that will allow this technology to go well beyond just the PC-centric households that we have today.

ANDREA: The set top box with DVD inside that allows you to play games, is Windows supported, will allow you to do your spreadsheets, will allow you to do some word-processing, and will allow you to do most of the applications on the PC. We are going to see different things in the marketplace, different terminals, and I am not sure that would be a place for a low cost PC. I think the low cost PC and the high end set top box will be difficult to understand.

TODD: Right now we are working with a lot of danger. If we were to get a lot of consumers on to the Internet today, their response would be, "This is ridiculous. Two and a half to three minutes to download a web page?" Unfortunately, the backbone behind the web is not prepared to take on that level of responsibility today, and other technologies do not hit until the year 2000, 2002. As Internet traffic continues to increase, we are going to find a more miserable experience on the Internet until we find another way to make it happen, i.e. things like direct PC, broadcasting, different ways to get the information.

Q: On a different communications topic, 1394, Sony was one of the first companies to introduce a consumer electronics product with a digital camcorder that included a 1394 port. What has to happen to make 1394 a valuable standard in the consumer electronics world?

TODD: 1394 will become more available on all of our AV equipment. Unfortunately, the digital camcorders, by themselves, are considerably more expensive than their predecessors, so the start out was at a $2900 to $3000 range. It does not make for a lot of volume at that kind of price. The real goal is to get the VCRs at the $200 to $250 range to have a 1394, and that is going to take a considerable amount, because there is not that much margin in today's cost of implementing 1394. But '97 will be a big year to get the interfaces. Ninety-eight will be the year that you actually use a lot of the interfaces.

ANDREA: The PC will start with 1394 quickly, everybody is working on it, and the consumer will like it.

Q: What is C-Cube's position on the DV standard, stating Sony will not license the hardware to Kodak, will C-Cube engage in development for DV?

DAVID: 40% of the consumer camcorder market in Japan today is DV, and this new 1394 camera that Sony has, I saw in the street in Osaka for $1700. We are very bullish. From all of the companies we are talking to that is clearly a trend.

THIS PORTION OF AUDIO NOT CAPTURED

TODD: . . . We will continue to push into other market spaces just to make the overall market grow. We do need the competition.

Q: The mini disk has had some strong shipments in Japan, and recently Sony has made a major effort to reintroduce the product into the United States. Could you give us a status report on that and what will help make that a successful product beyond Japan?

TODD: The mini disk has run into a lot of problems in the United States, from a format, as well as a quality issue. With the quality of the audio that comes from a lot of the mini disks, people are not willing to pay significantly more for that type of format, so we have had a difficult time trying to sell the higher price point with greater technology. We are going to try to reintroduce the mini disk with some better compression technologies to get more and more content, but it will be an uproad battle for even Sony in this area.

Q: What is the one most critical factor hindering the explosion of digital consumer electronics systems: content, media standards, delivery and infrastructure interactivity, cost, semiconductor components?

ANDREA: Content. People buy content. ISDN has been here for years. There was no point in paying for a line at higher price than another line. When people needed bandwidth they just purchased bandwidth. This is the same. People will buy content. If the content is good, and is better than what is available today, people will buy it. When we go to the shop and buy, we buy the latest game, or the latest release of a new product.

Q: As we move forward with these new consumer electronics products, will 0.8 micron technology suffice, or what types of technologies do we have to drive into the fabs in order to deliver these products?

ANDREA: The designs we are doing today in advanced consumer products is in 0.35, but you should not mix the advanced design and the volume production. What we are producing now is what we designed three years ago, or two years ago, so it is obvious there is a lag. The other point is there are not only digital consumers in this world. There are not only PCs. There are many other technologies where you do not need to push the lithography to these limits, and where the applications still need – and we are still building fabs for – 2 microns and above. It does not make sense to build an audio amplifier with a lithography of 0.1 micron, when you are limited by other factors. Needs to be looked at, the whole spectrum, of the application, of the market, and the lag. Today, these super-integrated chips in 0.5 are on the limit of competitiveness, and of margins.

DAVID: As C-Cube's line is focused not nearly as broad as SGS-Thomson's, we do not ship anything in 0.8. We designed a lot of things originally in 0.8, and they have all been converted. Today we are shipping only half micron and two, three layer metal. New generation designs are a 0.35 four layer metal, so even in very cost sensitive, consumer electronics markets, advanced process technology is very critical.

Q: We see a new playing field starting to emerge in a competitive sense. Traditionally, the Japanese companies have been dominant in both the systems, and even in the semiconductor side of this industry. But now we see many of the key core technologies coming out of Europe, out of the United States. How does the new competitive arena shape up with the new technologies emerging as they are?

DWAIN: Speaking from my semiconductor background, as well as my exposure in two major consumer electronic companies, it is a challenge. The move from analog to digital gets rid of an opportunity that we made use of –the concept of fidelity and differentiating on high end fidelity. Digital is digital, and you now have to switch to high volume, low cost, design time to market differentiations, and that will be an opportunity for change within the Japanese companies. Additionally, the move to networking brings in more software than previously had been involved, and we are addressing those internally.

DAVID: These consumer devices we are building are pretty amazing, ranging from one to five million transistors on some of the more advanced things we have doing. They are all about architectures, they are all about signal processing, they are all about real time software. This sounds like something we would have discussed for the computer industry, or other high tech industries, a long time ago. It is a design process that, for some reason, we tend to be pretty good at here, and of course, Japanese semiconductor companies have done tremendous things with memories and so forth, but typically not in the areas of signal processing architecture and software, and that is the name of the new game for a lot of these new consumer electronics products. We are going to have competitors out of the Japanese,– we have to. But right now there are going to be a lot of American and European companies supplying into Japan for the next few years.

ANDREA: The Japanese have been late in identifying the switch from analog to digital consumer. I think they are counter-reacting. DVD is the first case, which is between Phillips and Toshiba, Matsushita, Sony, and others in Japan. I would be surprised if we do not face a very strong competition from Japan in the next couple of years. The only point is how far our advantage today will allow us to keep going at a faster pace. They are getting ready, and they have everything to succeed–they have the strength; they have the intelligence; they have the end customers. So we will find them in the market.

TODD: We have thousands of television engineers in Japan, and when we go into the PC business, which is probably 70% software engineers, and the rest are some other digital engineers, we cannot expect a lot of help there. Sony has decided to launch that effort. The PC company is based out of the United States, and it then disseminates bits technologies, and its products, back around the world, versus the traditional model that we would have had, where everything is built in Japan, and pushed to the rest of the world. We actually have counteracted it by putting it where the majority of the talent is, and the capabilities, and then making that the source.

Q: There are many chip manufacturers that are developing what they call media engines, and bringing those to market, and, the first design opportunities they are looking for are PCs, and getting into the PC enabling, and enhanced video and graphics capabilities. What kind of hurdles do you think that they would experience in trying to leverage that technology and that development into a consumer product?

TODD: They are going to find that the software world that we deal with inside Microsoft, as well as the hardware world of Intel, have a squashing reputation that they have to compare with. They are going to have to make sure that the software development kits, SDKs, that they come up with, allow for system integrators to actually utilize their software, and do not expect to be able to provide an end solution. They are going to have to make sure that they are highly competitive and that they are well integrated into a Windows 95 environment, or they will be left behind.

ANDREA: On the PC, I think there is a window for these products. The window is before the next generation micros come out. Before the end of this century they will run at 1 gigahertz, so they will have enough power on the processor to suck all that is done today outside of graphics. There is a time window of a couple of years, and of course, the guy who is there will make money. The second thing is how to leverage this in consumer electronics. In consumer electronics it is a matter of cost. The reason we have been successful with MPEG is because we made the smallest and most competitive solution. If you are able to hardwire to the smallest solution, you will win in the consumer world. Some of these processors also need some backup from the host processor, so I think the consumer world would be more hardwired than programmable

DAVID: I think it will be tough for the media processor guys. They have a window. I agree with that, but they have to be as fast as Intel. With gigahertz processing at the end of the decade, they are doing DVD and software, plus they are multitasking with other things as well. They have to be as fast, and stay two years ahead of Intel all the time. Part of the problem with that is their strategy is based on programmability, and those media processors are typically not easily programmed.

Lan Switches, Round Two

Prabhat K. "P.K." Dubey, Ph.D.

President and CEO,
MMC Networks Inc.

Dr. Dubey has been president and CEO of MMC Networks Inc. since 1994.

Before MMC Networks, Dr. Dubey was the vice president and general manager at AT&T Microelectronics, responsible for the Wireless and Messaging DSP IC Division. Before that he worked for NEC Electronics as the business operations manager in the Microprocessor Strategic Business Unit.

Dr. Dubey holds a Ph.D. in physics from the Indian Institute of Technology, Delhi, and an M.B.A. from the University of Western Ontario, London, Canada.

Session #9: LAN Switches, Round Two

Dr. Prabhat K. "P.K." Dubey
President & CEO, MMC Networks Inc.

GENE NORRETT: Before joining MMC Networks Dr. P.K. Dubey was Vice President and General Manager at AT&T Microelectronics, responsible for the Wireless and Messaging DSP IC Division. Before that he worked for NEC Electronics as the business operations manager in the Microprocessor Strategic Business Unit. He has a Ph.D. in Physics from Indian Institute of Technology in Delhi, and an MBA from the University of Western Ontario.

P.K. DUBEY: Networking is a very exciting industry. I will share with you what is happening in the networking industry, giving you details of the infrastructure which is the key to the next two decades. The first thing to note is that the networking industry is a high volume and high growth industry.

This year we will see more networking nodes installed than modems bought. Dataquest forecasts 150 million PCs will be bought and sold in the year 2000. The networking industry is growing rapidly, according to Dataquest projections, some 42% per year. The second thing you will see is the networking industry going through a massive change.

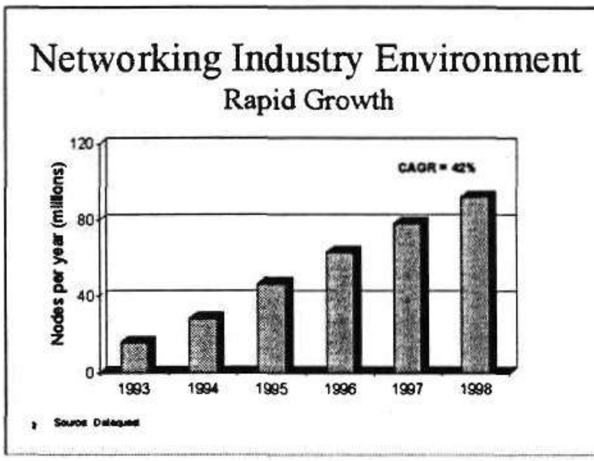

This diagram shows the way existing networks are built. Personal computers and workstations are tied to Synoptics hubs, which is a shared scheme, and backed by Cisco routers. Problems with this architecture developed with the large number of people on the network who want to communicate graphics and image-rich files with each other.

Think of networks as a shared highway. When the highway is used by a few drivers it is the greatest thing, but when millions of people use the highway it is a great bottleneck. Bumper to bumper traffic, exits, and accesses are blocked. Data networking is the same. Latencies are getting bigger, the effective bandwidth per user is now roughly 1/1000 of the wire speed. In other words the networks are very clogged.

In future Networks, you see switches have been installed. Switches deliver to the user bandwidth on demand on an asynchronous basis. Instead of sharing wire, a user is in command and control by using a switched connection as opposed to a shared connection. The switch industry has gone

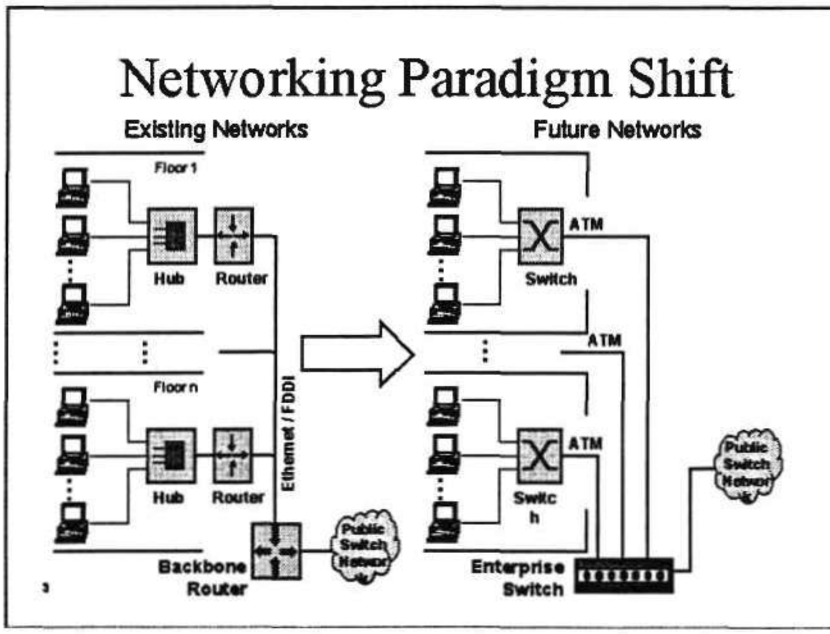

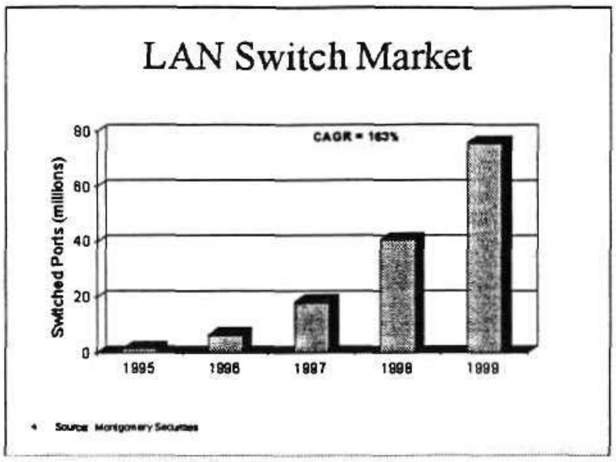

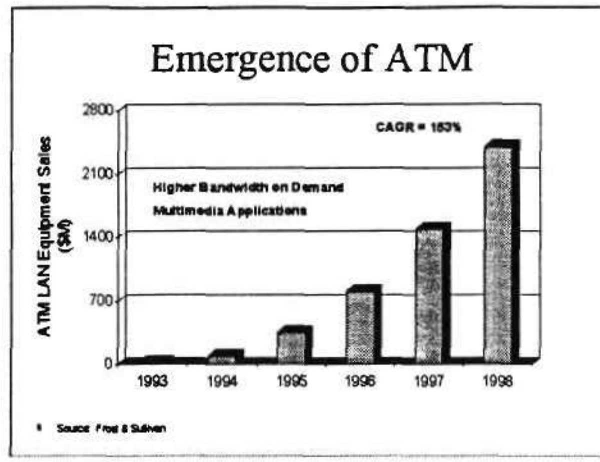

from almost a billion dollar industry to $7 to $8 billion sales this year. In the next five years, the bulk of the industry will be built with switches replacing hubs.

The LAN switch industry, from '95 to 2000, will grow by more that a 100% per year and will be a big industry by the year 2000.

Another trend in the LAN data networking side of the equation is the Ethernet going from 10 Mbps speed to Fast Ethernet, which is 100Mbps speed, to the emergence of 1000 Mbps, for the Gigabit Ethernet. The third trend you will see is the emergence of a new protocol called ATM.

ATM stands for Asynchronous Transfer Mode. Two factors driving it are higher bandwidth on demand and multimedia applications which are video-rich and voice-rich.

ATM is being installed in the backbone very rapidly.

The gigabit industry will also grow, you will see gigabit installations as well as ATM installations. Drivers for gigabit are higher bandwidth to the desktop, solving the congestion problem in the backbone, and ease of migration to the Ethernet world.

The high level trends are, the networking industry is a high volume and high growth industry, changing from shared network to switched networks, Fast Ethernet switches are rapidly being adopted, ATM switching in the backbone is emerging, and Gigabit Ethernet is going to emerge in the backbone towards the later part of the 1990s. The key to networking technology, to delivering bandwidth on demand, to delivering high bandwidth on demand, to delivering voice-, video-, and data-rich contents on demand, is switching technology.

Looking at the ideal attributes of switching technology four things come to mind. The first is scaleability, the switching technology should be

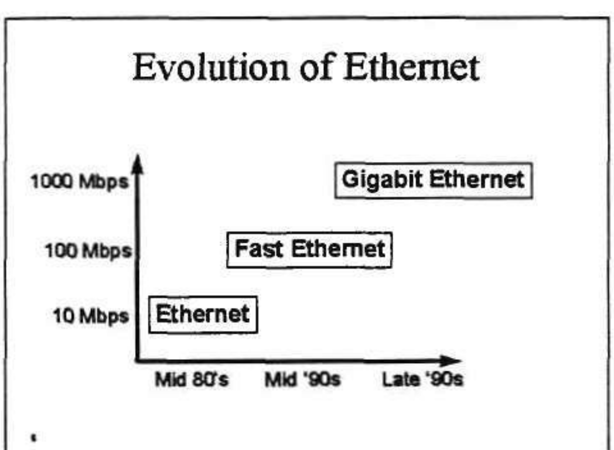

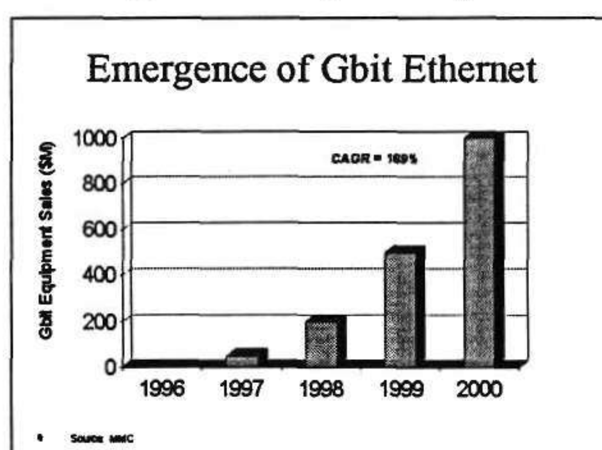

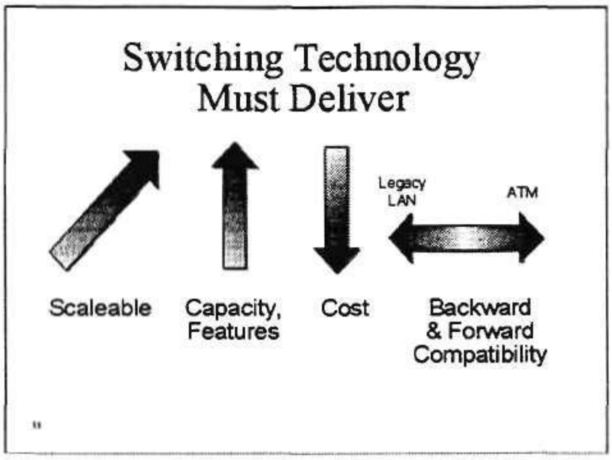

scaleable from low to high end, so that the equipment maker does not have to reinvent technology. The second trend is that as more people get on shared networks there is a need for switched networks with higher capacity switching engines.

Compare the switching technology to automobiles. The switches today are delivered without many features, the same way Ford introduced the model-T car, it was essentially an engine and transmission, without air conditioning, air bags, and so on. The same way the next generation switching technology must provide extra features in the switching box.

Cost must come down faster. Switching technology must preserve the legacy LAN investment as well as be future proof. Nobody wants to buy something now and throw it away in two years. Technology must be ATM and gigabit compatible.

In the switch box vendor environment, two dilemmas are evident. One, switching technology could not be bought from outside. A supplier built his own or did not participate in the market. Building your own requires 18-24 months and a lot of manpower.

Developing switching technology which is tested thoroughly and is bug-free, takes a lot of engineering manpower and money. When it comes to high capacity switching technology, it is fairly risky, so you may not succeed even though you spend millions of dollars, and twenty-four months, in that development. At the same time,

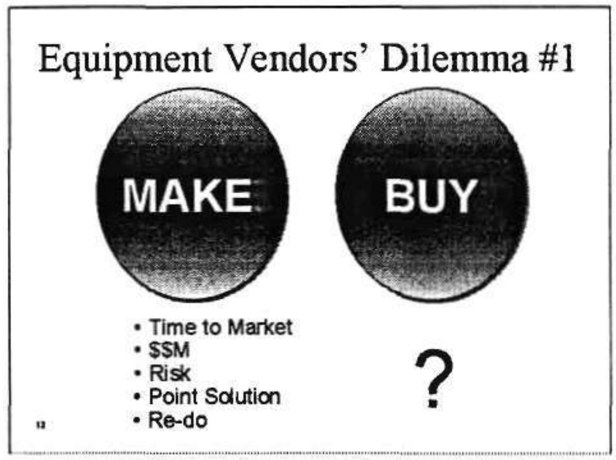

you get a point solution, which is not as scaleable, which means that you either get a low end, or a medium end, or a high end switching technology. When your marketing department says that 38% of the users require such and such capacity of features, you generally have to re-do all these things.

The second dilemma is that while the equipment makers are very concerned with developing switching technology, they find themselves faced with a different dilemma when dealing with the IT buyer and IT user. The end users purchase criteria is based on features. Back to the automobile example, we buy an automobile on the basis of styling, features, and comfort factors. Few buyers today buy an automobile because it has an engine implemented a certain way, or transmission implemented a certain way. With switching equipment, the purchaser's criteria are features, port configuration, up links, diagnostic features, maintenance features, LAN emulation, policing,

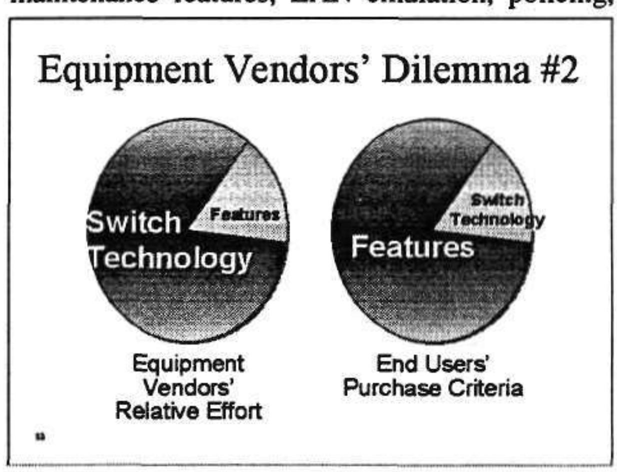

accounting, etc. The equipment vendor spends more time on developing switching technology, and less time on feature development. This is where MMC comes in as a vendor here.

You see very high growth in traffic and heavy congestion in the corporate network environment today. While switches are being touted as the elixir to all these problems, the shift to switches creates new challenges.

The challenges that a network administrator faces is to make sure that the client-server traffic, which is the dominant traffic in the networks today, is decongested. The bulk of the traffic is from user to server, rather than user to user, so the server traffic must be fed in a fat pipe. The backbone traffic, where everything aggregates, must be a fatter pipe. Adequate bandwidth to the desktop, depends upon the application itself. Switches need to be installed and managed to ensure they perform various functions such as moving/adding/deleting users, troubleshooting, doing performance analysis, and securities. Shared networks do have one important attribute which is the observability of traffic. Sharing a wire makes it very easy to observe the traffic flows and pattern by attaching a monitor and analyzer on a shared wire.

Troubleshooting is simplified. Switched traffic, by definition, is like tic-tac-toe. This tic-tac-toe behavior destroys the observability making troubleshooting a challenge. Other challenges are: the cost must be low enough; the right sort of equipment needs to be selected based on the horsepower of the switch; the flexibility in different ports for different users; and adequate features.

BW stands for the aggregate bandwidth or capacity. The number of desktops connected to the switch at the speed of either 10 Mbps or 100 Mbps. Most traffic to servers should be high speed type, either a fast Ethernet (100 Mbps) or a gigabit (1000 Mbps) connection. Because it is for both bandwidth and distance, it needs to be a full duplex connection, so whatever number of servers you have, you just simply multiply by two. Switches take you to the wide area network side, therefore, there must be some backbone connection. Take the number of backbone connections, either have ATM 155, 622Mbps, or a gigabit connection and to make it full duplex. Let us take an example of an average to network see what happens to the capacity requirement.

Ninety desktops are connected by the old Ethernet way, and ten desktops are at Fast Ethernet bandwidth. This network has six servers, which are full duplex Fast Ethernet connections with one ATM backbone connection. Plug in these numbers, you see that the switch with adequate capacity, which will guarantee non-blocking of the traffic, requires 3 to 4 Gbps capacity.

Legacy LAN Switches Today

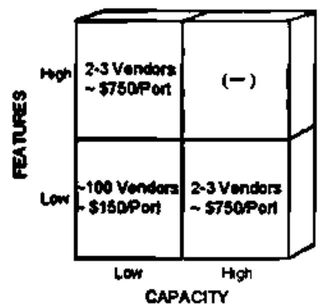

From this chart, you will see in the box labeled low-low some 100 vendors and very few vendors on the high feature or high capacity side. There is virtually none in the adequate capacity and adequate features side, in the upper right hand corner.

ATM LAN Switches Today

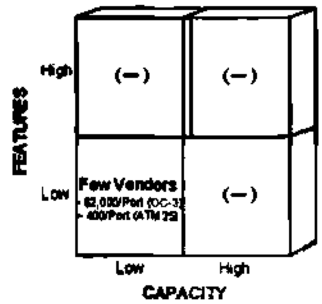

A similar situation exists with the ATM LAN switches today. There are very few vendors, their costs are very high, and it is mostly on the low-low side.

It is clear that the current LAN switches do not have adequate capacity or features to solve the network administrator's problems we have just been discussing. So what is the round 2?

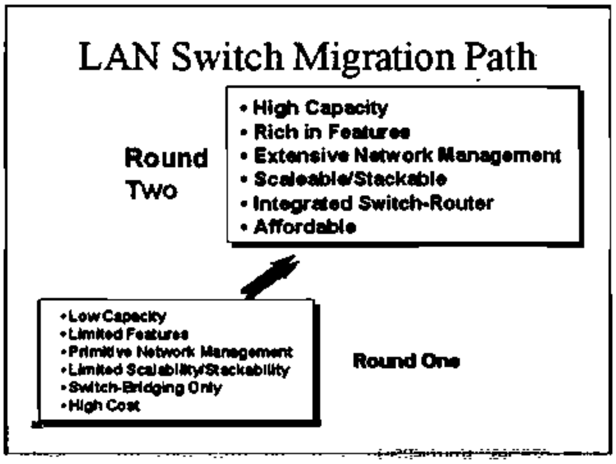

Round 2 of LAN switches require adequate horsepower, adequate high capacity, must be rich in features to allow troubleshooting, move/add/delete, performance analysis, etc. It must have extensive network management capability, it must be scaleable, and it must be affordable.

I believe that round 2 of LAN switches must have capacity scaling to 20Gbps in the corporate use environment, flexibility as different users will ask for different port configuration. V-LAN (virtual LAN) will simplify the network administration providing security, some firewalls and improving the efficiency of the network. Round 2 must have RMON (a remote monitoring capability), for troubleshooting and management. We are entering a period where there is a need for the observability delivered by RMON.

Now this is where MMC Networks comes in. MMC invented a switching technology called ViX technology, a switching routing architecture. We focus only on the switching routing area of the problem. We have expressed the ViX architecture in an Ethernet packet, or Fast Ethernet packet, or gigabit packet, and ATM cell switching, and these two platforms are now in production.

We are a company based in Sunnyvale, backed by venture capital, we have been profitable for three quarters in a row. To get additional information on our company, you are welcome to check our website at mmcnet.com.

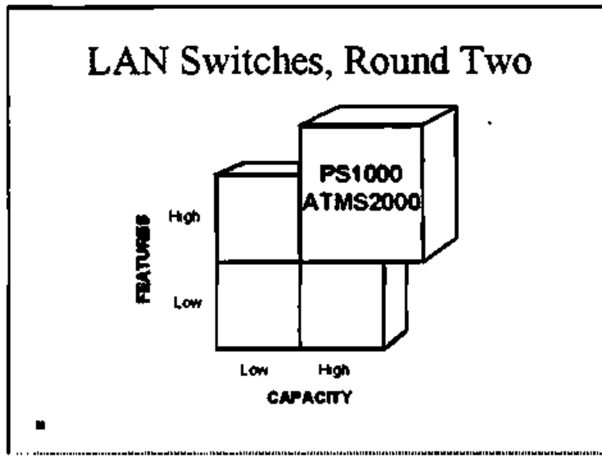

The two switching platforms that we have introduced fit in the high capacity, high feature box.

Our packet switch platform has a capacity of 4 Gbps which is roughly ten times the capacity of the normal switches available today. It provides ATM uplinks and every port is configurable. Many features are supported and it is cheaper; it allows our customers to differentiate their systems at the feature level and at network management software level. The ATM product is applicable up to 40 Gbps bandwidth which is more than adequate bandwidth in the corporate backbones. It is a complete ATM switching platform. Our switching platform plus 20 nanosecond static RAMs is the entire ATM switch hardware.

It is scaleable to 40 Gbps, the ports are flexible and we have completed numerous feature chips to make sure the LAN switches of tomorrow are here today, and it is a dramatically lower cost of switching. The cost of ATM connection, if you throw in the adapter cards and all the network management software, is roughly $2,500 for connection. Imagine buying a PC at $2,000 and paying $2,500 for a network connection. Clearly if ATM is to become widely popular, the cost must be brought down by a factor of 10. In the world of Ethernet, in 1987, once the costs were brought down to less than $300 for a connection, Ethernet took off like a rocket.

In summary, the current switches that are available today do not have adequate horsepower, nor do they have adequate features. Our company is enabling the round 2 of cheap, high bandwidth feature switches.

Q: Is the world going to be different five years from now? Does P.K. see Gigabit Ethernet ever being used on the desktop, or will it always be a backbone technology?

P.K.: Gigabit, starting out, is going to be a backbone technology. The majority of users in the next five years do not need 1000 gigabits for a desktop. I see this as a technology for backbone and for server links, rather than desktop.

Q: Your diagram showed ATM to the desktop. Do you believe ATM will come to the desktop?

P.K.: If multimedia applications to the desktop become mainstream, then you will see the quality of service which is the main differentiating factor between ATM protocols versus Ethernet protocols. ATM will become a mainstream desktop technology.

Q: What is the semiconductor content of a switch, ATM switch measured in dollars, number of devices, and silicon square inches per switch?

P.K.: Cisco has an ATM LAN backbone switch family call Lightstream 1010. It has 32 ports, meaning 32 users can get ATM 155 installed in the backbones with a 5 gigabit switch. This switch has a total semiconductor of about $3,000. If you have a bigger switch you will have more content, and if you have a smaller grade switch you will have less content

Q: Will more than one of the big four networking companies acquire the switching technology; will it remain somewhat independent or be dominated by one player?

P.K.: I am a contrarian at heart, most of the industry gurus are predicting the industry will consolidate. My view is that industry will deconsolidate. Many new things are happening at a much faster pace. I do not think it is possible for any one, two, or three companies to be a champion at everything.

Embedded DRAM Technology: The Next Major Technology?

Prakash Agarwal

President, Cofounder and CEO,
NeoMagic Corporation

Since 1993, Mr. Agarwal has been with NeoMagic Corporation as president, CEO, and cofounder. He has over 17 years of engineering, marketing, and general management experience in the semiconductor industry.

Mr. Agarwal began his career in 1977 at Intermedics Corporation, Freeport, Texas. Other companies where he worked during his career are General Electric in Lanham, Maryland; Storage Technology in Santa Clara, California; Silicon Compilers in San Jose, California; and from 1984 until 1993 was employed as vice president and general manager of Cirrus Logic's Portable Product division in Fremont, California.

Mr. Agarwal holds master's and bachelor's degrees in electrical engineering from the University of Illinois.

Session # 10: Embedded DRAM Technology: The Next Major DRAM Technology?

Prakash Agarwal
President, Co-Founder & CEO, NeoMagic Corp.

GENE NORRETT: Prakash Agarwal is President and Co-Founder of NeoMagic, a company that is focusing on the portable market. He has had a number of different, exciting positions in the industry, for the last 17 years. The last position he had prior to founding NeoMagic was Vice President and General Manager of Cirrus Logic's Portable Products Division. He has a master's and a bachelor's degree in electrical engineering from the University of Illinois.

PRAKASH: Gene asked me to talk about embedded DRAM technology, and especially the topic: "Is it going to be next major DRAM technology?" Is it really next major DRAM technology, or is it going to be next major semiconductor technology?

I will talk about why embedded DRAM technology is so important, the benefits, the challenges, the challenges ahead of us, and the NeoMagic experience for the last three years.

What are the applications? At NeoMagic we are focused on a few applications. There are many other applications. We will talk about that, and the future direction of embedded DRAM technology, what we should look forward to, and summarize it.

I claim that system on a chip cannot be implemented unless you are using embedded DRAM technology. I will focus on what is a system, and why embedded DRAM technology is so important to achieve system on chip.

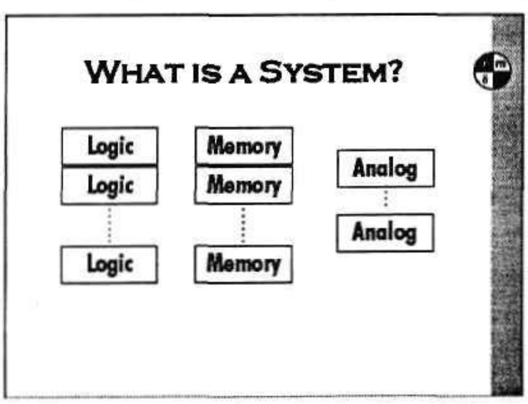

What is a system, or a subsystem? It contains logic modules, memory modules, some analog pieces, and industry, has been implementing more and more logic functions into a single chip, and has been, including more and more memory capacity from 1 megabit DRAM to 4 megabit DRAMs, to 16 megabit DRAMs, and beyond.

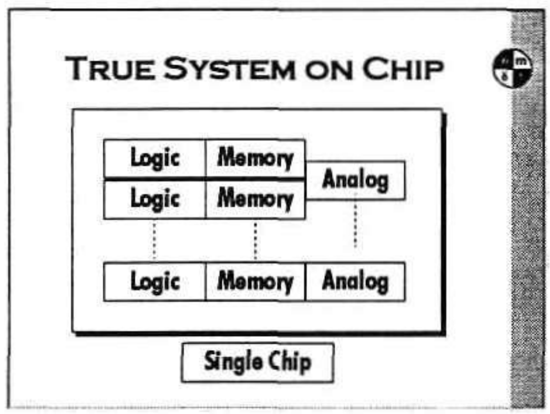

For true system on a chip, you have to integrate logic, memory, analog, all in a single chip. If you do not have that, you are not implementing a true system on a chip.

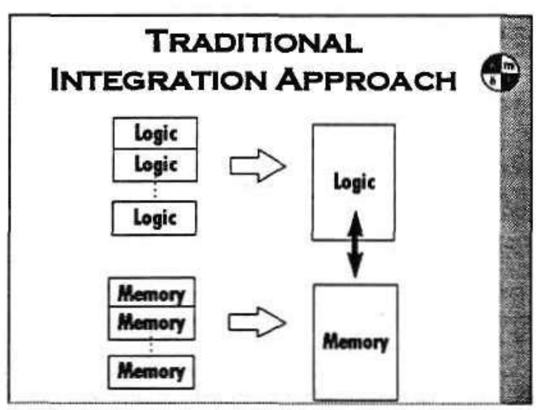

If you look at the traditional integration approach, you take a bunch of logic blocks, 100,000 gates, 200,000 gates, or 1 megabit DRAM to 16 megabit DRAMs. That has been the approach, where people have been incorporating more into a single logic chip, or single memory chip, and with that approach you do get the integration. You can reduce the number of components. You do get that benefit. But you do not really see a major performance enhancement, nor do you see much of the power reduction.

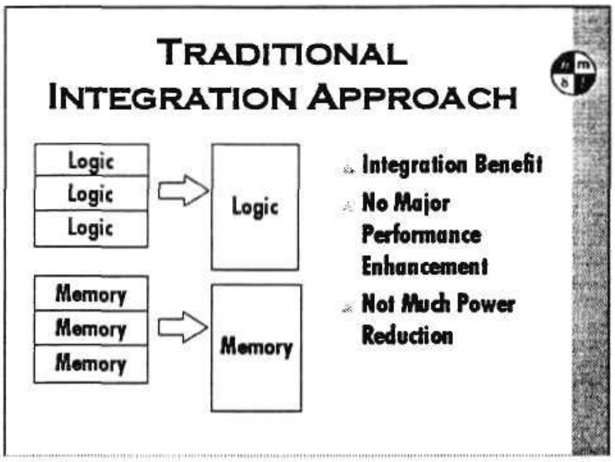

Embedded DRAM technology not only gives you performance, it also reduces the power consumption in a big way. Form factor: that goes without saying. Thermal: some of the experience we have had we will share that with you. At the same time that it reduces the EMI, system on a chip also helps from the economic side.

What improves performance? In any system or subsystem, the most obvious performance bottleneck is in the memory bandwidth. If you can solve the problem of memory bandwidth, you can improve the performance. You can have faster buses, or wider buses, or you could have both, and this is where the whole industry has been moving, from 8 bit, to 16 bit, to 32 bit, and even 64 bit. When you go wider you get the performance.

The problem is, as you go wider, or faster, it takes more real estate to implement the system. It increases the complexity of the board design. It takes more power consumption, and costs more, with more components, more pins, more real estate, more EMI, and the real solution is embedded DRAM. When you combine DRAM, you are no longer restricted by any of those problems, because you get the widest bus. You are not restricted by how many pins you have to dedicate. Is it 32, or is it 64? You can make it as wide as you want, because it is all internal.

What saves power? Low voltage operation, power management, slower clock operation, narrow buses, fewer I/O pads. Industry has made a lot of progress in low voltage operation, going from 5 volts to 3.3 volts, even getting down to less than 3.3 volts. Same thing goes on power management. If you slow down the clock, or have narrow buses, the problem you have is lower performance. Embedded DRAM gives you the widest bus with zero I/O pads switching simultaneously. Power especially in CMOS, is consumed at the I/O pads, because they are designed to drive very hefty printed circuit traces on the motherboard. When you eliminate that by combining DRAM and logic into a single chip, you have eliminated the most power hungry circuits in your chip.

Form factor: NeoMagic is focusing on the graphics side. Take today's comparative graphic solution, the 256-pin BGA, and two megabytes of 16 DRAMs. That is five chips. You can eliminate basically five chips with a single chip. The resulting chip is not only a single chip, but a smaller chip than the 256-pin BGA. That chip is only 176-pin.

Cost benefits: Any time you have a single chip versus multiple chips, you reduce the manufacturing cost. It is so much simpler for the manufactur-

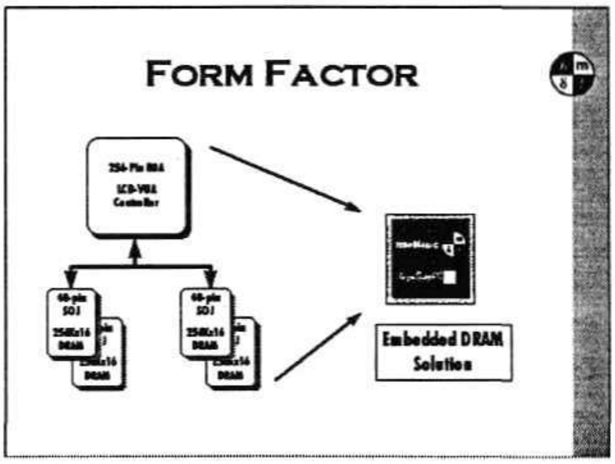

ing , the testing, and more importantly, granularity savings. Take graphics, for example, let's say you want to run 1024, 768, or 64 thousand colors. Let's say the next application requires 1.1 megabytes of memory, you are forced to go to two megabytes of memory, you have no choice, that is how standard DRAMs come. They come in half a megabyte, one megabyte, or two megabytes of memory. That 900 extra kilobytes of memory, which you do not need is wasted. When combining embedded DRAM to logic, we put in 1.1 megabyte, no more, no less, and that way, you save on the granularity side. You see the simplicity I talked about, fewer wiring layers, less inventory, and that translates to ease of use, and time to market. Time to market is so important in this business, if you can save three months off your design cycle, you can be ahead of millions and millions of dollars.

What are the challenges? Why hasn't anybody else done it, though people have been talking about it. There is actually more talk about it now, than when we started the company. We really have to understand the dynamics of memory and logic.

There are two different industries: memory is component-oriented, whereas logic is solution-oriented. For the logic side you have to provide complete solution, demo boards, BIOS, drivers, software.

Memory is manufacturing intensive, in order to be successful in the manufacturing business, you must lower your cost, and increase capacity. That is what some of the Koreans did in the last four or five years, became number one. Samsung became a number one player from nowhere. Logic is design intensive, the more you innovate, the more successful you are going to be on the logic side.

We are a fabless company, so we work with major DRAM suppliers, that also have the logic groups, too. But the memory groups do not understand logic, and logic groups do not understand the memory side of the business. Once you start combining them, design complexities exist because you need to run logic and analog on the commodity DRAM process.

Testing is a big issue, but DRAMs are tested on giant gang-testers, where you test hundreds at one time to bring the test time down, whereas logic is done one piece at a time. How do you resolve some of those issues? Burn-in requirements: DRAMs are such that, if you have any soft errors, and if a bit flipped, there goes your bank account, and you have to pay attention to that too.

The biggest fear in this industry has been dealing with yields because the DRAMs are repaired by laser. They are compacted so much that they have redundancy in their design so they test it, and then they do laser repair. With logic you do not. Once you bring that logic with DRAM, are you taking a yield loss because of that integration? You have to pay attention to that too. When we started the company, DRAMs were at their height, very capacity constrained; 16 megabit DRAMs were selling at $50 plus. How do you work out an economical, profitable model with your DRAM supplier where it becomes a win-win? Those are some of the challenges that one has to face.

NeoMagic experience: We have three products in production now. Our first product came out in December of '94, and since then we have added two more products. 2070 MagicGraph 128 is the first product. It is designed for notebook computers: graphics controller for notebooks – our bus interface between controller and memory is 128 bit wide. That is the beauty of integration. You no longer are restricted by 32 bit or 64 bit interface, and that gives us more than RAM bus type of memory bandwidth, without having all of the issues related with it. It has been a very successful product. This was done on 0.45 micron commodity, 16 megabit DRAM technology. It is about 20% bigger than the standard 16 megabit DRAM. The 128V: we took an M2070 and added video acceleration to it. With that product you have full-frame, 30 frames per second, MPEG-1 playing in software with a 100 megahertz Pentium. You do not need MPEG-1 in the hardware. That controller working with the Pentium gives you all of that. Same technology, 0.45 micron technology, and it is about 20% bigger than the first chip. The third chip, we added zoom video support, and TV-out. There are also analog blocks, RAM DAC and frequencies synthesizer, and a small RAM for TV support. The interesting thing is, and this is where the three years of learning has paid off quite well, it is the same technology, and we have 20% more gates in this chip than 2090, yet die size turn out to be the same.

The graphics market industry has been going from 16 bit, to 32 bit, to 64 bit which takes more board space. It used to be only 2 square inches, now at 64 bit, it takes 4.5 square inches. What is next for them, if we are not embedding DRAM? Going to 128 bit, means that your square footage may go up to 8 square inches. The power consumption has gone from 1 watt to 2.5 watts in the case of 64 bit, and over four watts for 128 bit. At NeoMagic, based on the embedded DRAM solution, power consumption is less than 400 milliwatts and board space is less than one square inch, for the 128 bit bus.

We have announced quite a few good customers: Acer, Compaq, DEC, Dell, HP, Hitachi, Mitsubishi, NEC, Sharp. These are all the people that are shipping the product now, and there are quite a few more customers we have in place, that will be announcing their products in next three to six months. The support we have gotten from our customer base, can be attributed here, what they are saying about our technology, our product, and our support.

David Altounian, who is a Director of Portables Marketing at Dell, who had been our first, and the number one customer: "NeoMagic's technology not only offers much higher performance, but also reduces power demand. No other technology has been able to do so." That is a true statement.

Ronald Chwang, President and CEO of Acer America: "By incorporating the NeoMagic graphics and video in our newest notebook computers, we have been able to provide state of the art multimedia capabilities while increasing the battery life over previous solutions." You must have heard about the Nuovo notebook, which they announced at PC Expo. This notebook has 10.5 hour battery life. We are adding about one to two hours extra battery life in that machine. A company called Veritest provides the battery test in the notebook, and when they ran their benchmark test on this machine, it broke after 8 hours because they never thought any machine could run beyond 8 hours.

Rod Schrock, VP of Presario Business Unit at Compaq: "The innovative technology of NeoMagic's video solution provides a breakthrough combination of exceptional performance and very low power consumption." No one talks about form factor because at machine level, what they had been able to do was put more functions on the motherboard, because the board space we saved for them enabled them to put PCI logic and other functions on the motherboard.

We started NeoMagic in July of 1993. Our first $3 million round was led by Kliener Perkins, the Sequoina capital, and was joined also by U.S. venture partners. The second round of funding of $6 million, in June of 1994, was led by Jafco, the Japanese venture capitalist firm. In December of 1994 we had the first functional silicon which we started sharing with our customers. In March of 1995, we introduced introduced MagicGraph 128. In July, Roger McNamee of Integral Capital, led the round, and we raised $8 million of funding, and that was the last round of funding we have done. Since Q1 of this year, we have been volume shipping to those nine customers. In June of this year we introduced the MagicGraph 128V, and ZV, and both of these products are in production. In July we announced a lot of customers, and we are approaching very close to one million units in shipment very soon. That is a very good milestone for a company that has only existed for three years, went after a unique concept, and in the first year of production, will be shipping close to one million units.

Besides graphics the embedded DRAM applications can be classified into three different categories: performance-driven, battery-operated, or cost-driven. You can apply this technology into all of these areas. Some of the semiconductor opportunities in those areas are 2D/3D graphics accelerators, audio/video, compression and decompression engines, video conferencing, DVD controllers, CPU and DSPs, and etc.

The mobility aspect of information technology is going to increase by 30%, because all of these new applications, like multimedia and video-conferencing. Multimedia is presentations, training, entertainment, education, and what that means, is portability. I am able to do this presentation from my notebook, because my notebook can do quite a bit of graphics, text, and video type of presentation, which makes me more productive and allows me to give a better presentation and training.

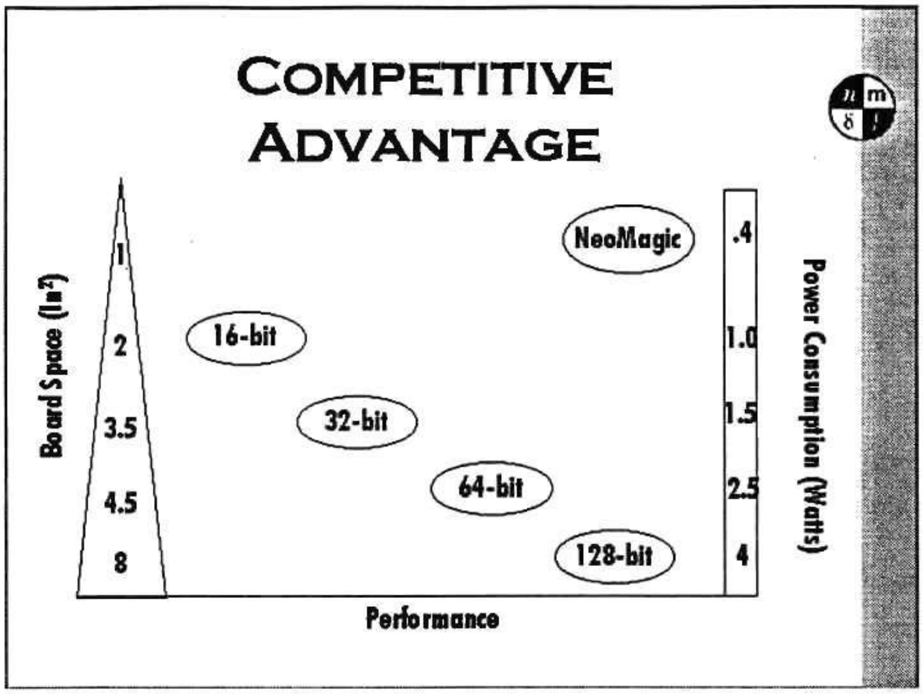

With embedded DRAM technology, we need to add a third metal, or even beyond that. The 64 megabit DRAM technology, that is coming will allow more memory to come in (because multimedia requires lots of memory and requires a lot of gates), so 64 megabit DRAM technology will enable us to put more functions, more features, and more memory into embedded DRAM applications. The DRAM technology is basically poly-driven, poly-silicon driven, it is not so much optimized for metal, so to get the speed out of your logic, you need to have tighter metal pitches. DRAM takes 23, 25 mass set compared to only 13 to 15 on the logic side. It has a longer production cycle, and, being in the PC industry and how fast the market moves, we really need a shorter cycle time. The key thing is you want to do all of that on a commodity DRAM process. You do not want to build a specialty DRAM process to do embedded DRAM, because you want to leverage the cost and technology of your commodity DRAM.

There is a lot of interest in embedded DRAM right now, probably due to what has happened in the commodity DRAM market. I believe this technology is going to be a major semiconductor technology, providing us with this system on a chip application. It should be focused from that direction. Those who have been focusing on long-term vision for the commodity DRAM process are going to be very successful.

Q: Is there a concern about embedded DRAM foundry capacity and prices?

A: There is always that concern, because industry goes through the cycle. When we started three years ago, DRAMs were in a very tight capacity. Prices were very high. The key thing here is it must be planned properly. If there is a value proposition, there will be solutions, and those who see it are busy solving capacity problems and pricing problems. Those who do not will be left behind.

Memory process drives semiconductor technology. Memory technology drives the semiconductor cost. If you combine those two, long-term, you will have the technology and the pricing you are looking for.

Q: What do you look for in a manufacturer foundry?

A: The key thing we look for in a DRAM supplier is do they share the same vision as we do because in the long-term relationship if we do not have similar vision, we will drift apart. Then we look at their technology and their road map to the future.

Q: There are quite a few questions related to manufacturing, and how long it took us to get the things going.

A: We have targeted all of our logic and analog blocks to work on a DRAM process. We have created a unique standards set library for that DRAM process. That takes us from foundry to foundry. We are working with two foundries right now: one in production, and one will be coming on in second half of next year, from four to six months.

The way we work with our supplier, we do not do a memory design ourselves, because we do not want to take the risk. We want to take the commodity DRAM block, and have them modify it for our 128 bit interface which took them from three to five months. You can go from one supplier to another supplier, technically, in about four to six months.

Q: What, PC system company has been most resistant to using embedded DRAM, and why? The question should be: Why not?

A: When we started working with system companies, they said, "Great story, if you can do it, we will use it. We do not think you can do it, and we are not going to risk our billion dollar product line, because it takes only one component on a notebook, to hold up the whole production." Companies like Dell wanted to take the risk. They said, "We believe you guys can do it, and we will take the risk," and they are the ones who are shipping hundreds of thousands of units every quarter, based on this technology. Dell has increased their notebook market share by 2% since they announced the machines based on the NeoMagic controller

You have to show the value proposition, and make it work for them, when customers announce machines, and other customers see the ads, see the benchmarks done by PC magazine, all of a sudden they say, "Boy, you know, I have to use the NeoMagic controller, too." Right now we see no customers saying, "We do not want to use embedded DRAM.

Ultra High Speed DRAMs: Unlocking the Potential

Dr. Fu-Chieh Hsu
Chairman and CEO,
Mosys Corporation

Dr. Hsu is currently Chairman and CEO of Mosys Corporation.

Before founding Mosys, Dr. Hsu was chairman, president, and cofounder of Myson Technology Inc. and before that was vice president and chief technical officer at Integrated Device Technology. Dr. Hsu was also founder of Knights Technology Inc. and was a project leader at HP Labs.

Dr. Hsu received his Ph.D. in electrical engineering from the University of California at Berkeley.

Session # 11: Ultra High Speed DRAMs: Unlocking the Potential

Fu-Chieh Hsu, Ph.D.
Chairman and CEO, MoSys Corporation

GENE NORRETT: The next speaker is Dr. Fu-Chieh Hsu, "Ultra High Speed DRAMs: Unlocking the Potential." Fu-Chieh has spend some time at IDT, along with Knights Technology, and was a project leader at Hewlett Packard labs, and he has a Ph.D. from Cal. He is here to talk to us about ultra high speed DRAMs.

Fu-Chieh: In my presentation, I will take a slightly different look at the DRAM and particularly the special niche application segment of the DRAM. MoSys's vision is to exploit one of the potentials that exist in advanced DRAM process technology but so far has been underutilized, the ultra high speed DRAMs.

What is speed? Speed has been talked about by all the various people, various groups, various application areas. Is it bandwidth? Is it latency, or is it throughput? Our view is that speed is all of these. Speed is access time. This is probably one of the oldest definitions for the 30 nanosecond SRAM (the "S" is the latency time), the speed is defined by the latency time. For traditional 60 nanosecond, fast page mode DRAM, or EDO, the 60 nanosecond is the row access time of the memory. Speed is also bandwidth. At this conference, over the last several years there has been a lot of talking about bandwidth. What people are talking about is actually the peak bandwidth, which is usually defined as the width of the data path times the clock rate, the frequency. The speed is also the throughput, which takes into consideration the reduction in the usable bandwidth due to the access latency.

With all the different attributes of speed, we think speed is really an application-specific concept, so with that, how to optimize the 'speed' is really dependent on a specific application environment. Speed is a figure of merit based on all those components that made up the speed concept – a weighted average of all those factors that are relevant, and which is application-specific.

Speed as defined by bandwidth, and this term is mostly used in the graphics or multimedia related applications because in those applications, the graphics pipeline needs to handle a lot of data traffic. Usually, in those systems the memory data bus has a fairly high to almost full utilization, and the efficiency of actual data transfer, depends on the access latency. The graph at the bottom of the slide shows that most applications in today's computing environments are multi-process, or multi-tasking. There are multiple tasks that have to be interleaved to access a shared memory system. With the advanced pipelining capability, the task may be queued for a relatively long period of time before it is actually executed, this is what we call a task latency. Between tasks, when the transition is going from task to task, this is the access latency, which defines the time from the time when the previous transaction is executed (or the previous task is completed), to the time when data transfer can start for a new task or a new transaction. In those applications the task latency may be relatively long and the utilization, or the efficiency, or the throughput, of the system is determined by the peak bandwidth as well as the access latency.

In another application, specifically the level 2 cache type applications for the processor, or the communication, or other similar type of applica-

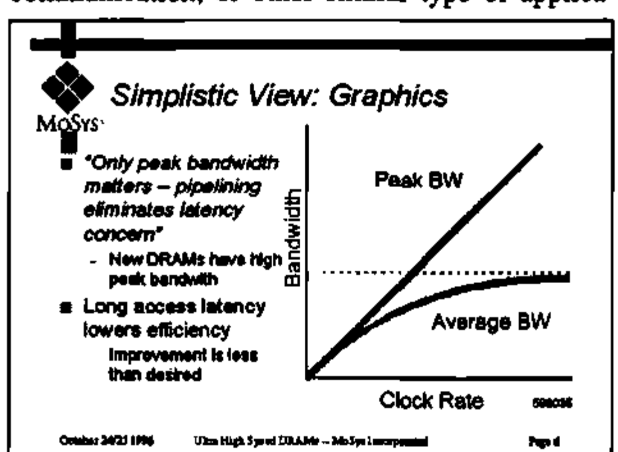

tions, the speed means latency. In those applications very short latency is required, and in general, the bandwidth requirement is moderate. The bar at the bottom of the slide shows a typical example of this kind of access. Ideally the task latency is equal to the access latency, and this latency has to be short to prevent stalls at the processor level. So far today, most of the systems with this type of application have relatively low bus utilization, typically 50% or lower.

Using a relatively simplistic view to look at the graphics or multimedia applications, one of the common perceptions is that the peak bandwidth is the most important parameter. The reason is that pipelining can eliminate practically all access latency concern. The task latency itself is relatively unimportant in this application, but the access latency is still very important. What it essentially does is reduce the usable bandwidth of the system, the memory the system can deliver. This peak bandwidth requirement, so far, has driven many of the new DRAM developments. As a result, much of the advanced DRAM memory today has very high peak bandwidth, such as a synchronous DRAM or RAM bus DRAM. However, the relatively long access latency lowers the efficiency of the data transfer on the memory system, so the actual net improvement is less than desirable.

On the cache application side, a simplistic view of the application requirement has always been that only SRAM's random access latency can meet the requirement, and that is the most important parameter. Today, these so-called "fast" SRAMs actually have a fairly low peak bandwidth capability, and this is primarily tied to the internal architecture that the data transfer frequency is determined by the internal memory access time. We believe improvement in even the traditional fast SRAM is needed to allow the much higher data transfer frequency, and with that, the new architecture is needed.

What is the driving force for speed? Certainly one of the most obvious parameters is the CPU performance. CPU performance continues to increase at a very fast rate, and then, to allow for balanced system performance, the memory has to perform in a proportional fashion. From the process technology point of view, certainly today's advanced process technology is capable of delivering 200 megahertz, or faster, CPU or microprocessors, or DSP engine, or any media processor, or any type of the processing logic circuits. The speed of those circuits will be at 300 megahertz, 500 megahertz, and before long, will exceed one gigahertz operation. That is certainly proof that the basic process capability is fully capable of supporting devices or circuits at those speeds. From the memory architecture point of view, innovation is definitely needed to match such explosive performance increases for the various processor technologies. One of the techniques already being used fairly successfully, is to decouple the access bandwidth from the internal memory access time. This is evident in the specification of the device of the synchronous DRAM or the RAM bus DRAM. However, improved access time is equally important. So far this aspect has not been addressed by the majority of industry suppliers.

Why DRAM technology? DRAM traditionally has been viewed as low performance, but low-cost technology. At MoSys we believe that the DRAM processes fundamentally are not different than the logic processers used by the advanced microprocessors or the various DSP, or media processors. The process technology is fully capable of operating at whatever frequency the processor is capable of. The DRAM, with its one transistor cell and the smallest area per bit of all the memory type, promises to be the lowest cost memory technology today, and perhaps for the next several generations. From the bandwidth point of view, there is really no limit on what DRAM memory devices can achieve, provided the right architecture is being used. From the access time point of view the DRAM internal access is always a two-step access, so by definition, it is always going to be somewhat slower than the high speed devices such as SRAM. However, this internal, inherently slower access time can be compensated by several techniques such as clever circuit implementation techniques used in the MoSys Mcache product line, or clever partitioning as used in the MoSys MDRAM product line, which statistically equals the performance of the synchronous SRAM.

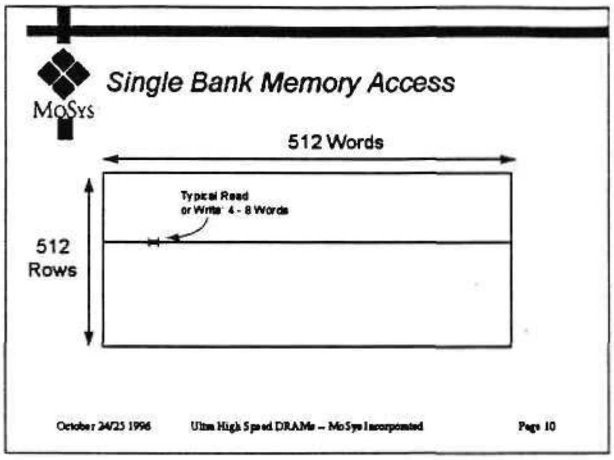

To look at the multi-bank concept, this shows a fairly traditional memory partition, a DRAM memory block, or memory bank, or memory array rows and columns. The typical access in most of the applications, whether it is graphics, multimedia, or processor accesses, typically involve read or write access of four to eight words, or maybe 16, but typically of that order. Most of the computer applications (the processor) now involve multi-processing, or multi-thread, or multi-tasking type of operations.

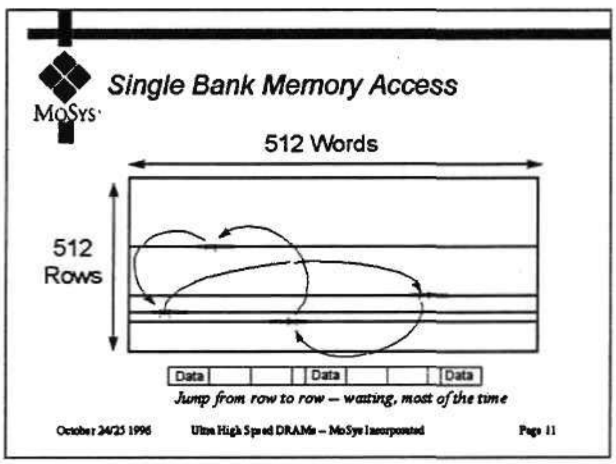

With a single bank type of memory device, the DRAM device, when the access is changing from task to task, the probability is such that the majority of the time the DRAM has a page miss. Even though the memory itself has a fairly high data transfer rate or the peak bandwidth, most of the time the memory bus is idle while the memory is performing precharge RAS and CAS access operations.

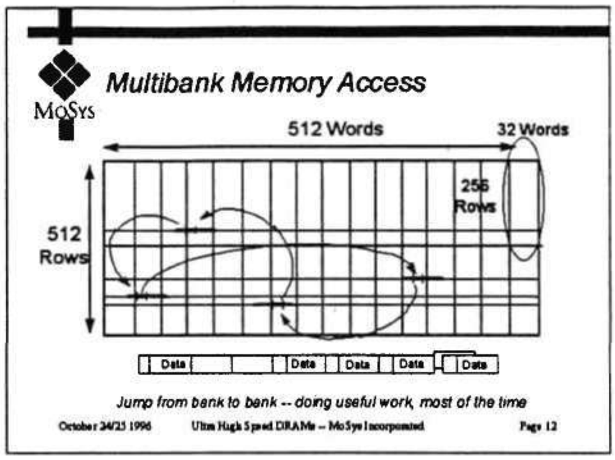

In a multi-bank memory architecture concept, the same memory array is partitioned into logically multiple, large number of small banks, so with that we can accomplish a statistical advantage. While all tasks are being executed by the memory, there is multi-tasking, the complex switching from task A to B to C to D, there is a high probability that when the task return to A, the page is still effective. Even in the cases where the data does not locate on "the right page," the multi-bank operation allows for very efficient overlap operation, thereby hiding most of the precharge and RAS operation inherent in the DRAM operation. With that, the memory subsystem is allowed to work most of the time, and is transferring data most of the time, achieving high efficiency, while still retaining and utilizing the same basic DRAM array.

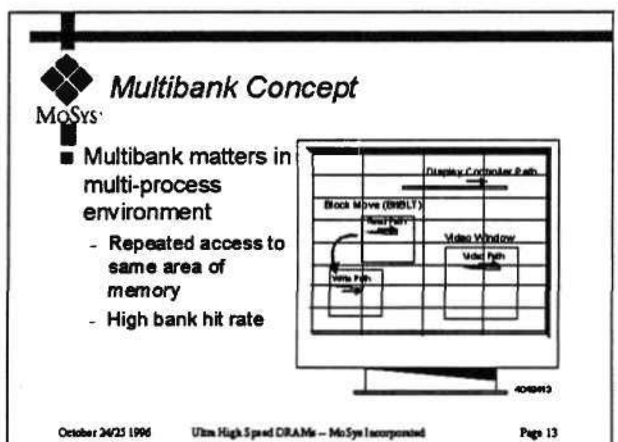

The multi-bank concept was specifically developed and engineered for the multi-processing environment. Repeated access to the same area of the memory ensures a high hit rate of the memory access, and this is resulting (statistically) in the memory with the multi-bank concept to deliver SRAM performance, because the majority of the access is the page access.

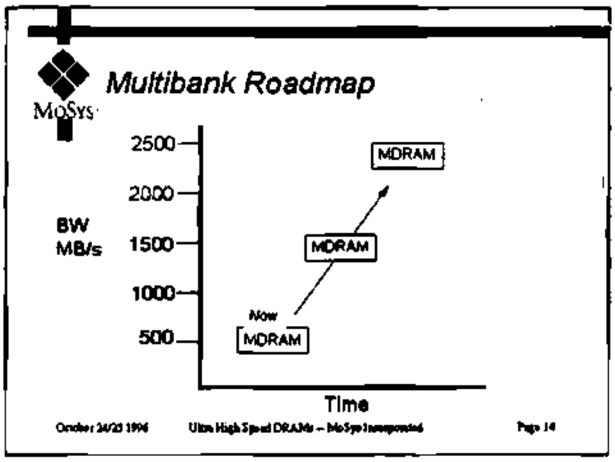

Today the MDRAM delivers the highest peak bandwidth memory at the 660 megabyte per second. We believe the multi-bank technology is fully capable of delivering the peak bandwidth at one to two gigabyte per second, or beyond.

A single chip, the MDRAM, has the bandwidth capability of one to two gigabytes per second over time, but one of the things that is most intriguing is that in the embedded application, there will be no limits on the I/O capability, or any bandwidth limitation for the multi-bank operation because every single memory bank operates independently. The performance capability potential on a chip is really one to two gigabytes per second per bank. This kind of capability provides almost unlimited bandwidth when it is used in embedded applications.

The multi-bank architecture provides a statistical advantage to take performance of a DRAM device and then effectively deliver the SRAM type of access, and data efficiency, by eliminating or hiding the majority of the row access and the precharge operations inherent in the DRAM devices.

MoSys announced the MDRAM product at last year's COMDEX show, and we started shipping the prototype in the first quarter of this year, with production shipping in the second quarter of this year. We shipped our one millionth unit within four months of our production starts. In the most recent PC World magazine, of the top 10 graphics cards being evaluated, the MDRAM-based solution won first place, second place, fourth place, and eighth place. We think that our MDRAM solution, being an application-specific, niche product, optimized for graphics and multimedia applications, has been relatively successful.

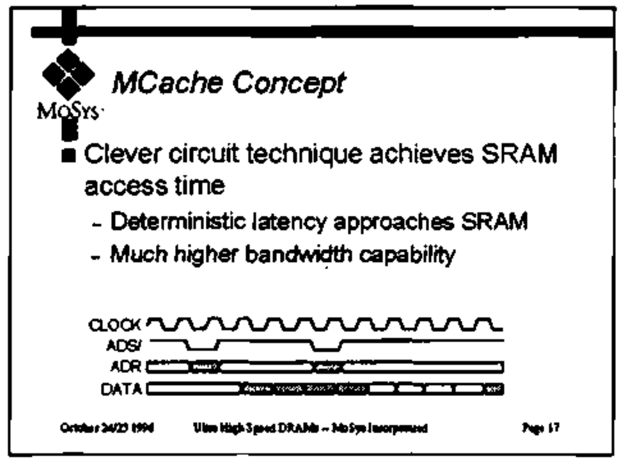

The second product family MoSys has developed is what we call the MCache product. This is the first DRAM memory that is directly a plug-in, drop-in replacement for the high speed, synchronous SRAM devices. On this product we apply both the multi-bank technology and clever circuit techniques so that we achieve the deterministic SRAM access performance with much higher bandwidth capability for future growth.

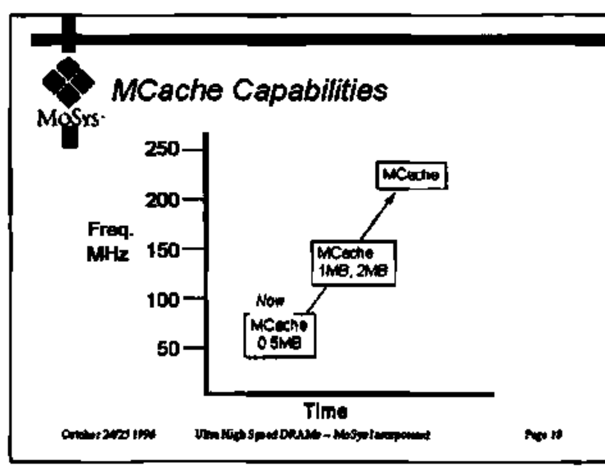

Today, the first generation MCache MPEG is offered with 66 megahertz and 75 megahertz operation for drop-in replacement for the Pentium class PBS RAM, at the 256K byte and the 512K byte configuration. Our MCache technology can be extended well above 200 megahertz operation and is suitable for most of the telecommunication, data communication, future workstation, processor, and the PC cache requirements.

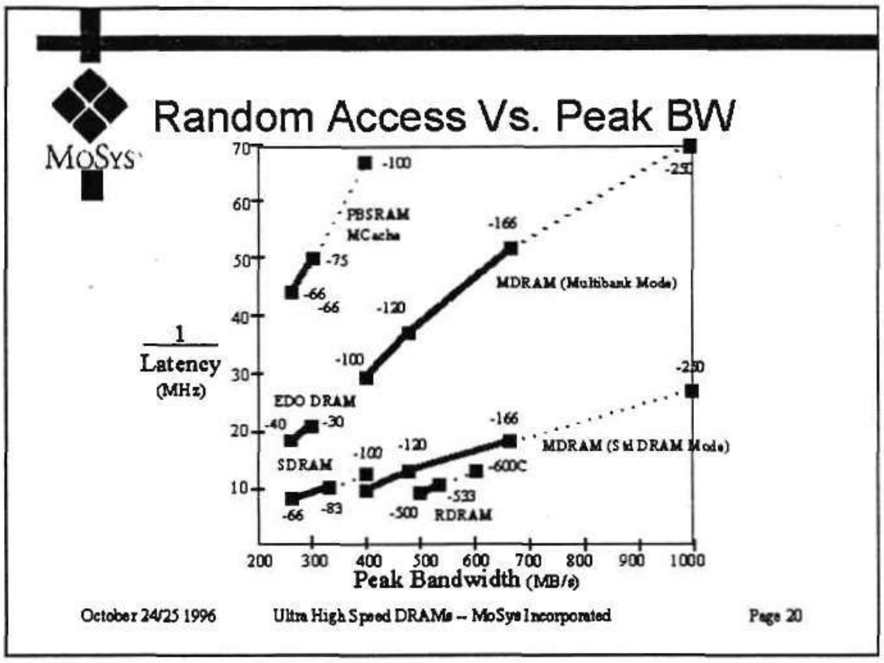

MCache technology not only delivers the same performance as a high speed synchronous SRAM, because of its inherent DRAM design, but it has much lower power than the comparable SRAM devices operating at the same frequency. The active power is one-third to one-fifth that of a comparable SRAM devices. In addition to the potential cost advantage and the integration advantages, this is a much more environmentally friendly device solution for the notebook or portable products.

This chart is by Steven Przybylski. It lists a comparison of all the existing and emerging DRAM devices and technologies. From the more traditional DRAM memory to the very advanced new memory devices, the memory system frequency has improved dramatically. At the same time, the row access time, or the basic access latency for most of the memory are relatively unchanged. Even though the peak bandwidth, or the maximum sustained bandwidth, can be quite high in most of those memory subsystems, the average sustained bandwidth can still be relatively low, even with very high peak bandwidths.

MoSys

DRAMs Compared: Przybylski

	EDO	Burst EDO	EDRAM	SDRAM	CDRAM	RDRAM	MDRAM
Density (Mbits)	All	15	4	2, 4, 16	4, 16	8, 16, 18	4 to 10
Width (bits)	1 to 32	4, 8, 16	1, 4, 8	4, 8, 16	16	8, 9	16
Interface Type	Async	Pseudo-sync	Async	Sync	Sync	Sync	Sync
Best Common Row Access Time (ns)	60	52	30	60	60	60	30
Memory System Frequency (MHz)	33	66	88-83	66-83	66-83	300	167
Memory System Width (bits)	64	64	64	64	64	2 x 8	2 x 16
Memory Controller Pins	~110	~120	~120	~120	~120	62	72
Sustained Bandwidth (Mbps)	152-267	213-533	356-533	193-533	193-533	217-960	666-970
Peak Bandwidth per Pin	19	35	44	53	46	155	148

Source: Steven Przybylski, June 1996

October 24/25 1996 Ultra High Speed DRAMs -- MoSys Incorporated Page 19

On this chart, I combined the data from the DRAM devices with the data from the SRAM devices. The horizontal axis shows the peak bandwidth, or the data transferring capability of the various memory devices, while the vertical axis shows the inverse of the latency, which gives a rough indication of the random access capability or bandwidth of the devices. The line in the middle demarcates the DRAM devices at the bottom, and the SRAM devices on the top. The figure of the merit on this chart is better on the upper right top corner, the corner indicate the best memory devices, having the best overall speed capability. The EDO DRAM, the synchronous DRAM, the RAM bus DRAM, and the MDRAM (the multi-bank DRAM), is shown in its native mode, or the standard DRAM mode. There can be a tremendous improvement in the peak bandwidth, but in terms of the random access bandwidth, there is only limited improvement you can achieve.

Using the multi-bank architecture changes the picture to make the effective random access (the latency), of the MDRAM equal and comparable to that of the standard SRAM. The other case is that the MCache product, being a direct plug-in replacement for the PBS RAM, is exactly the same as the performance capability of the PBS RAM. With this chart, from MoSys' perspective, although DRAM is inherently slower than SRAM, with appropriate architecture and circuit design techniques, DRAM can be very high performance memory, suitable for almost any applications existing today, whether it is served by DRAM or SRAM devices.

Speed is an application-specific concept. We certainly believe that DRAM devices should be application-specific. With that, we have the flexibility to optimize the DRAM circuits, the DRAM memory array to deliver the highest performance possible, for that particular application. Many other side benefits go with that. DRAM technology always has the lowest power of memory devices that operate at a similar speed. It has the lowest cost. It can deliver the same or better performance because the small cell size of the DRAM allows much more flexibility and capability in terms of higher level of integration, leading to possibly a single chip solution, and embedded solutions. The MDRAM and MCache products are only examples of application specific implementations that take advantage of the latent performance capability of DRAM technology. This application-specific concept, using a very high performance DRAM core, can be adaptable to any system partition and requirements.

Q: The previous speaker from NeoMagic indicated that using a commodity DRAM process was the way to go in building an embedded DRAM product. Do you agree?

A: Yes, we agree.

Q: How many MCache products have been shipped, and when were they introduced into the marketplace?

A: The MCache product was announced in April of this year, so we are just starting prototype shipment in the late second quarter and the third quarter of this year. We currently have three devices, 32K by 32, 64K by 32, and a 32K by 64 which covers the desktop applications. The 32K by 64 is particularly interesting. It is what we call the mobile MCache, or the green MCache. It requires less than one-fifth of the power of the SRAM, the comparable similar solution.

Q: Can you tell us a little bit about the manufacturing technologies required to build the MDRAM?

A: There are actually not very many requirements. Our first MDRAM implementation was done on the 0.7 micron technology with a 1 micron transistor. With this generation of technology, we have achieved 150 MHz (or 300 megahertz data rate operation), and so far, most of the technology migration we are accomplishing are primarily the cost down.

How to Win the Digital Communications Race

Hatch Graham

*Senior Vice President and General Manager,
Personal Communications Group,
TCSI Corporation*

Mr. Graham has been with TCSI Corporation since October 1995 as the senior vice president and general manager of the Personal Communications Group. He is considered a leader in business and technology within the broadband and wireless communications industries.

Before joining TCSI, Mr. Graham was corporate vice president and general manager of the Telecom Products Group of Stanford Telecom where he is recognized as the principal in Stanford Telecom's successful progression into commercial communications. He has served as the chairman of several communications symposium sessions, and has been featured as a leader in the electronics industry by publications such as *Electronic Engineering Times Magazine*.

Mr. Graham has successfully established major business agreements with industry leaders such as Motorola, AT&T, Hewlett-Packard, Westinghouse, Scientific-Atlanta, Zilog, DSC Communications, Tellabs, Goldstar, and Hughes Corporation. He has also personally developed many of the world's fastest and highest performance integrated circuit architectures for digital signal processing of radio and wireline communications. The applications of his ASIC designs range from cellular telephones to mobile utility meter reading to telephony over cable TV. Mr. Graham's patent activities include the invention of architectures for direct digital synthesis, programmable logic device, digital filtering, forward error correction, spread spectrum communications, cable TV, and cellular telephone applications.

Mr. Graham holds a bachelor's degree in engineering from Idaho State University where he now serves as a member of I.S.U.'s Engineering Advisory Council.

Session 12: How to Win the Digital Communications Race

Hatch Graham
Senior Vice President & General Manager, Personal Communications Group, TCSI Corporation

GENE NORRETT: We are pleased to have with us Hatch Graham, who is Senior Vice President and General Manager of TCSI. Not only is he involved in helping the customers, he has actually done a lot of inventing. His patent activities include the invention of architectures for direct digital synthesis, PLDs, digital filtering, forward error correction, spread spectrum communications and cable TV.

HATCH GRAHAM: I have a unique perspective of the changes in semiconductors for digital communications, because I belong to a company that has two businesses. One is providing semiconductor architectures and embedded software. The other is network and business management software for all the big PTTs, RBOCs, big service providers who are building this new wave of infrastructure. I think it will have much more of an impact on semiconductor manufacturers in the future, more than in the past. I think what I can present to you is a need for semiconductor manufacturers to be prepared to change, and change quickly, and maybe some of the components and types of technologies that you might keep in your arsenal to prepare for this sort of change.

We are going through a change where we have access to products, whether we are mobile or at home. We have different issues related to networks and how they provide services to us. It is becoming more and more convenient and we are able to get more and more information than we ever have before. Imagine the concern within the Pony Express Marketing Department as word of the telegraph spread. We saw the courier give way to the telegraph, the telegraph gave way to the telephone, and we had audible information. The telephone gave way to television and we got visual information. Now we have everything coming to us. Right now we are between the future and the past, we are integrating voice, data, and video, and we are getting pretty good at it. That is going on now.

We still have the mailman driving his Jeep to drop off the mail. Things like this did not go away, they just took a different form. We still have the need, but the means have evolved.

Communication is all about access to transfer information from a source to a destination. From the beginning of time the goal has been the same, only now we are packing in more information and again, we are getting pretty good at it, but there is a long way to go.

A few years ago, digital communication came around. Digital communications is expanding the limits. What the 'digital' does is increase capacity, quality, efficiency, and performance. The choice now becomes how to address consumer decisions so semiconductor manufacturers have to study the changes in digital communications to be ready for it.

A few of the benefits are information throughput, more data, higher speeds – achieved with semiconductor advancements and applied theory. Bandwidth optimization (I call it transmission density), are modulation techniques to pack in more information per bandwidth, per hertz. In a certain amount of bandwidth, there are different techniques that give a very robust signal, but if you have a very clean environment you can actually enhance the number of bits per hertz that you send and again, you have to weigh these trade-offs. Hybrid fiber coax cable TV has something called Ingress or impulse noise. There are

Graphic materials for this speaker were not available at the time of publication.

different ways to guard against that than in certain other applications, so these are the decisions that have to be made.

Quality, we all hear about forward error correction such as Reed-Solomon coding, and Viterbi decoding. These are basically ways to get better signal performance, or better bit error rate per a given signal to noise ratio. Spread spectrum is a different technique to guard against interference. Take all the energy in a certain bandwidth, and spread it across a wider bandwidth so if interference comes in and punctures a certain area, you have only destroyed a small percentage of that signal. There are ways to improve the quality and immunity to interference. Privacy, same sort of thing. Lower battery power needs, error correction helps you transmit at lower power because you can pull a weaker signal out. These are all benefits of digital communications and the architects that are working on these semiconductors architectures need to take these into account for the given applications.

Here I made a pretty bold statement. "The progress and commercialization of Digital Communications coupled with continued semiconductor advancements represent a change to commerce as significant (and similar) as prior milestones in communications."

I know we have not invented anything new. We are only integrating instant literal information, and instant visual or audible. We are making it affordable for the consumer to get information, whether they are mobile or not.

Integration of the three C's: computers, communications and consumers. Twenty years ago, 50,000 computers were sold, now 50,000 are sold per day. More money is spent on PCs than TVs. By 2000, over a 100 million cellular PCS phones are expected to be sold.

By the end of the decade, over a billion people are expected to be on the Internet. Everybody knows they are affected by communications today. Finally the consumer can be much more proactive and interactive. We have gone from the choice of two channels on TV to being able to make the choice of a little satellite dish and get direct broadcast satellite, a VSAT dish, or we can have hybrid fiber coax cable TV. There are all sorts of different techniques to give you a choice to be able to select what you want to see and where you want to see it. Service providers are competing to get you to select their service offering.

In the future, any product may have (and need) affordable, interactive access to a network. You as a consumer, either at your home or on the road will be able to buy any form of information, on any product, provided by many companies out there. I am not saying that every product is going to be integrated to offer everything that can be offered, but the capability will be there.

There are very complex value chains emerging, equipment system manufacturers, providers of the service, and service providers. Worldwide deregulation is causing competition among these companies like never before. They are doing what semiconductor manufacturers have done for quite some time, striving to become more competitive, improving customer care, and reducing time to market.

Integrated service management can be broken down into four parts: business management, service management, element management, and then managing the elements is the network management. Service providers now have an increased capability at the service centers to manage all of this equipment out there, change it in some cases, to provide you what you want, upon your demand. For example, if you want ISDN to your home today, you can make a phone call to your telephone company, they either have that service in your area or they do not. If they do, you wait a couple of weeks and they will come hook it up for you. In the future you will be able to call your telephone company, and without speaking to anybody push a few buttons and have your home provisioned to have 128 kilobytes from 9:30 to 10:30 Saturday morning when you need to download a big file. You will then be billed for that hour, and you will turn it off. That is the future.

This will be based on proactive consumer involvement, you pick what you want. That is what digital signal processing and these advancements in semiconductors will enable you to do. Another example is cellular phones. You could have the ability to push a function button on your phone and have premium service, meaning if you wanted to spend $3 a minute instead of 22 cents a minute, to have the best possible phone call, to be handed off between base stations at exactly the right time so you do not have a weak signal, so you have top priority, to allow you to have software downloaded to your phone to improve the equalization scheme to get better performance in your signal processing. All of these things are possible. All of these things can happen in the future. That is what networks are preparing to do, to bill you for services that you choose.

Business management includes financial control, marketing, payroll, service management, and billing. Billing is a hot topic. Companies do not like giving consumers something without being able to bill them. You need to be able to request the service you want and not be stuck with it when you do not want it. Element management is basically all of these networks you read about every day, broadband, PCS, cellular. Plain Old Telephone (something you have had for quite some time), as well as SONET, ATM, SS 7. These are all networks that make signal processing more convenient.

These networks really do some tough work. Fault management is used if there is a problem today. Error flags come up in service centers, there can be several hundred errors show up that have to be audited individually, logged in, and corrected. One part of the network going down can cause days of work.

The sophisticated software nowadays can take care of these problems. Companies have to revise their entire outdated infrastructure. The reality is they handle five million events per day, thousands of events over just a few minutes, so these are very complex issues. As this access comes to consumers, and as these services are used more, systems have to get better and better.

Involved in the challenge of migrating from legacy systems, called "spaghetti" code in software development, they keep writing code and writing code and putting in patches. Now we see frameworks and application frameworks where 80 percent of that software package is already done. Then they build framework packages and application packages to do different things like SONET, like billing, each one of these is a different object, a different package. In summary, for these networks it is an increase in service management capabilities, so they can change easier, so they can manage these services as independent accountable objects, and they can upgrade them independently without affecting the rest of their network. They can change equipment in the infrastructure using graphic user interfaces. They can see the geography and literally dive all the way down into the map to the actual piece of transmission equipment, look at it and correct a problem from thousand of miles away, or dispatch somebody to fix it, or provide additional services to the consumer.

Service access expands with each new network that comes up. Each of these access points, to each piece of equipment, is different. That is the key because for a semiconductor manufacturer to be able to have product and content in each one of these products you will have a slight variation. In some cases it will be dramatic and in other cases it can be something that you can set up and prepare for. The majority of these are point to multipoint systems, essentially a point that distributes information to several other points, effectively an operation maintenance center. Then base transceiver stations finally send signals out to handsets.

I broke communication access down to something that has five blocks of an access from an antenna to a consumer. There is a conversion process that takes a signal from a carrier frequency and brings it down towards the baseband that you operate the signal processing in. There are different ways of doing it: stopping in an intermediate frequency and picking up digitally from there, or take it all the way down to baseband. Modulation is QAM, BPSK, QPSK, even something called 'spreading,' and that is CDMA, spread spectrum. In some cases you pick up the signal from the conversion at a point that you can do digital signal processing in a DSP chip, or in some cases you need to dedicate circuitry. Error correction is the decoders and

Reed-Solomon coders. Compression is something that has really enabled digital signal process to help the consumer. Compression in video, MPEG, JPEG audio compression, VCELP, PSI CELP, different ways to compress a signal that you can hear so you essentially get more signals within a given capacity or improve the quality of the voice or the audio. Basically you have a stream of data at a certain point that you need to turn into something that can interface with the network or the consumer – I call that formation. Some people call that media access control, interface, I just blocked it out as formation. Remember that these are certain areas that just about every network, every point to multi-point network has to address and its content is all semiconductor based.

As these networks succeed, standards emerge. I used to claim that you did not need a standard to sell a lot of chips but to have high volumes there has to be something that is a target. It has to be set by some standard and even if it is created as a proprietary architecture, to be successful it has to finally give way and become a standard. Then there can be competition and there can be again a still target but as soon as these standards are set, they evolve rapidly. In some cases before the standard is even agreed on, there is the next evolution of that standard which affects those five blocks in different places.

Semiconductor manufacturers must plan to optimize this change process, these standards. All these new ways of having access to a consumer, all these new networks, are all going to be a little bit different. There has to be a game plan that says we can adapt to just about any new access type as rapidly as possible so that we can be in the market and be selling products before anybody else. There is always going to be an improvement on capacity, quality improvements, throughput – all of these things are going to be very important.

So what can change? Across the life of a standard, typically, this carrier that I talk about, the RF point, will stay set. If it does change, it evolves to become a new standard but through the life of that particular standard, that RF circuit can remain fixed. The modulation scheme, the process of timing recovery, and some of the carrier tracking specifics remain the same. Error correction needs to be a little bit more dynamic. There are always trade-offs between bandwidth optimization and correcting the signal but the real change happens in compression algorithms. This is where in every standard there is always a push to get more people per given bandwidth, to get better quality, but keep in mind that there is always a change needed in the compression schemes, there is always an advancement coming along, and it will be called a different standard that usually falls within a major category.

I show an example of digital voice for cellular. This is an example of compression evolution in cellular telephone. North American IS 95 is CDMA and Japan is personal digital cellular. Each one of the four primary digital cellular technologies has evolved in voice coding almost before they have been adopted. If you look at IS 136, its predecessor was IS 54 which was again North American TDMA. IS 54 came out and it had some issues with it, with the Raleigh fading characteristics, with the hand-off from base station to base station. IS 136 emerged as the brand new standard that replaced it for cellular, and it had improved quality through a digital control channel.

Once it was set, it evolved in voice coding from VCELP to ACELP, to a different rate of ACELP, and now there is a half rate in process but less than 5 million of these phones sold. GSM is another good example, going from full rate to half rate, to enhanced full rate, and as this cellular version of GSM splits off and becomes the PCS upbanded GSM version, again there is an opportunity for change.

Actually, this is a bit exact architecture that comes out with full quantization, the bytes being exact to the bus structures. Effectively it is a new standard. The next chart shows a relationship though in that yes, it is new but you will see that several of these standards have a relationship between them which makes things a little bit easier. You have to prepare for the change but if you know a little bit about the basis of the change, you can prepare.

These are all versions of a certain voice coding algorithm that was developed at the University of Sherbrook (Nokia was involved in this), and effectively the voice coding standard for several different cellular applications is very similar and comes from the same root foundation. In effect, if you build one of these boxes you are only about a month or two away from recoding to get to the other rectangular box. As a game plan, if you have the right semiconductor foundation, you can make changes quickly and address several different markets without having to repeat the same amount of work that you did the first time.

The first thing in preparing for this continuous change inside of these access channels is the adoption of embedded processors. Some of these things can be cores, they can be library cells and nowadays these cores are becoming more and more optimized for the application. For example, the DSP core from my company is basically dedicated toward signal processing type applications, especially wireless. It has certain features built into the instruction set that make it easier to do algorithm coding than if you did it from scratch. Once you integrate this and do the layout and target it to your technology, as a semiconductor manufacturer, you can use it as a cell. When new voice algorithms come out, you can literally make the changes in software and it becomes a ram-mask change or if you have memory outside, it is essentially no change at all. But this cannot handle the RF sections (what I called conversion) and it cannot handle some of the modulation portions, but in most cases, it can handle the error correction and the voice coding. In many cases, it can handle all of the modulation, too.

The second thing is something that is evolving in the base station sides of things, more than it is the handset, and this is what I call embedded accelerators. Do not give up on custom, or hard ASIC logic, because in many applications there is a need to have a certain type of function that could be dedicated to a hard accelerator, and leave a lot more room in these embedded processors. I show a chart that has many different functions and it has basically the number of cycles needed for a certain piece of a function. If you have an application, maybe an Internet access where you need to do FFP (fast forward processing), you might want to have a dedicated FFP accelerator married with a DSP core that picks up voice compression and other modem functions. This is another thing I think you need.

We just cannot seem to get away from A to D and D to A converters. These bring analog techniques and mix them up with digital techniques. There are many companies that have A to D and D to A converters, but in signal processing this is a must. To have either as a cell, or as a very valuable supplier in bringing a signal from the RF to the digital domain, you have to have good performance at low cost in your A to D and D to A technology. This is something necessary.

As more networks come in with more access to products, there will be more software coming from a network to a product. So there will be a need for more memory on chip. I believe there will be more memory within products. I also believe that there will be a need for reconfiguring. If you have more access types that share a common RF, conversion can be made the same but they have a different modem structure.

By loading different information into your memory and reconfiguring the processor or the actual instruction set, you could literally have multimode technology. This is coming. You could have a phone that addresses one sort of standard that immediately, on the fly (or on your demand), becomes adaptable to another standard. To do this, I list flash memory for example. Reconfiguring on the fly has to be something that is thought out. If you are going to dedicate your chip to an application, be prepared to throw it away as a standard evolves. This is a technique to basically prepare – as the cost of flash comes down, it will be much more attractive to have multi-mode products.

Mapping these functions to the needs again, you still have the logic I like to term as accelerators because I am becoming a believer in DSP cores and what the generic nature of a DSP can do for you. You have to have the A to D and D to A technology. The RF integrated technology is still an art that lies on its own technology platform, and will be its own chip. It is not yet integratable to the digital signal processing and CMOS. RFIC, RF components are obviously something that you should consider too, but I wanted to focus on the areas that can be integrated together.

In summary, there is a proliferation of networks. All of the big RBOCs, PTTs are all essentially revising their systems to offer you more services, and that means that the products you buy will have access to any kind of information, whether you are moving around or sitting still. There is going to be an endless change to the service needs, and there is a similarity to the access of this communications. I believe it provides an opportunity for semiconductor technologies to be balanced out, optimizing cost, performance, configurability, and time to market, as these new networks come up, and as these new access points are integrated into the products.

Q: What type of MCU cores, the microcontrolling units, are suitable in these applications?

A: In a lot of cases, depending on the architecture of the signal, the modem, there is a lot of set-up that goes on. When you receive a signal, whether it is wireless or across a cable, you have to acquire the signal, then you have to track, so the call processing that goes on is held in microcontroller units. These are not the same as the actual DSP processors so typically, in the cellular phone, there is a DSP core that is used for the physical layer, the modem, the error correction. That DSP core is crunching out that data. Then there is a microcontroller unit that handles the call processing. What microcontrollers units are suitable? Well these things do not typically need to be more that 16 bytes, they can be 8 bytes, but there are applications that are driving them to more resolution. There are several companies that offer suitable controllers, but these are not what is performing the DSP functions but rather the protocol of call processing and set up.

DSP Chips: The Hottest IC Segment in Sight

Joseph Grenier

Vice President and Director,
Semiconductor Device and Applications Program,
Semiconductors Worldwide,
Dataquest

Mr. Grenier is vice president and director of Dataquest's Semiconductor Device and Applications program. He is responsible for managing the semiconductor device research for Dataquest's Memory, Microcomponents, and ASICs programs. He also manages the research for Dataquest's semiconductor applications including the consumer, PCs, and communications programs in these areas.

Before joining Dataquest, Mr. Grenier was product marketing manager at GCA Corporation where he managed marketing activities 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 B.S.E.E. degree from the University of Detroit and an M.B.A. from the University of Santa Clara.

Session # 13: DSP Chips: The Hottest IC Segment in Sight

Joseph Grenier
Vice President and Director, Semiconductor Device and Applications Program, Semiconductors Worldwide, Dataquest

JOSEPH: I am going to talk about the coming DSP revolution.

My goal at this presentation is to raise your awareness of DSP, its many applications, and the importance DSP will play as an enabler in communications, consumer and PC multimedia.

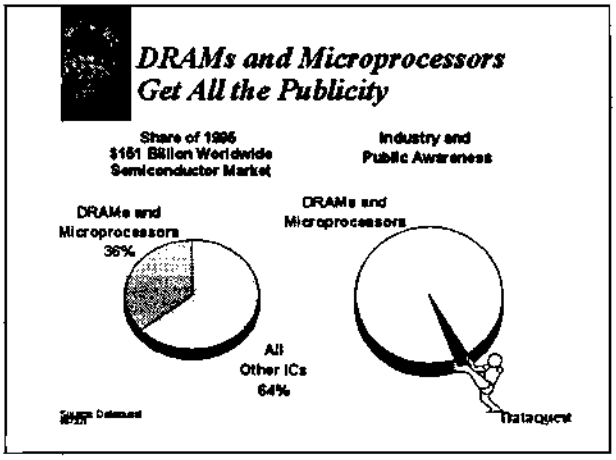

The pie on the left shows that, in 1995, $151 billion semiconductor market DRAMs and compute microprocessors got about 36% of the revenue. The pie on the right shows DRAMs and microprocessors get all the press and publicity in the media.

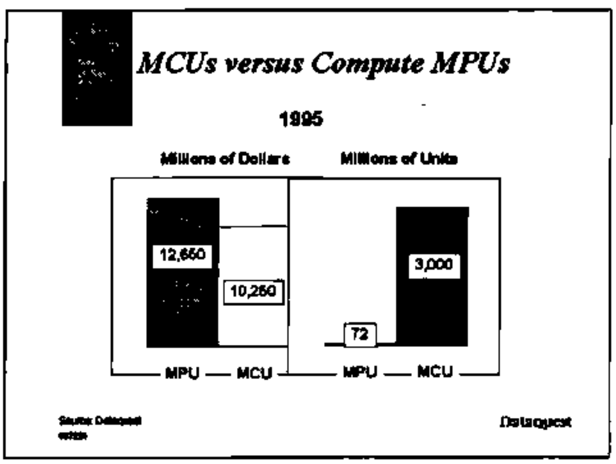

Microcontrollers or MCUs do not get much attention either.

I am going to show you some interesting facts on MCUs that you might not be aware of.

In 1995 the compute microprocessor market, which is mostly X86 type processors, totaled over $12 billion. The MCU market is over $10 billion. The three billion MCUs shipped in 1995 dwarfs the 72 million MPUs shipped. Motorola has forecast, by the year 2000, there will be more than 200 MCUs in the home.

This shows you the MCUs forecast. Today they are around $10 billion. There will be a slight pause, then a continual growth reaching $20 billion.

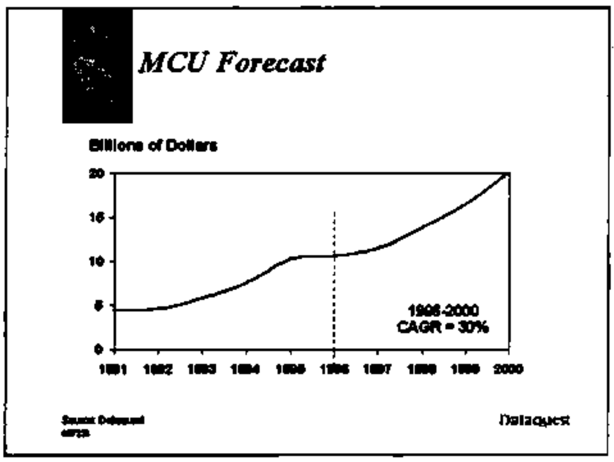

Let's get back to DSPs, the main topic of discussion.

Why don't DSPs get much attention? People have no idea what DSPs do. Even folks in the semiconductor industry have to struggle to define DSPs.

I suspect a lot of you are secretly wondering, "Yeah, just what does a DSP do?" This is what a DSP does.

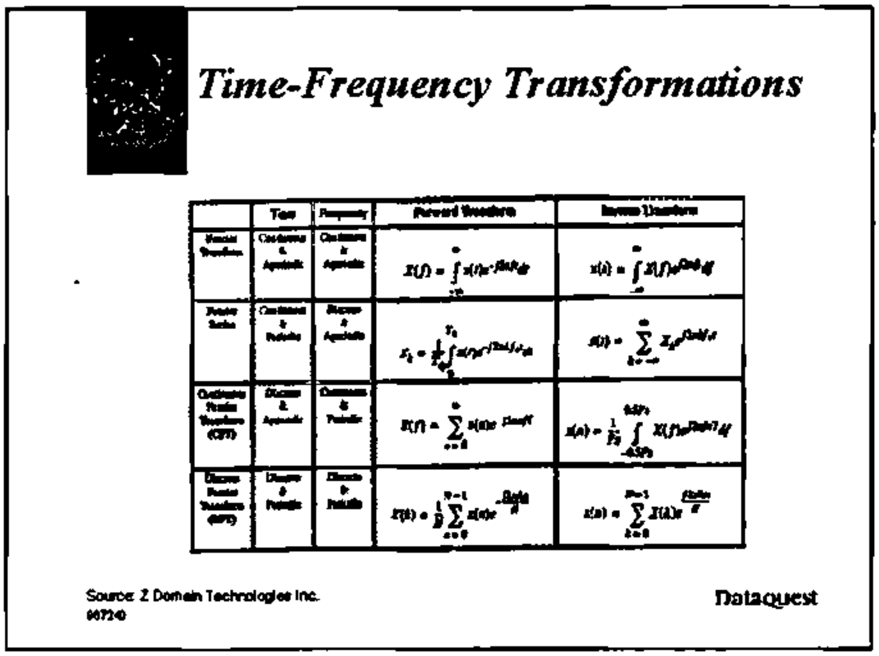

Unless you are a chip designer or a programmer this graphic does not help you very much. This tends to scare people away from understanding DSP.

A good example of an application that requires Digital Signal Processing is an automotive noise canceling system. Ambient analog noise is picked up and digitized with an A/D converter. A digital noise signal is then processed and sent to a D/A converter, where a process noise signal is converted back to an ambient audio cell or an ambient analog sound. That feedback should perfectly cancel the original noise signal. What you will hear if everything is perfect is silence. One characteristic of digital signal processing is that it is real time processing, no delay. To do real time processing DSPs have to be very fast.

The second characteristic of digital signal processing is that it usually deals with real world signals, analog signals that either originate or terminate in the real world. A third characteristic of digital signal processing is that it is very math intensive. In this application, creating another noise signal exactly opposite in phase and amplitude to the original random noise signal is very complicated, but this is what DSPs are best designed to do.

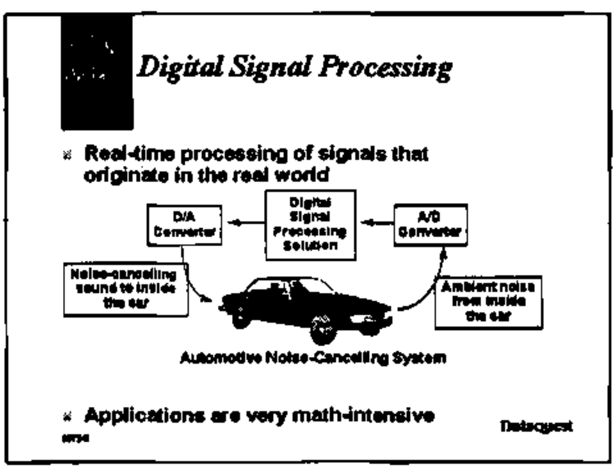

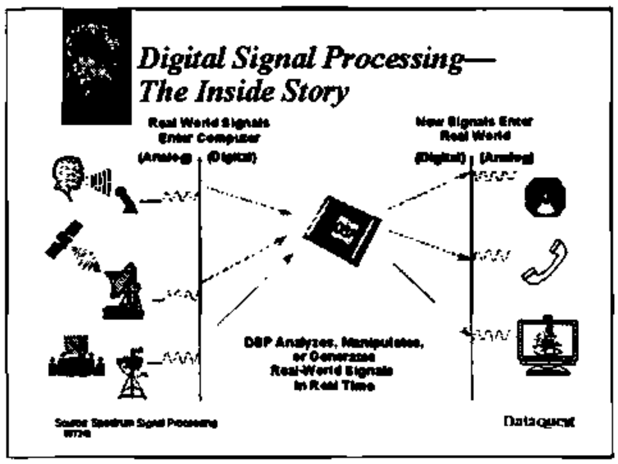

Here are some of the things that digital signal processing does. Analog signals that originate in the real world are digitized. The digitized signals are processed and then converted back to analog signals for presentation back to the real world. Not all digital signal processing deals with analog signals, this is a simplification to help you understand the process.

We refer to DSP chips and to digital signal processing solutions. What is the difference? There are a number of ways to implement a digital signal processing solution. As shown here, they range from discrete logic implementations to host MPU signal processing, to digital signal processing done with MCUs and MPUs, to a programmable DSP chip, to an ASIC with a DSP core, and finally to fixed function DSP chips.

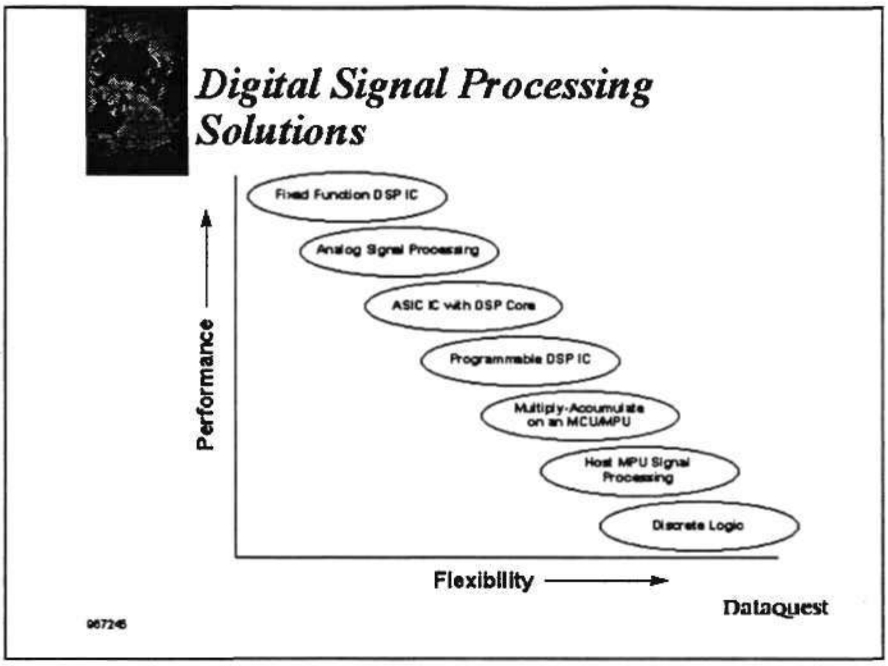

The automotive noise canceling system the box was labeled digital signal processing solutions. In theory, that particular solution could be implemented in a number of different ways. Each type of solution would have different cost performance characteristics and some may even be impractical.

A fixed function DSP is one that is designed to do a specific task. It cannot be reprogrammed to a different task. On the other hand programmable DSPs are general purpose DSPs that can be reprogrammed to do many different tasks. Only fixed function programmable DSPs are considered to be DSP ICs.

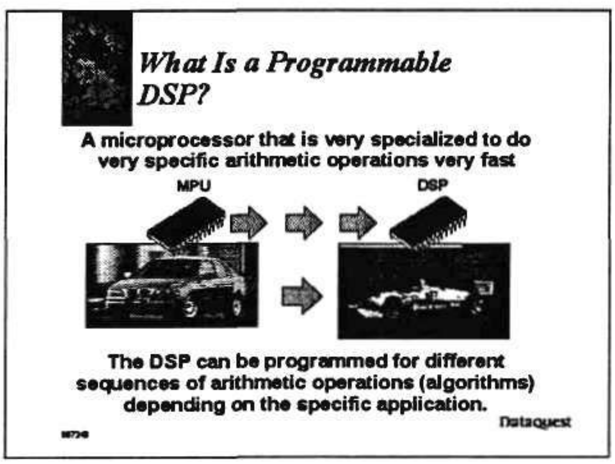

Programmable DSP chips are a very high growth area for digital signal processing. A programmable DSP is a microprocessor specialized to do fast arithmetic operations. Here is an analogy that you may find useful. A DSP is to a microprocessor what an Indy race car is to a high end luxury car. For instance, a high end luxury car is a general purpose car that can do many things, is comfortable, has many features and can go quite fast. An Indy car on the other hand, is very specialized, and designed for one thing – speed. This is a pretty good analogy, except when it comes to cost, the analogy breaks down. Indy cars cost more than luxury cars as they run about a half a million dollars a piece.

Another characteristic of programmable DSPs is that they can be programmed by the user for his own particular math-intensive application.

Now that we know what DSP solution is, let's take a look at some DSP applications. These applications require very compute intensive real time processing. DSP has long been used in military weapon systems. Notice the words 'vision', 'imaging' and 'pattern recognition' in several applications. DSP excels in these areas because of its ability to process real world images. Food inspection is a good example. Fruit on a conveyor belt passes video cameras at several feet per second. Images of the fruit are processed to be passed or rejected based on attributes such as color, size and texture. This system eliminates human inspectors with great savings and makes less mistakes than human operators.

While not the high growth areas or high volume products, I mention them first because they are some of the more traditional applications of DSP and they do indicate the range of interesting math intensive real world DSP solutions.

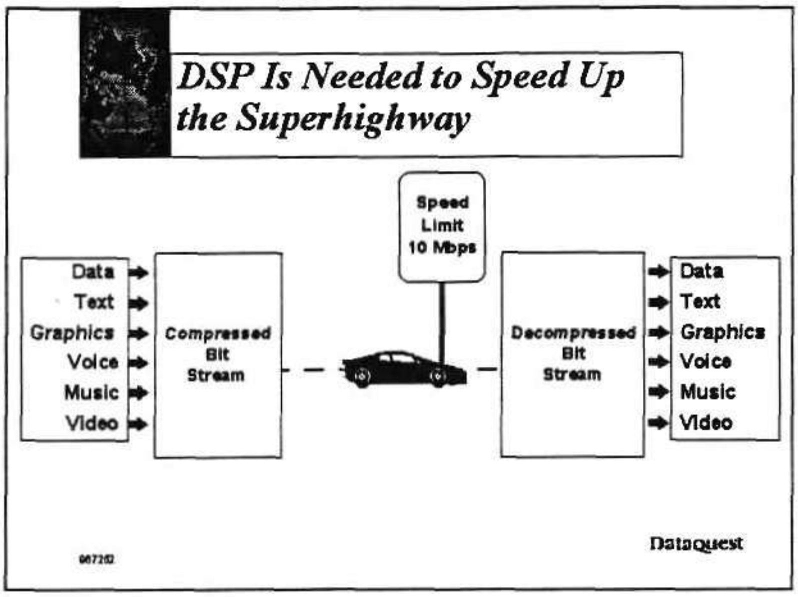

Communications is a leading application segment for DSP chips. Half of all the programmable and fixed function DSP IC revenue results from sales to communication applications. Modems, fax machines, and cellular phones are the leading communication application segments for DSP chip revenue. For instance, many of us have 28.8 modems, but soon 56 kilobit modems will be available. You will not have to throw your old modems away. To change from the 28.8 to a 56 KBs modem all you will need is software that will reprogram the programmable DSP. If your modem is by U.S. Robotics, they will probably make that software available to you.

Communication applications require real time compute intensive functions, such as voice and data compression and encoding, carrier modulation and demodulation, speech recognition, and text to speech conversion. Notice the words 'voice' and 'video' throughout these applications – more examples of signals that originate in the real world.

Here is a partial simplified block diagram of a cellular phone, indicating why DSPs are used in some of these communications devices. The speaker's voice is digitized with an A/D converter, the digitized signal is compressed, the compressed digital signal is encoded and encrypted for the transmission protocol, and then sent to the RF system for transmission. Upon receipt of the signal the process is reversed.

Here we are dealing with analog signals that originate in the real world. Digitized signals have to be processed with high speed, in real time. Compression, coding, and encryption algorithms are very math intensive. The DSP solution has to be low power to prolong battery life, and low cost because cellular phones are becoming consumer devices. Currently phones may use one DSP chip for encoding and decoding, another DSP core on an ASIC chip. In the future, in order to reduce cost, probably all of these baseband functions can be combined on a single chip.

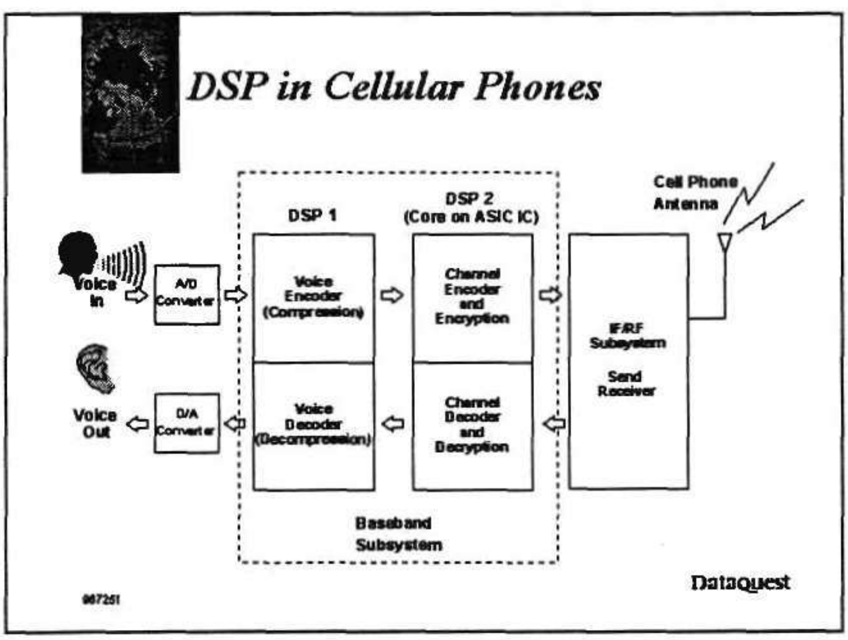

ll traffic on the information super highway will be digital. Not only will it be digital, but it will also be compressed with digital signal processing. Why compression? Because bandwidth is both limited and costly on a super highway. Compression is required to pack the digital traffic more densely to take better advantage of system resources. For instance, video data requires more compression than any other. Transmitting uncompressed video data requires a huge amount of bandwidth, about 216 megabytes per second. That same video data compressed requires six megabytes of bandwidth.

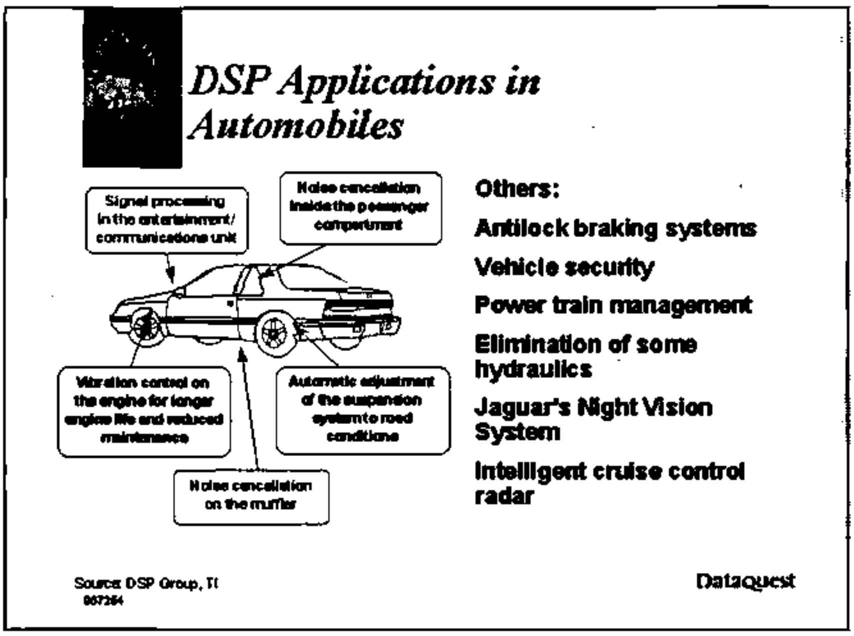

Digital compression is a key enabler of both tomorrow's communications, the super highway, and communications in what I call the 'wireless' super highway, such as direct broadcast satellite TV and direct broadcast radio. Digital compression will be vital in the success of PC multimedia, video conferencing, interactive video training, games, and digital consumer products, such as digital camcorders. Digital compression requires fast, low cost, real time number crunching capabilities to execute the compression routines, or the algorithms. Today's DSP chip technology along with industry compression standards, or algorithms, like JPEG and MPEG, can make this happen.

Here are some applications in PC and multimedia that will require digital processing. Some involve audio and video digital signal processing, which will require digital compression in intensive math calculations. 3D audio requires some pretty hefty mathematical manipulations to create the 3D sound we hear. Continuous translation is real time, continuous translation of speech to text and is another example of intense real time mathematical processing of an audio signal.

PC fax modems are already a very high volume application of DSP chips. Another example is disk drives, traditionally MCUs have been used in hard disk drives but for the new high end drives MCUs are beginning to run out of steam. Seagate has recently introduced a 2.5 GB hard drive, which uses a custom DSP chip, integrating a DSP core with logic and flash memory. In this application DSP provides better performance. The MCU, using past designs, is a DSP processor that is ideal for executing the complex position algorithms that are required to quickly position the rewrite head over the track.

Today's autos already have several microcontrollers. DSPs will also find their way into cars, in those applications that require fast real time processing. We have already mentioned noise canceling systems, some cars already have DSP controlled automatic adjustment of the suspension system for smoother rides. Hydraulic systems are used for power steering. With the new DSP designed for motor control, we may see the pump and belt systems replaced with direct drive motors. Another example is Jaguar cars. Jaguar, in conjunction with TI, has developed a night vision system that uses a near infra-red camera to allow drivers to see objects in the dark that would not be visible using normal headlights. The range of the near infra-red system is far greater than that of standard headlights. DSP is used to process the images detected by the IR camera.

In consumer electronics DSPs will be used throughout the new wave of digital electronics which is now emerging. DSP will be used wherever there is a need to process audio, video and other images like photographic images. You have already heard the terms AC3, MPEG-1, MPEG-2, JPEG, these are basically just algorithms for the digital processing of audio, video and images.

While not a particularly high volume application, restoration of Hollywood films is certainly an interesting example. The restoration of Snow White took 18 weeks, 60 workstation operators using 40 workstations, working three shifts, seven days a week. Using digital signal processing the restoration could be done in much less time with much less money and operators.

DSPs are also finding their way into arcade games as motor control chips, into power tools, and home appliances, replacing MCUs in some cases. Sega's new Indy 500 arcade game uses two analog devices, Shark DSP chips. Shark is the fastest floating point programmable DSP on the market, capable of 120 million floating point operations per second. The Indy 500 game is capable of rendering about 300,000 polygons per second.

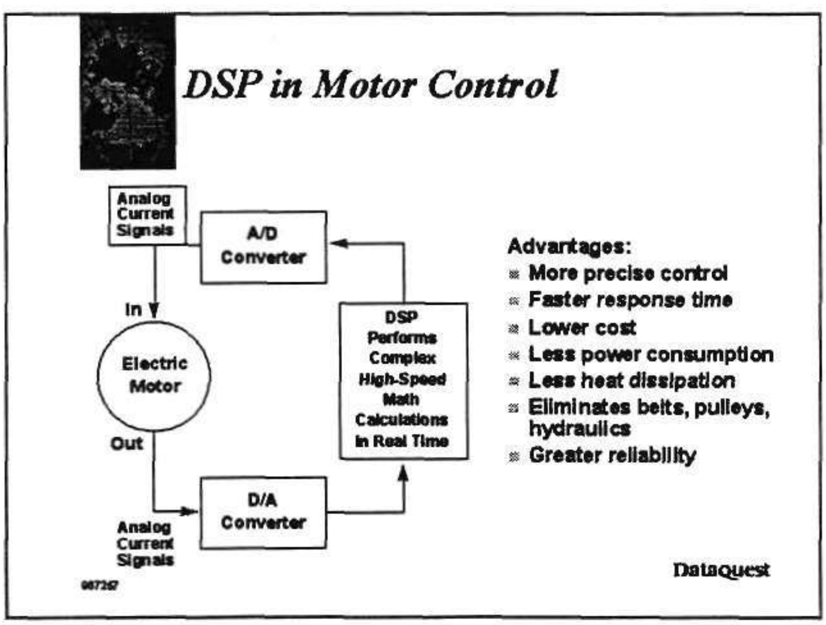

Here is a simplified block diagram of a DSP controlled motor. Real world analog current signals are digitized and processed by the DSP. Advantages include more precise control, faster response time, the use of lower cost motors, less power consumption, less heat dissipation, and elimination of belts, pulleys and hydraulics. DSP motor controlled systems will find their way into automobiles, heating, ventilation, air conditioning systems, blowers, compressors, and other industrial uses. Office products like printers and copiers are other applications. Home applications will include refrigerators and clothes washers.

Systems that use DSP motor control can reduce power consumption by as much as 40%. When you consider that over 10 billion motors are shipped every year and motors account for a huge fraction of the earth's energy budget, then such savings will make a huge impact on the overall energy requirements of the population. MCUs are and still will be used for motor control, but they are running out of steam as they cannot keep up with the increasing mathematical demands required by advanced motor control.

Clothes washers with direct drive motors controlled by DSPs are already on the market. In the Bay area, PGE will give you a rebate if you buy one of these machines, but they are horrendously expensive, even though we say the motors will be cheaper. Not only is power consumption reduced, but these washers run quieter, use less detergent, and can ramp up more smoothly to fast spin cycles, reducing instances of imbalanced distribution of clothing.

What is driving all these diverse applications? Collectively, they are all being driven by the three main characteristics of DSPs. In the past DSP applications were limited because low cost and low powered DSPs were not available and what was available was expensive. When DSPs were first introduced, back in the early '80s, they cost about $200 per MIPS. Today, the going price is about a buck per

MIP. Thus the availability of low power, low cost, and low cost per MIPS DSPs are causing the explosion in the DSP market. Each of these characteristics are creating new markets.

Who Are the Programmable DSP Players? (Millions of Dollars)

		1995	Percentage Change	Market Share (%)
1.	Texas Instruments	635	69	38.0
2.	AT&T	510	86	30.6
3.	Motorola	178	18	10.7
4.	Analog Devices	117	37	7.0
5.	Zilog	59	NA	3.5
6.	NEC	49	22	2.9
7.	Fujitsu	46	28	2.8
8.	Toshiba	28	75	1.7
9.	Hitachi	13	47	0.8
10.	IBM	10	NA	0.6
	Others	25	38	1.5
	Total Market	1,670	62	100.0

067261 Dataquest

This chart, courtesy of TI, shows our DSP product portfolio. Unlike microprocessors which are generally high performance, high cost products, DSP manufacturers are evolving their product lines to meet a very wide spectrum of applications ranging from low cost, low performance, products costing a couple of bucks, to high cost, relatively high cost, high performance products such as the TI products here and Shark, the analog device, which is about $200. It is this spectrum of products that are needed to meet the diverse applications that we have talked about so far.

Here are the top ten suppliers of programmable DSP chips. The market is dominated by only a few players. Heading the list is TI with 38% market share, followed by Lucent with 31% share, and Motorola and Analog Devices. Together these four companies account for more than 86% of the market and all are U.S. companies. I want to emphasize this is programmable DSP, it does not include fixed function DSP chips. TI, earlier this year, announced that they would be making a 1996 capital investment of about $2.3 billion and 90% of this investment would be for DSP capacity. Clearly TI intends on keeping or expanding their #1 position. Another point is the relatively low growth of Motorola compared to the other three top players.

Programmable DSP Application Segments

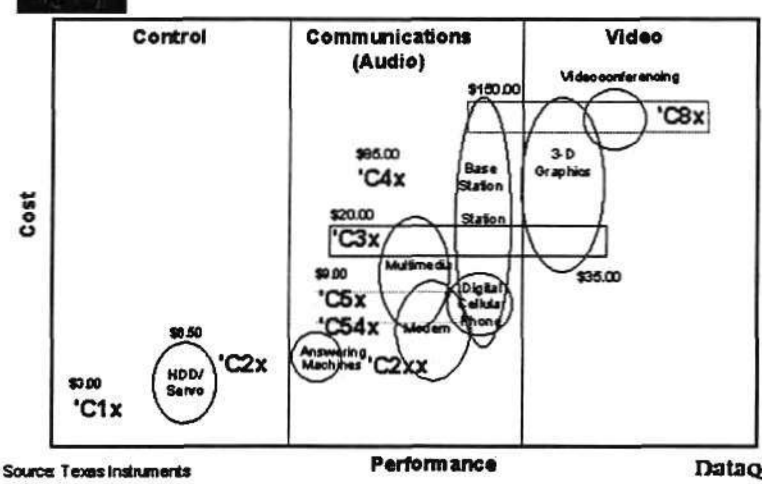

Source: Texas Instruments
067260 Dataquest

The 1995 market grew 62% to reach more than 1.6 billion. What does the future look like for programmable DSP? Here is the forecast. We expect a programmable DSP market to reach nearly $7 billion by the year 2000. This represents a compound annual growth rate of about 32% over the forecast period. For 1996, we expect a market to grow 33% and to reach over $2 billion. The 1996 growth of 33% is less than the 62% of 1995, but this does not indicate a slowdown of growth. It

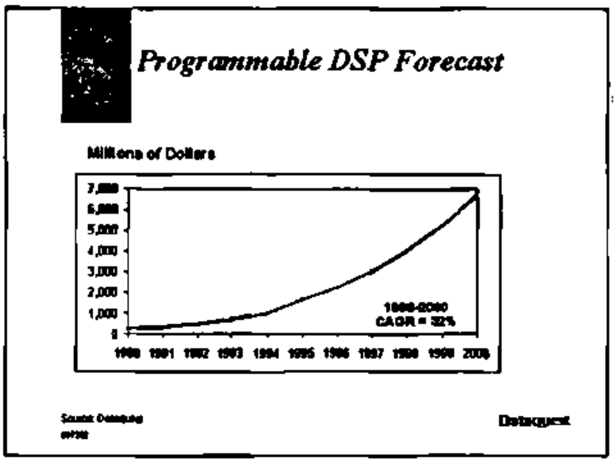

is the law of large numbers beginning to apply. When coming from a small base like DSP has been doing, you cannot continue with these high growth rates.

This chart shows the microprocessor market for a 10-year period when MPUs were just starting to ramp up. This is the period of 1982 to 1992. DSPs are just now beginning their growth. The chart also shows a DSP market for a 10-year period, starting in 1990. Note the striking similarity between the early growth years of both MPUs and DSPs. The compound annual growth rates for the 10-year growth periods are almost identical or at least forecast to be identical for both products, whereas MPUs literally exploded off this chart

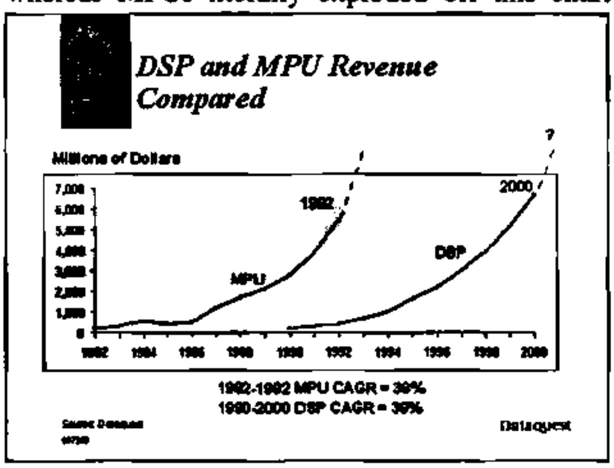

after 1992. What would DSPs do after the year 2000, when its initial growth phase is over? Would DSPs show a like explosion? As someone once said, "If Intel made DSPs there would be an Intel DSP on every motherboard." There probably will be anyway, but it will not be provided by Intel.

DSP versus MCU

	1995 DSP Revenue	1995 MCU Revenue	1995 MCU Ranking
1. Texas Instruments	635	189	17
2. AT&T	610		
3. Motorola	176	1,838	1
4. Analog Devices	117		
5. Zilog	69	71	20
6. NEC	49	1,554	2

Dataquest

Here are the MCU market leaders compared with the DSP market leaders. First, the DSP market is considerably smaller. Second, the market leaders in DSP and MCU are different, no company is a strong player in both categories, with the possible exception of Motorola. However, Motorola's growth last year was only 18% and had lower growth than the other DSP players. Was this because Motorola was making all of their investments in MCU capacity? NEC is the number two company, and MCUs are not a major player in the DSP market. Most of the market applications for DSPs and MCUs are not overlapping. There are some applications that already do overlap – hard disk drives and motor control. It will be very interesting to see how the DSP and MCU markets evolve in the coming years and how they may begin to significantly affect one another.

Java versus Intel: And the Winners Are ...

Raj Parekh

Chief Technology Officer, Vice President and General Manager,
Embedded Products Group,
Sun Microelectronics

Mr. Parekh, 43, brings more than 20 years of engineering and management experience to Sun Microelectronics. As chief technology officer, he oversees the division's research programs, international R&D, architectural innovations, and Internet-enabling components to support Java and other network technologies. As vice president and general manager of the division's Embedded Products Group, Mr. Parekh will oversee the proliferation of 32-bit SPARC and Java processors into the emerging embedded network marketplace.

Most recently, Mr. Parekh was vice president of engineering and chief technology officer for Sun Microsystems Computer Corporation where he was responsible for the computer system strategy, technology, and international design and development. Before that, he 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's degree in electrical engineering from the Polytechnic Institute of New York. He received a B.E./B.S. in electrical engineering from the L.D. College of Engineering, India. He also holds U.S. patents for a bias control circuit for a substrate bias generator, EPROM reliability test circuit, and a segmented channel field-effect transistor.

Session # 14: Java versus Intel: And the Winners are ...

Raj Parekh
Chief Technology Officer, Vice President and General Manager, Embedded Products Group, Sun Microelectronics

JOSEPH: Our first speaker is Raj Parekh. Raj is a Chief Technology Officer, Vice President and General Manager, Embedded Products Group, Sun Microelectronics. Mr. Parekh brings more than 20 years of engineering in management experience to Sun Microelectronics. As Chief Technology Officer he oversees the divisions research programs, international R&D, architectural innovations and Internet enabling components to support Java and other network technologies. As Vice President and General Manager of the division's Embedded Products Group, Mr. Parekh will oversee the proliferation of 32 bit SPARC and Java processors into the emerging embedded network marketplace.

RAJ PAREKH: Every time a new market emerges, a new paradigm shift takes place, it creates some opportunity for new winners and new losers. If it is not taking place it is very hard to take the incumbent out. When the transistor was invented, the computer became reality. Without the tuning capacitor the radios and TVs were not reality, but as soon as that happened it became reality. At this point in time with Internet, World Wide Web and the emergence of Java it creates the paradigm shift. As it creates the paradigm shift it creates the new opportunity.

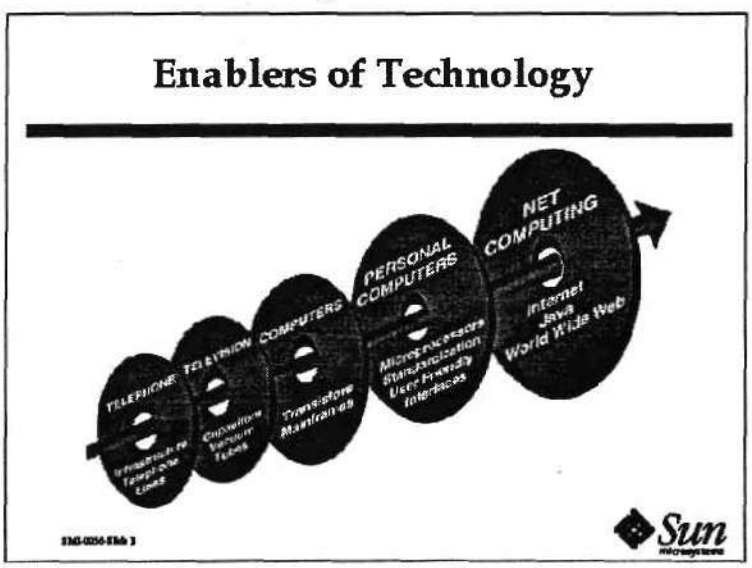

One of the problems of this presentation is, I have tremendous respect for Intel. The title itself of my presentation says I must say otherwise. So I will be very careful about not hurting Intel because I do have a lot respect for them. With this new change of paradigm, let's look at Intel's view of the world. It is extremely straightforward, it is very easy to understand. They are going after the PC, the desktop, and the laptop. That was very good for them, they did an excellent job in standardizing as well as creating a very high gross margin business for themselves. As far as Intel is concerned it was an excellent model for them.

However, as we look at Java's view of the world, it is much broader. Of course we are worried about the desktop, we are looking at the Net computer, but the overall breadth Java can provide is far beyond just the desktop. It is into the networking world, the switches, hubs and routers, the communication world, the telephones, the entertainment industry, it is in the automobile and it is in the consumer. All of those industries can take advantage of Java and throughout my presentation I will be going into more detail about it.

With this size and shape of the world, what kind of a potential market does Java exhibit? The emergence of the market is measured in millions and millions of units, far larger than Intel, or any CPU, MPU, individual DSP, or any other product can really exhibit. Because of the wideness of the market, from the pagers to the games to the set top boxes to the net computer and enterprise and control side, it becomes an extremely interesting and dynamic market.

What are the key success factors? How can we really penetrate such a market? We think that one of the key aspects is the universality of the application. Universal applications run on any computer. It is highly modularized, so you can actually merge multiple different applications together. Telephoning is a well understood market and PCs are also a well understood market.

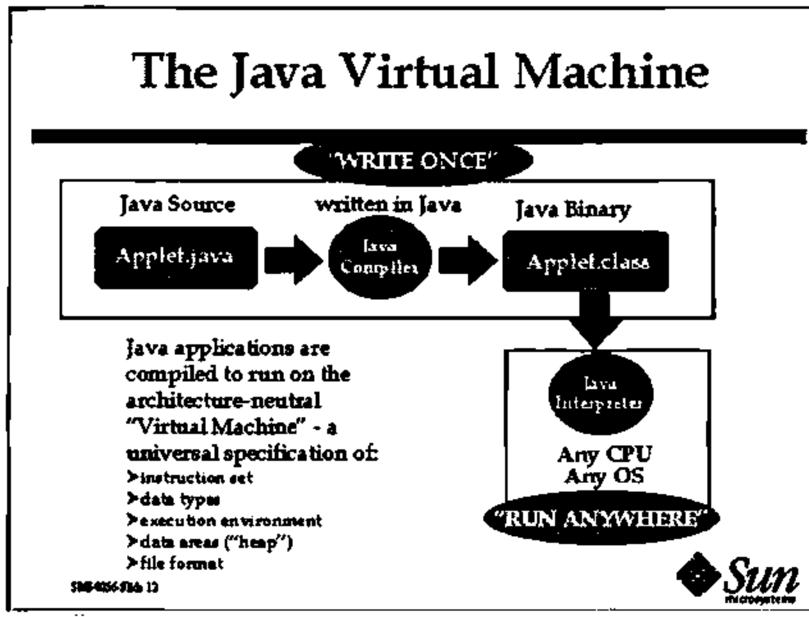

That is the combination of the different technologies and Java has become an excellent way to provide that combination. The issue is how do we take multiple applications, integrate them together, and make it available from multiple different companies because that is the only way we are going to get a tremendous amount of innovation on one side and very effective price and cost control on the other side.

How can we remain evolving and not get stagnant with a single architecture or a single way of doing it? How do I take it advantage of it? These are the success factors for Java. We are hitting all of them and that gives us much more confidence that it will become extremely pervasive into the market.

Java's way of doing business is rather interesting. We are not going for single source, we are going for multiple source in almost every geography of the world. There will be some sources who can produce this product. They can produce the product in the area of expertise they have, rather than making just one generic part and trying to compete with each other. It allows people to do the innovation and make their products available all over the place.

Because of the nature and modularity the chip has achieved you can actually operate at much lower frequency and optimize power quite significantly so that if it is in your pocket the batteries will last literally forever. In case of a desktop where you demand tremendous performance, it can really scale up extremely easily because all those thoughts had gone into it when we created the architecture in the first place. In the business model the multiple source creates a rather interesting way to deal with it. It also creates the modularity. That certain segment of the market can be handled by certain key suppliers. Other segments of the market can be han-

picoJava I core vs. x86

Cost	multiple sources	single source
Power	power management unit flexible for vendor	fixed power
Business model	modular	fixed design
Future potential	ability to innovate	fixed architecture

Sun

dled by somebody else.

The picoJava we have designed, which is a core, will be available for licensing. Sun will create the product based on that, but many other people will also create the base on this. The span of this product is from the simplest pager to the most complicated Net computer and everything in between in multiple different dimensions. The reason we were able to do it this way is this is not a general purpose processor. Think of it as a DSP for Java. This processor executes Java netively, all the instructions of Java netively and we took care of the legacy core issue by adding the ability to execute C in this processor.

The primary design cycle and the primary dimension for this product is the world's best execution of Java. In this particular core we are talking about 60,000 gates. We do not measure ourselves in how many millions and millions of transistors there are. We do not take pride in making it bigger or more complicated, we take real pride in making it the simplest, easiest way to use with flexibility so that people can scale it up, down, run different applications, and integrate other functions with it. That is the only way to create a minimum cost product. PicoJava has achieved that particular goal for us.

All of these products are in design. This is a very impressive list and today they are not all designed by Sun. They are designed by many of our partners. These products will start rolling out in the calendar year 1997. Some of them may come out even before Sun can have a product and that is perfectly acceptable to us. They can create next generation and the generation after without help from Sun. That is also acceptable to us. While we continue to provide the technology, they can also continue to innovate.

If one company is falling behind, another company will be waiting to take over that particular marketplace from them. It goes beyond the Net computer, the Java Processor will touch your life, in one way, shape, or form, in the near future. Some of the partners we have include NEC, Samsung, LG Semiconductor, Mitsubishi. NorTel has already announced their intention to make the chip, use the chip, and create products based on this. As you can see these are not small companies, they are all industry giants. They are extremely serious about this product all the way up to their chairman level, everybody understands what they are doing and why they are doing it. Many other companies in the near future, in the next six months, I believe, will announce what they are going to do as well.

Nothing happens unless we have a very significant performance advantage. There are two ways to upgrade the interpreter as well as the just-in-time compiler and this is a different trade-off. One makes it more general purpose but has a slower interpreter. The just-in-time compiler is somewhat faster, but it takes more memory and it is very specific to an individual processor. With picoJava running at the same MHz clock rate and with 60,000 gates, we should be able to run a lot faster. Nevertheless, at the same clock rate, we are experiencing, at this point, a five times advantage over Pentium, running the fastest thing known called 'Jit Compiler.'

Many of the species are coming from the language, that it is platform independent, it is dynamic, it is highly secure and object oriented. That is why many companies are extremely interested in Java as a language. Under the Java language there is a server at work where Write 1 takes place, there is a climb side to it, where being able to run anywhere comes into play. The communication between the server and the client is very fixed, it is very dedicated. By virtue of that, what happens is that anybody can replace the CPU operating system and not break the application. That is an extremely powerful piece. What we have done is the communication of Java applet class or Java byte code, which is highly compressed, we take that directly into picoJava and execute. Period. Nothing is needed in between.

That is why it creates the lowest cost solution, not only at the chip level but also at the system level. The memory requirements and other requirements start to go down. With this power of the language, what it has created is a rather interesting momentum in the field. I do not know whether I should call Netscape an operating system or not, but look at the other numbers. The number of programmers working on Java already exceeds 200,000. At its peak, Windows programmers were noted at close to 400,000. Java is already halfway there and growing very rapidly. Fifty-seven colleges and universities are offering Java classes and claiming that most of them are so full that students are on waiting lists. One hundred and fifty books were published in one year, so somebody is reading.

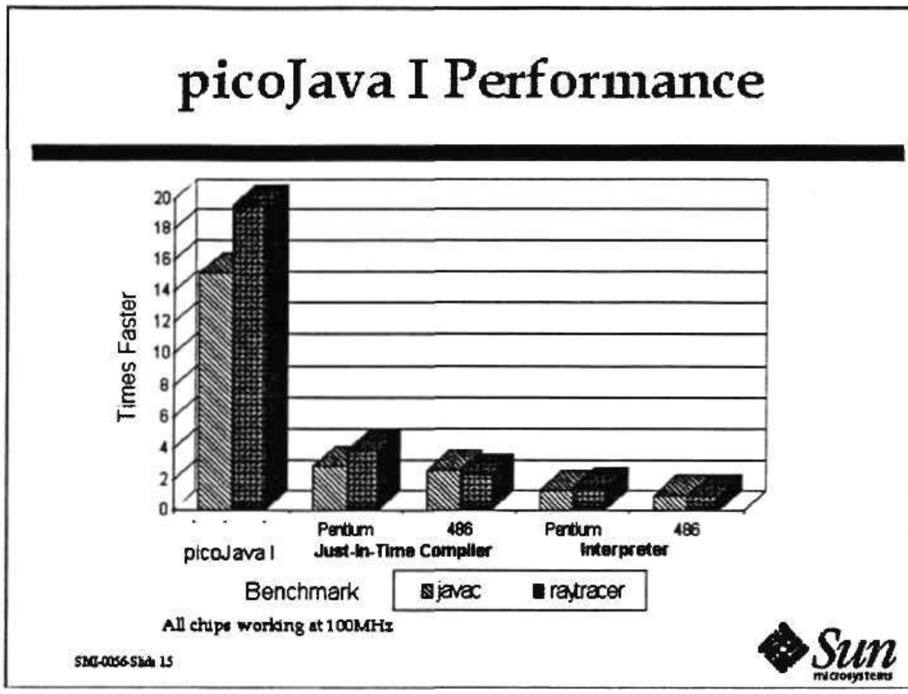

This is rather amazing, 80,000 web pages running Java already. That is the power of the language. We take that power and say we are going to do mundane things, we're going to just focus on a chip technology, which actually helps this Java language by delivering much higher performance, much better costs, 60,000 gates versus millions of transistors. You can give that to any foundry and they will tell you that it makes a huge difference – the ability to control power and the size of the chip as well as change the configurability so that it can adapt to different markets easily.

Who are the winners? Intel or Java? I don't know, but I believe that the emergence of Java is breaking the paradigm that you have to have one operating system, or one CPU architecture. That particular phase has changed and with that change, I believe consumers are the true winners of this particular paradigm.

The ability to seamlessly integrate different technologies and different applications, of evolving particular universal applications, has become an extremely powerful way to deal with it. On top of that we have created a business model which makes a huge difference for us.

You have to ask yourself a very simple question. Is Internet important to you? If so, can you live without Java? Once you go to Java can you live with low performance? Or do you have an infinite amount of budget? If you are on a budget, if it makes sense to reduce the cost and increase the performance, then a Java-based product becomes a very natural selection not only from Sun, but from many other people as well.

Q: How does Sun make money from Java and picoJava?

A: We will sell licenses, chips, and operating systems and each one of them has a profit associated with it. However we also sell systems and servers and every time there is a thin client a fat server is needed. We have a large share of fat servers, but we really believe that if the market truly opens up there will be many winners, including Sun, and that strong belief is driving us and we believe that we will be fine as a company.

Q: What is the danger that Microsoft will seize control of Java?

A: The issue is that there is something called a compliance test suite that they license. We have a compliance suite, which is available openly on the Net and every application developer writes to the compliance suite which is a common API common protocol. One of the conditions of the license is that they must also remain compliant to that open standard. By virtue of that, we think that no one company, including Sun, will have total control over Java.

Q: Will you be licensing the picoJava core to independent semiconductor manufacturers?

A: At this point in time we are licensing picoJava to the people who are also creating systems solutions. We want to make sure that the initial round of picoJava licensees are indeed the companies who have both system and semiconductor operations, and they will use it in real applications. As time goes on we will also make it available in a selective way to the right semiconductor companies.

Q: Does your solution for making money in Internet reinforce the concept that only hardware will make money?

A: We think that the solution makes money, not hardware, not software, not system integration, not manufacturing, no distribution channel individually makes money. How we attack the total solution for the customer is where the difference is and our focus is exactly on that. That is also reflected in our licensing, that we are looking for partners to also look at all the aspects and deliver excellent solutions.

Q: Sun claims that Java is architecture independent yet you say that it runs best on your own, so what is going on?

A: Java truly is CPU independent. The Java byte code is universally defined and always consistent. Any CPU, any operating system can take it, through the interpreter or a just-in-time compiler, can word into the native operating system, operating commands for the CPU and execute that. The difference with picoJava is that our natural instruction set of the chip itself is the same, or very similar, to the Java byte code so we do not need a translator in between, that is the difference.

Our SPARC product which is made by Sun will need an interpreter, as will the X86 CPU, but picoJava is the only product which does not require an interpreter in between which is why the performance is in a different matrix. Everyone is comparing their CPU through an interpreter or just-in-time compiler and saying "I am 10% faster or 20% slower." What we are talking about is several times difference as long as you are running Java application because we actually execute it in the hardware itself.

Q: Who will ship the first Java processor, what will the price be and what will be the first application?

A: I would like to have one of the other companies who is our licensee answer this question. I would say that typically a smaller dedicated application will be the first one in this arena whether it is PDA or a set top box, but I do not know exactly. In terms of price, 60,000 gate devices are less than $10.00, so the cost is fairly low but typically people will put the complete application on a chip rather than just have a Java core. We are looking at a solution on a chip and the price varies based on the solution. In my opinion, the price will be $10 to $20.

Q: Is Intel a Java user? Or can they be?

A: Our policy is not to disclose any user unless they prefer to disclose, so I cannot say whether Intel is user or not. Can they be? Of course they can. It will augment their situation, they have a very good product in certain areas and I am sure they would like to expand in other areas and Java would be a very interesting vehicle for them.

The New Paradigm for Online Services

Mark Walsh

Senior Vice President,
AOL Enterprises,
America Online Enterprises

Mr. Walsh is senior vice president of America Online Enterprises and a corporate officer of America Online Inc. He is responsible for developing and executing the company's entry into business-to-business services.

Mr. Walsh has been actively involved in information, database, and entertainment services throughout his career. Before joining AOL, he was president of Genie, the consumer online service owned by GE. He was also president of Information Kinetics Inc., an interactive information and database publishing company focusing on the job and employment marketplace. IKI was funded by major U.S. venture capital firms. Before IKI, Mr. Walsh was the vice president and general manager of Interactive Services for CUC International Inc. He was in charge of all aspects of CUC's interactive information businesses, which included shopping, travel, automotive, dining products, and membership services. He originated CUC's involvement in virtually every online, interactive TV, and screenphone platform in the United States. Mr. Walsh has also served as a managing director of a direct-marketing firm specializing in cable television; was head of telemarketing and special projects for Home Box Office; and, directly out of college, was an on-air news anchor and news director for a CBS Television Affiliate.

Mr. Walsh served as the chairman of the Interactive Services Association, a 330-member company trade association focusing on the consumer market for interactive and online services. He has been a frequent speaker at information industry, direct marketing, cable television, and newspaper publishing conferences and conventions, and has been widely quoted in a variety of business, trade and industry publications. He is on the Editorial Board of *Online Access Magazine* and *Interactive Television Report*.

Mr. Walsh is a graduate of the Harvard Business School, where he was editor in chief of *HarBus News* (school newspaper), and Union College in Schenectady, New York. He is a former member of the board of the Information Industry Association and a recipient of its 1992 Leadership Award. He is a former board member of the New York Theater Workshop, an "Obie" award-winning theater group, based in Manhattan.

Session 15: The New Paradigm for On-Line Services

Mark Walsh
Senior Vice President, AOL Enterprises, America OnLine Enterprises

Moderator: Mr. Walsh is the Senior Vice President and Corporate Officer of America OnLine. He is responsible for developing and executing the company's entry into business-to-business services. Mr. Walsh has been actively involved in information, database, and entertainment services throughout his career. Before joining American OnLine, he was president of Genie, the consumer on-line service owned by General Electric. He was also president of Information Kinetics, an interactive information and database publishing company focusing on the job and employment marketplace.

Mr. Walsh served as the Chairman of the Interactive Services Association, a 330-member trade organization focusing on the consumer market for interactive and on-line service. He is also on the editorial board of *Online Access Magazine* and *Interactive Television Report*.

MARK WALSH: We are a company that has enjoyed explosive growth in the last two or three years, and because of that growth, a lot of industry analysts, fund managers, and other large investors from Wall Street and the global investment community, have made AOL stocks a bellwhether holding. Those stock holders were very excited at $70 a share which was our price in June, and are less excited at $25 a share, which is our price today. It is still about a thirty to one multiple over our IPO price in '92, but, as is true here in the Los Angeles, Southern California area, you are only as good as your last movie. In the investment community you are only as good as your last stock price

I am the business-to-business guy at America OnLine and the cynic would say that is like being the electric car guy at GM, in that GM knows it is going to build electric cars someday, but it still likes building gas guzzling cars. AOL is so far, the world's most successful consumer orientated on-line service. Because of that, I sometimes find myself at meetings within the company, selling a little bit harder and describing the future a little bit more dramatically that I would in the outside world where the obvious is more intuitive. What is obvious and intuitive, to a lot of the world right now is that a platform called America OnLine has penetrated many homes and desktops and in fact is being used by a significant portion of our members for business purposes.

I was asked to talk about a new paradigm is emerging. I felt that perhaps AOL was the best model to spend a few moments on because AOL's migration, and experience by having the Internet impact the whole marketplace is a very good example of the experience that the marketplace is going through right now.

J. Robert Oppenheimer was the chief scientist of the Manhattan project, the group of scientists in World War II that invented the atomic bomb. After six months of work in Los Alamos, Oppenheimer and his team had successfully created the first atomic bomb, and had a prototype called 'Fat Boy' which they were going to detonate in Alamogordo in a test. A bus circulated the Los Alamos facilities gathering all the scientists together to take them out to Alamogordo for that initially detonation at 6:00 am on that fateful morning. Oppenheimer noticed that the mathematician, Dr. Brock Johnson, was not on the bus, and he said, "We have to have everybody," and he went to the bungalow, knocked on the door, and when Dr. Johnson came to the door in his bathrobe, Oppenheimer said, "Come on, you've got to come with us. We've all worked so hard for six months. This is an amazing creation that we've successfully completed. It's a whole new age. We're going to change the world. Come on out with us!" The mathematician said, "No, you don't understand. If I got the math wrong, none of you is coming back."

We got the math wrong in this business, all of us did. I know we got the math wrong, on the low end, at how successful this would be, how fast. Nobody thought it would be this big, by now. My chairman, always felt it would be a large consumer medium with millions of households, but I think even he, in his most private moments, admits that this hockey stick curve growth of membership and usage is amazing, stunning and unexpected.

Now we have the math on the high side on how much commerce, transactions, shopping, stock buying, EDI, and all that stuff would happen, how rapidly. It was always the business of the future, and always would be, was the old saw about the interactive commerce business. We kept expecting that interactive shopping, both on the corporate and the personal environment, would take off but it really has not. It certainly has not in the calendar we expected. It seems to be taking off now, but that delay is how we got the math wrong.

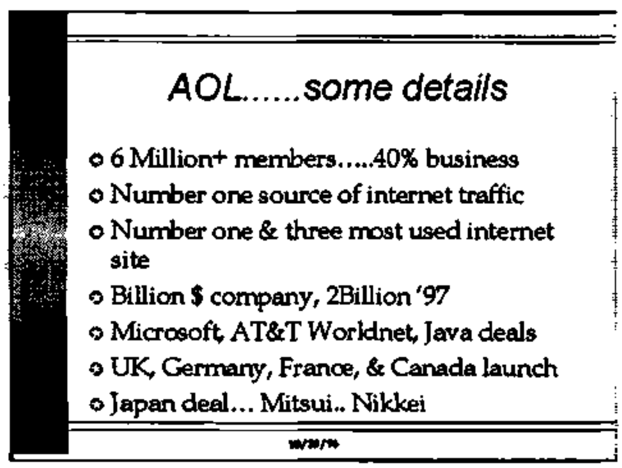

Some quick details about my company, we have almost six and a half million members today. Forty percent of our usage is business related. We are the number one source of Internet traffic today, globally. We own the number one and three most used Internet sites. That is an arguable statement that Netscape would contend is wrong, and they are certainly welcome to contend that. We are a billion dollar company, just finishing our fiscal year July 1 at a billion dollars and our run rate for '97 will be two billion. We have successfully completed some interesting technology deals with Microsoft, AT&T, Worldnet, Java, and a number of organizations. We now are a global company with AOL services, in the UK, Germany, France, and Canada. We have linked a deal and announced it four or five months ago, to launch AOL Japan. Our partners there in a joint venture are Mitsui and Nikkei.

Internet access on AOL, over thirty percent of our gross usage is Internet. In fact another ten percent or so comes through hybrid forms that are actually accessing websites, but it looks like you are still on America OnLine.

How many people here belong to America OnLine today? Can I just see short show of hands? Maybe 10%. How many people have ever belonged to America OnLine? About the same. How many people here have not received an AOL disk in the last thirty days? I am always interested in this statement, I ask people how many disks they think AOL mailed in the last nine months or so. Typically people say five, fifteen, thirty, forty-five million disks. That is the typical range, from forty-five to fifty million to maybe fifty five million, and the answer is a hundred million disks in the last nine months.

We currently offer HTML authoring capabilities for our members. The new version of our AOL client uses Internet Explorer IE, for both Apple and Mac. We also offer Netscape's browser runs over our link. The Winsock compliance is now complete. We get about 70 million hits a day on our home page. We do over 7 million pieces of Internet mail. That is mail out from our service daily.

We are today a feeder, a supplier, a supporter, an access provider, a standards creator and enhancer for the Internet. We are not a foe; we are a friend of the Internet. I think what you find, from the most professional and savvy media observer or analysts down to the least observant or least aware citizen in the U.S. or the globe, that there often is this positioning of AOL as a closed proprietary service, and the Internet is something separate. In fact, we are the Internet to a lot of people. But the Internet has forced us to change our paradigm a lot.

The Internet is a pipe. It is not a place today. It is an access. It is not a destination.

There is today, no there, there, as Gertrude Stein once said. There is no thing called the Internet that you buy. There are no brands, no organization, no editorial, no sifting today, and I keep saying today because it is happening faster than even we expected. We got the math wrong again.

Many of you who are from the East Coast or from major cities might know the restaurant guide called Zagat's Restaurant Guide. It is in about twenty major cities. I am an old friend of Tim Zagat, who is either the luckiest or the smartest man on the face of the earth. I do not think he is both, but what he figured out in New York in the early '80s was that instead of having a book that managed the 35,000 restaurants in New York and Manhattan specifically, there is probably about 1,000 restaurants that really mattered. He was a wealthy lawyer, and he got all his rich friends to go to the various restaurants and gave them a form after they went to rate it. He brought all the forms together, made a little book, and then sold it back to his wealthy friends. One of the greatest closed loops of information generation and purchasing that I have ever seen.

Now he has a lovely business because in fact this is the self-referential element of rating. More importantly, nobody cares about 35,000 restaurants. They only care about the thousand or so that matter, and in fact, they really only care about 40 to 50 restaurants in Manhattan, where the majority of the impact, and frankly the majority of the revenue, as far as the high end is concerned, is generated. But there is no "Zagatization," there is no segmentation between various sites and the web yet and people buy brand names. The Internet is a vacuum of brands and nature abhors a vacuum.

The greatest example of people buying brands, that I ever personally experienced, was in southern California. I used to work for a company called CUC International, which, among other things, had an on-line auto buying club. On that club you could log on, and get a quote for the value of your used car. To get the data of used car valuation, I went to a used car evaluation company here in southern California. It is called Kelly Blue Book. I mean no disrespect if a relative of Mr. Kelly is in the crowd. It is a great business, but I met with old Mr. Kelly who was the son of the old, old Mr. Kelly, that started the Kelly Blue Book. Many of you in the audience will know Kelly Blue Book. It is the book the dealer uses to say your trade-in is worth $7,000, or the purchaser of a used car might say its worth $9,000 book value.

I visited with old Mr. Kelly all day. We did our deals so you could get a used quote on-line, and then at the end of the day, I sat down with him.

The entire operation that this man ran, was equal to one of the sectors of this room, bunch of cubicles, people typing, a few computers. I said, "You know, you have a fantastic business. You put out this book that values Miami cars versus Minnesota cars, that goes to all the auctions, gathers all this data of price points of cars by year, by model. It forms a giant database, does regression analysis of the valuation, puts out the mean valuation for good and bad condition. Incredible. All this data, putting out the book, where, how do you do this business?" I was saying, "where" do you do this business?

Old Mr. Kelly looked over his shoulder conspiratorially and said, "You know Mark, we just make it up." Completely straight-faced. I said, "Well what if you are wrong?" He said, "Well we are never wrong, but if we are, we just change the number and then we are right again." I said, "What is the difference between a Minnesota and a Miami car?" He said, "Ah, we knock fifty bucks off the Minnesota car." I said, "That is a great business." You make it up. Then you sell it. The essence there is the brand. The brand is what matters. You not buying what old Mr. Kelly made up, but you are buying the brand name, the referential validity of what that brand does.

You have all heard this joke. I will go through it quickly. The two campers, bear coming at them, one drops his pack, starts putting on his running shoes. The other camper says, "You can't outrun the bear." The other camper says, "I don't have to outrun the bear, I just have to outrun you."

The reason I suggest this is that there are few paradigm shifts occurring in the Internet service space today. There are lots of people running with, ahead, or against the Internet. If the Internet is a bear, we think there are a lot of campers out there, because the Internet is consuming so many people today, people that wear ties, people that do not, people that are retired, people that work, kids. It is a tsunami.

What is next in this business? We think that our brand will become larger and larger. We expect to have 10 million members by the mid to end of 1997. Some of our competition will stay with us or perhaps grow as big or larger than we will. We think that the connectivity side of the on-line business and the content side will become more and more separate, much as it did in the cable, broadcast television, and other media businesses, even the movie business. We think that Internet weaknesses will become more and more glaring.

We think that the Internet's going to break. The perception of the consumer is that the Internet is slowing down. Whether it is true or not becomes immaterial.

My example of perception versus reality is Andy Grove and the Intel Pentium chip. It had a calculation flaw that occurred once every few billion times. This audience is probably stupendously sensitized to its flaw, but the world out there did not care. Yet on the Internet, the awareness, the existence, the validity, and the danger of this flaw became a cause célèbre. Suddenly all of us were worried about our Pentiums failing as if this were going to truly change the way we played F-14 Strike Commander on our computer. Now, I knew this became an issue larger than Andy Grove could ever contain, when I opened the New York Times one morning to the sports section and the headline of the article said "Nets (the New Jersey Nets) -- Nets Play Like Pentiums, Still Win." I knew he had a real problem because the word "Pentium" was becoming synonymous with "crappy." Perception, reality and the Internet's perception now of its flaws becoming more glaring, is starting to become truth.

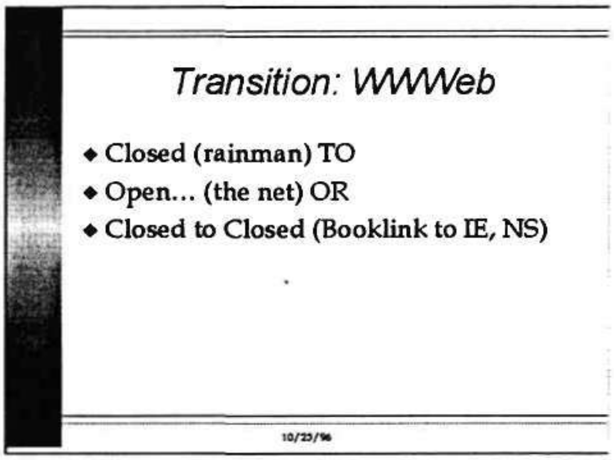

My company is a company in a series of transitions today: transitions that are both painful, exciting, expensive, wasteful, exhilarating, and many more adjectives I might suggest. We are transitioning to the world wide web and to becoming the world's largest Internet access provider. From a closed set of production tools called Rainman, that make the screens you see on America OnLine, we have migrated to an open system of tools and sites we point to and aggregate for our members called The Net. In fact, we went from a closed to a closed system. We bought a browser company called Book Link that was the original browser we used in AOL about a year and a half ago. Now we buy and use browsers from Microsoft Internet Explorer and Netscape.

I would contend they are as closed as the old closed proprietary on-line systems. Open is good, yet the battle is to own the proprietary code for browsers. I think the war going on is not about being open, it is about being closed. A transition in our revenue model is from an hourly, metered service where you pay a flat fee to get in, and then after you use up your free hours you are on the meter. We were motivated to keep you on-line for a long time. Now, since flat-rate pricing has clearly taken hold in the consumer, and frankly

Transition: Other revenue

- No advertising... TO
- Lots of advertising
- Not to much Commerce.... TO
- Lots of commerce
- No Daytime... TO
- Lots of daytime

10/23/96

the business marketplace, you will see AOL become far less metered, far more oriented towards flat-rate pricing, become more oriented to surcharging access for certain parts of its system.

Other revenue AOL used to have no advertising on it, now you are going to see lots of advertising. AOL did not used to have much commerce but you are going to see a much higher stress of buying things on AOL – trading stocks, buying roses, buying records, buying business information, buying cars, buying insurance, buying travel, a wide variety of parody products.

We did not used to have much daytime use, in fact my whole team has been hired to do nothing but drive daytime use, by building private AOLs for companies and corporate users.

Our connectivity – we used to have outside suppliers, Sprint, BBN and others, who were our networkers. Forty percent of our revenue dollars went to outside connectivity companies. We had to buy our own company, ANS, to get into our own hosting business, to control our cost structure. The marketplace forced us to.

Transition: Content

- Reseller... no real VAR... no equity TO
- Owner, w/equity, joint ventures on the web, etc.

10/23/96

A transition in marketing – probably many of you have seen on television, our new many millions of dollars branding campaign. We are much more involved in sponsorship, in more mainstream media, to brand our company's products, not to give you a sampling of them.

A transition in geography – we used to be simply U.S.-centric, look at the name, America OnLine. We have launched global joint ventures where we have kept the three initials AOL (just like AT&T used to mean American Telephone and Telegraph), but people knew it globally as AT&T. We think that will occur with our brand as well.

A transition in content – we used to be a reseller. We took news feeds or other sports news and we added no real value. We had no equity in this content. We were just a repackager, a recommunicater. Now through our joint ventures, through our venture capital fund, called Greenhouse, we are owners and equity owners in a wide variety of content plays.

A transition in audiences: we used to be solely consumer. We were the cool alternative to CompuServe and Protégée. We were very much nighttime oriented; in fact, the joke was that loneliness and boredom were our best salespeople. To the new transition of a balance of consumer and business to business orientation of content and users, we are simple, we are reliable, we are integrated, we are cheap, we are multi-platform, and that is becoming more and more appealing to businesses. In fact, we think that simple wins in the business marketplace.

Just because something is hard, does not mean it is good. Just because it is easy, does not mean it is childish. Lotus Notes is like every button on your VCR except play and rewind. You know they do a lot, you can never figure them out. I routinely go into executives offices, at a wide variety of levels, in a wide variety of companies. I walk up to their telephone, and I say to them, "I dare you to tell me what more than four buttons on this phone do," other than the digits, obviously. "I dare you to explain what any of these buttons do. I will give you four but on the fifth, I bet you cannot do it," and I have never been beaten. Nobody ever learns their phones, nobody ever learns and never uses staggering amounts of the robustness of what we buy.

Productivity is still elusive and difficult to prove in many of these products. Most importantly the business marketplace is adamant about not being Beta-maxed again. They are adamant about not investing in platforms, in hardware and connectivity, in content, in services, in packaging, where you will tell them in a few years, "Oops! Got to upgrade. Sorry." It is not going to happen. That demand for flexibility, we think, puts my company and some others, in some interesting places. AOL is in a marketplace that is morphing daily. Our partners are morphing daily.

We have an interesting dance of love and hate in every meeting we go to where we are a partner, a supplier, a customer, a competitor, a litigant, all these roles in one room. The same is true when we visit Disney, or Viacom, or Netscape, or Lotus, or IBM. All of us are very wary of each other. It is an interesting time. There is no proven model. However, we know that this business, be it business-oriented or consumer-oriented must be a religious conviction.

If your company considers on-line or interactive services a marginal play to protect some franchise you already have, or to extend and protect something you already do, you will fail. Dow Jones had one of the first, largest, on-line services called 'Dow Jones News Retrieval,' now it is marginalized. Reader's Digest owned something called 'The Source' at one point in the eighties, the second largest on-line service in America. It is dead, bought by CompuServe. Genie, owned by General Electric, at one point the second largest on-line service in the world, with tremendous consumer reach and business reach with its corporate parent, is sold and dead. Sears and IBM owned Prodigy, dumped over a billion and a half dollars into it, sold it to a Mexican company for a lot less than a billion and a half dollars. Many think it has lost its ability to survive in this niche. H & R Block owned all of CompuServe Information services, the largest on-line service provider in the world. It is now struggling to find itself, and in fact, announced record quarterly losses by its CEO just a week ago. MCI News Corp. bought Delphi. Delphi was the third largest on-line service in the world. They announced great plans. A year later, they sold it for nothing.

All the regional Bell operating companies in the U.S., and many worldwide have tried to get in the on-line business. They are trying again with Internet access provision deals but all of them, at least in the past, have been quite sour.

In this business you've got to eat it, you've got to sleep it, you've got to breathe it, and you've got to live it, and you've got to die by it. We think that marginalization for larger companies is very difficult to do.

What is really going on in this business? I am reminded of a great line I heard from Albert Einstein, who said that there were probably only five or six people on the face of the globe who truly understood his theory of relatively. "But," he said, "I am not one of them." If somebody claims that they know what the future of the interactive services business is going to be in the business or the consumer space, you should brand them a liar. Nobody knows. We are making this up and I am the first to admit it and everybody that is in the business, is in the same boat.

What is going to happen in the future? More pipe, more bandwidth, more content, we will get way deeper into the education business, which is a fantastic market, which is always slower to adopt than it should be. You will see the adoption of

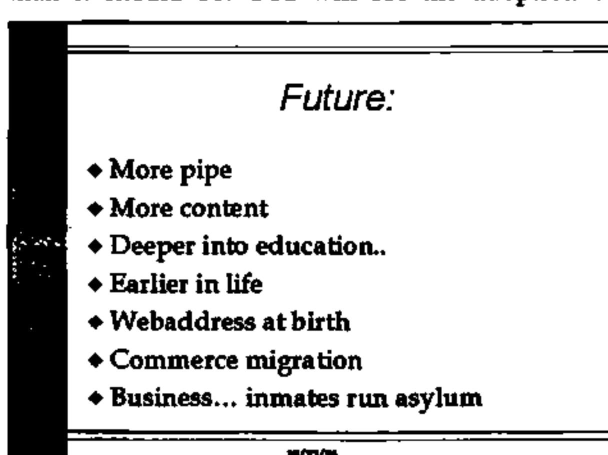

these products earlier in life, earlier in your business career, earlier in your physical life on the face of this planet. I predict that by the end of this decade and this century that you will receive an e-mail address or a URL at birth, like you do a social security number. You will have a lifetime URL.

Migration of commerce will occur fast. When it finally happens, it will occur real fast. I would suggest the MIS function of America's and the global corporations, that are important to all of us as customers, is under severe duress to prove its value, under severe duress to address this rising tide, which rises all boats, called the Net.

Continuing changes in my brand and the Internet – more women are using it, which is very encouraging. At AOL today, over 30% of our users are female. That is up from 5% not that many years ago. You will see more business on America OnLine and more business on the Net. You will see more purchases from the business platforms; EDI will finally work. My joke is that EDI will work because the manager will go home and be using Intuit to pay his or her bills and say, "How come I can't do this at work?" EDI will finally hit and consumer use will drive it.

You will see far less platform dependence, and frankly, far less chip dependence, we think, in future applications, be it the network computer, set top devices for televisions using coax, the X2 modem doubler from U.S. Robotics, or ISDN, or ADSL. Whatever pipe is going to happen, it will mean less chip and less platform dependence for services and programming like AOL. In fact you will see more programmed services. I know it

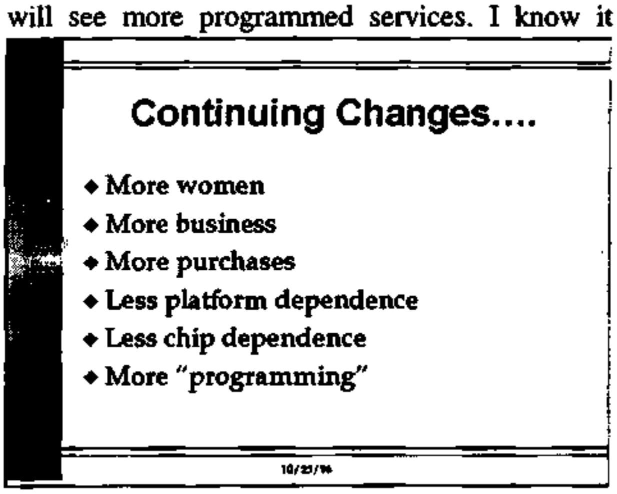

sounds suspiciously like television, but it will be a lot like television.

F. Scott Fitzgerald, the author, was once asked what it was like to become poor because he died penniless. He said, "You know first it happens very slowly, then it happens very suddenly." That is what has happened here. We burbled along in this thing for decades. Many of you probably were using the Internet in the '70s, and the '80s. You called it other names back then, but it hit, and it hit big, and I think that whatever you think it is going to be, you should double or triple it, because we are seeing evidence that this is an incredible paradigm shift.

I would conclude with the following phrase that I think sums it up – shift happens. Shift happens, and it is happening more rapidly than anybody would have ever predicted because we got the math wrong. With that, I thank you for your attention and take some questions.

Q: Do you have any suggestions for recycling these hundred million disks that you have sent out? You know nobody throws them away, they just accumulate in our offices and bookcases.

A: There is actually a web site on a hundred and one things to do with your used AOL disks. It is very inventive. A guy I worked with said he knew we were becoming a national brand when he went to a cocktail party and somebody was using AOL disks as coasters.

Q: As more people go on-line who will become responsible for upkeeping and growing the Internet?

A: There is no simple answer. It used to be the government. Now it is not. Everyone in the consumer world thinks it is free. It is not. The Internet is truly a network of networks, as you know, and it is as weak as the weakest point. What we think is going to happen is that there will be new sub-Internets, but they will call themselves the Internet. We think there will start to be performance guarantees routinely issued by Internet providers, because they are effectively not going to give you access to the Internet, they are going to give you access to their Internet.

In fact, At-Home, which is the interactive service that TCI is touting it will offer to cable subscribers, be it at business or at home, is really building the Internet for its own use. So the Internet, as we know it, I think will break and die, and you will see vulcanized versions of the Internet that are run by brand name companies who guarantee performance. So the short answer to your question is, who is going to be responsible for upkeeping and growing it? Companies with three letters in their names, like AT&T and GTE and IBM and AOL.

Q: How important are the new 56K modems to AOL and other ISPs?

A: We are one of, I think 30, signatories to use U.S. Robotics 56K modem as soon as it is deployable and replicable. We are galactically excited by this kind of development. The single most common complaint by people who cancel from our current membership is "too slow," the single most common complaint.

Many of you probably do not realize, but AOL is available now over your LAN. Any TCP-IP connection can get into AOL. You can go in your corporate LAN, boot up an AOL disk, and choose TCP-IP link, and off your LAN use AOL, and when it i that fast, it is great. Instant on, everything comes down, all the graphics instantly. It is a much more appetizing visual experience, and a much more satisfying way to use it. You do not feel like you have to wait. So anything that makes it faster we are whole-heartedly in support of.

Q: Do you see Java playing a role in AOL's future?

A: The answer is yes. I am not educated, informed, or at liberty to say much more, but yes, Java apps, shock wave apps, anything that makes the experience richer, robuster, and more visually appetizing is in our interest.

Q: What is the life of a subscriber? What percentage of new subscribers come through your relationship with Microsoft and then his most important question, any regrets?

A: Well you can't not do a deal with Microsoft and be in this space, is what we concluded. The percentage of our members that come through the Microsoft bundling deal: we are on the Windows '95 desktop now. It is, so far, relatively unimportant, but we think this Christmas push that is just starting now – our acquisition ramp is starting to go real strong again – will be far more important. There are no regrets, just like there are no regrets that I breathe to keep myself alive everyday. It is intrinsic in behaving in this space.

Q: What percentage of time spent by AOL users is in chat?

A: North of fifty. A lot of those chat rooms are not Mistress Helga's Leather Room. A lot of chat is extremely meaningful. I will give you a quick example. I was in Washington D.C., a little over a year ago, in Union Station at 7:30 in the morning the day of the Million Man March. I think I saw all million men come off the trains to go to the mall for that amazing day. At around 2:30 that afternoon, I went back to my office, and at 7:00 that night I logged onto an AOL chat room in an area call Net Noir. It is an investment AOL made, and a company that runs an area on AOL that is African-American-centric information, chat, bulletin boards, purchases. I am in the Net Noir chat room, and it is about that day's event, the Million Man March. Now a typical chat room on AOL will have anywhere from 10, 25, maybe 30 people in physical screen names in that chat room, talking to each other through text, and we are talking unformatted text on a computer screen. The world's most emotion-free medium. In this chat room in Net Noir that night there were 1,700 people – live. The text was spewing up the screen. They did a state check, and all 50 states were represented. In fact, many nations outside of the U.S. were represented. There were many people that were at the march talking; many people that were not at the march asking what it was like. In this most emotion-free medium, I seriously had one of the most emotional experiences of my life. It was staggering to read, how these people were relating to each other what had happened, sharing the joy, sharing the pain, sharing the moment. You do not get that in a lot of other ways that we interact. So this idea of chat, I think, got a bad moniker when it shouldn't. It can be amazing for a wide variety of reasons.

Q: Why has your cost of gaining new subscribers increased so much? Cost of the free disk, or high level of people dropping AOL, so low net increase?

A: The cost of the free disk has gone down. The cost of gaining new subscribers has increased for several reasons. The first is that as you go deeper and deeper into an audience of early adopters, which was our typical sweet spot, you find the less and less convinced, i.e., you find people who are more "show me" orientated than the original group.

Secondly, the idea of flat-rate Internet access by local ISPs, although numerically not a huge problem for us, philosophically and cyclically is a bigger problem. What we found is that a lot of people perceive their AOL bill to be way larger than it actually ended up being. If you ask people what a lawyer cost, they routinely suggest that a lawyer would be way more expensive than lawyers actually are. There is that huge gap of perception and reality, once again of expense, and I think we suffered some for that.

The third reason is that there is a bit of a "wait and see" attitude. We think, in this Christmas and next Christmas, it will start to come back stronger, with "What is the next chip? What is the next modem speed? What is the next service? Do I have to commit now?" Somebody paid $800 for a Bromar-brain four-function calculator. We kept seeing calculator prices going down and VCR prices going down and I think there is a "when do I get on the carousel" perception for one of the next sets of waves of consumers. What will they buy? How cheap will it be? Will this network computer from work really exist? Will I have a set top device? And that confusion I think sometimes hampers our acquisition efforts.

Q: I no longer purchase floppy disks as I receive a steady supply from AOL. I promise to subscribe if you could ship future disks blank to simplify my use of them.

A: I wish that we did not have to pay for all these disks, but it is quite funny to see people's reactions to these disk mailings. Most of them ask, "How come you can't purge me off your list?" I have never responded. Some of them say, "I already belong and I keep getting disks," which is true. It is actually cheaper to mail than to try and figure who does and does not have this stuff we think the disks might have run their course. MSN didn't mail any disks. Gee why? They had the operating system on their PC. They didn't need to, but CompuServe, and Protégée, and WOW from CompuServe, or our product from GNN, and a bunch of others are mailing disks. So the differentiability of our disks has been reduced to almost zero. In fact, maybe the way we differentiate ourselves is by mailing blank floppy disks and say, "Go to the Net and download our new client," which is available at, by the way AOL.com.

The idea of disks as a trial medium to get the uninitiated to try this product out worked well. It was a pretty smart stroke by Steve Case, the guy who runs my company. The next step is to make us into a lifestyle, or a business, or a content choice. You choose to buy the *Wall Street Journal*. The stock price for IBM is the same in *USA Today* as in the *Wall Street Journal*, but which would you rather be caught looking up a stock price in? It is a choice of branding for the same data. You will see multiple brands from my company to make ourselves a choice of a product that is branded. The migration by us to a branding strategy is, I think, our next big paradigm shift.

The Technology Investment Outlook for '96 and '97

Charles A. "Chip" Morris

Managing Director and President of T. Rowe Price Science & Technology Fund
T. Rowe Price Associates

Mr. Morris is a managing director of T. Rowe Price Associates, which he joined in 1987, and the president of the T. Rowe Price Science & Technology Fund.

Mr. Morris graduated, summa cum laude, with a B.S. degree in finance from Indiana University and earned an M.B.A., with honors, from the Stanford University Graduate School of Business. He received the Alexander A. Robichek Award as the No. 1 finance student at Stanford and has also achieved the Chartered Financial Analyst (CFA) accreditation.

TRANSCRIPTION FOR THIS SESSION IS UNAVAILABLE

Asia/Pacific and China: Looking Ahead to 1997

C.D. Tam, MBE, JP

Senior Vice President and General Manager,
Asia/Pacific Semiconductor Products Group,
Motorola

Mr. Tam is senior vice president and general manager of Motorola Asia/Pacific Semiconductor Products Group. One of his objectives is to bring the most advanced semiconductor technologies to the Asia/Pacific region. He spearheaded the establishment of advanced integrated circuit design centers in Hong Kong and Taiwan; two fully automated assembly and testing plants in Hong Kong; three advanced facilities in Hong Kong (Silicon Harbour Center—A/P Semiconductor group headquarters); Tianjin (operational 1993); and Singapore (Innovation Center—operational 1994; including IC design center). Mr. Tam is also the leader who champions and encourages numerous revolutionary devices in Motorola, among which are the "DragonKat" microprocessor, which received the Hong Kong Governor's Award for Industry in Machinery/Equipment Design in 1989; and the "Poc Set" chipset, an innovative system for pen-based palm-top computer with communication capabilities.

Mr. Tam began his career teaching for two years and joined Motorola Semiconductors Hong Kong Limited in 1968 as an applications engineer. He gained a wide spectrum of expertise in the company through different positions in engineering, sales, marketing, and product management. Mr. Tam was general manager for Motorola's Asia Pacific Marketing Operations and has accelerated the establishment of a strong regional organization through decentralization. His responsibilities were further expanded in 1982 with the opening of a facility in Hong Kong that houses an advanced testing plant. In 1984, he was appointed vice president and general manager for Motorola's Asia/Pacific Semiconductor Division, the first Chinese to achieve regional responsibilities in the corporate history. In November of the same year, he was elected as one of the Ten Outstanding Young Persons of the Year in Hong Kong. He was promoted as corporate vice president in 1988, and further promoted as senior vice president in 1991. In the 13 years since Mr. Tam headed the Asia/Pacific area, sales have climbed to over 16 times the 1980 level. In 1992, two significant events marked the milestones for the Asia/Pacific region: the division's 25th Anniversary, and the elevation of the division to group status. In addition, the company received measures of accords over the years.

Mr. Tam is active in Hong Kong community activities and is a board member and committee member of several electronics and business associations. In recognition of Tam's contribution to the community, he was granted the Young Industrialist Award by the Federation of Hong Kong Industry and the Courvoisier Award for Business Excellence—Industry, both in 1988. In 1992, he was appointed by the Hong Kong Governor as Non-Official Justice of the Peace and later honored by her Majesty the Queen as the Most Excellent Order of the British Empire (MBE). Again, the same year, Tam was awarded the Executive of the Year, Hong Kong Business Awards sponsored by DHL and SCMP. Again in 1995, he was elected as one of Asia's Top Ten Executives by Electronic Business Asia.

Mr. Tam graduated from the University of Hong Kong in 1966 with a B.Sc. (honors) degree.

Session # 17: Asia/Pacific and China: Looking Ahead to 1997

C.D. Tam

Senior Vice President and General Manager, Asia/Pacific Semiconductor Products Group, Motorola

JOSEPH GRENIER: Our next speaker is C.D. Tam of Motorola. Mr. Tam is Senior Vice President and Group General Manager of Motorola's MCU Technology Group worldwide. He spearheaded the establishment of advanced integrated circuit design centers in Hong Kong and Taiwan, two fully automated assembly and testing plants in Hong Kong, advanced facilities in Hong Kong, Tianjin and Singapore. Mr. Tam is also the leader who champions and encourages numerous revolutionary devices in Motorola, among which are the DragonKat microprocessor, which received a Hong Kong Governor's Award for Industry in Machinery and Equipment Design in 1989, and the "Pock Set" chipset, an innovative system for a pen-based, palm-top computer with communication capabilities.

Mr. Tam joined Motorola Semiconductors Hong Kong Ltd. in 1968. In 1984, he was appointed Vice President and General Manager for Motorola's Asia Pacific Semiconductor Division, the first Chinese to achieve regional responsibilities in the corporate history. Among his many awards is the "Most Excellent Order of the British Empire" awarded by the Queen, and in 1995 he was elected one of Asia's Top 10 Executives by Electronic Business, Asia.

C.D. TAM: I am honored to be here speaking at this Dataquest conference to share with you all, some of our outlook in Asia Pacific, particularly that of China.

Motorola Inc. revenue was $27 billion U.S. last year, and 37% of this revenue came from the USA, which is obviously the biggest country in term of sales ranking for Motorola Inc. Which country occupies the No. 2 ranking in our corporate revenue? That is the combination of Hong Kong and China, which represents 12% of our entire revenue. Eight years ago our business in China was virtually zero, for the corporation. Even this year, when there has been some slowdown in some regional business temporarily, the business in China continue to grow.

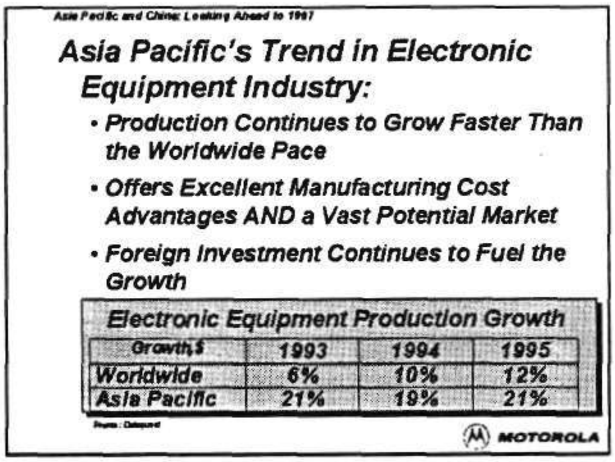

So let us get on to Asia Pacific. It is not a secret that electronic production in the world has been shifting towards Asia Pacific for many years. Historically this was driven by the lower manufacturing cost in Asia Pacific. What I want to talk about today is about a real paradigm shift. The real reason for producing continuously in Asia Pacific is to also serve the increasingly large domestic consumption markets created by improving the living standard of a large population base of over 2.5 billion people.

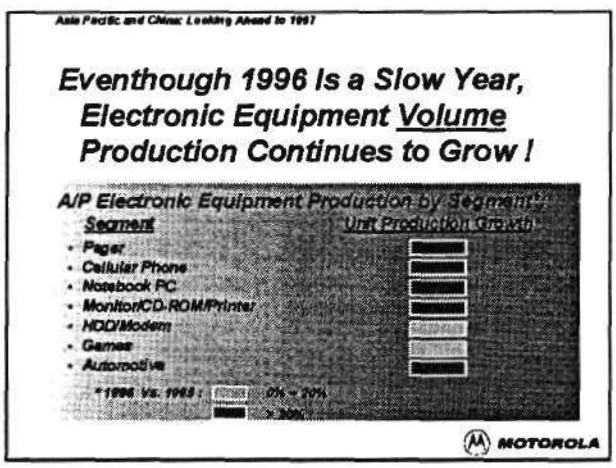

Asia Pacific Semiconductor TAM

	Annual Growth Rate	
	90-95	*1996*
• *Non - DRAM*	*28%*	*9%*
• *DRAM*	*51%*	*-50%*
• *TOTAL*	*35%*	*-10%*

First Adjustment After 11 Years of Growth

MOTOROLA

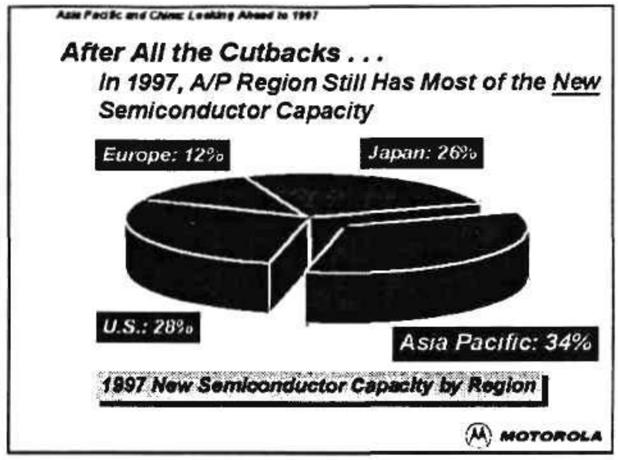

Even though 1996 is supposedly a comparatively slower year, electronic equipment unit production continues to grow in Asia Pacific. This year many of the major production segments such as pager, cellular phone, notebook PC, monitor, printers, continue to grow quite well.

What about semiconductors? After many years of 35% annual growth, the Asia Pacific semiconductor consumption market does have a recession this year, primarily caused by the falling prices of DRAM. This adjustment is the first time in 11 years. However at non-DRAM semiconductor demand continues to grow at a 9% rate, even through this so-called recession.

In 1996, for semiconductor suppliers, there are problems. First of all, at the beginning of the year there were high component inventory at most of the customers. The second one is because of very fast build-up of fab capacity, ASP does drop significantly. DRAM of course, led the pack with an 80% drop. Some of the more mature commodity prices also drop. But on commodity semiconductor price actually only decrease at about the learning curve rate. This semiconductor recession is not a general economic recession, as the economy in Asia Pacific continues to boom. This is an excessive capacity-induced recession and therefore collectively we should be able to correct it in due time.

Definitely, there is some slowdown in the rate of wafer fab capacity investment in Korea, Singapore, Taiwan for the short-term. Raising funding, in the private sectors, has become a bit more difficult than before. There is still plenty of money around in Asia Pacific but the investors are more critical and more choosy.

Is it really that bad? Not exactly. There are delays and scale back, but some of the delays are basically 6 month push-out. Construction of new facilities is still ongoing but at a slower pace.

The investment momentum of the past two years still means Asia Pacific would have the highest percentage of new wafer fab capacity coming on stream in 1997. We estimate this to approach almost 34% of all the new capacities next year. The government in Asia Pacific, such as Taiwan, Singapore, and Korea, would get into the act in order to maintain the country-level competitiveness during this slowdown. For example, in Taiwan, the Taiwan government recognizes that since the private semiconductor sector, through their own effort as well as strategic alliances with foreign company, are closing the wafer processing technology gap in manufacturing with U.S. and Japan.

They surely do not want to see this momentum being lost. So a little series of financial incentive has recently been announced during the Taiwan Semiconductor Conference last month. As an example, in total 17 intelligent industrial parks are planned to attract foreign investors and to establish the island to become a high-tech island, before the year 2000.

These intelligent industrial parks will be categorized as science-based, technology-based, software, aerospace, biotech, incubator, and research-based. Special operation zones have been planned to allow the enterprises stationed within the zone to enjoy zero tax, freedom to hire foreign workers (because Taiwan actually ran out of workers), direct transportation with mainland China, etc.

In fact there will be new incentive programs for companies within these intelligent industrial parks, five years exemption of tax, shareholder investment credit, acceleration of machine depreciation, joint venture with funding from Taiwan government, the interest rate 2% lower than the typical interest rate from bank.Taiwan has sufficient capital resources at the government level because of the U.S. $19 billion foreign exchange reserve. When thing start to slowdown quite a lot of governments in Asia Pacific are jumping into the act. There are not letting the country lose momentum in the electronic and semiconductor arena.

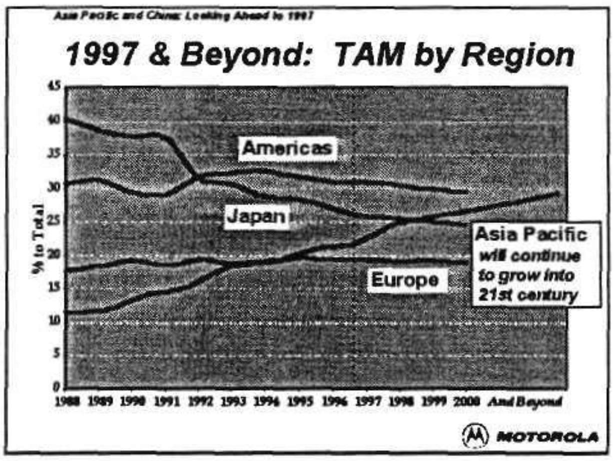

From Asia Pacific semiconductor manufacturing to consumption, look at each of the geographical market trends. America will maintain the No. 1 spot in the semiconductor consumption market ranking, in the run up to year 2000. The question is therefore: When will Asia Pacific overtake Japan? Because of the weakening of the yen, there is some thought that the overtake time will be extended to the year 2000. However there are other thing which are happening within the region, such as some domestic markets in Asia Pacific, like China, that will mandate in-country production. That could move the year to 1999.

This is a major paradigm shift that will happen within Asia Pacific in the next several years. There is sufficient to say that Asia Pacific would reach the number 2 ranking, give or take one year difference.

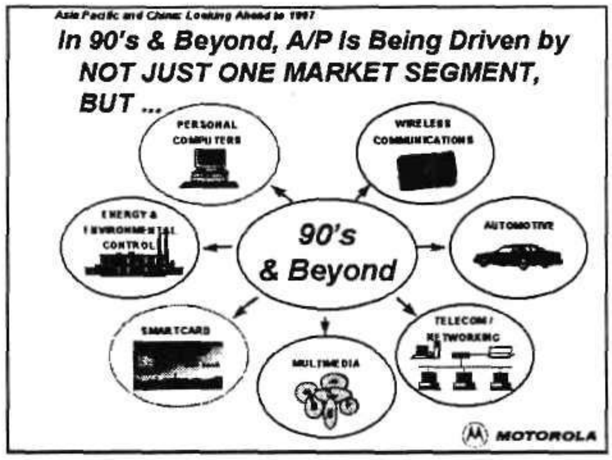

Semiconductor growth in the '60s, was very much driven by military, U.S.-led. In the '70s, driven by mainframe computer, again U.S.-led – primarily IBM. In the '80s, the semiconductor market was driven by consumer products, and globally led by Japan.

In the '90s, Asia Pacific is being driven by not just one market segment but seven personal computer, wireless communication, automotive, telecommunication networking, multimedia, smartcard, and energy and environmental control. Every one of these market segments has increasing semiconductor content.

Asia Pacific also sees another phenomenon. The creation of domestic consumption markets have accelerated dramatically. I will use the pager as an example. In 1990 (about 6 years ago), there were 23 million pagers around the world, over half of which were in the USA. By last year the pager use had grown to 95 million and Asia Pacific was almost 60%. In the span of 6 year, the Asia Pacific

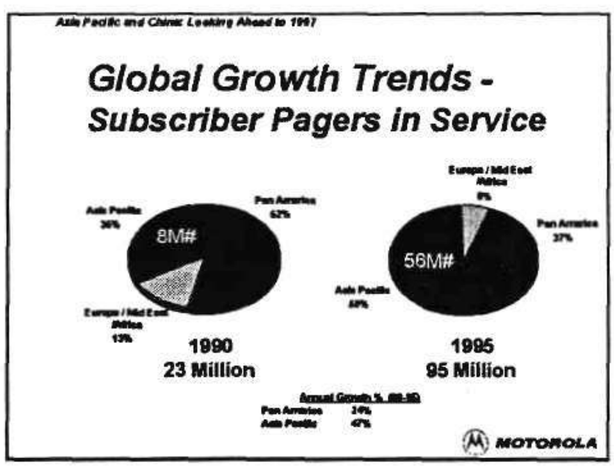

market is already bigger than the USA market. The USA market took 15 years to be created, and many of these pagers are in China. Think about the implication of the creation of a domestic market in Asia Pacific within a five to six year cycle, to a size that is bigger than the USA market.

Obviously, the China factor cannot be ignored because there are 1.2 billion people. As soon as living standard continue to improve, the buying power becomes higher and higher.

There is another thing that has been going on called the Hong Kong/China transition and its impact. There are two typical misunderstanding about China. The first one is that China does not have enough money. Second, China state-owned enterprises are inefficient and therefore there is a wrong conclusion that the entire China electronic industry is inefficient. Thus the impact to the world will be small.

Let me explain some of the things that have been happening in the last eight to ten years. The world does not understand the Hong Kong/China factors. Since the opening of China in the early '80s rich overseas Chinese have been channeling investment through Hong Kong to China. In fact, over the past 10 years, 80% of the foreign direct investment in China came from overseas Chinese, in Hong Kong, in Taiwan, in southeast Asia, channeling through Hong Kong.

The 20% foreign direct investment in the past 10 years came from USA, Japan, and Europe. For example, Motorola's half a billion dollar investment into China, is part of the 20%. China has found a way of getting capital and investment when they need them the most, which was about 10 and 5 years ago. Now, of course, everybody more or less invests in China. So that takes care of the so-called "China does not have enough money to build some industry."

Hong Kong has been moving low and mid-range manufacturing across the border. Today four million people in China, across the Hong Kong border, are working for Hong Kong companies in factories there. Efficient factories and management were introduced in the south by Hong Kong and Taiwan companies, and this has moved across the country. Thus the Chinese have a whole new breed of factories, which are different from the less efficient state-run factory. This is what creates the large export surplus for China. Obviously most of the profit does go to Hong Kong, Taiwan, and overseas Chinese. The result is that per capita GNP in China significantly jumped, in quite a number of locations.

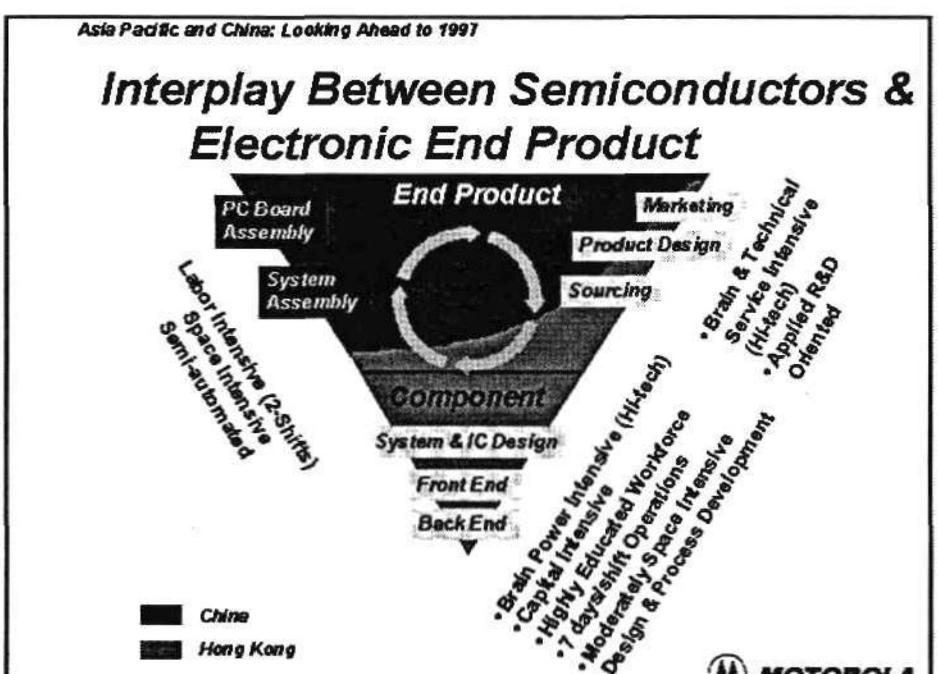

Thus a domestic consumption market emerges, so it is not surprising that they will buy more pagers than USA. The Taiwan government has a foreign exchange reserve of almost $90 billion U.S., and would use part of them for their building up the industry into the future. The Hong Kong government has $70 billion U.S. foreign reserve and do not know what to do with it. That is the difference between the Hong Kong government and the Taiwan government. So all the previous investment into China, by overseas Chinese, actually comes from the private sector. Some of this $70 billion is now getting used, more or less like Taiwan.

This is the interplay between the semiconductor and the electronic end product in the Hong Kong/China border. Most Hong Kong companies keep the marketing, product design, and material sourcing in Hong Kong, but have all the end product manufacturing, like PC board assembly, system assembly, just across the border. The most sophisticated components are done in Hong Kong, or Taiwan, Korea, and Singapore.

Now that China has a dependable manufacturing base new policies are put in place, like the Taiwan, Korea, and Singapore governments have been doing for many years. The difference is that China has a large domestic market to use as leverage. For China, in the '96 to 2000 five-year plan, their top priorities are: reinvigorate agriculture, development of some of the more interior poor parts of China, inflation control, and reform some of the state-owned enterprises by learning from the foreign companies.

The five-year plan for electronic industry focus is on strategic development in semiconductor, communication, transportation, distribution channel, software, audio/video, and medical.

There are three gold projects: golden bridge, golden card, and golden tax. The last one is more important as quite a lot of people try not to pay tax in China. They want to automate it, and computerize it. In the ninth five-year plan for China, for the electronic industry, their goal by year 2000 is to achieve an annual growth rate, at about 20% plus. The total electronic industry will reach about $75 billion. As a comparison, last year the Taiwan IT industry production was about $20 billion.

They would like to export about $25 billion out of the $75 billion production, so the remaining $50 billion worth of production would translate to production of sales of about $63 billion domestically. What do all these numbers really mean? There is a certain implication. What it really means is that at least $63 billion worth of electronic end products, or semiconductors, they will not be importing. They will be trying to build it themselves. Either through 100% foreign-owned company through J/V or through China domestic enterprises.

So import of finished end-products from other countries, say Japan, would be much reduced in the future. Since end product production will be in China, semiconductor consumption will also shift there. That is why, even though the yen has been weaker, there is a popular belief that Japan factories moved to Asia factories. But they have not, because of the domestic market requirement, and government mandate. These are important paradigm shifts to observe.

There is always the impression that Chinese worker cannot make sophisticated products or engage in production that requires great attention to detail, so they may not be able to build all these things. I want to share with you some of our experience in Motorola, in China. In the Motorola factories in Tianjin, GSM and soon CDMA digital cellular phones are being manufactured. Motorola makes some of the smallest and most sophisticated cell phones in the world. There is no problem making them in China.

What is the scenario for the China electronic industry? Foreign investment is being encouraged, for certain leading edge technology areas like semiconductor and other designated sectors. They are preferential policy for these sectors. For example early pioneers can have 100% foreign ownership and still have access to China domestic market. If you do not want to invest 100% yourself, say for advanced integrate circuit manufacturing, China will be willing to put in money and operate like a foundry.

Investing in mature industries, like color TV, then there will be an export requirement, while you can sell domestically. The more mature the segment, the higher the export percentage for new factories built to serve those mature segments. For the mature product, if you want to access the China domestic market, you also have to export. The semiconductor consumption is not just for product built for China consumption, but also those built for export. This will have an impact on countries like Malaysia and Thailand. Malaysia and Thailand do not have a big domestic market as compared to China, so there will be manufacturing resizing going on in the next 5 to 10 years.

Continuing with the China electronic scenario, government will also try to support the creation of large conglomerate groups. This is like the Korea model. Basically the China government wants to help create or merge something like 30 large conglomerate groups with about 10 billion RMB annual sales by the year 2000.

The ideal complement: China offers lowland, low-cost labor, potential market, and in certain high technologies China will also make capital investment. Foreign companies contribute some technology, capital, and job opportunity. There is a word of caution. You really have to know China, China's culture, and the negotiating techniques, and build a strong relationship before you can really achieve the ideal complement. I am not trying to say it is easy, but it can be achieved.

At the current economic growth rate the total GDP of China will exceed about one trillion dollar by the year 2000. How do you really go after China with such a large geographical spread and huge population? I will share some of our experiences with you. You don't; you pick targets. You

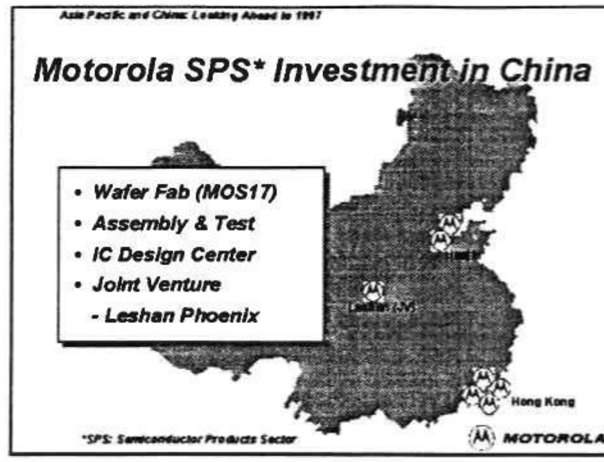

look at per capita GDP (measure in purchasing power parity, which is PPP measure) and go after the cities and provinces that have high per capita GDP, on a PPP measure, that means Shanghai, Beijing, Tianjin, Guangdong province and so forth on the chart. In fact, some cities like Beijing, Tianjin, and Shanghai are almost as significant as an entire province.

Where do we stand in Motorola semiconductor investment in China? By the middle of next year Hong Kong will be part of China. This map shows our semiconductor sites in China. In Tianjin, the northern part of China, our assembly and test facility alone produce about 500 million integrated circuits a year in assembly and test. The product will include microcontrollers, so it is not simple logic or analog devices at all. The facility actual reached 6 sigma quality level after startup in less than one year. Why? Because even a technician in the production line has a bachelor's degree, as you can hire directly from university. No other country, and Motorola has many plants, has every

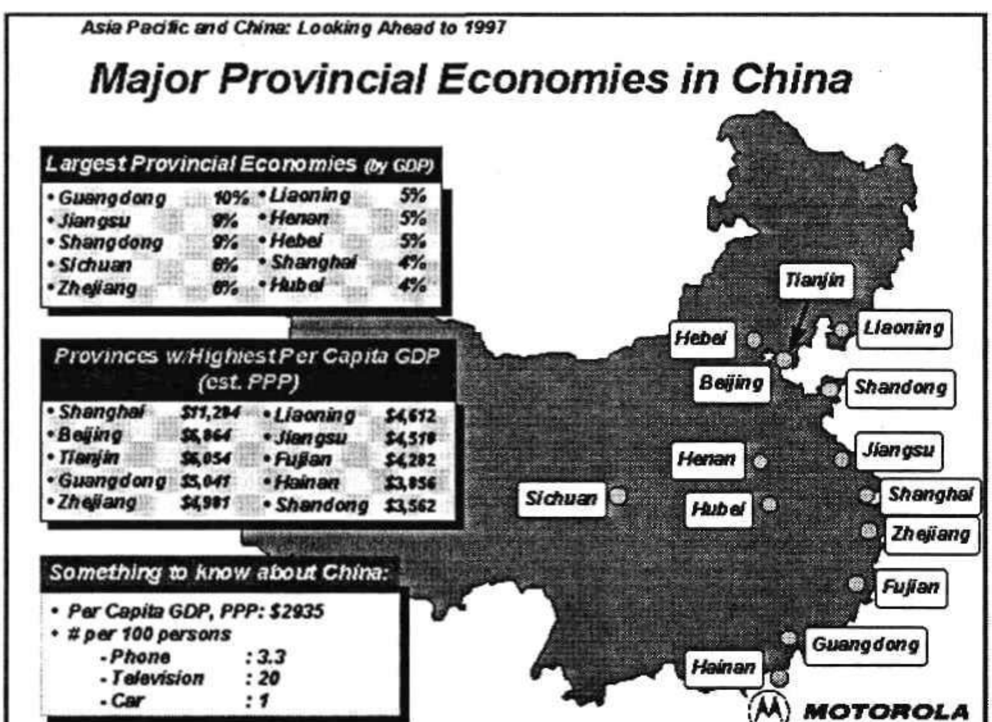

technician having a bachelor's degree.

Our wafer fab called MOS17 is under construction. It is also in Tianjin and it will be the first eight inch wafer fab in China, at the 0.65 micron level, and with some staff changes it can handle finer geometry. The type of product that can be made and will be made are smart mouse type of devices, MCU, as well as IC mouse.

I hope this give a quick overview of Asia Pacific and China, even though this year is a recession year for the semiconductor industry, and again recall that most of the product outside of DRAM are still growing, and there is tremendous market possibility.

Q: How will China control inflation while achieving 20% profit growth rate?

In the last few years, the China government has been putting a lot of control into the inflation. China inflation, this year, has been running at about 11% and lower, and the electronic growth is in the 25% growth rate. In fact, they are running at 31% growth in the electronic industry. They have been able to control the inflation rate. The economists in Motorola have been tracking this for years, because we have, of course, a lot of vested interest in trying to understand China, and they have been able to do something which most other countries have not been able to do.

Q: How do you suggest U.S. companies balance their desire for profit versus China emphasis on long-term development before profit?

A: The Chinese have always been capitalistic. That is why the overseas Chinese are so rich, because the Chinese have always been a capitalistic culture. They do not think that long-term development means low profit. It is a negotiation. If you understand that you always have to give and take something, to be able to share the profit, okay. When you achieve synergy, then both side will be profitable, and that is the secret. You do not go after a project that you end up having to do 10 years of investment, and never expect much return. Why should you? The Chinese government do not expect that to be the case. There is no harm for them to ask, just like there is no harm for you to say no. You are dealing with Chinese who are very capitalistic. They call the factory "profit-buildings." Most of the people that you will be dealing with just run the company, but they really want to make profit, so a balance is possible.

Q: Please comment on non-transferring business practices.

A: There are two questions here. One, I presume he means is there a lot of bribery? Does it require kickback,? The answer is no because if you do things above-board, you do not see bribery and kickbacks. Remember that people get shot there if they are caught. You have to always do things above-board.

Now, the other thing called "non-transferring practice" comes in. There are hidden taxes, or charges. These are all legitimate, except you do not know it right up front. You have to look for more than your contract. Ask somebody who has been there, and if you ask for those to be excluded, they will exclude it, and you do not have to pay it. I am just sharing some information out of experience.

Finally, the most important thing is the relationship. If you build a relationship, you will find doing business with the Chinese are pretty easy, because they want to do business with people they know. They have this mentality, "I am going to buy from somebody anyway. I better buy from somebody I can trust, and can have a relationship with." Once you master this, it is not difficult.

What Do the USERS of Technology Say?

Scott Winkler

Vice President and Research Director,
Platforms and Operations Software Technology,
Gartner Group

Mr. Winkler brings 18 years of computer industry experience to Gartner Group. His research area at Gartner Group is in platforms and operating systems.

Before joining Gartner Group, Mr. Winkler was director of systems marketing at Sequent Computer Systems where he was responsible for communications, operating systems, systems software, and hardware technology directions. Before that, he held a variety of management positions in sales and marketing at IBM, including responsibilities for mainframes, communications, personal computers, and UNIX systems.

Mr. Winkler holds a bachelor's degree in electrical engineering and computer science from George Washington University.

Session # 18: What Do the Users of Technology Say?

Scott Winkler
Vice President and Research Director, Platforms and Operation Software Technology, Gartner Group

JOSEPH GRENIER: Manny Fernandez in his opening talk mentioned the huge increase in dollars that will be spent over the next several years by the information technology industry. To expand on that, we have Scott Winkler from the Gartner Group. Scott is Vice President and Research Director, Platforms and Operation Software Technology at Gartner. Mr. Winkler brings 18 years of computer industry experience to Gartner. His research area is in platforms and operating systems, and involves some 65 Gartner analysts. Before joining Gartner Group, Mr. Winkler was Director of Systems Marketing at Sequent Computer Systems, where he was responsible for communications, operating systems, systems software, and hardware technology directions. Before that, he held a variety of management positions in sales and marketing at IBM, including responsibilities for mainframes, communications, personal computers, and UNIX systems.

SCOTT WINKLER: I would like to shift from the supply side to the demand side, and talk to you about the users of technology. At Gartner Group we have relationships with many thousands of user organizations, and provide them with advice and counsel about their information technology investments. In doing so we are able to collect some very interesting information about who is buying what, why, when, and how. I would like to share some of that with you.

There are two types of users we deal with at Gartner Group and the users themselves fall into the category of end users in the business units of organizations, and the Chief Information Officers, the so-called CIOs, who are centralized buyers or standards setters. There has been some flux between both camps with respect to power in the buying decisions, and technology setting decisions in large organizations. Let's start with the CIO.

The chief information officers of large to medium sized organizations are in a squeeze – they are being asked to do more with less. They are being asked to take on more responsibility, and sometimes with dramatically less authority. We see a trend which is not working in their favor. The half life of a CIO these days has decreased to about 18 months and is still decreasing. The likelihood of them holding onto their jobs two years hence is very low, the turnover is very high. There is a vast supply of CIOs in the world and smaller demand for them. Running scared as they do they are often in a position, in large to medium sized organizations, and they are looked upon around the executive table where someone turns to them and says, "You and your organization are the reasons that the business cannot advance. We are trying to do new things, yet our technology is in the way." That is the position no CIO wants to be in, but unfortunately, too many of them find themselves there. Their number one goal therefore is to enable their businesses to get things done. To absorb new technology though, is not anywhere near their number one goal and in many cases absorbing new technology is the last thing that they would like to do.

I would like to cover five points about what they have told us in this tour and it is all very interest-

The CIO View of Technology

- Not so fast!
- User perception *is* reality
- Keeping up with the consumers is tough
- Prove to me the benefits of early adoption
- Internet: business enabler or cyber water-cooler?

Gartner Group ©1996

ing, however it is not very positive. They want to slow down the pace of technology, but they do not really mean it, so let's explore what they are saying and what they mean by it. They will talk about speed, saying "not so fast," and deal in technological obsolescence. They will speak about the users. The CIOs that are speaking about users are talking about business organizations, departments, divisions, business units. What they believe to be true is the reality of the CIO and they are influenced by the suppliers of technology. They are finding it very difficult to compete with the consumer business. Now position the CIO who has end-users who, during the day, are business users, and during the evenings and on weekends, are consumers, and therefore they are greatly influenced. The CIOs are asking for prescience, they are asking for us at Gartner Group to help them understand which technologies are going to give them benefits. If they are going to be the early adopters, they do not want to adopt those technologies that are going to fizzle out or become fads. They want to know about Internet, and they want to know why Internet is a technology enabler for them and why it is not a productivity drain. What they see too much today is a lack of productivity associated with Internet as opposed to a boost.

Let's start with speed. The latest technology for a CIO is nowhere near the highest priority. They frankly do not care – unless that technology can be translated into a business advantage or an opportunity for one of their business users to do something they could not do previously. CIOs want release schedules to be tuned to business priorities. For example, they do not want technology to be released by technology companies when it is ready. They would much rather it wait and be packaged up in more consumable pieces and therefore receiving a steady stream of technology over a year is much less desirable than receiving one big package of technology on an annual basis. The CIOs play this back to us a number of times telling us, "We just cannot afford to upgrade very often so therefore we need to make bigger jumps less often than small, smooth, incremental jumps." Now, of course, they have heard for years about the up-ensuing and up-coming technological marvel of things like object-oriented technologies, that will allow them to make smooth upgrades as they go. Frankly, they are very skeptical. They do not believe that will ever come to pass and the more the technology industry tells them the component architectures will make upgrades more smooth, they tend to focus on the unmanageability of thousands of components.

Not So Fast!

- The latest technology is not the highest priority
- Release schedules must be tuned to business priorities - do not let features "dribble out"
- We know obsolescence is you business model; impress us by subjugating it once in a while

Gartner Group

It is not a winning game and those technology companies who have reoriented themselves to the business user and said, "I get it, and I am going to give you a big upgrade every year, versus a small upgrade every month" have won the reverence of these users. The final point is a quote, I have this is straight from a large CIO in Europe, who said, "We know obsolescence is your business model. Please impress us by subjugating it once in a while." They were speaking directly to the technology vendors and recognizing that there is a push and a pull here. That is, it is important for technology to keep moving. There are many times when things are working wonderfully and yet they become obsoleted by the next generation and they sit back and ask themselves, "Why did I move? I moved to stay current. I did not move because I had a business reason to move." They are willing to accept some of this because the CIOs, although not all of them are technologists, if they survive, tend to be people who understand the technology industry and recognize it. There is some balance there, and if more moves going forward are just for staying current than for getting real value, the users will then revolt.

The CIOs view of the end user is also something that is in play here. They believe that the users have the money, and therefore they have the power, and therefore their perception is the reality of the day. Five, six, and seven years ago, there was a user/end user revolution. The CIOs lost a dramatic amount of their power and many were washed away in a major trend. The users believed they could do everything faster, cheaper, more productively, and much more in tune with what they needed to accomplish at their local level, than their large, monolithic, centralized IT infrastructure could deliver for them. While they were correct in assessing that their centralized IT infrastructures were not getting the job done, they were incorrect in believing that they could do it all themselves. So there was a vast sweeping change to the end user, to the departmental level, ergo, the growth in business of the personal computer and the local area network, and all the accouterments that go with that. In the end, the user organizations found that they took on more than they wanted to, they were creating their own, small IT departments. They had tremendous organizational inefficiencies. They did not know how to own or operate their IT, and frankly it wasn't their business calling, and within the last three years we have seen the pendulum swing.

We have seen it swing back to an era of organizational, centralized authority to some degree, control to a lesser degree, influenced to the highest degree. So the new generation CIO is one who works for the end users. He does what the users want him to do, to operate the systems, make sure they are backed up, make sure they are available 24 hours a day, 7 days a week, and to integrate those systems both internally and for the rest of the world. Whatever the users desire, the users are going to get.

There are very few organizations we have met, less than 5%, where the CIO can make things happen and make decisions stick, technically, regardless of user input. In the end, what is happening now, is very different, even though we have gone through a re-centralization, in that this re-centralized world is one where there is much less centralized power. One thing that the CIOs have told us universally is that the users are influencing technology, because technology decisions are coming through the consumer channels.

Thirty-two bit operating systems for personal computers is very topical among these users, whether they be Windows 95, or Windows NT on the desktop. Still, the vast majority of business users are running PCs with 16 bit Windows operating systems. It will be two to three years before we see it turning over to being a majority position of 32 bit operating systems in the installed base.

User Perception is Reality

- Users believed they could do IT:
  - Faster
  - Cheaper
  - More productively, and
  - In tune with business priorities
- Now users want the CIO to:
  - Operate our systems
  - Integrate our systems with the corporation and the world

Gartner Group

However, of those who are moving quickly from Microsoft Windows to Windows 95 or Windows NT, we still find a small portion of them with a true business reason to do so and who are building a business application on top of the new platform that exploits the new platform or buying one that exploits the new platform. This is fundamentally a situation where the organizations are both staying current as well as responding to user demand. Because the users run their PCs at home, and what runs on the PCs at home is Windows 95. The level of noise that is heard in the CIOs office about, "Why do we have this archaic operating system called Windows on our production desktops when we have this new operating system called Windows 95 running at home. I upgraded in two hours. Why can't you?" This is just too much noise for the CIOs to stand. So what this creates is something of a backlash effect, where the CIO feels under duress from companies that go end-around, and market directly to their end users. This is not going to stop. The two companies that are leading the charge in end user marketing are both Microsoft and Intel; clearly the ones who have the most to gain by such an upgrade. It creates a question for the CIO about which of those trends are the ones that are necessary and which ones are just trendy.

The ones that we see being adopted less consistently are multimedia and things like Internet at the desktop for business users. Although these are technologies that have blossomed in the consumer space, in the business space the CIOs are much less aggressive on the acceptance of multimedia for lack of business justification. Therefore it is our expectation that these types of technologies are going to continue to be led at the consumer level and only creep into the business space when there is an application that justifies them.

CIOs are not afraid to be first. They believe that being first will provide them benefits. However they must generally have an over-justification for early adoption. There are too many examples of adopting technologies first, only to find out that they did not provide benefit or that the benefits they provided were short-lived because the technology didn't last in the marketplace. The second point is one that CIOs tell us most often, that skipping technology generations lowers their risks and saves them money. Therefore they must have an understanding of what they are going to get before they are going to accept it and they are slowing down as we speak. We have found in the last six months CIOs are once again slowing down their willingness to accept new technologies into their organizations because of a lack of fundamental justification that they can understand.

There is also a growing reticence to accept anything new. Release 1.0 of any operating system, or any major subsystem, or new hardware designs, is going to be avoided by the vast majority of organizations, even if they call it something like Release 3.1. These organizations are becoming much more savvy and understand that the first release of anything is to be avoided. In fact what we are finding here is an over-justification and an over-reaction to this concept. To this day we are finding organizations who are being conservative about the adoption of technology that was proven to be effective many months and sometimes many years ago. Once the technology takes on a reputation of being new or buggy, it takes a very long time for it to be washed out of the market. We are finding that organizations are being slow and conservative and willing to wait. The number one influencer of these organizations is still the popular press and the trade press. They continue to read and look at the sensational stories of what does not work and be less influenced by those stories of what does work. We expect that is going to continue as well. Ultimately what they are trying to separate are fads from trends, trying to understand that if they invest dramatically in one direction, their investment will pay off over the long term.

The one area of the most heightened interest today is the so-called area of groupware, where for some time there were organizations adopting groupware strategies, which three years ago meant they were adopting LotusNotes. The most common reason that Notes was adopted was because everyone else was doing it, and it became a fad long before it became a business value for many of the organizations who adopted it. Some who adopted it paid a fairly large sum, not only in the software license fees, but for upgrades to their Local Area Network infrastructure, and upgrades to their personal computer infrastructures, then turned around and asked themselves just six or nine months ago, "Did I go in the wrong direction, because of the explosion of Internet and the Intranets that are available for similar types of activities?"

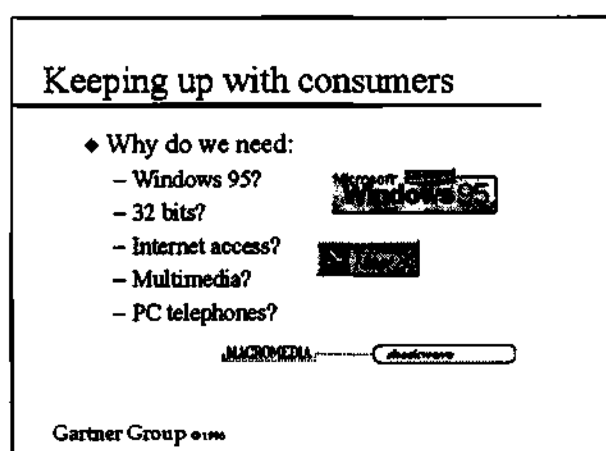

We separate those users into generally two camps. Those who had a business application for which Notes was the right answer, and for which Notes and its capabilities was a good investment. Therefore they are doing something for their users today that they could not have done without an enabler like Notes. Those organizations are doing fabulously for the most part. They are growing, and they are living off of that investment.

Alternately, those organizations who tossed in the technology because it was trendy did not do much with it other than using if for bulletin boards and casual conversation, and have found that they could have gotten away a lot cheaper if they had not reacted so quickly. It is that type of lesson that gives everybody a pattern to match to in the future, asking themselves once again the question, "What am I trying to accomplish? Have I got a business goal that I am trying to achieve with this technology investment, or am I just upgrading because it seems like the right thing to do?" We expect that that will continue.

Why be first?

- CIOs (generally) believe:
  - Early adoption must be "over-justified"
  - Skipping technology generations saves money and lowers risk
  - Release 1.0 of anything will be buggy - even if you call it Release 3.1
  - Only some fads develop into trends

Gartner Group ©1996

The big question of the day revolves Internet and Intranet. In many ways the extraordinary explosion of Internet and Intranet news and technology has been negative for technology adoption. It has slowed people down dramatically. Many existing plans of CIOs have been altered, have been slowed down, have been dropped, and have been changed because of the emergence of Internet. Part of the slowdown that many of us have felt this year is attributed to the boom of the Internet. It has both effects concurrently, both a booming effect as well as a slowing effect. These organizations take a step back.

A lot of money that was going to be targeted for solving some business problems is now being retargeted to solving the infrastructural issues for Internet and Intranets in large organizations. It is a shift of emphasis and a shift of investment.

However, a higher priority still remains, the worst computer virus of all time, the year 2000 problem. As we speak to CIOs, they are now understanding the magnitude of the problem. Many have been through a test cycle where they have looked at their existing IT infrastructure and asked the question methodically, "What will happen to our applications, and what will happen to our business processes, when the clock turns over in 1999 to 2000, or actually before that, and what impact will it have on our business?" The result is so profound, the year 2000 problem is so enormous, the spurious results, system crashes, and unavailabilities that will occur in many businesses and organizations, are so significant that there is nothing that can have a larger priority. Therefore we are now beginning to see, after two, three, four years of talking about it, CIOs shift their investments dramatically towards solving their year 2000 virus problems. This is going to slow down purchases in other areas. It is going to be an inhibitor to some degree, for the next couple of years, something that we have to grow over.

One good news factor here, though, is that the year 2000 problem has received enough publicity now that smart CIOs are able to go to their Presidents, Chairmen, Boards of Directors and say, "You understand the year 2000 problem, don't you?" If they do not, they can bring in outside experts to help them understand the magnitude of the problem. They can help them understand that it is not just their problem but everybody's problem. Nary an organization in the world is exempt from a year 2000 issue. What is happening is we are seeing incremental funding being given to IT buyers for solving the year 2000 problem.

I had one CIO tell me, "I have been working for years to try to get a justification for a new network infrastructure here. Every time I go to the well to get this new network infrastructure funded I don't get anything coming back. I am trying to cobble together systems to meet user demands and no one user project has enough money to justify a new network, there is no funding mechanism inside the company to get all these monies together to support the new network."

This CIO used the year 2000 as a catalyst, an opportunity to visit the highest levels in the corporation and the Board of Directors, and cry, "Fire, something bad is going to happen here, if you don't fund something new, and fund something that's going to fix this problem," and he tucked the funding for the new network infrastructure in the year 2000.

We have surveyed CIOs and asked them publicly, "How many of you have done this?" Publicly none of them agree that they have done it. Privately many say that they are doing this, that they are using the year 2000 awareness to help fund technological infrastructure investments. So while they are fixing the problem they're buying and building some infrastructure, which is a good thing.

Now finally there is two points to be made. When new technology comes along this causes pause and re-evaluation of what is going on. So there is both a good news effect to new technology and a bad news effect, as it relates to the human experience of planning. I am going down one road and I'm going to spend some money, then something comes along that's new and different. I was not aware of it and it is a big discontinuity. It is very exciting. It will cause me to stop and wait and re-evaluate. We would perceive that in the industry as a bad thing, because people are not spending their money. In the end what it does is cause uncertainty and delay. What is necessary to avoid those uncertainties and delays is to provide more of a continuity, as opposed to a discontinuity.

However, there are exceptions and the discontinuity itself provides a clarity of value when the new technologies come along. It is so clear to the user organization what they are going to get, when they know it, when they can see it, and they are going to be able to do something they will jump all over it. I'll go back to my groupware example. There were many organizations in the world, and there are many today, who are disabrogated, who are global and worldwide, and are having an extraordinarily difficult time running meetings, doing data interchange, doing knowledge interchange, for whom the concept of groupware is an extraordinarily clear value. They know what they are going to accomplish. It is very easy for them to see and then they reorient their companies to be able to take advantage of the technology itself.

I will also point to another provider of technology, at one of the higher levels in this chain, at the application level, and that is in the manufacturing space and process control space for SAP Germany – their R3 product. Once again, a product that costs a fortune. Many millions of dollars of implementation go into an SAP R3 project. However the clarity of the value to some organizations is so precise. It is so clear what they are going to get, from being able to run their businesses more efficiently and deal in a process of new manufacturing that they are willing to spend anything to get there quickly. That has been shown to be the case over the last couple of years.

There is also the opportunity to bring forth killer applications, but they come along just every once in a while. We have met a lot of technology providers who hope that their product will enable a killer application, but in the end the killer applications are not so clear as they are after the fact. We are seeing Internet at a business level, business to business, growing. There are very few companies who are not afraid to avoid Internet and be able to see their customers, but to them its a marketing expense.

Internally the Intranet phenomenon is one that is information decimated oriented. You see a tremendous investment today, by organizations who

A Balanced Internet Strategy

- Nearly all enterprises will re-evaluate its technology plans in deference to the influence of internet/intranet technologies
  - Some existing plans will be altered or slowed
- Some additional resources will be put into IT to capitalize on, or exploit internet
  - After Year 2000 costs are budgeted
- CIOs remain concerned about fad, over-investment, and misuse

Gartner Group

recognize that if they have something to say to all of their employees, the easiest way to say it is to publish it on the internal Intranet private home pages of departments, corporations and organizations. This is a very powerful phenomenon. It substitutes a very simple and clean methodology, the home page, for a very difficult problem, internal communications. We expect the vast majority of these organizations to take advantage of that simple and easy method for promoting information throughout their organizations.

Discontinuity

- New technology directions lead to:
  - Re-evaluation
  - Pause
  - Uncertainty
  - Delay
- Rapid adoption requires:
  - Clarity of value, or
  - The killer application

Gartner Group ©

We also expect them to take a big step backwards after that and say, "Okay, how else might I use this?" and not be so aggressive in assuming that their entire IT infrastructure is going to be replaced by Intranet technologies. There are some interesting promises, platform independence, productivity, etc. But we would not expect CIOs, from what we have learned from them over the years, to be very excited about exchanging the technology they have today for something new and different. If they are going to get somewhere from here, it is going to be slowly and gradually, as opposed to quickly and abruptly.

Q: "Are CIOs moving towards fewer vendors in favor of several strategic solution-providers? What are the pros and cons of this trend?

A: Yes. CIOs are moving directly towards fewer vendors. Those vendors who are providing a more integrated solution are doing better. We are not moving back towards the time when there was one supplier for large organizations but we are moving from an era of multiple suppliers to fewer hardware vendors, system software vendors, applications vendors, and services vendors. That is four. It is not 45, as people might have thought it would be two to three years ago. The pros of that trend is management of vendors and management of complexity. The cons are concentration of power in the industry. The users understand the pros and cons but they are less concerned about the concentration of power. What they tell us is they want lots of competition, but they want somebody else to foster the competition. For them, they want simplicity.

Q: Internet access is becoming ubiquitous in the business office at large companies and e-mail is the killer application. What makes you believe CIOs can or should delay this inevitable move to business Internet access?

A: I would agree that e-mail is becoming ubiquitous. However, we have a different view of Internet access at large organizations. We see Internet access at large organizations being tamped down upon. There are many companies who made Internet available to many or all of their PC users, who have backed off tremendously, who refer to it as the cyber water cooler, a place where people can spend time and waste time. They are worried about what are they doing surfing the net, when they are supposed to be getting a job done.

Now, that might be an over-reaction in many cases. However, even Scott McNealy of Sun Microsystems, at a Gartner Group conference last year, said that he shut down external access to the Internet at the end of every quarter, so they can get their books closed. Now if Scott admitted that the Internet was a bit of a productivity drain we would recognize that most CIOs think that the Internet can be a productivity drain. E-mail can be a productivity drain. No one is turning off e-mail. There are lots of heads being scratched about what to do about nuisance and annoying e-mail though, both from the external world and all the SPAMs coming, but also just internally. As e-mail becomes part of culture, how do people spend less time in their day going through what turns out to be less interesting e-mail?

Panel:
Mainstream Multimedia: Which Long Distance Approach Will Win?

Moderator:
Nathan Brookwood

Principal Analyst,
Semiconductor Application Markets Group,
Semiconductors Worldwide,
Dataquest

Mr. Brookwood serves as principal analyst for microprocessor research as a part of Dataquest's Semiconductor Application Markets group. He directs research on microprocessors used in computational 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 a Pentium-based computer with 130 mips of processing power. Before that, he worked for Intergraph Corporation and planned the strategy for their 10.0-mips CLIPPER microprocessors. Mr. Brookwood has also worked for Convergent Technologies, where he directed the marketing of a line of proprietary X86-based workstations, and Prime Computer, a then fast-growing (and now extinct) vendor of 1.0-mips minicomputers.

Mr. Brookwood received a B.S. degree from the Massachusetts Institute of Technology and attended the Program for Management Development at the Harvard Business School.

Panel:
Mainstream Multimedia: Which Long Distance Approach Will Win?

Wes Patterson

President,
Chromatic Research

As current president of Chromatic Research, Mr. Patterson brings to Chromatic over 25 years of industry experience, most recently as executive vice president and chief operating officer of Xilinx Inc. He joined Xilinx in 1985 as senior vice president of Marketing, was promoted to executive vice president of Marketing in 1986, and was appointed chief operating officer in 1990.

Before Xilinx, Mr. Patterson spent four years at VLSI Technology Inc. as director of Standard Cell Operations. Before joining VLSI, he worked in various management and engineering positions at Motorola Semiconductor and Honeywell Information Systems.

Panel:
Mainstream Multimedia: Which Long Distance Approach Will Win?

Prakash Agarwal

President, Cofounder and CEO,
NeoMagic Corporation

Since 1993, Mr. Agarwal has been with NeoMagic Corporation as president, CEO, and cofounder. He has over 17 years of engineering, marketing, and general management experience in the semiconductor industry.

Mr. Agarwal began his career in 1977 at Intermedics Corporation, Freeport, Texas. Other companies where he worked during his career are General Electric in Lanham, Maryland; Storage Technology in Santa Clara, California; Silicon Compilers in San Jose, California; and from 1984 until 1993 was employed as vice president and general manager of Cirrus Logic's Portable Product division in Fremont, California.

Mr. Agarwal holds master's and bachelor's degrees in electrical engineering from the University of Illinois.

Panel:
Mainstream Multimedia: Which Long Distance Approach Will Win?

Carl Stork

Director,
Windows Platform,
Microsoft

As director of the Windows Platform at Microsoft, Mr. Stork is responsible for the development of the Windows operating system platform and for strategic relations with hardware partners for the system software products. His group drives the evolution of the hardware architectures supported by Windows and Windows NT and also develops and implements programs to permit the hardware industry to support the Windows operating system family. Mr. Stork is also responsible for the annual Windows Hardware Engineering Conference.

Mr. Stork has served in a number of positions since joining Microsoft in 1981. Most recently, he served as director of Windows Platform Definition and Business Development. He was responsible for Microsoft's strategic relations with technology partners, including the technical and marketing programs that promote computer platforms for Microsoft Windows. Mr. Stork has also served as director of Windows NT Business Development, director of CD-ROM Marketing, European Business Development manager, and technical assistant to the president.

Mr. Stork holds a master's degree in business administration from the University of Washington and an undergraduate degree in physics from Harvard.

Panel:
Mainstream Multimedia: Which Long Distance Approach Will Win?

Michael A. Aymar

Vice President and General Manager,
Desktop Products Group,
Intel Corporation

Mr. Aymar is vice president and general manager of Intel Corporation's Desktop Products Group.

Mr. Aymar, who joined Intel in 1976, has extensive experience in micro- and minicomputer systems. Before his current assignment, he held positions in system and software engineering, microprocessor marketing, and VLSI design automation, and general management assignments in development systems and mobile personal computers. Before coming to Intel, Mr. Aymar was employed by Hewlett-Packard.

A graduate of Stanford University, Mr. Aymar received both B.S. and M.S. degrees, cum laude, in electrical engineering in 1970.

Panel: Mainstream Multimedia: Which Long Distance Approach Will Win?

Martin Reynolds

Vice President, PC Technology Program,
Computer Systems and Peripherals Worldwide
Dataquest

Mr. Reynolds is the vice president of Dataquest's PC Technology program, which analyzes desktop, notebook, and handheld personal computer products technologies and trends. The analyses focus on product costs and viability of technologies, which are factors that relate directly to the success or failure of products in the market.

Before joining Dataquest, Mr. Reynolds was employed by Computer Intelligence InfoCorp where he developed the Technology Roadmap Services. Before that, Mr. Reynolds spent nine years with Memorex Telex as Director of Engineering, where he was responsible for the design and development of IBM 3270 communications and personal computer products, system architectures, and technological planning. Before joining Memorex Telex, Mr. Reynolds lived and worked in the United Kingdom, where he designed and developed advanced electronic cash registers, wrote software as a consultant in the banking industry, and was involved in the early development of IBM-compatible personal computers. He has been involved with industry groups, including the Video Electronics Standards Association (VESA) and the Micro Channel Developers Association (MCDA), and has served on the Microprocessor Report editorial board.

Mr. Reynolds, originally from London, England, was educated at Oxford University and holds an M.A. degree and a B.A. degree in engineering science.

Session # 19: Mainstream Multimedia: Which Long Distance Approach Will Win?

Moderator:
Nathan Brookwood
Principal Analyst, Semiconductor Application Markets Group, Semiconductors Worldwide, Dataquest

Panelists:
Carl Stork
Director, Windows Platform, Microsoft
Robert Maher
Vice President of Engineering, Cyrix
Wes Patterson
President, Chromatic Research
Prakash Agarwal
President, Co-Founder and CEO, NeoMagic Corp.
Martin Reynolds
Vice President, PC Technology Program, Computer Systems and Peripherals Worldwide, Dataquest

NATHAN: As we go forward over the next couple of years, we have recognized a need for incremental processor performance, especially to handle several new applications – multimedia, 3-D, and so forth. There are a lot of different approaches that vendors have taken to providing those extra cycles.

We are going to try to sort out which of those approaches might have certain advantages in particular environments, and which of them might actually over time be more or less useful.

We have brought together experts in microprocessor design, media processor design, graphics accelerators, software, and PC technology, to help you with that and we will introduce those people as we go on.

Why MMX? Basically I am just going to look at the challenge that the industry faces, and how Intel fits into that, to the extent that the industry and Intel might have slightly different perspectives on it.

The challenge for the PC semiconductor industry is really very straightforward. Dataquest has said: "Go forth and grow at 17-20% per year between now and the year 2000, and if you do not do that you will answer to us." So in order make our forecasts come true, vendors are looking for growth opportunities.

The challenge facing Intel is a little bit different. As the major industry leader with the most R&D dollars to spend, and the most clout to set a lot of standards, they have more responsibility to drive that growth forward. They also have a selfish greedy reason, typically whenever anybody buys a PC, $200 dollars or so of that purchase price goes into Andy Groves' pocket, and he has every intention of keeping it that way.

The other final factor here is that Intel has made a massive capacity bet on the growth in the personal computer industry going forward. Their challenge now, especially as die sizes shrink, is how to keep their fabs full. I think you will see an awful lot of Intel's actions on many different fronts associated with creating demand for PCs because as long as Intel can find – or Intel the X86 community can find – 10 million give or take new buyers every year, then everything goes fine and the semiconductor people will be happy, and all the component suppliers will be happy, and the software suppliers will be happy. If that should stall out for any reason, in other words, you go home at the day and you say, "Gee this machine is fast enough, I do not need a new one," we are all in a lot of trouble.

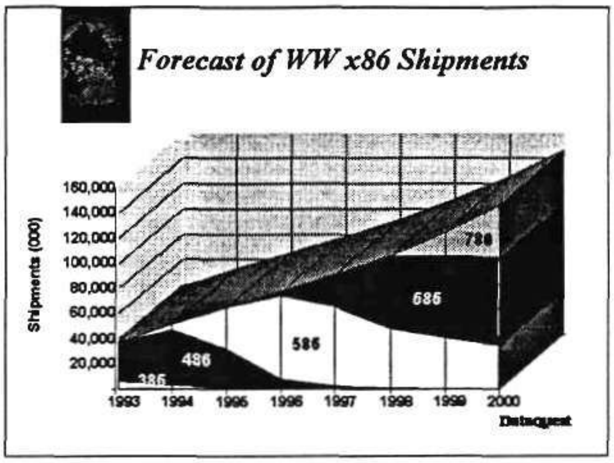

This is the Dataquest forecast for worldwide shipments of X86 over the next five years. We are going to be breaking through 140 million units by the year 2000, to support the growth of PCs over that interval. I have also tried to give you an idea here, of this transition changes from the current Pentium to the Pentium Pro, or sixth generation, and to the Merset, or seventh generation products, over time. Basically in two years the Pentium Pro will be the mainstream processor, or sixth generation products, I should say, in order to make that a cleaner kind of definition. We really do not see the seventh generation having any significant impact in this decade.

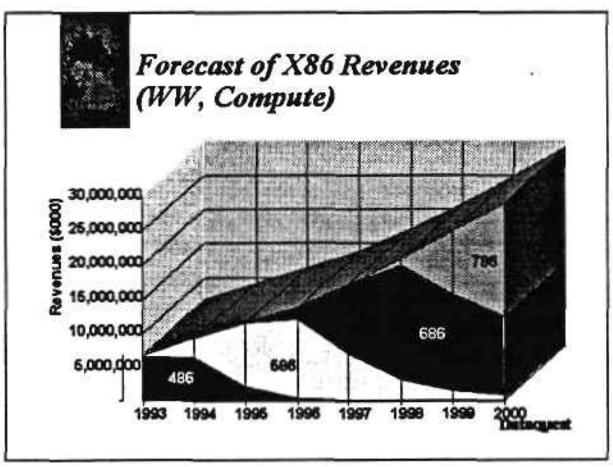

This is on the basis of units, and if you translate that to dollars, you can see we are going to be approaching $30 billion in X86 revenues, by the end of the decade. The total microprocessor market in this period grows to about 33 billion, and there are people in embedded applications, and in risk compute applications, who are fighting over $4 or $5 billion, it is really a very lopsided structure.

One of the reasons is it is real hard to get into the X86 marketing business as several companies have shown.

Intel needs to have lots of new compute-intensive applications show up, that use lots of data and have broad market appeal. Their problem, historically, has been that compute-intensive applications have very limited market appeal.

The candidates for this include 3D games, 3D entertainment in the form of movies, MPEG decoding, video teleconferencing, and then finally, making the machines easier to use, speech interpretation, and so forth.

When they look at these increased requirements, they look at the platform and ask: Where are the roadblocks to going ahead? They try and incorporate features, such as adding things to the silicon so that they can maintain their ASP, and expand and have the products fit a new market. Sometimes they add features that might allow external acceleration, as is the case with a lot of server products. They also look at the technological, architectural, economic, and infrastructure barriers, and try and influence that.

MMX falls into the category of: How can we make the computers better at working in image and audio streams, and video streams, where one of the characteristics is, you are working on a lot of relatively small pieces of data in parallel? You know, every Pentium chip that has ever been built has a lot of area dedicated to 64 bit data paths, to work floating point, and nobody really uses floating point, except for a few scientists and engineers. It is really, and of course it created a little bit of a problem for them, when they discovered that sometimes you do not even get the right answers.

One of the major things that MMX does is enable you to use 64 bit wide data paths for things that might apply to typical applications, like bringing in dual audio streams, or 4 or 8 pixel streams. It is an example of how to make the processor work smarter, rather than harder. We are not quite sure how this is all going to play out, but it does not look like it costs very much from the silicon perspective to add it, so why not.

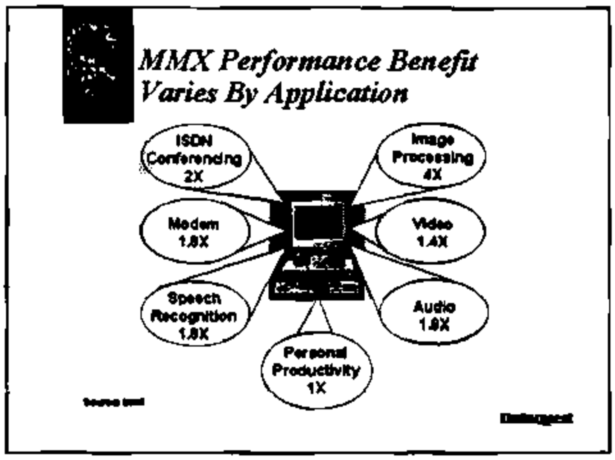

This chart shows you that Intel estimates that the MMX extensions can provide anywhere from a factor of 4 to a factor of 1.2, depending on the nature of the application. (I have added the category of 'Personal Productivity Applications' to this chart). From that standpoint, if you are building a new product that is not oriented to any of these new media streams, then MMX is not going to have any impact on you at all.

The balance sheet for MMX: it brings several advantages over the alternatives including what I would call ubiquity. That is a common architecture that everybody can count on, or ISVs can count on, appearing in all microprocessors, from all vendors, starting in 1997. It will be in Intel's products, AMD's products and Cyrix's products. It is a common target for ISVs to develop and incorporate into new software packages, although clearly, there is still going to be 160 million PCs at the end of this year that will not have MMX, and so ISVs will have to be selective in how they choose to implement that.

The disadvantage is that there are bandwidth issues and data size issues that keep this from being a superb solution across the price performance spectrum. In general we feel it will be better for low-end products, although it is going to be introduced initially in high-end products.

How do we see this playing out? We think it will enable new applications, it will cause ISVs to rethink what they can do, and how to make their applications faster. They will create new products, which will create new demand. Some of those products will undoubtedly fit very well with MMX, and people will go, "Boy, I wish they had thought of that sooner." Others will basically still be too slow, and will create opportunities for external hardware accelerators and therefore, opportunities for media processors, and PGA accelerators, and so forth, going forward. There will be a lot of opportunities created by MMX, and by the new facilities it introduces into the X86 architecture.

Our next speaker is Carl Stork. Carl has been at Microsoft for 15 years and is the director of hardware platforms for Microsoft.

CARL: Well Nathan, I will be sure to point that out to Bill. I know he will appreciate the input. I knew that when I came to the Dataquest semiconductor conference that it would be a learning experience. I am not really a hardware guy, but what I have learned the most is what you really think about Microsoft. It is enough to make you perhaps think, "Oh boy, am I at the wrong company," but I can tell you, having worked at Microsoft for a long time, while it is certainly a company full of smart and ambitious people, it is people that really are trying to do the right thing, and improve computing. One does not like to hear it all the time, but we go on, and hopefully continue to build better software and better products.

One other thing that I just wanted to point out, I actually read the title of the presentation in the notebooks and on the agenda, and it said: "Mainstream Multimedia: Which long distance approach will win?" I could not figure out if this was about, Sprint versus AT&T, or what exactly that question meant, and we will see if any of the panelists can answer that for me. I am going to talk about multimedia and the operating system and the issue of MMX versus media processors versus fixed function accelerators, how I see that playing out. I have a lot of slides and we will go through most of them very fast. I will be taking a little bit more of a software view. I handed out copies of some of the slides, to some of the people in the room, and there are more in the box here, if you did not get one and want need them for any reason.

First of all, I do want to sort of mention the history, if you look at what has happened in PCs

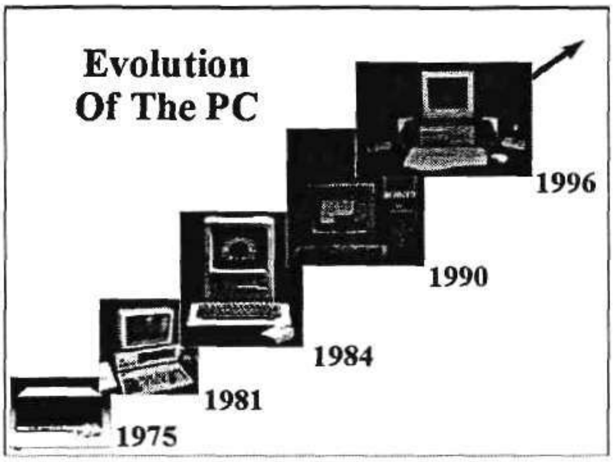

over the last, depending on when you want to begin counting, 20 years or 15 years, I certainly remember my first IBM PC with 16K memory, and a cassette recorder, and a monochrome non-graphic display. Over the course of the last 15 years, PCs clearly have become much more multimedia. Color is a standard part of PCs now, graphics have at least a VGA resolution level was a standard part of PCs, things like CD ROM drives and audio have become very high volume as well.

Multimedia is becoming an expected part of the PC platform. Initially PCs were used for information processing, for doing word processing, spreadsheets, databases. Over time, PCs have become much more interactive, and even information consumptive, meaning that you sit back and watch something, you sit back and just experience something. The software, the hardware, our whole model of working with PCs, really has been focused on the data processing or information creation side of things, up until fairly recently. The area where we have a lot more work to do, is in the areas of communication, interactivity, and entertainment. Multimedia clearly plays a big role there.

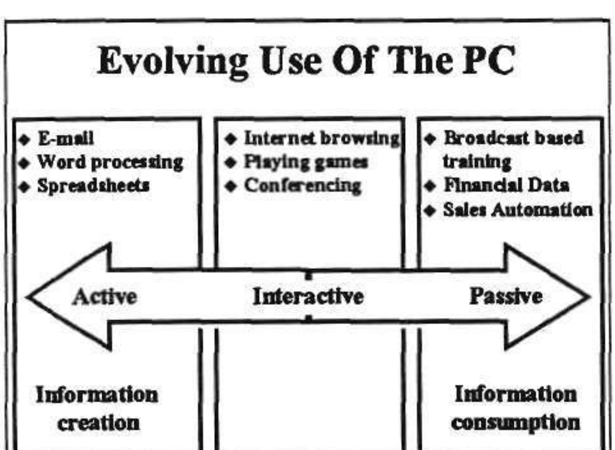

Another trend that is happening outside of the PC space, is that more and more media are becoming digital in the consumer or electronic space. That started with audio CDs, although while the way audio information is encoded on a CD is digital, up until recently, they have been analog devices in most people's entertainment centers. The device itself is digital, but it had an analog audio output, and it really was an analog component for all intents and purposes. That is starting to change now, as we see connections like toss link, and SPDUF appearing on high end CD drives, so that the audio information is now leaving a CD player or video disk also in digital form, and is being processed in higher quality audio components.

This trend is clearly continuing with devices like DVD, DSS broadcast, which is digital obviously, and then in Japan, these DVC (which stands for Digital Video Cassette) camcorders and cameras now exist that use the 1394 interconnect. We are on the verge of a lot more digital multimedia in the consumer electronics space, and I think that is going to open up a lot more interaction between consumer electronics and PCs.

Microsoft launched an effort that we call Simply Interactive PC, or SIPC, which has the goal of making PCs become great for entertainment and communications. There are a lot of elements to this including driving ease of use, and an appliance experience, getting rid of the whole process of booting and making these machines always on. The WDM driver architecture, I do not think I am going to spend too much time on here; digital pipes, meaning a way for it to get digital information in and out of PCs easily. USB is a fairly low speed digital pipe, USB is the Universal Serial Bus.

The IEEE 1394 standard is a very high speed digital pipe, that in today's version of the standard, goes up to about 400 megabits per second and there is work on extending that well above a gigabit per second. That kind of bandwidth is go-

Panel:
Sync or Swim: The DRAM Bandwidth Challenge

Neal Margulis

Senior Vice President,
Research and Technology,
S3 Incorporated

Mr. Margulis is senior vice president of Research and Technology. He joined the company in December 1989 and has held various engineering and marketing positions. He was appointed vice president, Business Development in October 1994; vice president, Home and Advanced Systems Products in January 1995; and senior vice president, Research and Technology in March 1996.

Before joining the company, Mr. Margulis held various positions at Intel Corporation, a semiconductor and system manufacturer, in the Microprocessor Group. His positions included design engineering and most recently chief applications engineer for high-performance processors.

Mr. Margulis earned a B.S.E.E. degree from the University of Vermont.

Panel:
Sync or Swim: The DRAM Bandwidth Challenge

Moderator:
Jim Handy

Director and Principal Analyst,
Semiconductor Memories Worldwide Program,
Semiconductors Worldwide,
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 B.S.E.E. degree from Georgia Tech.

clock rates, their caching strategies will improve, their pipelines are being tweaked even as we speak. So if you got a processor and never touched the MMX features, you would still see overall performance improvements of 10 to 15% in many cases.

CARL: It is my expectation that the marketing messages around MMX will not be very clear to end users. The message will be, 'MMX is better than non-MMX. You want it.'

Q: How does your upside-down PC differ from a Sega/Nintendo or other advanced game station?

MARTIN: The fundamental premise is that it includes basic X86 compatibility. That means that all those X86 games that are out there, and things like Microsoft Works, are available to run on this system, so fundamentally, it is much much more compatible than any of those products.

Q: But from an architectural standpoint in term of the way the graphics and sound are implemented?

MARTIN: It is a PC. It is a small subset of a PC, with a much smaller processor.

Q: The upside-down PC?

MARTIN: Yes. Big multimedia, small processor.

Q: Can you describe Chromatic's license and manufacturing strategy?

WES: Our idea has been to focus our efforts in the areas where we can really add value with architecture and software. So we have licensed the chips to LG and to Toshiba – they are building the chips and selling them. That gives us access to their best process technology and the lowest cost. I think it is what Martin was saying about 3D, that when you are pushing performance like this, we need the best process technology we can get. It is typically not available from foundries, so we have to make our living by providing the software that makes the media processor do its functions.

Q: The decision making process for multimedia peripherals of PC manufacturers is segmented. Audio, graphics, modem, each occur in different groups. Given that the decision making process is fragmented, how can you convince a PC system vendor to use a media processor?

WES: That is a real issue, and it involves a lot of extra work in the early stages, it involves a lot of coordination in companies across organizations that are not normally in close communication. It is one of the challenges.

Q: Martin, you deal with a lot of PC system vendors, what is your experience in this regard?

MARTIN: This really has to be rolled up through the product manager, the person in the company who is responsible for the development, implementation, and success of that particular PC product. He still has to go fight all these battles, but that is where you have to concentrate the effort.

PRAKASH: What we see is that there is no issue with the graphics, video, and audio. On the communication side there is a lot of reluctancy to combine that together, because communication is just a pipe. It solves a pipe problem and as new technologies are coming up they get very reluctant to get that integrated as part of the media processor. I don't think that we will see a lot of that in coming years. But communication will be part of the same solution.

MARTIN: The problem with UMA is you end up with all that frame refresh bandwidth polluting your main memory.

ROBERT: There are advances though in memory compression techniques that are being employed and that will really grow over the next couple of years to help augment that bandwidth problem.

WES: It certainly would help but from everything we see bandwidth is the whole issue in high quality displays, high quality multimedia, and we can use every bit of bandwidth we can get, so to think that we are going to share that with the processor, we could also be easily starved for bandwidth. It does not seem to be the right way to get the performance that people are looking for.

Q: One of the arguments for going back to the UMA scheme is that the granularity of memory chips is going to force frame buffers to be ever bigger in order to provide enough bandwidth for access to the frame buffer. Is there any solution in sight for that other than a UMA solution?

PRAKASH: AGP solves part of the problem. Even if you are talking about 3D applications you could have a lot of your texture memory requirements put in the system memory. So the way that we see it, AGP is the best way to do the UMA type of implementation. The way that UMA was defined it was not going to work and I do not think that we are going to see that again.

WES: With the evolution of demands that we see for media processors there is no particular challenge in being able to utilize the 64 meg frame buffer memory and 256 meg is stretching it.

Q: Given the economics of embedded DRAM, which tend to make it fairly expensive compared with commodity DRAMs which we know are becoming very inexpensive, do the benefits of integrating the DRAM with the graphic controller offset the cost penalties that putting those two guys together impose?

PRAKASH: We believe embedded DRAM is a cost comparative solution to non-integrated solutions for many reasons. If you look at the cost performance, or cost per megabyte of memory bandwidth, there is no comparison. More importantly those who own fabs can appreciate that we are talking about five chips. In a two megabyte graphic solution, we are talking about four memory chips right now and one controller. At least 30% of any chip is the overhead of the I/O padring. So you have 30% overhead in five chips versus 30% overhead in a single chip. If you add the silicon area used in every wafer you will see that the embedded DRAM approach will be a far more cost effective solution than processing five chips. There is no comparison there.

Q: Isn't MMX good for anything?

ROBERT: The experience we have had at Cyrix is with the MPEG-2 algorithm which is part of the DVD specification and without MMX we would not be able to do that. I believe that is true for Intel as well. Even in their climate and the future products they are coming out with in 1997, they would not be able to do that type of processing without MMX. It has some great benefits for parallel processing in terms of audio and video.

CARL: The types of instructions we are talking about are the SIMD instructions (which stands for Single Instruction Multiple Data), which are very specialized and you have to have specialized algorithms that that will apply to. It is not going to improve the performance of an operating system or word processing document or spreadsheet. So there will be certain classes of algorithms that will benefit fairly substantially, which is the numbers that Intel showed you, and those are probably the highest payback uses.

Q: Does this mean that Office '97 is not going to support MMX?

CARL: I cannot see how, for most mainstream things like PowerPoint, Excel, or Word, MMX is going to make much of a difference. However, the MMX enhanced processors will run Office faster than the non-MMX enhanced processors, because of the other improvements in the processors.

Q: Clearly what is going to happen is processors are going to get better on many dimensions over the next two years. One of those dimensions is MMX, and if you use that you might see benefits in selected areas. They will also get higher in of 1998, and for most mainstream software applications, it has to work on the installed base. You cannot target a very small subset of the base.

Q: Intel, at one point, had talked about where MMX gets introduced, the bottom end or at the high end. I think that the decision is coming in, that it is being introduced at the high end, as a performance extension at the top. Initially it will be a high end feature, not a low end feature, which is a little bit strange because the multimedia functions that MMX enables, will not necessarily be higher performance than hardware that MMX might displace in those systems.

MMX will increase total system performance because there are some architectural improvements in the CPU, so the operating system and the word processor, and so on, will run faster as well.

Q: What is the impact of multimedia on system memory requirements in terms of size, bandwidth, technology, and that concept called unified memory architectures?

PRAKASH: Like I talked about you have text, graphics, video, audio. So far, all of that media has been sitting with different memories, processed by different controllers, and is accessed independently. As they start integrating it, you need to process and access at the same time, whether it is done at the CPU, the controller level, or processor. You need very low latency, high memory bandwidth buses to do that. An embedded DRAM type of approach allows you to do that. It is very critical that your memory architecture supports that.

Q: Do you think as these applications get bigger that they will require more memory than you can get on a single chip?

PRAKASH: Memory technology is going from 4 meg to 16 meg to 64 meg up to 256 mega bursts sooner than we imagine.

Q: Will you be able to put all the logic that you want on there?

PRAKASH: Absolutely.

MARTIN: One of the things that we see coming, perhaps by the end of the decade, are LCD panels with resolutions of perhaps 3000 by 2400. That is about 8 million pixels with 3 bytes per pixels. We are looking at about 24 megabytes in the frame buffer. Are you sure about handling those, Prakash?

PRAKASH: 24 megabytes of what?

MARTIN: Of memory in the frame buffer.

PRAKASH: What type of panel are you talking about?

MARTIN: For high resolution 3000 by 2400 LCD panels

PRAKASH: We are talking about notebooks here and PCs.

MARTIN: I am talking about PCs, yes. We will see that.

PRAKASH: You have to look at the mainstream applications. Right now portables are 800 by 600. In '97/'98 they will be 1024 by 768 and the focus would be about 16 bit per pixel, 64,000 colors. Beyond that we will have 1280 by 1024. If we are really talking about 3000 by 2000, you probably will need some different technology.

MARTIN: There are alternate solutions, you can work the two together.

PRAKASH: We are talking about mainstream multimedia here, not something which Hollywood or somebody else will use.

MARTIN: That is right and in 1990 if you told me you were going to have 32 megabytes in your PC in 1996, everybody would have said no.

Q: You have all ducked the issue of unified memory architecture. Are we going to see that rear its ugly head again?

CARL: The AGP, the advanced graphics port that Intel has announced, may open up an opportunity to have a shared memory architecture between devices on the AGP port and system memory. UMA as it was defined by the VESA committee we do not expect to see.

It is easy to make us think that perhaps everybody wants this, when in fact they do not.

We do not see interactive television as something that is compelling in an entertainment sense. That is where the upside-down PC comes in. This is a device that we see having more multimedia silicon, if you like, than processor silicon. An Intel-type processor, with a multimedia processor from Chromatic, or Samsung, or Philips, that drives things really aimed at entertainment and not productivity. So, HDTV, we think, would be a function of this kind of box. DVD, games (the same games that you can play on your big PC), and some limited Internet. Primarily this is a device for passive viewing. That means it has to cost $500 or less. It has to connect to a TV set, and it probably has to include exotic features, like AC3 sound decoding. We think this could represent the third wave of home computing. This becomes a device like a VCR, that sits in the stereo rack, that drives all these advanced entertainment features. It could be a very high volume application, as we get into the next decade.

It is very difficult to compete in 3D graphics if you do not have a fab. Ultimately, the quality of your 3D graphics depends on how many, and how fast the transistors are that you can deliver. That depends on how much you can charge for the part, therefore, it is ultimately dependent on your cost. We think it is very important to have a fab, or a zero cost adder when you are building these chips, to compete. The biggest challenge is that Intel is coming to this market. Intel is not a great graphics company, they have already failed twice at it. I am not sure they will fail a third time around.

Talisman, Microsoft's new 3D architecture, is exotic to watch and very compelling. It has five very large chips, a media DSP, a compositor, a polygon object processor, a smart DAC, a big buffer, 4 megabytes of DRAM and it is very expensive to build today. We believe that a lot of the ideas from Talisman will find their place in the market, but at the moment it is fundamentally a well supported unknown. With Microsoft behind it, that is a lot of pressure.

Can MMX displace media processors? We think MMX performance will roughly double every year and a half or so. We think media processors can stay ahead because media processors are much easier to scale. When processing these video streams, as a processor, you know where you are, you know where you are going, there is no need to have good predictors in there. You know where your next piece of data is coming from, and you know where the next piece of data is going to. That means it is relatively easy to crank up the clock speed, and as long as you have the bandwidth into the media processor, you can deliver scaleable performance. Ultimately, MMX will win if there is enough performance. If the MMX enhanced X86 processor can do all of these things, then there is no point in having a media processor.

We think 3D can take as much as you can throw at it.Video compression is another major challenge and other applications will keep media processors ahead of the MMX curve. The processor then, has to find something else to do, we think the big one there will be speech recognition. To do speech recognition without an algorithmic breakthrough, you need a bigger processor, bigger hard drive, more memory to make it all work. So we think media processors have a way to go yet.

In the long run, we think that the media processors will ultimately become the video controller of the PC. In this part, we see integrated memory to deliver very high bandwidth. We expect to see AGP, Intel's high speed bus, delivering high bandwidth for large objects, brought from main memory to the graphics subsystem. Integrated DAX for audio and videos becomes a very simple subsystem, and may be the place for 1394. In the long run, I see media processors as the dominant solution for multimedia in the PC.

Q: What benefit is MMX to low end systems, if it comes out first on high end microprocessors?

CARL: Nathan, in his presentation said that one of the advantages of MMX was going to be ubiquity. From the point of view of a software developer, MMX will not have that advantage for quite a few years yet. Because even with a very rapid ramp, in MMX-enabled processors by Intel, at best they might have 10-15% of the installed base by the end of 1997, perhaps 30% by the end

This chart shows how the PC market split in the U.S. in 1995. Of all computers sold, 18% were portables. Those computers go primarily into the business market. That leaves 49% of all computers were desktops, sold into the business environment. The balance therefore, is 33% of all computers sold to the home. One of the things we saw over the last few years, is business buyers, when the economy goes a little sour, do not necessarily back away from buying PCs. They like the productivity, the benefits they bring, and indeed, the PC market has been fairly recession-resistant over the last few years.

On the other hand if home buyers feel insecure about their jobs or their future, when it comes to Christmas, it is not PCs, its ties and slippers again. If the economy turns down, we can expect to see the PC market take a hit, but I cannot predict when that will happen; when it does, we will see this downturn.

Next is the overall size of the PC market. We are a little less aggressive on the content of Pentium class products in 1999. A lot of the reason for that is many of those processors are perhaps not going into what we would consider to be a PC.

Something else that I would like to draw your attention to is universal serial bus. This is a PC technology, any of you designing devices or semiconductors for communication, where you need a simple serial interface that can carry high speed data, need to look very hard at USB. It is going to be very cost effective and absolutely ubiquitous within the next few years.

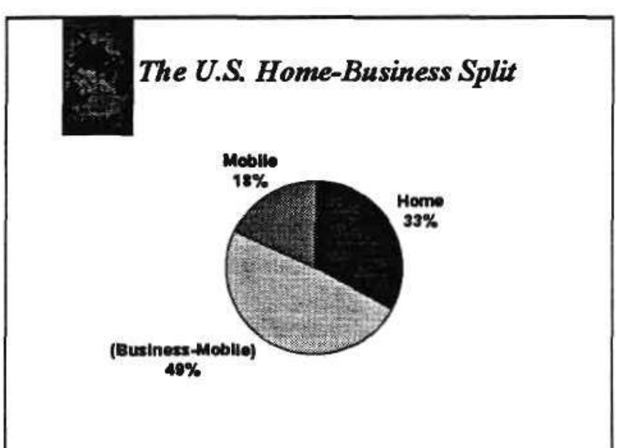

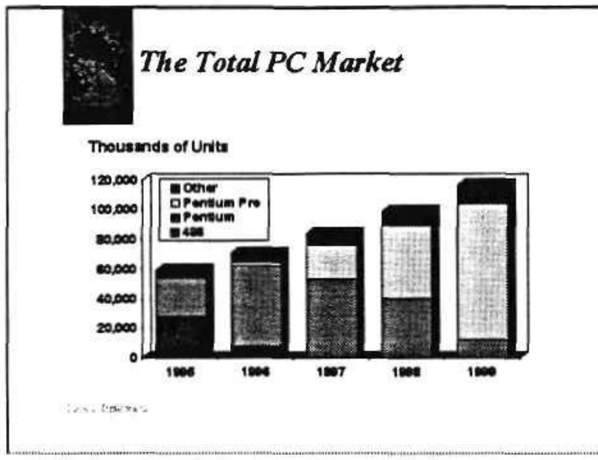

Another question that keeps coming up is the $500 PC. We are not that aggressive on the $500 PC. There are two points there I want to bring up. Users are very much in control. If you, as an IT manager, deploy Internet computers around the office that do not have the features these people expect, six months later, you will see PCs sitting by the network computers that you bought them.

Deploying network computers to productivity workers is fraught with risk. For the home users, people say, "Well, if we could give people a $500 computer to take home, and surf the Internet, then we could get all these people connected, without them having to buy an expensive PC." That is another flawed argument because most people that have a PC at home today are not connected to the Internet. Although there are lots of them, it is still a relatively low penetration. The assumption that we can get more people on the Internet, with an inexpensive device, is unnecessary and probably wrong. We in the PC industry are very blinkered by the fact that the Internet is so wonderful.

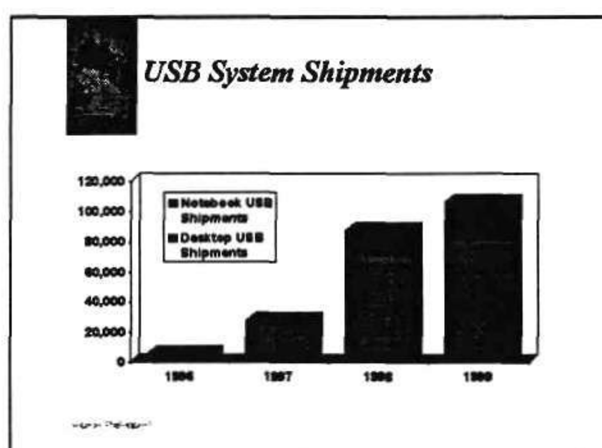

rate, to increase the memory bandwidth. Once you bring the memory inside, you can make that interface as wide as you want, or you can clock it even faster because it is all inside the chip.

It does offer the battery life, especially in CMOS, because most of the power consumed is in the input/output paths. In the multimedia applications, whether the user is doing anything or not, the screen, the display, has to refresh at 75 Hertz per second. That means if you are using 1024 x 768, 8 bit, one megabyte worth of data has to be accessed from the memory, and sent out to display, 75 per second. That is a lot of data being accessed, lots of input/output parcels switching, burning lots of power. If you bring the memory inside, you eliminate those I/O paths, and you can cut down the power.

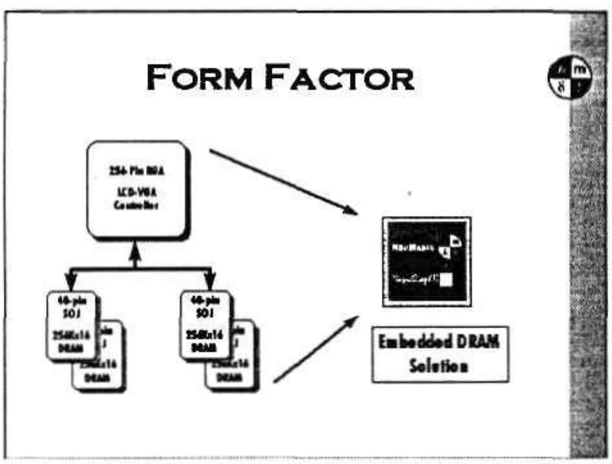

Form factor: any time you take four chips, five chips, or memory chips, and bring them to a single chip, you solve the form factor problem. Cost: you have heard that any time you take multiple chips, and put them in a single chip, it helps on the cost side. Memory comes in configurations of half a megabyte, 1 megabyte, 2 megabytes, or 4 megabytes. If your application only requires 1.1 megabyte then you only integrate 1.1 megabyte in you chip, no more, no less. The simplicity, ease of use, ease of manufacturability, and ease of testability, all translates into time to market.

The key thing here is, you get the benefits of performance and battery life with no compromise. With traditional technology so far, you get performance at the expense of battery life, form factor, or cost. You cannot achieve all those functions at the same time, unless you are using embedded DRAM technology. In our experience, embedded DRAM technology is the only semiconductor technology which can offer you all those benefits at the same time.

On the graphics accelerator side the performance of the competitive solution has been 32 bits wide or 64 bits wide, and that relates to memory bandwidth. NeoMagic MagicGraph is 128 bits wide, which gives much higher performance and the benchmarks at the system level show a difference of anywhere from 45% to 80% higher performance.

Power consumption: the competitive solution takes 2.5 watts of power. Our solution is less than 400 miliwatts. At the system level, that translates into anywhere from one hour to two hours extra battery life which is very significant. Form factor: five chips versus one chip. So, which long distance approach will win? Multimedia accelerators based on embedded DRAM technology.

NATHAN: Our final panelist is Martin Reynolds. Martin is Vice President of Dataquest, and Director of the PC Technology Directions Service, which keeps an eye on what is happening in all of the component technologies that go into personal computers.

MARTIN REYNOLDS: I would like to start off by answering a couple of questions about the PC market that keep popping up. The first one is: When will it slow down? Everybody wants to know when the PC market is going to slow down, and while I cannot predict the future, I have a chart here that might help you understand what the driving factor is going to be.

level of multimedia performance, at the right cost point. That should help to drive some new applications for the PC and help to expand the market.

NATHAN: Our next speaker is Prakash Agarwal. Prakash founded NeoMagic, which is a company that has integrated DRAM and a graphics controller onto a single chip, and has found a home for itself in many portable applications.

PRAKASH AGARWAL: The question is: Which long distance approach will win? I think what we mean by that, talking about mainstream multimedia, is which type of solution will win in the long term? Before we can answer that question, I think that we should answer: What is multimedia?

Multimedia is not application, it is a combination of the following technologies: text, graphics, video, and audio. When you start combining all these different media, that becomes multimedia. The application of these combinations of technologies is presentations, training, entertainment, and education. What is important to notice here is they all require portability.

How often do you get presentations at your desk? You go to conferences, you go to conference rooms, which means that you need a tool or a system that is portable, that you can carry with you. Same thing for entertainment. I do not believe I am going to watch an MPEG movie at my desk. I am probably going to use the TV, VCR, or DVD. But if I have a portable computer, then I can carry a movie, a portable DVD, that I will watch the next time I am going to Japan. Take that 12-hour flight, watch two or three movies. Again, that requires portability.

Education. If students have computers that they can carry to their classrooms they are going to benefit a lot. That also requires portability. I believe multimedia requires a lot of portability. So which solution will win? The one that is simple, and solves the portability problem.

Portable computers have not offered as much performance or features as their desktop counterparts. Their battery life is not as good as is needed and they are too bulky.

We find that consumers are not happy with the performance and features of their notebook computer compared to their desktop computer. The reason they do not bring their notebooks is because they are not happy with the battery life, and with the size of it. I am driving this presentation from an HP notebook that is a 133 MHz Pentium, 800 by 600 TFT, 1.4 gigabytes of hard disk drive, 16 megabytes of memory, 256 kilobytes of cache. With this computer, I can do anything I can do with my desktop computer. As that starts happening, people will have one computer as a primary computer, and if you can bring the multimedia functions into that computer, that is when you will start seeing multimedia become really interesting.

If we look at the overall solution we talked about, Carl gave you the perspective of the operating system, Robert talked about the CPU side of it, Wes talked about the multimedia processor. In my opinion, if you look at that architecture, there will always be a multimedia chip. I do not believe, as multimedia functionality gets more complicated, that the CPU alone would be able to do that. You need an integrated multimedia chip, where you know you have access to memory, and can do a lot of those things.

It has to offer the performance, it has to offer the battery life, form factor, cost, and more important, simplicity. The reason a lot of people have problems, integrating graphic, video, audio, and 3D is because, right now, it is being offered by different vendors, different software drivers. When you start looking at plug and play, it becomes a big problem. Once you start integrating them together, whether it is a multimedia processor, or multimedia accelerator, you are solving that simplicity problem.

Embedded DRAM technology, which NeoMagic has pioneered, does offer performance because the biggest problem in dealing with the multimedia applications is memory bandwidth. Once you bring the memory inside, you are no longer restricted by how many pins you need to dedicate on your chip to talk to memory. People have gone from 8 bits to 16 bits to 32 bits to 64 bits, or they have gone from 50 MHz to 250 MHz on the clock done up to now clearly has an advantage. They are high performance, they can do high concurrency, but they also have the highest cost. They have the highest obsolescence rate because they are fixed in functions, and the cost of integrating these various function with the hardware and the drivers, are non-trivial. In fact, some of the customers that we are working with have found this to be a tremendous obstacle in putting a lot of multimedia functions into a box, and making them cooperate and work together.

There are other bottlenecks in the system, beyond what the chip itself can do. The PCI bus can become something of a bottleneck, drivers can bump into each other and cause conflicts, and the lack of a real time operating system on the host can also be an issue.

Media processors are, as the name implies, processors. They are programmable, in our vision they are dedicated to PCs, so they are capable of dealing with the legacy issues, graphics, and audio, for example. They are very high performance, they have to be able to deliver multimedia functions concurrently, at very high quality. They are not DSPs, although there are a number of DSP-like architectural features. They are optimized for the PC. They work on the PCI bus, they deal with these legacy issues, as we said.

They are programmable, so they are not just a collection of multifunction ASICs put on the same piece of silicon. In our view, they are certainly not general purpose CPUs. The chip we have designed was developed to work in cooperation with the X86. A lot of the software that we are writing runs over on the host processor because, frankly, it is done more efficiently there.

The advantages of media processors are they are the lowest cost solution to deliver the full range of multimedia functions. They are also very high performance, and you can see in this graph that the performance is well above what you can realize today or in the next generation, from host-based cycles, or from DSPs. Because they are programmable, they are very flexible; that means that we can respond to changes in evolution – 56K modems is one example. It also allows the users of these devices to differentiate their products. Because we have a software-based solution, we can do a fair amount of load balancing between what takes place on the media processor, and what takes place on the host. If the host is MMX equipped, then we can use those resources to further improve the concurrency or the quality. The biggest issues for media processors are the same ones that affect any new technology and it takes a lot of software to provide the critical mass. We do not have a media processor solution until we have 2D graphics, 3D graphics, the basic audio functions, the modem, and MPEG video.

The requirements for multimedia include programmability because that is how we get the flexibility that we need, and they also include high performance and low cost. Media processors are the best solution to this set of constraints. The size of the media processor market is expected to be about 9.5 million units next year, growing to over 55 million by the year 2000. We are the first to market, but we are not going to have this market to ourselves. NEC, Philips, and Samsung have all announced products in this market space.

We think that media processors, working together with the host CPUs, are going to deliver the best

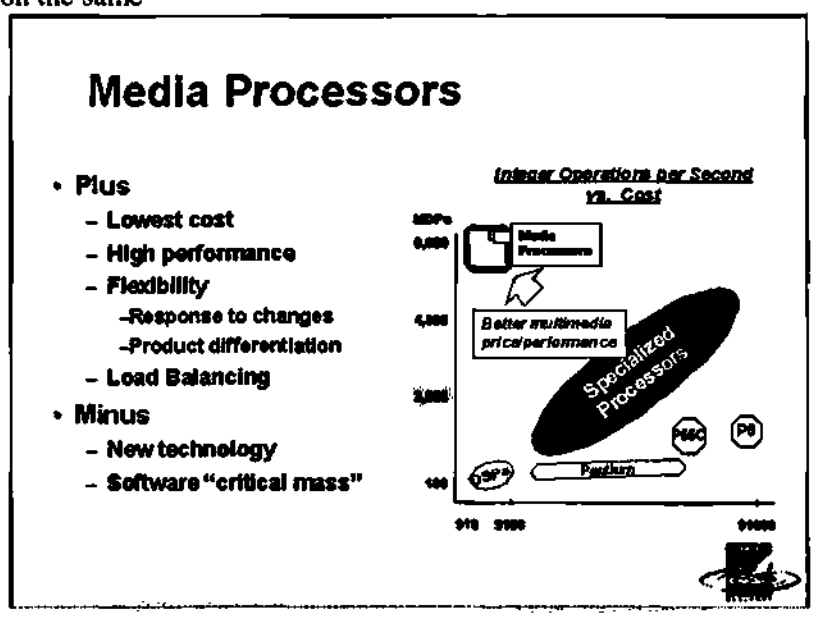

that concept, or working with its partners to ship that concept, a few months ago.

WES PATTERSON: The chart about the evolution of the PC was pretty well covered in Carl's presentation. The thing that is important for the opportunity we are discussing here is that this change is in the function of the PC, from being in effect a computer, to being something that is looked at as a source of entertainment and communications, is that the data types are changing. In my experience, there is really no precedent for this. The computer industry has evolved over a long period of time, there have been changes in performance and architecture, but they have all been aimed at, fundamentally, the same data types. When we move to entertainment, there is a new kind of data, and that makes newer architectures appropriate because there are better ways to do computing on video, graphics, and audio data types.

Obviously, Intel recognizes this. This is their chart, they show the split in computing demands between so called "conventional computing" shown in the lower bars, and the upper part of these bars show their estimates of multimedia demands. If high quality DVD is a reality next year in the market, then the multimedia demands will be far in excess of what you see on this chart.

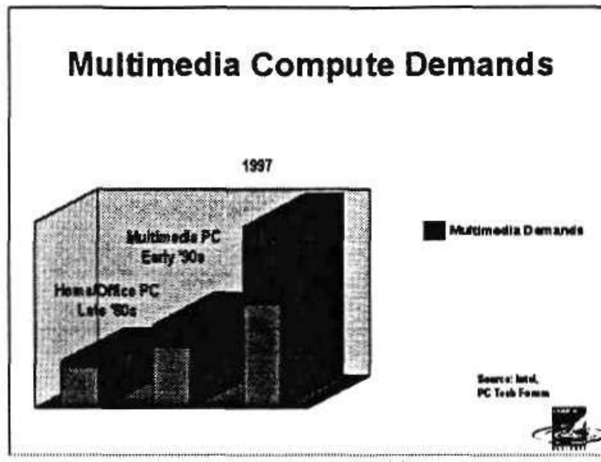

There is really no upper bound for multimedia computing requirements because no processing capability translates to better quality. Whether it is audio, video, graphics, or communications, we are going to see a steady pressure to put more computing cycles on these problems. For audio, video, and graphics, that will show up as better quality. For communications, it will show up as higher bandwidth.

We do not think that we are anywhere near a point of saturating, or completely meeting the needs of the user for computing cycles. We can see a requirement, looking forward just a few years, for something in excess of 10 thousand MOPS, to solve all these problems with reasonable degrees of concurrency.

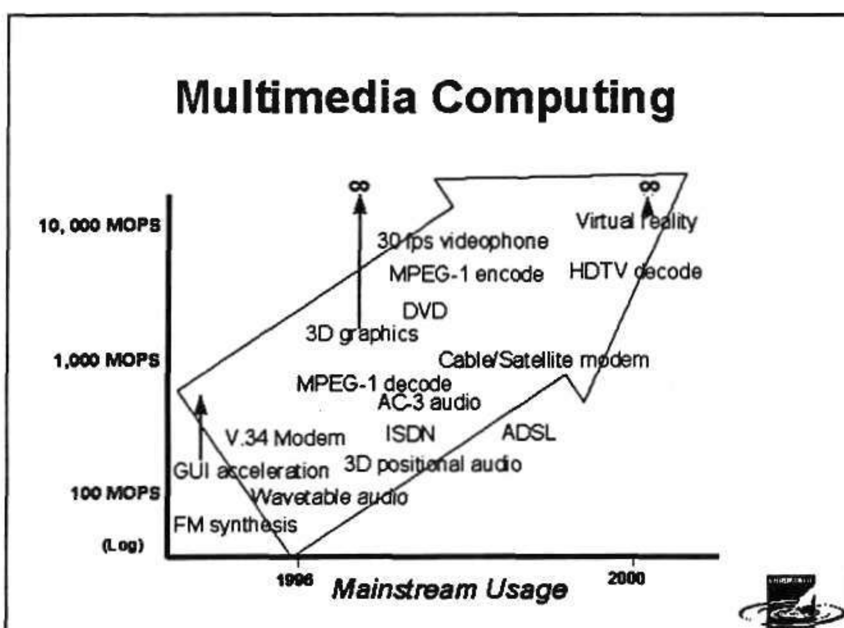

In the PC, there is going to be traditional computing, and Intel and others seem to own that space pretty clearly. For the multimedia portion, there are at least these three alternatives. We have talked a little bit about what Cyrix, and Intel, and AMD are doing with extended instruction CPUs, there are dedicated peripherals in their media processors. Some multimedia functions require real time response, and the host side solutions do not have good access to a real time operating system.

Dedicated peripherals – the conventional way that this has been

GX processor, where we integrated the memory controller. With that memory controller onboard, we have bandwidth as good as 3D graphic cards.

NATHAN: So, you think that bus, that very wide bus, that connects the CPU and the main memory is also a very busy bus, and it is got a lot of traffic that has nothing to do with multimedia going across it. Wouldn't it be better, if we could take the stuff that is easy to offload from that bus, and put it on something separate, like a media processor?

ROBERT: If you needed to, yes, absolutely. But with some of these new applications that are coming out, with the bandwidth that is capable in these processors, we can still meet all the requirements to do things like MPEG-2 decoding.

Obviously there are applications that consume the CPU, there are not enough cycles available, memory bandwidth is a problem. In the future, as those applications become more mainstream, CPU architectures are going to be focused, they are going to be looking at addressing that. So, I think that the bandwidth issues that you see will be alleviated with some new architectural enhancements. There will be new media processing capable engines added to the processors to address that.

The CPUs today are typically 6 to 8 million transistors. Over time that is going to grow to huge numbers, and the frequency is going to go up, and that will give us more capability to address, to do multimedia processing, as well as have cycles left over to do other things. If an application or number of applications outpace what a CPU can do, then there is a real need for a media processor to help out.

NATHAN: Do you have any theories as to when the processors will be fast enough?

ROBERT: Today, with the applications that are out there, there are not a lot of problems, other than in 3D, where there is not enough cycles left over to do the job.

NATHAN: Now, one of the interesting issues with MMX is that Intel initially tried positioning this as an alternative for low end 3D. From what you know of having implemented MMX, do you see it providing a big assist there?

ROBERT: In terms of 3D, I do not see it as being a big assist. The large problem there is there's a pipeline involved with 3D. There is a floating point part, and there is a rendering part. The floating point part is typically done by the host processor, but in order to do the floating point, as well as the rendering both by the processor, that is a real tough job. So I certainly see in the 3D arena, where there is clearly an insatiable appetite for more processing cycles in that rendering stage, where that media processing capability will most likely be in a separate unit, maybe external to the processor, maybe not.

NATHAN: So fundamentally, if we had kept the 3D story out of MMX at the beginning, it might have been a slightly cleaner story, not that it is Cyrix's fault.

ROBERT: Yes, that is Cyrix's position, right now.

NATHAN: One final question for you, Robert. You have implemented both the MMX type extensions and taken more liberties with microprocessor architecture, adding features to do graphics and sound, in a slightly more direct fashion in your GX product. Between those two approaches, which one do you think has more merit from a technical standpoint?

ROBERT: It really depends on the market you are going after. I really see the two as being quite separate, both have obvious benefits and disadvantages. We solved some of the bandwidth problems that are inherent with the host CPU doing a lot of the work, with our GX processor, integrating that memory controller. With MMX, we clearly get a higher computational throughput, to do the computational algorithms, that are part of those media applications. Combined together that could be a real powerful product.

NATHAN: Our next speaker is Wes Patterson. Wes is the President and Founder of Chromatic Research Company, that introduced last year, the concept of a media processor, and began shipping ware implementation between different PCs. There is also a hardware abstraction layer that lets a given piece of software run on multiple platforms.

DirectX queries the hardware to find out what the hardware can accelerate, and if it can, it makes use of the hardware function. If it cannot, then DirectX emulates the hardware function and software. These DirectX emulation pieces can be written to take advantage of things like the MMX multimedia extensions that Intel is introducing next year.

The question that I am trying to answer is: What do ISVs want? If ISVs select a certain approach and write applications that are dependent on a certain approach, that will almost certainly determine which approach wins. What ISVs really want is a common platform. They want a single interface, a single view of the system that they can write to, and not have to rewrite their application or their game for every single platform. They definitely want rich, pervasive multimedia functions. They want high level APIs, they want hardware independence. They also want the ability to understand and adapt to the performance of a given platform. I want to be able to write my game so it looks decent on whatever level of performance you have in your system. Now, that may mean that in order for the playability and the reactivity of the game to be good, I need to have much simpler 3-D graphics, I can throw out some shading, I can throw out detail, there is a lot of things I can throw out, or maybe I do not do as good audio or I do not do any audio, if you do not have a certain level of capability.

I can adapt what my application or my game does, based on the capabilities of the hardware. ISVs also want enough CPU cycles left, on the main CPU, so the game is playable. They do not want to use the majority of the CPU cycles to do audio, or 3D rendering, or MPEG decompression, or modem function, if they want to be able to have a good gaming experience.

Lastly, ISVs do not want to place a bet among the different approaches. In most cases ISVs will be working through libraries and other software packages that actually implement functions. Then ultimately, it will be up to customers to decide what combination of hardware capabilities, price, and so on, sell the most.

NATHAN: Our next speaker is Robert Maher from Cyrix Corporation. Robert is the VP of Engineering at Cyrix, he is also the designer of the MMX extensions that Cyrix is including in the next generation processor that they intend to launch early next year.

My first question is: How do you see MMX fitting in, overall?

ROBERT: In terms of MMX, and the advantages it gives, it does provide an added benefit for the PC platform. Number one, it provides a baseline capability for the PC to do some of these multimedia applications.

Without MMX on some of the next generation or the processors that are going to coming out in 1997, we would not be able to do MPEG-2. Sound-blaster type audio is easily handled with MMX instructions. There are a few applications like video encoding and 3D audio that it does not cover, but it really does take us to another level.

NATHAN: You were careful to say 'sound-blaster' type audio, is that to suggest FM as opposed to wave table?

ROBERT: Right now, FM is the mainstream, but it will help with wave table as well. Wave table is more memory dominant, it is going to be more memory intensive.

NATHAN: One of the big problems for multimedia solutions is bandwidth, and bandwidth between the processor and the memory system. One of the arguments for putting this stuff outboard is, you can get a lot more bandwidth between an outboard processor, and an outboard accelerator, and button memory than you can between the CPU and the main memory. How do answer that?

ROBERT: I do not necessarily agree with you. The largest bandwidth, the biggest bus on the personal computer is the bus that interfaces to the processor. There are issues with latency, which is a big problem, but we have solved that with our

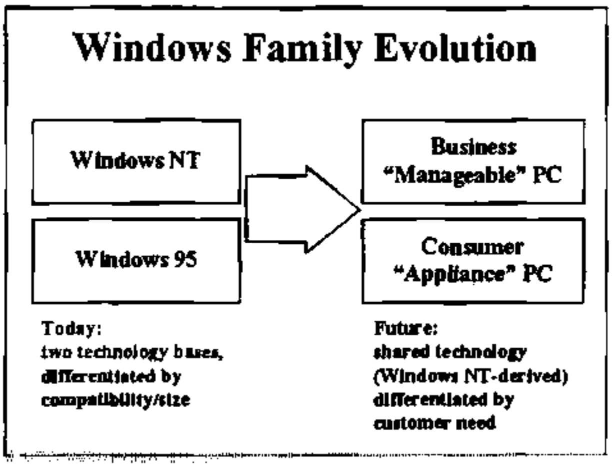

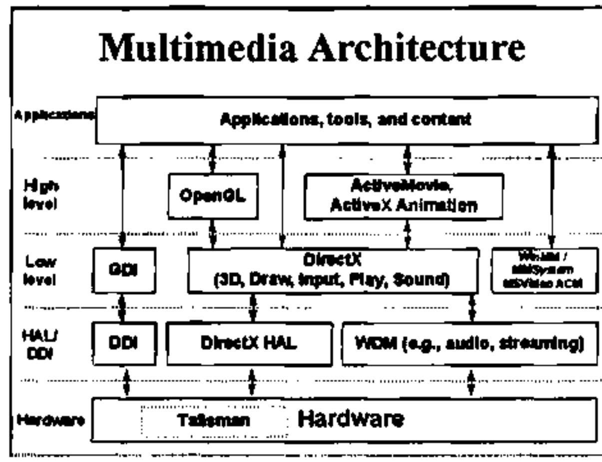

ing to be needed to get in and out of PCs. Higher quality audio and video: audio quality in PCs today is poor, the signal to noise ratio, the speakers, etc. We see that as an area where the audio processing is likely, over time, to move outside of the physical PC box, and the PC is likely to emit audio in digital form to another device that does the processing, where there is cleaner power.

I threw in one obligatory Windows family slide, just to say, " Hey, we have two members of the Windows family, as far as desktop systems go, Windows '95 and Windows NT," and we see two versions of Windows extending out for some period of time, with the Windows '95 stream aimed at consumer small office, home office kinds of uses, and Windows NT aimed at corporate and large organization uses.

We will, over time, share more and more technology between the two operating systems and we will be taking more technology out of Windows NT into the Windows '95 product stream. In the long term, there is a technology convergence, but for the next several years, Windows '95 is not the end of life for Windows, and Windows NT is not going to take everything over in the next year or two years. These two products will coexist in the marketplace.

What are all the different multimedia software pieces? For Microsoft, they fit into four major groupings. The DirectX, APIs and driver interfaces, which provide relatively low level abstraction of the multimedia capabilities; Open GL, which is a high-level 3D modeling language; Active Movie and ActiveX animation, which deal with media streams, and then some research investments that have had the code name of Talisman, all built around the Windows base. The key thing that most ISVs are going to look at when developing titles that are multimedia is: What are the APIs, and what are the standard interfaces that I write to.

This chart attempts to show the multimedia architecture that is part of the Windows family (both Windows '95 and Windows NT) with a hardware layer, with hardware interfaces, DDIs, the "hows" that are part of Direct X, and the interfaces that the Win32 driver model will provide. Above that there are low level APIs including the Win 32 GDI, the DirectX APIs, and then the Win MM subsystem. Some applications will talk at this level, others will use higher level interfaces including things like Open GL and Active Movie. The bottom line is that, in general, most Windows software does not talk directly to the hardware; most applications talk through APIs that abstract the actual hardware.

Active Movie is a media streaming API that is useful for things like video capture, video playback – it is a configurable connection architecture. The first version of Active Movie ships with Windows '95, service release two ships with Internet Explorer, and it is being incorporated in Windows NT as well.

DirectX is a low level API for games and things of that nature that abstracts the hardware at a relatively low level. It provides a consistent set of multimedia functions, so that applications can be written to a single APR regardless of the hard-

Panel:
Sync or Swim: The DRAM Bandwidth Challenge

Avo Kanadjian

Vice President,
Marketing, Memory Products,
Samsung

As vice president of Marketing for Memory Products, Mr. Kanadjian is responsible for marketing Samsung's extensive lines of high-volume memory products, including static and dynamic RAMs. He joined Samsung in July 1996 with 13 years' experience in memory marketing.

Mr. Kanadjian held key positions with Toshiba America and was previously the director of Memory Marketing. He is considered to be an industry architect of innovative strategy. Before Toshiba, Mr. Kanadjian held a position with Xerox as Memory staff engineer.

Mr. Kanadjian did his undergraduate studies at Academie de Lyon, holds a B.S. degree in electrical engineering, cum laude, from the University of California at Irvine, and is fluent in French and Spanish.

Panel:
Sync or Swim: The DRAM Bandwidth Challenge

Tom Dye

Director, Austin Design Center, Chief 3-D Graphics Architect, Cirrus Logic Inc.

Mr. Dye is the director of Cirrus Logic's Austin Design Center and is the company's chief 3-D graphics architect. Along with managing the Austin Design Center, his primary responsibility is to oversee the development of innovative 3-D architectures and algorithms for PCs.

Mr. Dye joined Cirrus Logic in early 1994 when the company acquired technology developed by his start-up company, Design Technologies Inc. Since joining Cirrus Logic, he has been instrumental in bringing to market products based on the company's Laguna graphics platform, including the new Laguna3D multimedia accelerator.

Before joining Cirrus Logic, Mr. Dye was director of multimedia at Dell Computer. Other positions in his career include director of Advanced Technology for Nth Graphics of Austin, Texas; graphics product manager for Motorola Corp.'s Computer Systems Group; and graphics section manager for Texas Instruments.

Mr. Dye earned a B.S.E.E. degree from Michigan State University.

Panel:
Sync or Swim: The DRAM Bandwidth Challenge

Michael Nielsen, Ph.D.

Director of Engineering and Chief Engineer,
Digital Media Systems,
Silicon Graphics

Dr. Nielsen is director of engineering and chief engineer at Silicon Graphics. He was responsible for system architecture and product definition of the Indy-follow-on workstation, and was also chief designer of the high-performance unified memory system for this workstation. Before this he was involved with flat panel interfacing to SGI workstations.

At MIPS, Dr. Nielsen was director of architecture and a MIPS Fellow responsible for microprocessor and system architecture. At Digital Equipment Corporation, as a consulting engineer, he was chief architect and designer of the DECstation series of MIPS-based workstations.

Dr. Nielsen has Ph.D., M.S., and B.S.E.E. degrees from Stanford.

Session # 20: Sync or Swim: The DRAM Bandwidth Challenge

Moderator:
Jim Handy
Director and Principal Analyst, Semiconductor Memories Program, Semiconductors Worldwide, Dataquest

Panelists:
Avo Kanadjian
Vice President of Marketing for Memory Products, Samsung
Neal Margulis
Senior Vice President, Research and Technology, S3 Incorporated
Tom Dye
Director, Austin Design Center, Chief 3-D Graphics Architect, Cirrus Logic
Dr. Michael Nielsen
Director of Engineering and Chief Engineer, Digital Media Systems, Silicon Graphics

JIM HANDY: I'm Jim Handy and I am hosting this panel called "Sync or Swim: DRAM Bandwidth Challenge." We are going to be talking about the need for improved bandwidth in DRAMs.

With DRAM bandwidth, what kind of speed do you really require? Some needs in DRAMs are very random. Cache line refills and general CPU accesses require very low latency. High bandwidth requirements and little concern about latency apply to some graphics operations and video output. Different types of DRAM architectures satisfy different sides of that. Random addresses need very fast first accesses, streaming does well for video output, but there is no such thing as a solution for all kinds of problems.

We see a lot of system approaches for getting the bandwidth out of the system that can take advantage of lower costs, lower DRAMs, and we also see more varieties of DRAMs being aimed at these problems.

Names like Wide Words, UMA, Interleaving, that apply to the architecture, and the DRAMs will be going through in the discussion.

Finally we have the people who are subject to everything that is going on in the world. They are subject on one side to the constraint of ever-increasing needs from the user standpoint, and on the other side from DRAMs that are becoming more dense with very low pin counts to get the data out. A myriad of solutions rather than a single uniform solution to solve all these problems.

These people are Avo Kanadjian, who is VP of Memory Marketing for Samsung. He just recently joined Samsung from Toshiba and has worked for 13 years in memory marketing. Avo has also worked at Xerox as a memory systems staff engineer.

Neal Margulis is Senior VP of Research and Technology at S3. He has been at S3 since 1989, during the initial phases of the company, and before that he worked at Intel in applications and support for high end microprocessors.

Tom Dye is the Director of the Austin Design Center for Cirrus Logic and is their Chief 3D architect. He has been responsible for the development of the Laguna 3D multimedia accelerator, was a director of multimedia at Dell Computer, and also worked in graphics PC architecture positions at Nth Graphics, Motorola, and Texas Instruments.

Dr. Mike Nielsen is the Director of Engineering at Silicon Graphics. He is responsible for the systems architecture and definition of the Indy-follow-on projects and has been a chief designer of the unified memory system there. He also was at MIPs Computer, responsible for systems architecture, and was at DEC as a chief architect and designer of DEC station MIP space workstations.

AVO KANADJIAN: I want to say that more memory is always better. For example, the-state-of-the-art notebook that has only 16 megabytes of memory would operate 41% faster with 32 megabytes of memory. That is the result of a recent benchmark that we did. The idea of a synchronous DRAM was suggested to us by system engineers who wanted to improve performance by having accesses happen during the precharge period, so we developed an internally double-banked memory architecture. We have inputs and outputs that are controlled by the master clock, we designed the part and brought it to production. The only thing that is not synchronized is the fact that a lot of the controllers, the ASIC core sets, are not out in the market, in order to take advantage of the performance improvements of a synchronous DRAM, with the exception of companies that have control over the development and production of their own chip sets.

In 1996 approximately 10% of Samsung's DRAM production will be synchronous DRAM. As these new chip sets become available, in early '97, we anticipate that, by the end of '97, as much as 50% of the total DRAM production will shift over to synchronous DRAM. By 1998, that percentage will increase to 70% and over. We think that synchronous DRAM is a revolutionary solution, but once it is implemented it will have incremental improvements in performance, such as putting four banks instead of two, improving the processes, and getting higher performance parts altogether. Synchronous DRAMs will definitely be the solution for a lot of the new applications that are coming on line, such as multimedia, multimastering, and AGP. We are ready, from a production point of view, to transition the market from standard DRAMs to synchronous DRAMs.

JIM: Next we have Neal Margulis.

NEAL MARGULIS: S3 is a graphics company, but when looking at memory and different memory choices, we really have to consider system memory first because that sells on the order of 6 to 10 times as much as graphics memory. On the system memory side, we definitely agree with the observation that the industry is moving towards SDRAM for main memory. It is higher performance, with clock bursts, and is easier to design with.

One of the other reasons we see this as a good step for the industry is it is very well suited for SMA (shared memory architectures), not to be confused with UMA, which tends to denote separate chips for graphics and core logic. With SMA we have redefined the term to refer to an integrated memory subsystem including an integrated controller that controls one set of system memory.

S3 is pioneering SMA as the architecture for cost effective entry level segment one, PCs, and some of the emerging markets for PCs. SMA has a higher level of integration and cost savings than AGP, where the entire frame buffer is moved into main memory. SMA provides integration, cost, and flexibility advantages. The memory can stay in one place, it does not have to be bused around the system to different memory subsystems. Especially for multimedia data which tends to be streaming-oriented, there are performance advantages to dealing with one pool of memory.

Because the graphics controller is integrated into the main memory controller, you can do a better job with system memory arbitration and reduced latency. For graphics memory subsystem, S3 believes that EDO has another year of dominating the percentage of the systems that go out. SG RAM needs to perform at a significantly higher megahertz rate to match EDO's performance, be-

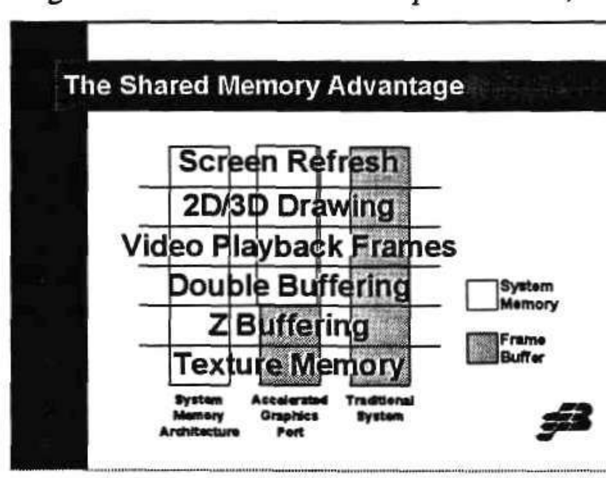

cause EDO has lower latency. The products we are bringing to market now are an EDO version and an SGRAM version. We do not feel that strongly about it, but the overwhelming majority of our customers are choosing to deploy their systems based on our chips with faster EDO memories, rather than SGRAM. In the future, we do see the path to above 100 MHz requiring synchronous memories. Over 1997, EDO will far outship synchronous memories.

A trend that S3 finds a little disturbing is this universal appeal of AGP and the universal dismissing of SMA as AGPs from Intel, such as, "It is the way to do high performance, and it is a great thing, and SMA, because of some anomalies of the past, is an inherently bad thing." If you look at a 3D graphic subsystem, AGP is really one step towards what a SMA architecture provides. In AGP, the texturing and Z buffering is done out of system memory, while things like 2D and 3D, and screen wrap refresh, are done out of a private frame buffer. With SMA it is all done out of system memory and with the traditional systems it is primarily all done out of the frame buffer.

High performance 3D, texturing, and texture fetches are more demanding than screen refresh. It is almost an anomaly to say that AGP will not have the CPU suffer from accesses, yet SMA will be bad because refresh happens from it. It is just simply not true, they are very technically similar, and we believe that, for the entry level SMA will win out. Then for the high end, traditional systems with larger frame buffers are the highest performance way to go for 3D graphics.

JIM: Our next speaker is Tom Dye from Cirrus Logic.

TOM: I wanted to talk about our memory selection criteria, because ultimately people look at that as the ultimate judgment. It is difficult to look into the future, when you have a new memory architecture coming down the pipe, or a number of new memory architectures, and your design is a year out, and you have to pick the right one. You will lose if you do not pick the right one. We have noticed that the VRAM is not applicable any more for the mainstream market. There are a few high end niche players that use VRAM for specific applications but it is a little too costly for mainstream.

Over the long run, the second half of '97 into '98, the EDO will not compete with new memory architectures. EDO is a good burst RAM, it has the capability of doing fast burst page fetches that are nonsequential, while SGRAM and SDRAM do not, so there is some advantages. As you bring the frequency of a synchronous memory up above the point the EDO will run, the synchronous memories have an advantage in bandwidth.

SDRAM, as we see it, is good for system memory, but it does not compete well in the mainstream. SGRAM competes much better with its special features, the block clear, and the mask per bit functionality, and allows a simple lower cost 3D accelerator to work fairly well. Another observation that we have at Cirrus, is that there will be a period of time where the system memory evolves from EDO to SDRAM. The belief is that the system memory will then have to switch to a protocol-based memory, or something that allows a much higher frequency, to get the bandwidth, because of the difference between going with more pins and more switching outputs versus less pins and higher frequency. It is our belief that the frequency will win out in that scenario.

My personal belief is that SGRAM will continue to evolve and SDRAM will continue to evolve. Manufacturers will have to pick up transmission theories required to have lower voltage swings, to be able to run these memories at higher frequencies. For 3D graphics, Cirrus Logic picks the RAM bus implementation or DRAM, for its latest 3D graphics accelerators. There are some definite advantages for mainstream controllers in the PC with RAM bus. Basically the RAM bus protocol allows you to have nonsequential accessing of memory. When you read textures, or are doing the mapping, you are usually not placed into a sequential fetch. The fetching is usually random, and what you need to do is stay in a page and be able to randomly access the memory. With an SDRAM or an SGRAM you basically do a row selection, and then are required to read that data sequentially in order.

A lot of 3D graphics do not apply to that rule – triangles are drawn from multiple directions, which means that you have to be able to fetch data in different orders. The faster, not wider, issue is going to be a trend in memories as RAM bus has pioneered the transmission line memory at 600 megabytes per second. That is a good clip and requires swinging a voltage of less than a volt to keep the noise down and the speed up. I believe that will be a trend in memories.

Another feature that is good for RAM bus and 3D is the protocol concept. In a protocol you can embed a packet of information which is then transferred over to the memory, and the memory can decode that packet and do special features, which everyone would like to see in future versions of memories. There are a lot of functions that would be better placed in the memory. As time goes on the protocol technique means we do not have to relay out a board, or change a package to add a pin, to get another function.

What is it that motivates this criteria selection of memory choice? One of the first things that we look at is the cost performance ratio. The cost performance ratio is critical for us, and there are certain application markets where the cost is more important than performance, and vice versa in other ones. 3D is a good example of where we need the most performance. You may have to pay a little bit more to get it.

The second criteria is availability. A lot of our system manufacturer customers will not go off single source solutions. They just will not take that risk. They expect to see at least two or three memory vendors with compliant solutions. Another is market acceptance. The story that RAM bus is difficult to sell to our customers is true. Initially, the onslaught of customers were skeptical about RAM bus, it scared them, the frequency was too high. SDRAM and EDO are solutions that are very comfortable. They run the same voltage swings with the same address data technique. In recent months, we are starting to see the acceptance of this higher frequency memory device, and it no longer becomes an issue. Market acceptance is an important criteria that will ultimately drive cost.

The design and implementation costs are important. Certain memory devices require more design costs and have a longer lead time to market, and the benefit of that (you would hope), is higher performance. Another is future growth path. I believe that a protocol-type memory is the type of memory that will ultimately win in the future.

JIM: Our last speaker is Mike Nielsen, from Silicon Graphics.

MIKE: If you look at SGI's product line there is no one solution for us. We have everything from the very low end systems like the Nintendo 64 up through super computers. In each of the various bands that we target, from embedded consumer applications, low end workstations, mid-range workstations, high end workstations, there are different trade-offs in selecting the DRAM technologies. For low end systems like Nintendo 64 where there is a very limited amount of memory but still high bandwidth requirements, a technology like RDRAM has been selected because it provides that trade-off.

An entry workstation, like the O2 product announced earlier this month, is a single memory system architecture with 2 gigabytes of memory bandwidth that is applied to both 3D graphics, screen refresh, CPU, and all the I/O. All these things contend in this one memory system, but on the other hand, it is a memory system which is an order of magnitude higher bandwidth than you would have built if you had made many different separate memory systems. Things like rendering for 3D graphics and texture map fetching, far outweigh activities like screen refresh. For that class of workstation, where the main focus is data visualization, there are minimum memory requirements, typically 64 megabytes of memory. The aggregate bandwidth that you get out of the minimum number of synchronous DRAMs is more than sufficient to provide the bandwidth that that total system needs.

A mid-range system is based on RDRAM because it is a graphic subsystem, and systems can be configured with one or more of those subsystems so there is a different trade-off in terms of memory capacity versus bandwidth that is needed. A high end machine, where there is a tremendous amount of memory capacity in the system there is maybe over 1024 bits per pixel, and 500 megabytes of texture memory. Again, it is a case where synchronous DRAM provides the aggregate bandwidth that is needed.

When we use synchronous DRAM, we do not take advantage of things like burst features. We do access them, supply a new address every cycle, because rasterization is not a sequential process. You are rasterizing a 2D space. We may not use synchronous DRAM in main memory in the same way that a mainframe or a PC would, but it is still a high bandwidth solution for providing the graphics memory. As we go forward, DRAM latency and bandwidth is the limiting factor. With the rate of improvement in ASIC technology, a talented design team can put just about as many gates as you need to saturate the memory system and the fundamental performance becomes limited by what the DRAMs can do.

We are always looking at what is going to be the next high performance standard. In order to keep improving performance at the rate that we need, the memory has to follow this 18 to 24 month doubling in bandwidth. We expect synchronous DRAM, which is just now coming to market, is certainly going to be dominant in the two-year timeframe. They then either need to continue to evolve in bandwidth, or sync-link, enhanced versions of SDRAM, a variety of technologies are being proposed, evolutions of RAM bus. There will be a choice for the systems that are being designed for the year 2000 timeframe in the technology that is going to be the most cost-effective in each of the various market bands.

We favor mainstream DRAMs because of our focus at the higher end of the workstation market, where we have so much DRAM in the system that we are very sensitive to the price per component. In a very small frame buffer of a couple megabytes, if the DRAM is a bit more expensive per part, that may be a good trade-off. For tens, hundreds, or thousands of DRAMs in a system, that price premium adds up to a substantial amount of material cost. That is why we favor driving standard commodity DRAMs.

To choose sourcing there are the customary motivations such as quality, reliability, availability, and for a company like SGI which does ship worldwide, local content issues. Besides the traditional issues, we are trying to pick the next emerging standard. We select a DRAM technology before it actually physically exists. We need to have a vendor who is willing to work and take feedback. In the typical standardization process in the industry they have focused on protocol and pinouts, but not necessarily on electrical specifications. There is a tremendous amount of effort in working with the vendor to really hammer out those issues that make the DRAMs useable at a system level, not just as components. There are long lead times to change DRAMs, they are very difficult parts to produce. Partnering with vendors willing to invest a tremendous amount becomes a crucial aspect to us.

All current generation shipping systems in SGI are based on synchronous memory, either RAM bus or synchronous DRAM. For all next generation systems, that trend is going to continue for the products that are now being conceived, out into the '98-'99 timeframe.

JIM: The first question that I have is about the video RAM, I noticed that Tom said video RAM is dead. Nobody else mentioned video RAM but Samsung, although not a supplier of video RAM, does supply windows RAM. Where does such a beast fit in today's world? Where does a part like that fit in?

AVO: Samsung does make and sell 2 meg and 4 meg video RAMs. When we look at the spectrum of applications, this year the 256K by 16 EDO DRAM has the lion's share for graphics applications. We estimate about 65% of that market are serviced by the EDO. The remainder of that market is split between video RAM, window RAM, and RAM bus technologies. Next year we anticipate the SGRAM, which is best suited for the graphics applications, for its synchronous version, to grow and become the most significant share. I estimate that to be about 50% by the end of next year. Video RAM will become less of an attractive solution. Window RAM, which is the only remaining dual port cost effective part, will maintain its market share, and we estimate that to be about 10%, including designs wins.

JIM: I heard you say some time before this presentation the reason why people want that is for high performance, because of the fact it is dual ported, is that correct?

AVO: That is correct, this is for the high end.

JIM: Mike says that for high end systems you do not need any of this exotic stuff because you have such huge DRAM that you get off-band what you want to out of just having wide words, is that correct?

MIKE: In general.

JIM: Will you be inclined to use the windows RAM in any new designs at Silicon Graphics.

MIKE: We are not currently looking at it.

JIM: So, where does it fit in?

AVO: It fits into the PC industry, in particular for high end graphics. There is a new design win that you will see come out in production next year, and it will sustain the run rate, I estimate through 1998.

JIM: Do you two see any particular place that it fits in PC graphics?

TOM: There is one particular place, and that is high end CAD/CAM application boards, add-in boards. Any design that requires an external DAC on the board, is a high end resolution type product. A windows RAM can work very well and replace a VRAM in that socket, because a windows RAM is designed to be more cost effective. I believe that there is a market for that particular device. The size of the market is really what we are looking at, and I believe in the mainstream graphics and video market, the windows RAM will play.

AVO: I do not believe video RAM and window RAM will play and because of its cost efficiency the window RAM will win out, over the next couple of years.

NEAL: This year VRAM out-shipped window RAM, S3 sold well over a million controllers that supported VRAM with our 968 product, which uses an external RAM DAC and was targeted at the high end of the market. We recently started shipping the Verge VX which also supports VRAM, and includes an integrated 220 MHz RAM DAC, so for the top 15% of the market, the 160 by 120 market, or the 24 bit per pixel market, at 1280 by 1024, VRAM is still a key product. It certainly is not the volume but it is very profitable for our board level add-in card customers.

JIM: It sounds like Silicon Graphics is using synchronous DRAM for both the low end and the high end, and for mid range you are using RAM bus. If you look at high end PCs as being the mid range represented on this panel, then windows RAM, or the video RAM, would be an attractive solution for the high end PC, which is the mid range of what is being represented here. Would you agree with that? Response: As an added cost.

JIM: What is likely to be an architecture for the mid range in the future? Like synchronous graphics RAM. Is that likely to be something that solves problems at the middle of the spectrum but not at high end and not at the low end?

AVO: I firmly believe that SGRAM, because of the block write and the write per bit, is quite attractive for many graphics system design engineers.

NEAL: I would agree that for the mid range and high end SGRAM will win. I am not so convinced it will be because of those two features as much as the configurations memory vendors are offering in the 256K by 16, and the 256K by 32, that happen to be SGRAM, not SDRAM, but in any case that configuration will win. Single port is clearly the trend.

MIKE: Single port is the trend, the wide configurations are tending to show up first in SGRAM and that is why they are attractive. For the machines that SGI builds there is at least a byte per component, so the write per bit features are not really that useful. Even the block clear is marginal because generally speaking we are filling with textures and things like that.

NEAL: Block write – you have taken the fastest thing that you can do in SGRAM, or in a memory system, clear a big chunk of memory and you have made it faster, so that is not where the bottleneck was to begin with.

AVO: Every time we try to eliminate SGRAM we get arguments against that, but since you are customers I will not argue this point.

JIM: Also the SGRAM is burdened with extra features so there has to be some trade-off that makes it worthwhile.

TOM: It is a marketing bullet. A lot of folks put the highest fill rate, based on a block clear, on their marketing sheets and for some reason that tends to sell.

JIM: It sounds like there is no easy answer. There are going to be certain places where SGRAM is right, just like there are certain places where video RAM and windows RAM are right.

MIKE: The features are fairly minor, so it comes down to the premium. A negligible premium is attractive, but at a 2X premium, it is not going to be used.

JIM: The issue of granularity probably has no easy answer. Silicon Graphics has shipped a fair number of systems based on a unified memory architecture machine. Is that another place where you have to make a decision system by system on where UMA makes sense?

NEAL: Clearly there are some barriers to doing SMA. UMA got a bad rap last year because people were shipping 8 meg systems, Microsoft did not want to run Windows '95 in 7 meg which was a viable argument for last year. This year and next year, when people are shipping 16 meg and 32 meg systems standard, Microsoft is perfectly happy running Windows '95 in that amount of memory minus 1, 2, or 4 megabytes for the frame buffer. The only issue left is in the PC space, Intel has 90% market share in the core logic. To do a SMA graphics controller you need to wrestle away market share from Intel.

JIM: But a one or two megabyte frame buffer is insufficient for 3D, isn't it?

NEAL: Sure, but for the additional texture memory needed, you have a more flexible approach with SMA than any other possible approach.

MIKE: How much memory there is for the system is always an issue. On the other hand, at least in my product band, the memory configurations are so much larger to start with that the amount of memory used for 3D graphics, which for us is typically something like 16, 17 megabytes, is still not significant relative to a 64 meg entry machine. Our machines are configured with 128 megabytes of memory. The other point is with shared memory architecture, you only need to dedicate that memory when you are actually using it. In our case, when one application is running it may need hundreds of megabytes of texture memory, then five minutes later it is not running and all that memory is available to the system. One of the attractive features about shared memory architecture is that you dynamically use the memory in the machine as needed for the current app, rather than using dedicated memory to handle your occasional peak demand.

JIM: So UMA would make a lot of sense, even for a very high end system.

MIKE: The distinction between the mid range and high end Silicon Graphics machines is not only their inherent performance, but also their modularity and expandability. It may be very common to have two or three graphics pipelines added on to a high end machine. In that case you really cannot unify that directly with the CPU and I/O memory. There is a distinction that when you get past a certain class of machine, where you want multiple displays, multiple pipes, you have to go to a different architecture. But for our main volume systems where there is only one display, one pipeline, it is the most appropriate architecture now.

JIM: What is the appropriate type of DRAM in a UMA type system?

NEAL: SDRAM is lowest cost per bit.

TOM: There may be some issues with that too. A core logic device trying to run all the buses present in the PC, plus memory in an interface to the local CPU bus results in a pin limitation without going to a very large package. Any time you run into a pin limitation and you need high bandwidth memory, then a RAM bus is a very good solution.

JIM: I know Mike will agree with that because Silicon Graphics uses RAM bus.

MIKE: Yes, but on the other hand, because of the price band we are selling into (we typically have ASICs that are up in the 600 pin count), so supporting a 256 bit memory world is not a problem for us. From our market segment, SDRAM is the most cost effective solution.

JIM: S3 is the only company here that does not have anything to do with RAM bus, at least not publicly anyway, and I wonder, Neal, if you could explain what your stand is on RAM bus.

NEAL: Publicly we have no plans for it.

Q: When does the panel think that Intel will increase the PC system bus speed, prompting the move to SDRAM, sync-link, or RDRAM, as the main memory of choice in the PCs?

AVO: End of '97. This is my opinion, not anything based on any discussions whatsoever.

NEAL: The VX already supports SDRAM and 66.

JIM: I guess I should ask Avo what is going to make people start using SDRAM with the Triton VX chip set.

AVO: The answer is simple. There is no price premium. One of the conditions for using synchronous was no premium for synchronous over standard DRAM, and once you have the right chip set you immediately achieve a 20% or more improvement in performance. Very simple.

TOM: The other issue is the way the data types are being bused around the system they are very bursty in nature. In SGRAM, SDRAM solution for packets that are sequential burst, that are cache fills, are very effective.

JIM: The way Dataquest sees it there is no customer pull from the purchasers of PCs, requiring the PC manufacturer to switch over. Until that exists there is no reason for PC manufacturers to go from something that is extremely widely supplied to something that is supplied only by the top DRAM vendors. There is always going to be a cost premium on something supplied by a limited set of vendors.

Q: Can anyone comment on the merits of sync-link versus RAM bus DRAM?

TOM: RAM bus is here in production today simply because sync-link is at least two, if not four, years out. Something needs to emerge beyond the bandwidth levels that are currently proposed for SDRAM, but it is a long ways out.

AVO: I agree with that, we are producing RAM bus and sync-link is really still at the conceptual level.

JIM: What is the road map to incremental improvements in SDRAM performance for the future?

AVO: The incremental performance improvements will be achieved with multiple internal banks – we are going from two banks to four internally. That will give an incremental performance improvement. We will also have faster speeds because of finer processes. Just as the DRAM evolved from page mode, to fast page mode, to EDO, the synchronous will have the same incremental performance improvement.

Q: When will you change voltage levels? At what frequency?

AVO: We are concerned about maintaining consistency in voltage, because it also has to do with the ability to go from 16 MB technology to 64 MB technology, so that transition has to be planned today. I believe that synchronous started with 3.3 V and will stay with 3.3 V even at the 64 MB level.

Q: But will you get to the 143 MHz at 3.3 V.

AVO: The answer is I think so.

MIKE: I think that there is a defined road map for how SDRAM evolves, both through JEDEC in terms of pinouts and number of banks, and from every single vendor there is a roadmap of moving to SSTL interface from 125 to 133 MHz and above and road maps for 143 MHz, 200 MHz parts. There are parts that will be 2X the bandwidth per pin, in the late '97 timeframe they will be ready for markets. The question is really: What trade-offs do you have to make when you get to 267 or 300 MHz and does SSTL 2.5 interface come in or does it turn out to be some other interface at that point?

JIM: Designers resist change in the way they do things, and when they are comfortable using SDRAMs, they will stay with them for a while, which is a more painful direction for them to move, towards a higher frequency (RAM bus or sync-link), or towards a logic level they are unfamiliar with.

MIKE: It will be difficult from a customer perspective because it will be harder across generations of machines to have interchangeable parts. It is most likely that every generation will have incompatible parts, as we go forward at the high end.

TOM: That is one of the advantages of the protocol memory.

MIKE: Nevertheless it is getting harder at the rate things are scaling to have an interchangeable platform similar to that in the PC industry. SIMMs have gone back for many years as a least common denominator. From that point of view, the voltage changes that will happen and protocol changes that will happen, whether it is in SDRAM or other forms, is a case where every two plus years, there will be a mandated voltage interface change, just from a semiconductor process point of view.

AVO: I think that you will have two generations which will coexist. Today we have five volt and three volt that coexist, and they address two different types of applications. The market decisions that you make are more concerned with standards and availability. Others are more aggressive in their specifications and can produce add-on cards that can be put into installed base systems, giving more performance over time.

TOM: The typical engineering community has to start facing the facts of higher and higher frequencies. Transmission line theory and a lot of the ability to run those frequencies will be commonplace. It will have to be accepted.

JIM: Neal, your position is that EDO is right. You have also said EDO is something that you are going to plan to use through 1997.

NEAL: People are pushing EDOs to 80 MHz and because of the advantages of lower latency and random access, 85 MHz EDO significantly outperforms 100 MHz SGRAM.

JIM: Does that mean you will probably be staying with three volt synchronous DRAM longer than people would imagine?

NEAL: For the SMA market, clearly that would be the case. S3 is a volume supplier so we cannot shift too far out of the mainstream and keep the company running.

JIM: You have to make sure you are in sync with the people who make the majority of the units too.

NEAL: Sure.

Q: In the NeoMagic presentation Prakash Agarwal shared a Tai photo where 50% of the area was embedded DRAM. Is embedded DRAM a real trend? If so, are current DRAM manufacturers positioned to participate, and third, what are the potential implications for logic venders?

AVO: I think embedded DRAM is something that we definitely need to explore. There are definite performance improvements by integrating the memory with the peripheral circuitry. For it to become a cost effective solution, some serious challenges need to be addressed in terms of manufacturing yields. When you put that much memory with circuitry on board, test vectors and yields become very big challenges. Over time, we are headed in that direction, but there are some lessons to be learned along the way.

JIM: Does anyone else have any comments on the NeoMagic approach? Where it fits in?

NEAL: Embedded DRAM is a premium type product today. There is no way to justify it on a cost per bit basis. However some segments of the market, power sensitive notebook computers, are willing to pay a premium for that. It is just a question of how big the segment of people willing to pay premium per bit is.

Q: Is 3D the wave of the future for PCs?

NEAL: The design cycle for embedded DRAM controllers is longer, and despite the fact all that bandwidth is available, it basically increases your design cycle, probably another six to nine months.

JIM: I was talking about how bad the DRAM size was for 3D.

NEAL: There are two issues, one is lining up your feature set for something that is going to take so long. The current NeoMagic chip, despite all the theoretical advantages of embedded DRAM bandwidth, it is a 32 bit controller, it under-performs products that we had on the market two years ago. However, it sells well in certain segments because of the power. For 3D, the market is the feature set, the performance, all those are moving so fast it is going to be really tough to intersect the right features of 3D with the right amount of memory on board. Then you still have a tremendous issue of the onboard memory uses. If you use it for the frame buffer and take the texturing from somewhere else, then your biggest bandwidth consumption, the texturing, does not take advantage of the onboard DRAM. There is a real system partitioning issue once you start talking about 3D, and you will not have the 8 or 16 meg required for 3D on one chip.

Q: What would be a reasonable smallest size DRAM you could use on a chip for 3D?

NEAL: I think someone mentioned they would probably go with just the frame buffer on chip, say 2 meg, then try texturing over AGP, which is a very medium range to low end 3D approach.

JIM: Once again, it just appeals to low power types. Would you agree with that?

MIKE: Yes. here is also the issue that with the processes currently available, you can get good DRAM and lousy gates, or you can get good gates and lousy DRAM, and 3D wants good gates and good DRAM. From that point of view the intersect does not happen until there is a generation of unified process. I have not seen a road map for that.

Q: Any comments from Cirrus?

TOM: I agree. 100%.

Q: If Intel backs RAM bus, can the sync-link effort catch up?

AVO: No, there is no way the sync-link effort can catch up, because RAM bus is here now and is a reality. It will find some applications both in graphics, it is already designed into a lot of graphics and main memory applications. Sync-link is a concept that will be talked about for some time, and hopefully we will come to an agreement in terms of what the real advantages are, and whether or not it provides any benefits above and beyond what we have today.

Q: I need a 64 bit data interface but today I have to pay a penalty in pins, connectors and granularity. Is there a way to have 64 bit performance and no cost penalty?

NEAL: That is the whole beauty of shared memory architecture. You pay for the pins once, for system and graphics, and use as much memory as you want for graphics, have the flexibility you want for MPEG one minute, for 3D textures the next, and for high res 2D after that. So SMA is exactly addressing that concern.

TOM: There also is an advantage for RAM bus

JIM: Yes, I was thinking that any of the existing architecture ought to be able to do that.

TOM: RAM bus is available today in that form factor to do that.

Q: What will four banks in an SDRAM be for performance? To increase performance.

MIKE: It depends on your application, more is better.

JIM: This is just a straight interlinking question.

TOM: They are probably talking about concurrent operation. Where you can precharge multiple banks, more banks than two.

MIKE: In applications, particularly like 3D graphics, where you may have to simultaneously access texture memory, depth buffer memory, and color buffer memory, having more banks, so that you can keep as many different data structures in page mode at once, is very desirable.

TOM: It will help with core logic that controls AGP buses. The AGP is highly pipelined, which allows you to go to different pages and pipeline the operation, so if you can precharge multiple banks out of your DRAM, it will give a more effective bandwidth.

JIM: This sounds like an argument for MoSys. They do something similar to that, don't they?

TOM: Yes.

JIM: What is keeping MoSys out of the mainstream?

MIKE: It comes back to cost. There is a die area penalty for banks, and that is why the SDRAM road map is moving slower in terms of adding banks than I might prefer. It is ultimately cost-driven. Two worked well at the 16 meg density level, four works well at the 64 meg level. Beyond that you start to pay a premium in terms of the overhead of banking.

TOM: There is a certain number of banks where you have already reached your 90% effective rate and that additional percentage does not really buy you much as far as the cost differential.

JIM: It is almost like cache design. How many ways of associativity do you need in something like that. If Samsung increases your SDRAM mix, then you will have a higher proportion of TSOP packages, but a lot of OEMs do not like using TSOP packages in their systems, so how is this going to affect SDRAM acceptance?

AVO: I very much disagree with that statement because our industry is shifting right now from SIMMs to DIMMs. When you go to DIMMs it basically is by 64, by 72, by 144 organization. All of those modules use TSOP. A supplier in our industry that has not increased and invested in additional TSOP capacity does not know what is going on. My feeling is TSOP helps us streamline the number of parts that we supply into the market. All the portables use TSOP, all the new DIMMs use TSOP. We would like to see TSOP take over, so we can do away with SOJ the way we did away with DIPS.

JIM: Apparently there is a package called TSOJ. What is the implication on TSOJ?

AVO: TSOJ is a package that has been requested by some customers who are interested in higher reliability. They think that because of the shape of the lead, TSOP is not as reliable in some of their high density boards.

Response: TSOP is also used for simplified stacking of DRAMs.

AVO: That is right, so it provides those two solutions.

Q: Is stacking going to go into the mainstream?

MIKE: I doubt it.

Q: Who would use it?

MIKE: High end workstations and servers.

JIM: I thought that high end workstations were not under size constraint, and isn't stacking only for . . .

MIKE: No, fixed size box customers always want more memory capacity in the same volumetric box. From the point of view of driving buses in terms of minimizing parasitic capacitors, from connectors, traces, DIMMS and such, it is advantageous.

AVO: I do not agree with that statement. I think most people who need high density and high reliability, if they use stacking, use it as a band-aid approach for as short a period as possible, then shift over to the 64 megabit density. Reliability becomes a concern when you are stacking parts.

JIM: That same thing happens in SRAMs where every time a new generation comes out, the next generation is available as a module.

Response: It is most attractive at the middle of generations.

Q: Will the 64 meg generation be 100% SDRAM in 1998?

AVO: 64 meg density will be close to 100% SDRAM in 1998, because performance requirements are steering it in that direction. Upwards of 70% of the market will require SDRAM in 1998, and of that, practically all of the 64 meg density requirements will be synchronous.

JIM: For a long time Dataquest has been saying the 64 meg is going to ship in appreciable units in its lifetime that are synchronous parts – maybe 5%, maybe less. But that is looking over the entire life cycle. The bulk of what ships now is asynchronous, and probably is not going to take too long to shift over to being the bulk of synchronous.

TOM: It is hard to say exactly what the system memory of choice will be in '98. It very likely will be SDRAM. One large manufacturer of all the core logic can turn that around on a dime.

AVO: But whenever there is a performance improvement that justifies it, like 20% improvement in performance, it is better than upgrading a CPU.

JIM: Can you convince PC buyers that is true?

AVO: We have been doing a lot of benchmarking, and a lot of visual demonstrations, and we have hired independent labs to demonstrate that.

TOM: So you produce RDRAMs and SDRAMs now?

AVO: Yes we do.

TOM: And you have done performance analysis of both those types of memories?

AVO: We have not done the benchmarking yet on the graphics area. What I was addressing is the synchronous in the main memory transition.

JIM: There has been a big study done on how much main memory is the right amount for Windows '95, and surprising enough, Samsung has found more memory is better.

AVO: Interestingly enough, it was shown that cache was the number one upgrade to make, because cache memory in itself gives a 20% improvement in performance. Increasing the memory from 16 to 32 gives you the next greatest improvement in price performance, and third is CPU upgrade.

NEAL: Whenever we do that independent study, changing your graphics card to a faster one always comes out first.

JIM: They told us it was our disk drive.

NEAL: It really does not matter how long system level things take in the background. What really matters when you are interacting with the PC is those milliseconds where the graphics really count the most, and the perceived speed of the system, which is most highly dependent on the graphics.

Q: Which of the large DRAM manufacturers, domestic and foreign, do you think will not be ready for the 1997 to 1998 introduction of SDRAM? And why?

AVO: I think that most, if not all, DRAM manufacturers will be able to make the transition. From a process point of view, it is not a very difficult shift. From a test point of view, and yield point of view, there are some challenges ahead of them. The technology itself is not exotic. Some of them clouded their thinking with things like burst EDO before they started developing their synchronous SDRAM, so the time delay between different manufacturers has to do more with their road maps, and less with their ability to technically bring the synchronous DRAM to market.

Q: When does the 64 meg cross over in price?

AVO: It is my estimate that the 64 meg density will become a volume runner as we introduce the 2 meg by 32 organization. 2 meg by 32 granularity for main memory applications is very critical. Today's 64 meg production consists of 16 meg by

4, 8 meg by 8, and 4 meg by 16, which addresses the high end spectrum of applications, all the way to notebooks. The 2 meg by 32 granularity is essential because two pieces make up for 2 meg by 64 which is 16 megabytes of memory. We need that kind of granularity and configuration in order to give the upgrade path to the largest user of DRAMs, which is the PC segment. That organization will be available in the second half of '97 and price parity between the 64 and the 16 we are estimating at this point will be in the first half of 1998.

JIM: I think that Toshiba already has that part. Something that is interesting about crossovers, and I always like to point this out to anybody who talks about price per bit crossovers, is that the last time we had a tidy crossover, from one generation to the next, in price per bit, was at the 1K to 4K bit transition which was 1975-76 or something like that. Ever since then, two generations of DRAM prices have tracked each other for anywhere from three to six quarters. The 16 meg and 4 meg first reached price per bit parity around the middle of last year. They are still at price per bit parity and it looks like they are going to go out this year with price per bit parity, so a crossover is never a really tidy thing. If you ask me when the 16 to 4 meg price per bit crossover was, I would say it happened sometime around 1995, 1996, 1997. We suspect the same thing will happen with the 64 meg, and with all successive generations, just because the two parts can be used in the same socket for a long period of time.

Panel:
Packaging Performance Breakthroughs: Is the Industry Ready?

Moderator:
Näder Pakdaman

Principal Industry Analyst,
Semiconductor Equipment, Manufacturing, and Materials Program,
Semiconductors Worldwide,
Dataquest

Mr. Pakdaman is a principal industry analyst for Dataquest's Semiconductor Equipment, Manufacturing, and Materials program in the Semiconductors group. He is responsible for research and analysis of semiconductor equipment and trends in IC manufacturing technology and capacity with a specific focus on the lithography and process control segment.

Before joining Dataquest, Mr. Pakdaman was at IBM T.J. Watson Research Center and IBM East Fishkill. His responsibilities included fast optoelectronic testing of ICs and qualification of advanced optical lithography systems.

Mr. Pakdaman has B.S. degrees in mathematics and physics and an M.S. degree in electrical engineering from Purdue University. He was a doctoral candidate at Columbia University in applied physics before he joined Dataquest.

Panel:
Packaging Performance Breakthroughs: Is the Industry Ready?

Dr. William M. Beckenbaugh

Vice President and Director of Sector Technology,
Advanced Interconnect Systems Laboratories,
Motorola Semiconductor Products Sector

Dr. Beckenbaugh is vice president and director of the Advanced Interconnect Systems Laboratories in Tempe, Arizona. In this position, he is currently responsible for Motorola's Sector Technology packaging research laboratories with R&D projects related to advanced design and manufacturing technology for new silicon and GaAs packaging systems.

Before joining Motorola in 1985, Dr. Beckenbaugh was a research leader at the AT&T Engineering Research Center in Princeton, New Jersey.

Dr. Beckenbaugh received his Ph.D. in physical chemistry from Rice University in 1972 and held postdoctoral positions in physical chemistry research of the Swiss Polytechnic Institute (ETH) Solid State Physics Laboratory, Zurich, Switzerland (1972) and Ohio State University, Columbus, Ohio (1973).

Panel:
Packaging Performance Breakthroughs: Is the Industry Ready?

Barry Lieberman, Ph.D.

Engineering Manager,
Intel Test Tools Operation,
Intel

Dr. Lieberman is presently engineering manager of the Intel Test Tools Operation. He is a member of a Sematech Project Technical Assessment Board chartered with developing new solutions and infrastructure for wafer sort and wafer probing.

Dr. Lieberman has 20 years of experience in the semiconductor and related industries and has worked both in the technology and product development areas. During his career, he has worked in the fields of electron beam lithography, magnetic bubbles, multichip modules, microprocessor development, and test technology development. He has published 12 technical articles and holds one patent.

Dr. Lieberman earned a B.S. degree from Caltech in 1968; an M.S. degree from UCLA in 1970; and a Ph.D. from UC Santa Cruz in 1975, all in physics.

Panel: Packaging Performance Breakthroughs: Is the Industry Ready?

Michael J. Cadigan

General Manager,
Packaging Group,
IBM Microelectronics

Mr. Cadigan is currently general manager of Packaging Group, IBM Microelectronics and IBM East Fishkill senior location executive. He is responsible for the development and manufacture of semiconductor packaging products at IBM locations in East Fishkill and Endicott, New York, in addition to his senior location executive duties in East Fishkill.

Mr. Cadigan joined IBM in 1979 as a quality engineer at Endicott, where he assumed his first management post in quality engineering in 1982. In 1985 he became technical assistant to IBM's Systems Technology Division president. He was named panel operations manager in IBM Austin, Texas, in 1986. He became manager of Electronic Card Assembly and Test (ECAT) operations in IBM Charlotte in 1990, and ECAT plant manager in 1991. He was named general manager, IBM PC Company, in Austin in 1994, and RISC System/6000 director of Worldwide Manufacturing, Austin, in 1995. He was named IBM senior location executive in Endicott in September 1995.

A native of Buffalo, Mr. Cadigan holds a bachelor of science degree in mechanical engineering from SUNY Buffalo.

Panel:
Packaging Performance Breakthroughs: Is the Industry Ready?

R.P. St. Clair

Engineering Manager,
Systems Development Engineering,
Memory Test Division,
Teradyne Inc.

St. Clair is currently in the Memory Test Division of Teradyne as engineering manager of the Systems Development Engineering organization in Agoura Hills, California. In addition to system responsibilities, this group has a charter to participate in the development and evaluation of any new technologies that may be suitable for probe cards or contactors. These technologies can then be used with the next-generation VLSI and Memory test systems for both high-performance and high-density applications.

After receiving his electrical engineering degree, Mr. St. Clair was employed by another ATE company before joining Teradyne in 1969. Since joining Teradyne, he has served as applications manager and final test manager on some of Teradyne's DC Parametric and DC Functional test systems as well as Pulse Parametric test systems. He then assumed the role of engineering manager on both Memory and VLSI test systems.

Panel:
Packaging Performance Breakthroughs: Is the Industry Ready?

Igor Khandros

President and CEO,
FormFactor Inc.

Dr. Khandros is the president and CEO of FormFactor Inc., founded in 1993. FormFactor is bringing to market products and technology for silicon back-end packaging interconnection.

Dr. Khandros' career history includes managing a packaging interconnection group at IBM Research Division, and the cofounding of Tessera Inc.

Dr. Khandros earned a Ph.D. in material science from Stevens Institute of Technology.

Session # 21: Packaging Performance Breakthroughs: Is the Industry Ready?

Moderator:
Nader Pakdaman
Principal Industry Analyst, Semiconductor Equipment, Manufacturing, and Materials Program, Semiconductors Worldwide, Dataquest

Panelists:
Ray Bryant
Manager of the Profit Centers, IBM Micro Electronics
Barry Lieberman, Ph.D.
Engineering Manager, Test Tools Operation, Intel
Reed Bowlby
Interconnect Technology Planning Manager, Motorola Semiconductor Products Sector
R.P. St. Clair
Engineering Manager, Systems Development Engineering, Memory Test Division, Teradyne Inc.
Igor Khandros
President and CEO, FormFactor Inc.

NADER: The panel members have come to us from different sectors, both as end users and providers of packaging and test technology.

The presentation by Ray Bryant from IBM, will set the stage, Barry Lieberman from Intel will follow. Reed Bowlby is a replacement for Bill Beckenbaugh who both work at Motorola. R.P. St. Clair. R.P. is the manager for the memory group at Teradyne Inc. From FormFactor, we have the CEO and president Igor Khandros.

Where do we need to go? Where do things need to develop? Who is going to take ownership? Who is going to actually manage to keep the cost performance where it is required, for this market to grow?

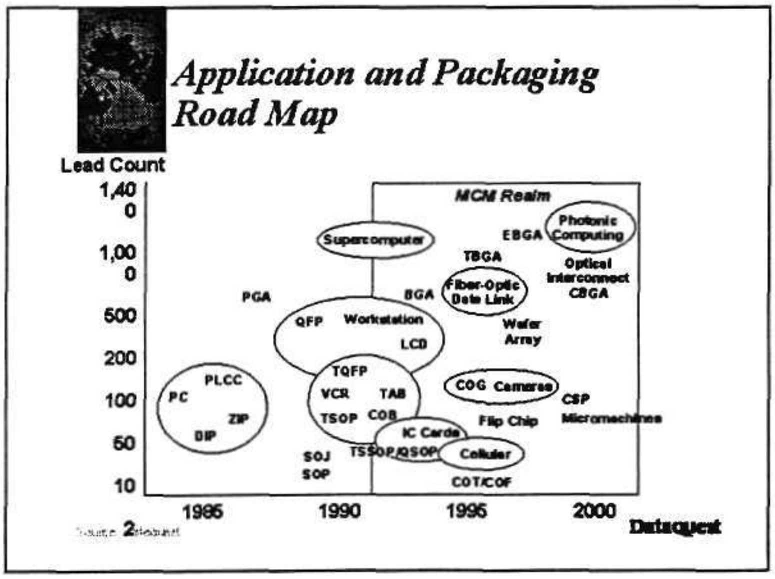

This is a general point of view packaging application road map. It takes different types of packaging and puts them together where there are areas of coverage and what kind of applications will be driving them. In the upper right hand corner are the very high speed, high performance, very high lead count packages. With the seller, you obviously do not have the same lead count but you do have the requirements and the growing need for increased performance. We are seeing more and more applications as semiconductors go into different types of applications, and different types of ICs, both the demand for lead count and performance are increasing.

This block shows where the MCM realm is, and where some of the other packaging types and packaging formats actually would either be alternatives or would be used in conjunction with MCM.

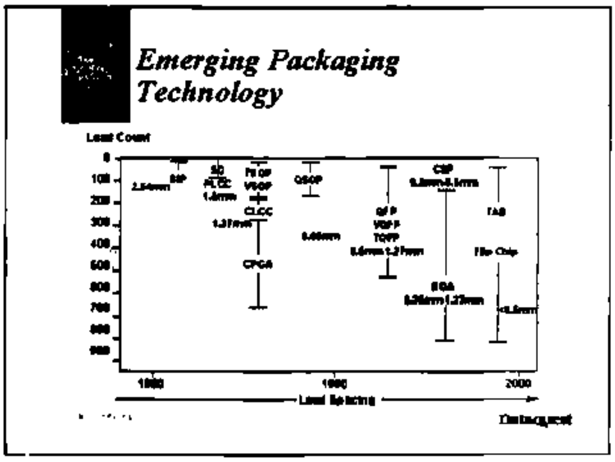

What is shown here is the lead count increasing from top to down. The lead spacing is actually going down. There is 2.5 mm lead spacing going all the way to the extreme low 0.5 mm and with the lead count increasing. The market, in terms of applications, is actually here. The low lead counts may be as high as 200 but are probably below 100. This is where you will actually be heading. With such a phenomenal upswing in technology, in terms of the number of leads which has direct impact on cost and direct impact on test (actually the technology and the design issues that are involved going from what packaging was an afterthought), the chip is built and then handed over to someone else and you have the luxury of test and the afterthought of packaging.

What will happen in this realm? Who is going to take ownership and how are we going to manage the industry and the infrastructure for that?

Looking at the SIA road map on the same thought

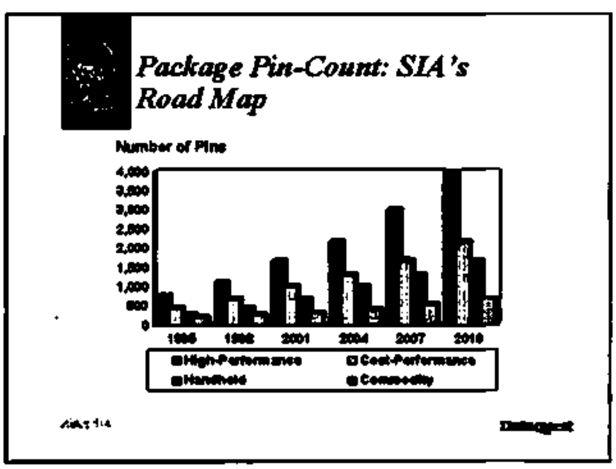

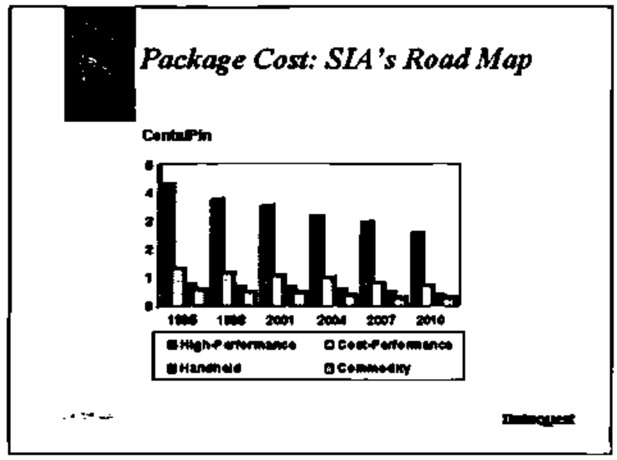

line you see the number of pins, an exponential growth, especially for the high performance. This would include the ASIC and the microprocessors. For the commodity we do not see the same type of growth. There is an increase in the number of pins and increasingly complexity.

What does that mean? Where is the cost and the cost centers, where are the bottlenecks and who pays for what? In a very generic sense you can take the cost per pin, and multiply by the numbers that you saw in the previous slide. Some time in the next decade, we will talk about $120 packages. Obviously if we stay on the same road map, somebody has to make sure that that cost fits not only the performance but the viability of the business. We can do this same thing shifting gears into test. If we remain in this neighborhood of about $14,000 per pin, and we go out to the 4,000 pin microprocessor or an ASIC, we will have a $70 million piece of equipment.

A major semiconductor manufacturer mentioned,

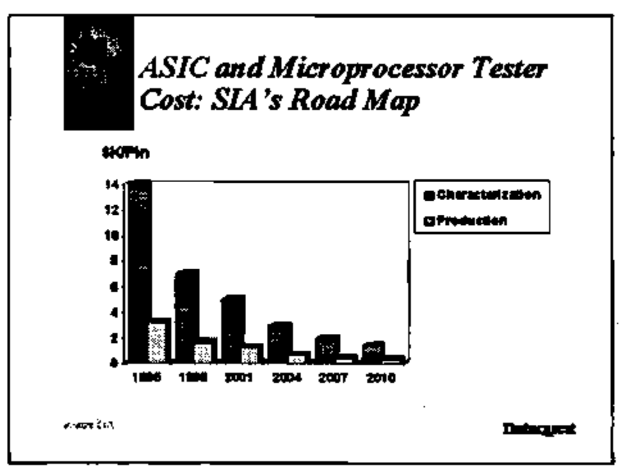

that their cost for test equipment in the past couple of years has been in the same realm and regime as lithography steppers. So how will we go through this exponential growth drop?

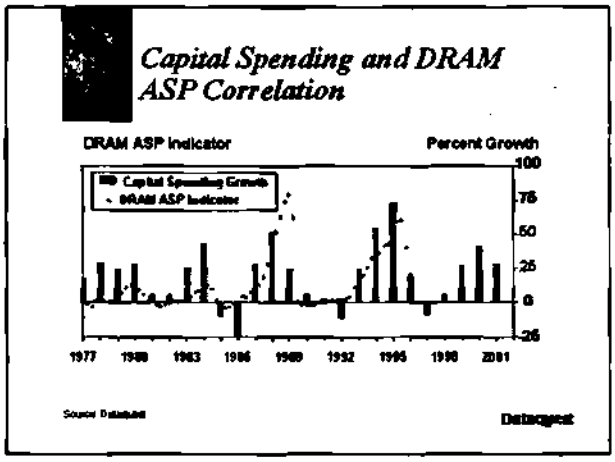

Being from the semiconductor and equipment and capital spending group, I wanted to show you capital spending figures and the impact of DRAM pricing. Prices of DRAM go down as technology transitions happen. Every time we head for a technology transition, we are doing the spending. Every time we go down, and prices of DRAM go down because of higher availability, we go through a period of slowdown in capital spending.

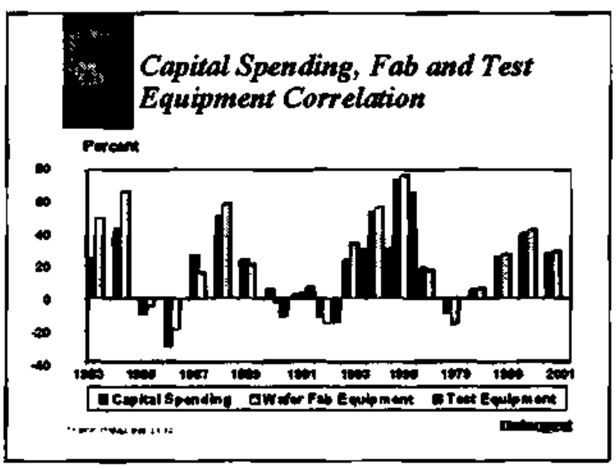

Capital spending, wafer fab equipment, front end equipment , and test equipment actually follow the same trend. We did not see the same growth in '94 that we saw for fab equipment and capital spending. It was at the expense of test equipment that the front end equipment actually out-paced the growth for capital spending. There is definitely a pattern here which reflects itself in front end equipment and the test equipment. A migration of technology and volume and the capital that is needed to put the capacity in place. On that note, I would like to have our next panel member come up and give up a perspective, from IBM's point of view. Ray Bryant is going to discuss his perspective, both in terms of the user and also as a provider of high end packaging to the semiconductor community.

RAY BRYANT: Nader has taken you through some of the specific numbers: the trends in the SIA road map, the size and increase in the I/O, and what the cost is going to be doing. I want to put in behind that some of the specific items that I think are going to be important, so that when we look ahead at what packaging has to do, over the next three to five years, and the goals that each of us must look at, we can have a discussion on that in the panel.

Clearly there are three major drivers that are taking place in packaging today. One is the increasing function that the SIA road map and the product menus would show over time. Decreasing cost is a key piece of that. Those two are very similar to the front end, where process materials, semiconductor equipment, any specific aspect of the industry try and achieve that goal. One thing that is starting to show up as a significant difference though, is this transition that is taking place in the way we are interconnecting the die to the package, and the package to the next level of assembly. This transition is going to cause some discontinuities that provide an opportunity for any particular provider to either do well or not, depending on how we collectively manage that change.

The SIA road map shows increases in power, increases in speed, increases in I/O. There is clearly a directional trend taking place. One of the areas not focused on beyond the immediate metrics, is what the pitch changes are going to be driving, the pitch of the chip level, the pitch at the substrate to module level. What that may provide as a cost reduction opportunity to achieve these long term changes, and also what the increased power requirements that the device is going to drive will force us to do at the package level, and force some changes in the way the business is run between the silicon designer, the package provider, and the system menu user, who puts this out to the market.

There is also taking place a significant proliferation of package types. Nader showed all kinds of different package families, with different changes over time. They are being added to almost every week it seems and it certainly is not clear which one, or which few, may become significant volume drivers in the industry.

One thing that is very clear is that old packages never die off. We are continuing to use package technology from 20 plus years ago. What has typically happened is they just shrink down to a slightly finer pitch or a little bit higher functionality, but the number of options out there, are not going down. They are going up significantly. One of the challenges to the packaging industry is how do we take this list of options, which are very nebulous and which are very ill-defined in any specific application, and help the silicon designer and the package user at the system level find a solution which is the best choice to met their cost constraints.

We talk about decreasing package cost, the silicon integration going up, when you go from 0.5 to 0.35 down to 0.25 micron and below, at the litho level, is driving an increase in functionality that the package has to be able to support. It also drives increases in mechanical and thermal problems now, as you start to take that package and shrink it down, and have the silicon play a more dominant role in the mechanical aspects of the package. When you look at packaging as a space transformer, from the silicon function down to the card, it is clear that you have to have a robust package. You need an ability to support that, and it has to go through a much wider range of applications and environmental stress than ever before. One of the drivers that has to be looked at in this cost equation is the cost per I/O is shrinking. Otherwise, you will see package price going up dramatically beyond what anyone can support.

How do we take cost out of the total system between the silicon, the package and the system user? It has to be eliminated. It has to be taken out of the system completely. This reduction in package size and increase in functionality is driving significant cost challenges at the board level. In order to put fine pitched devices down in fine pitch packages, we are driving either increased layer counts or tighter lines, tighter spaces, finer drilled holes, or photo-defined holes at the board level, which is also part of that total cost equation that we have to step up and figure out how to eliminate.

The interconnect transition gives a set of challenges which may be the biggest driver over the next two to three years. The infrastructure today, in supporting pin products and lead frame products, is now being stretched dramatically, when you look at ball grid arrays and flip chip packaging. What are the drivers which will be used to handle this transition? What is pushing ball grid arrays down to even finer features? It began with roughly 1.5 mm pitch balls in plastic packaging, now moving to 50 mil or 1.27 mm, and going down to 1mm pitch, again driving complexity at the board level, which has to be eliminated, and a continual interest in taking it down to this gray zone between BGAs and chip scale packaging, down to 0.8 mm and below. This is going to require an increase in complexity to go to a built-up multi-layer process or to some type of photo-defined via process when you get to the board itself.

When you look at the dico straight process which is used in a number of areas, or a photo imaginable dielectric with an SLC type process, both of which can support these, but they are again a migration of cost from one level of packaging to the next, that we all have to eliminate across the board?

The way this will be achieved, in my opinion, is to look at a more collaborative effort between the silicon producer and the designers; and between the packaging community that puts it into a form that the end user can get function from; and the system user who looks at his total thermal mechanical solution to support that, and the board technology required to handle the complexity they have over time. This is a case where these opportunities and the way they are managed will be the difference in success between providing a product which can meet those cost challenges or not having as competitive an offering, because you have only focused on one of the aspects – on silicon, on the package itself, or on the board and the end user. The suppliers and producers who have all three of those, effectively will be the ones who succeed when we look at packaging in '97, '98, '99 and the year 2000.

BARRY LIEBERMAN: I have labeled my talk: "Packaging Performance Breakthroughs: Is the Desktop Market Ready?" as opposed to "Is the Industry Ready?" because, from Intel's standpoint, the big volume market still remains the desktop with the mobile market coming up behind. Of course, Intel is a consumer of package assembly and test technologies, and it is an enabler for us to provide, basically, computing power to our customers. I want to compare, what we would call the back end of manufacturing, package assembly and test, to the front end, namely wafer fab. How these two parts of the semiconductor manufacturing process has, what it is relative contributions have been, to the performance treadmill we have been on, basically since the late 1970s, when the first wide scale use of microprocessors occurred on the desktop market.

Look at how the clock speed of semiconductors has increased. With the introduction of the 8088, it had a clock speed of about 5 MHz. The Intel product family, has current clock speeds of 200 MHz with Pentium and Pentium Pro. This has basically been a 40-fold increase in 15 years – a phenomenal increase.

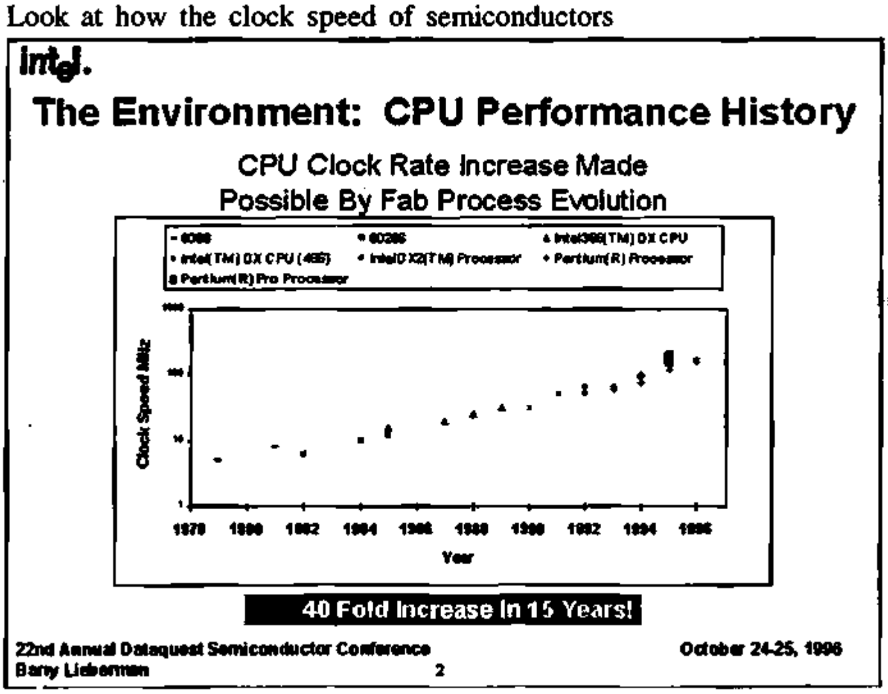

What has made this clock rate treadmill possible? The scaling of wafer fab processes and the scaling of transistors, which then run faster and faster. That is the first primary component of the performance treadmill we have been on in this industry. It is not only the fact that the wafer fab people have had so much success scaling silicon technology, but that success has come at a price. That price has been a large capital investment the industry has made, in wafer fabrication technology – the billion dollar fabs, the whole infrastructure of the capital equipment industry that drives this whole thing.

If we ask a metric which is more directly related to the performance the user sees, called the Y scale relative performance, this too is also noticed on a log scale. This performance trend, over the past 15 years, combines both the wafer fab process which has allowed higher clock rates, and it has also allowed transistor densities, because as you shrink the transistors not only can they run faster, but on the same square centimeter of silicon, you can pack a lot more in.

A third thing that is covered here is architectural innovation, which has driven this relative performance ratio. Here we can measure, treating the 8088 again in the 1979 time scale as a relative performance rating of one (back in that time about the only performance measure was the so-called MIP), and then going on up to now with the Pentium Pro at an increase of 1350 times. That has been driven by the clock rate, the fact that you are now doing five and six million transistor

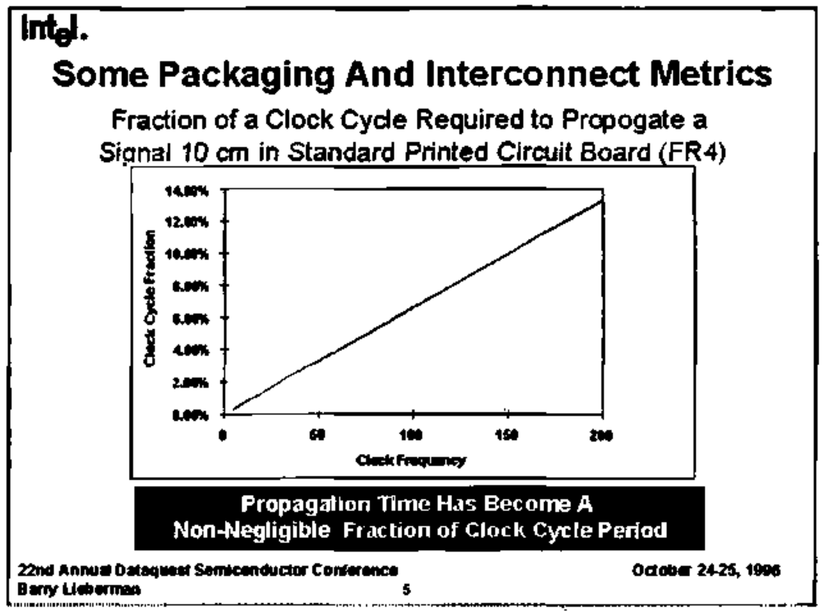

parts. The buses are wider; there is internal cache. If I take out the clock rate, if I divide by the clock rate scaling of 40, I am left with another 35 or 40 that has come through these other innovations, namely architectural and transistor density.

Because the relative performance metric has changed, the earlier parts are measured in MIPs and the later parts are measured in metrics such as spec int. ratings. there was a time at the Pentium where both were being cited, so I used that to normalize the curve. It is a straight line on a log scale. You can imagine if you plotted it on a linear scale, what it would look like. It would be just an exponential shoot-up: 1350 times in 15 years. What other industry can claim that kind of product enhancement over such a short time?

The next thing is packaging and interconnect metric as opposed to a silicon-driven performance metric. I want to plot a historical trend in the number of required interconnects, again using the same X axis, namely the evolution of the Intel microprocessor family. On the left scale, are the number of package pins starting at the 8088 which came in a 40 pin dip package, and that is the solid line. The dashed line, which refers to the right scale, is the percentage of pins on each generational product that has been used for power distribution as opposed to I/O. There is a couple of very interesting things to notice here. The first is that these are not log scales. The ability to drive package density, or package interconnect density of which number of pins is a metric, we can only claim a linear improvement, as opposed to a logarithmic improvement. The second thing is, as the absolute number of package pins has increased, notice what has happened to the percentage that are devoted to power and ground. I think this fact is not well appreciated sometimes.

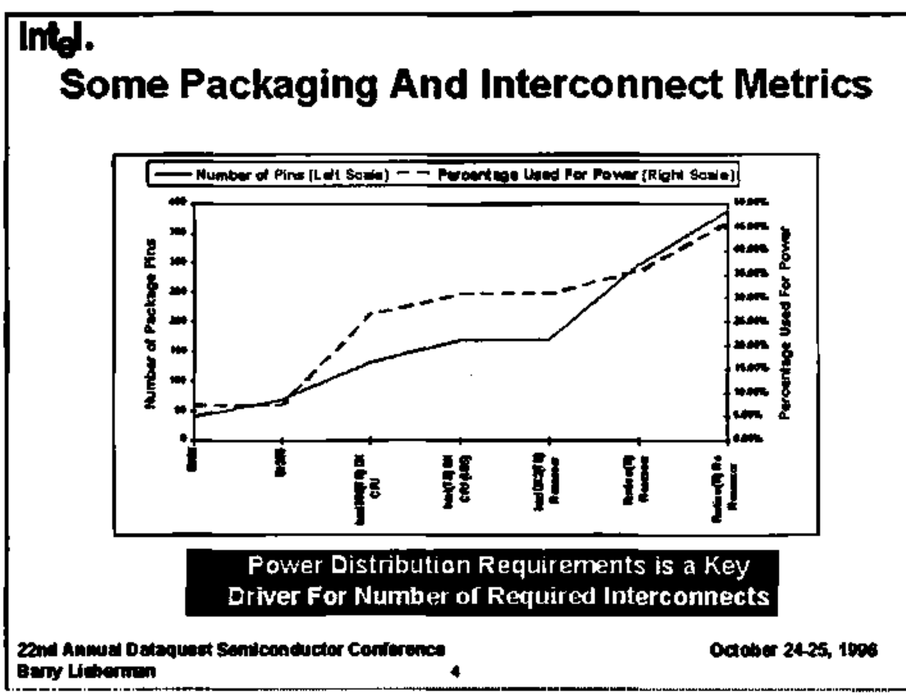

When we get into these large microprocessor dies, and we look toward the next generation of die, the power distribution problem of supplying clean power to the die is becoming a critical problem. In the case of the 8088, there may have been two power and two ground chips, and a Pentium Pro processor with, in round numbers, 400 I/O pins of which 200 of them are devoted to power distribution.

Why is that? One thing that maybe also is unappreciated when you scale the silicon is, even though, as the scale of the transistors get smaller, even though the power per transistor goes down, because you are packing so many more transistors in, the voltage of course goes down, which means your noise margins get much poorer. On a 3V logic its noise margin is a bigger problem than on 5V logic and as we tend toward 2V logic it is going to become even more of a problem.

The other thing is the absolute current that you are supplying is increasing. The overall power, because of the transistor density, is increasing. The scaling of the silicon process is slowing down this trend. It would probably not be physically possible to build these parts if not for the fact that the power per transistor is going down. At the same time the net trend in power and current is up, to the point that we are going to be talking about 50 watt parts on the desktop. The challenges of providing clean power to these parts so that they work properly in getting the heat away, becomes a major challenge.

Here is another metric I would like to discuss. If I look at standard circuit board material, namely FR-4 (FR for fire retardant), which is still in the desktop market because of the cost issues, and primarily what is used as an interconnect medium for chips, and I ask, How long does it take, what fraction of a clock cycle does it take to move a signal between a microprocessor and a cache memory. That may be as close as we can get given the packaging density we are able to achieve. In the days of the 8088 when we were down in the four or five MHz realm, the clock cycle fraction required for transit time of signals is basically irrelevant. It is negligible.

In 1992 at the time of the so-called DX2 chip (the 486 or the Intel DX2 chip), basically that chip represented a significant departure in architecture because of this problem. That was the first chip to use the so-called clock doubled architecture where, because of the propagation delay on the board and the problem of transit time of signals, it was decided to run the silicon core faster than the I/O. That has been the trend ever since, and the reason for that is this interconnect metric problem.

What kind of challenge do we face given this as the situation? Packaging interconnected test technologies has not yet been a major contributor to the metric of the 1350 full increase in desktop computing performance. Packaging has contributed little to that increase, maybe nothing at all. It has all come from the more silicon-centric side of this industry, fab process and architecture, and if anything the increase in packaging density has been less than 10-fold. A lot of the packaging enhancements, at least on the desktop market, have not been driven by performance but more by the mobile market and the need to package chips for laptops.

We face major challenges as an industry going forward. These problems are possible to fix, but there is one other issue. From Intel's standpoint, it has to be fixed in a way which is not a niche solution. Intel has always maintained that there is a "sweet" spot for the mass market, which is the box that has to cost $2,000 or less. Solutions which drive the cost of the silicon up, and then cause the integrated cost of the CPU or the computing solution to get well above that $2,000, is not considered a solution. This is a cost sensitive problem.

Test represents a similar challenge. Test is basically a measurement technology. In any measurement technology you can define something that we call a Process Over Tolerance Ratio. This is when you want the precision of the measurement tool or tester to be ten times better than what you are trying to resolve in the measurement. Otherwise your measurement tool is no better than what you are trying to measure. An interesting thing about this industry is this is possibly the only place where we are forced to use measurement tools that use probably a generation of silicon that is one behind the silicon that we are trying to test; if we are trying to test the latest CPU generation. That is another big challenge.

So what do we do about all these things? Speaking from the side that is a consumer of these technologies as opposed to a supplier of these technologies, we have heard about things such as multi-chip modules to solve the transit time and the flip chip, other kinds of high density ways of packaging, known good die. If we are going to package die in a multiple configuration, known good die is clearly the technology required in order to keep the costs down. We need the same ability to not add cost as we move through the manufacturing process. Known good die has to solve two problems: it has to solve the burn-in problem, and it also has to make some measurement of speed of the device.

There are the new test methodologies to solve the P over T problem and of course there is the cooling. Those are like five technologies that are ripe for exploitation to solve this problem. It has to be done in a manner that is cost effective, and allows a company like Intel to make tens of millions of these parts. A lot of these technologies have been emerging for too long and they have not yet shown that they can become mainstream.

The interesting thing that characterizes the companies that focus on supplying back end solutions, relative to their fab counterparts, is that they are small companies. If revenue is plotted against R&D expenditures for various companies, there are companies engaged in selling back end solutions, namely package assembly and test, and the leading fab suppliers. Even though the regression line of basically a 10% R&D expenditure against revenue seems to provide an excellent regression fit, we can only identify one company in the back end, which runs at a billion dollar revenue rate. Almost all the other companies, being the fab suppliers, are clustered at a much smaller size.

There is a whole host of "Mom and Pop" companies that supply technology which have been ignored in terms of tracking industry trends and this has to stop. The key message here is back end suppliers tend to be smaller, the industry infrastructure is poorer, and in terms of road maps, have not been as well defined as the fab-centric front end manufacturing. If we do not stop ignoring the back end infrastructure, we will do it at our own peril.

So in conclusion I would like to say that certainly in packaging performance breakthroughs, the markets are ready. We need better solutions. The supplier base probably has to elevate effort, and places for improved funding have to be found. A better infrastructure has to be developed where road maps are being shared. A better coordination between the various facets that all have to play together to move to this next step.

REED BOWLBY: My comments are going to reflect and reinforce some of the comments that have been made by Ray and by Barry. We tend at Motorola to look at, not only the high performance types of parts, such as the Pentium and Power PC, but there are a lot of applications that use not so high performance kinds of parts.

They are microprocessor or microcontroller types of things, but they are very much driven by density. The number of microprocessors that a person is likely to own, looking out to the year 2000, is a pretty significant number. They will be in various places around the home. They will be in the automobile and in communication type of products, and also in the office.

There is a family of devices (we call them

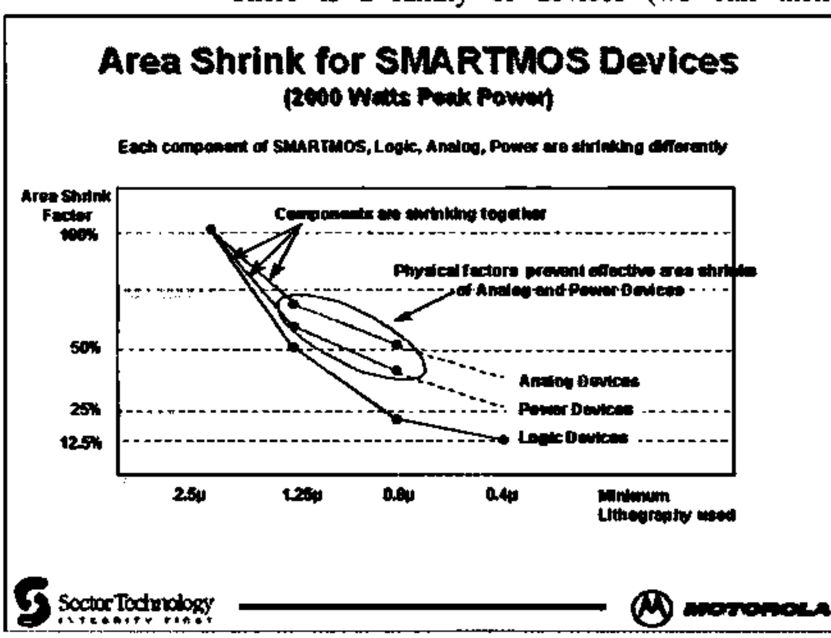

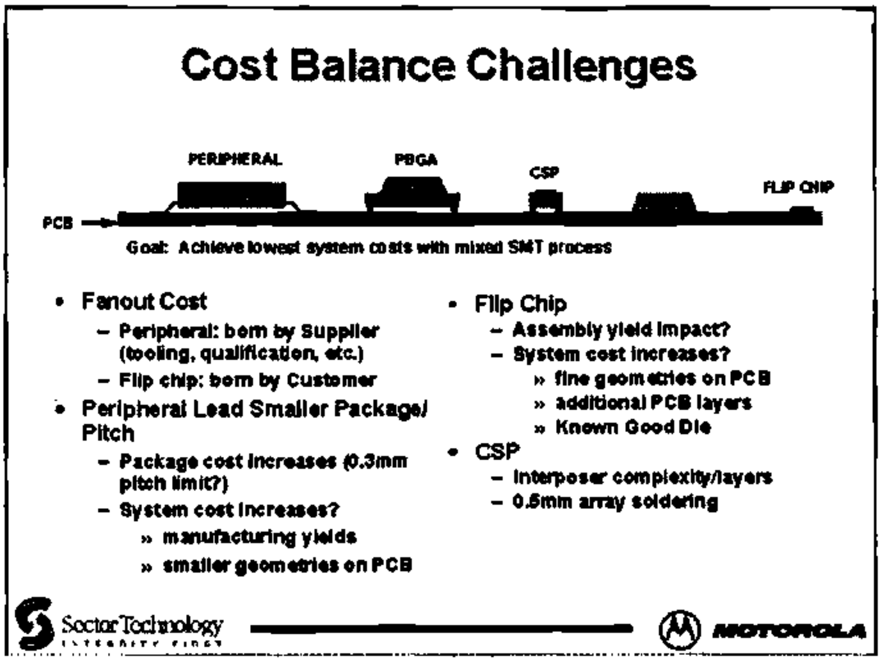

SMARTMOS) that have logic on them, they also have power and analog for controllers. If you look at what is happening in terms of the amount of reduction, these different technologies are not reducing at the same rate. As we try to make higher density devices in these kind of package, or device areas, we are faced with some dilemmas on how to locate the different functions within a chip, as well as the way the interconnect is accomplished.

Some people say that the best package is no package at all. That makes a lot of sense. We are talking about very high performance, high clock rates, fast transition times, so putting bumps on a chip, turning it over, and putting it directly on the using substrate makes a lot of sense.

Flip chips have been around for well over thirty years. So why isn't it the dominant one right now? I believe that it is a cost issue. In spite of performance drivers that may use these things and in some cases there are space drivers that require them, but the cost of attaching bumps to a chip and attaching a chip to a PC board is still quite expensive. Some of the limiting factors, are all related to cost. With flip chips we want to use a lower cost substrate which would be a laminate of some kind – FR-4, BT, FR-5. Because of the differences in CTE we would be forced to underfill those which adds processing steps. It gets even more complex in that when someone wants to use multiple sources, to get the same function from different sources, is going to find it almost impossible to have the same materials (that is paservation materials) on a chip. It complicates the assembly process. If we underfill before we test the system and we have a problem then we have complicated the die removal process, or the rework.

Flip chips are typically 250 micron down to maybe 200 micron pitches on the bumps. That requires an HDI type of PC board which raises the cost. When people design in a device into their chip module or other kinds of module, they want to be able to buy that device forever. We know that semiconductor companies shrink the chips. They can be designed so that the footprint does not necessarily have to shrink with the chip, but it is still problematic. Guaranteeing a footprint from one vendor over a long period of time is problematic, and it is almost impossible when you try to use multiple sources. That is a part of the design that is very difficult for the user to control.

In this thing I call the Cost Balance Challenge, on the left side we show a peripheral leaded device, such as a quad flat pack. The actual fan out then is born pretty much at the device level. When we talk about flip chips, we have these very fine pitches and now the fan out is borne at the PC board level. Again the impact there is increased cost because having to use photo defined vias or build up kinds of technology, is expensive.

The industry has not figured out yet how to define chip scale. I look at it as being just a finer pitched ball grid array package, now that the size of the chip is immaterial. Some of the early pitches that were looked at in chip scale were the 0.5 mm pitch and now, because of finding the escape routing problems with the PC boards, 0.8 mm and 0.75 mm. Even that is a dilemma because in

JEDEC the preponderance of opinion seems to have only one pitch, between 1mm and 0.5 mm and they would prefer that to be 0.75. If you look at it as a grid problem, it is a logical grid because you can populate or depopulate and still stay on a 0.25 mm grid. On the other hand, in Japan, there is quite a bit of use developing the 0.8 mm, which does not fit on our grid. But they make a very good case in that, with 0.8 mm pitch they are still able to have perhaps one or two routing channels for traces between the vias and between the bump pads on the PC board. It is an issue that is not going to be resolved soon.

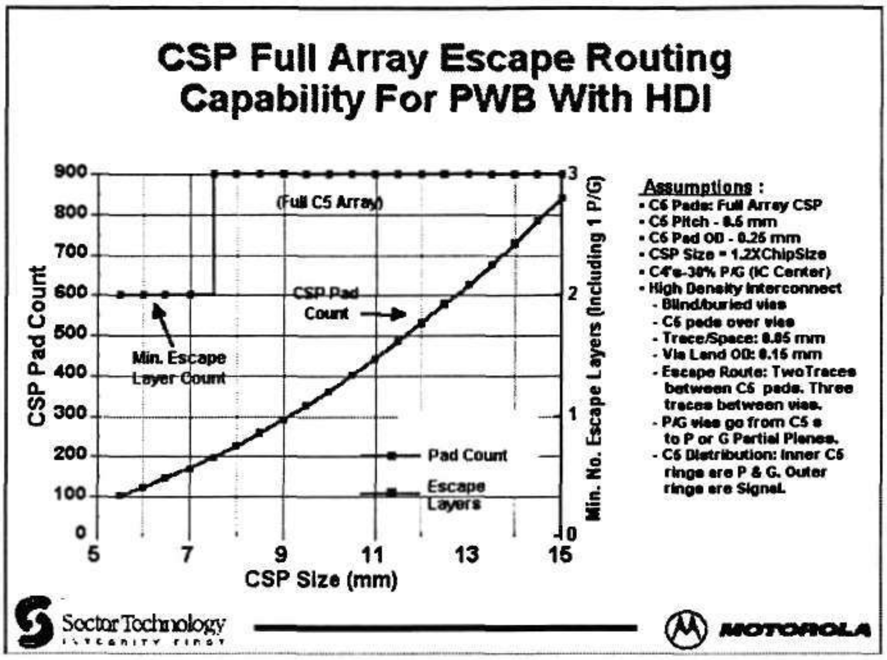

Here is a example of escape routing and I chose here to look at an HDI board. Here we are talking about 0.5 mm pitch, and it is a full array. It is not a depopulated package. We are using typically about 1 or 1.5 mm vias in the capture pads or in the neighborhood of 5 mils, I believe. Looking at this chart we can see that across the X axis the substrate size that the chip is mounted on is in 1mm increments, and the number of I/O that are possible on a full array. The line indicates the number of layers that it takes on a board to escape all the I/O. This is an HDI case and you can route almost any number with three layers of metal. Use the same conditions, on the leading edge PC board, which still uses blind and buried vias, about 16 mil capture pads, and the same size for the pads that the balls are mounted on as well. In that case we would require nine layers of metal in order to route out a 15 mm package. If we take it down one level further to commodity level boards, where we have through hole vias, and even assuming that you have the through hole via located in the pad which will be a great assembly problem, we are looking at reduced size pads, again 16 mils. We find that by the time you get down to the third layer of the board, you have run out of routing channels.

This all suggests that, if we want to balance the cost of escape routing on the PC board with the density we can put on the package, particularly at 0.5 mm pitches, we are going to be forced to look at, not full arrays but depopulated arrays that have two, to three, and possibly four rows of balls around the periphery of the package. We tend in the semi industry to look at reducing cost and increasing density at a device level without a whole lot of thinking about how it impacts the system. We have embarked on this study program, in developing a philosophy at Motorola, that forces us to take a harder look at the impact on the user when we are developing these new packages.

R.P. ST. CLAIR: I would like to talk about the ATE perspective in these various issues and I would like to start out with where we are today, where do we need to be tomorrow and how are we going to get there?

I would like to take look at the memory and VLSI aspects of this in particular in both probe and final test. at one time, the VLSI and memory devices were diverging as far as test requirements. VLSI would require a higher frequency and a much larger pin count than memory. with more and more memory devices being tested in parallel and with the frequencies increasing to match some of the VLSI products the problems become fairly common between the two.

Starting out with where we are today, in probe at memory, we are testing 16 devices in parallel. Frequencies typically do not run more that 60-80MHz. One of the reasons the frequencies are not that high, is because the needle technology, as it is known today with the tungsten cantilever needle, is essentially going to run out of gas when it comes to testing more devices in parallel. Because of the overlap of the needles that are required in order to contact all the pads in anything other than a two by some number of devices in an array.

In final test, we see that presently people are testing 32 devices in parallel, once again memories. Frequencies are at or above 100 MHz. The packages today for memory are typically SOJs and TSOPs and a mixed bag for VLSI, we have BGAs, PGAs, and QFPs. Like the limitation on the probe side, most of the contactors used in the automatic handlers do not have the electrical performance required to carry them much beyond where they are now. In addition there is also, apparently, some mechanical reliability problems with contactors in general which need to be solved in the future.

If we look at where things are going to be tomorrow, we will be at the 32 in parallel level for probing, for memory devices. When people start talking about being able to do that many devices in parallel, and since memory have fewer pins, people all of a sudden start talking about making contact on a full wafer. At one point in time the parallelism was driven by the probe card producers and the testers could always keep up with the number of parallel, as compared to the technology in probe. If a probe card manufacturer, is able to contact a full wafer then the ball is back in the court of the ATE manufacturers to up the parallelism one more time.

Another problem that we see with the wider parts, is that when we test 32 in parallel right now we would like that to be 16 wide memory devices, as an example, and that takes 512 I/O pins. As they go to 32 wide and 64 wide that goes to a 1000 and 2000 I/O pins. That is just the I/O pins, not counting the address clock pins and that is up above what the LSI testers are having to do for pin counts these days. When you talk about cost of tests, adding that many more pins to a tester is expensive. If, by testing that many more parts in parallel, like 32, you can offset those costs, then the cost per device is probably going to be lower. As they get wider, that does not help, if you keep the same number of parallels.

With these devices we will need a new probe part technology. The needle technology, because of the number of needles and frequencies that we are operating in, are typically not going to make it. When we have things like synchronous DRAMs, sync-link DRAMs, MPUs, etc., that are running in the 150 to 800 MHz range, we must maintain a 50 ohms environment directly down to the pad of the die under test.

When we look at final test we see the same types of things happening. At 64 in parallel (which is where they will be, and 128 after that) if you look at some of the tray handlers there does not appear to be an architectural limit on the number you can do. Frequencies like a probe approaching 800 MHz and need a true high frequency contactor that is both reliable and able to carry the signal fidelity down to the lead of the device under test.

One of the challenges we see is the drive for smaller packages, for high performance. The smallest and highest performance is achieved with no package at all. Then if we look at where we are being driven as an ATE industry our customers want us to double the pin count, double the frequency, improve the accuracy, keep the footprint the same, and reduce the cost. These are all wonderful things and how do we get there?

If we can get rid of packages and mount die directly onto our channel cards, this is one way to double the number of pins that we have in a tester without increasing the footprint, which is one big step. This would probably require known good die. If we are going to mount bare die onto a printed circuit board, then we need known good die. When you have known good die you need a die carrier. After the wafer is tested the dies are singulated, put it into a carrier, then it follows the rest of its path through production into final test in this die carrier, then it is removed and the de-

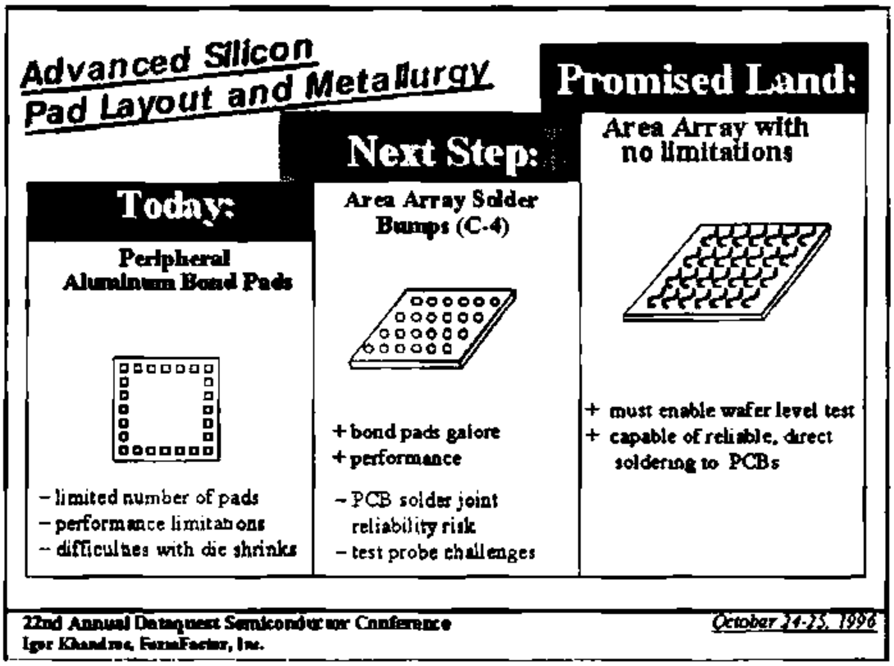

vices are shipped. This is going to increase the cost of the die, not decrease it.

Something has to be resolved in that particular area as well. One of the things with wafer level burn-in, is the ability to eliminate some of the steps required by the die carrier. If we can achieve a wafer level burn-in at test, and whether that means finding some other way to detect the same failures or induce the failures other than sticking the devices in an oven, then something may happen there. Then of course we have the boardattach problem that has been described. Presently we use FR-4 in the circuit boards, and without the fine lines on the surface as well as the small pitch on the vias we will have to go to filled vias, and other types of technologies.

One of the ways that we are proposing to prepare for this, is to make wafer probing part of the core competence of our company. We have to be able to know as much about probing as our customers. As customers demand more and more integrated solutions, they look to a common source to solve their problems, and it is usually where they spend the most money.

We have to have the ability to have wafer probing as part of our core competence. We also have to work with prober companies in order to develop the interfaces that do not require our customers to compromise testing on their parts. In addition, we have to participate in the development of probe part technologies.

If our customers are going to look to us for help and if we are confident in probing, then we need to ensure that we can maintain a 50 ohm environment and high signal fidelity all the way down to the pad of the device under test.

Likewise in final test we also need to make that part of our core competence for the same reasons that we do in probing.

Once again we have to work with the handler people in order to ensure that the interfaces do not prevent the customer from testing the part the way it needs to be tested.

Finally we have to participate in the development of contactor technologies for exactly the same reasons that we have to do in probers.

IGOR KHANDROS: The whole theme for back end industry should be: Let's get exponential! Let's get all that linear trend over with. So I will try to talk about the trends that we see, and our possible recipe, for how to get exponential.

This is a humble attempt to plot where the industry is, and where the industry is going. If you look at the pad layout trends: today we have peripheral pad-out. We have fine pitch peripheral sets of pads and it is difficult to shrink die. Every time you try to do that you go through all the pain of worrying about bonding it front pitch and other things. You have performance limitations for the reasons Barry addressed. You have a number of pads limitations.

The next step that we see happening imminently in the industry, in microprocessors and fast

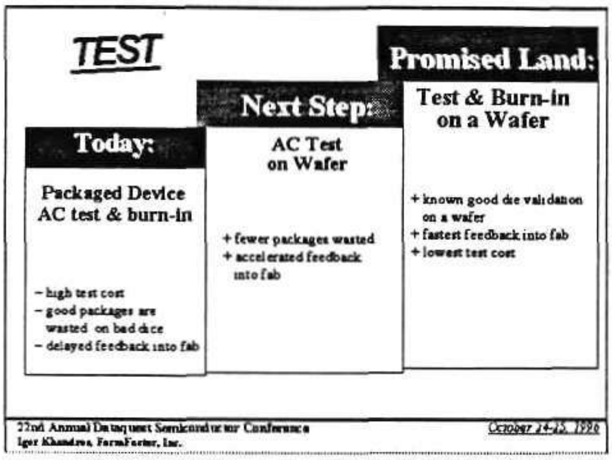

SRAMs, is transition to C-4 technology. The challenge now is to map that old technology onto high volume, low cost infrastructure of the industry, and you obviously get enormous amount of pads. You have performance advantages, especially power, and how power is handled. Reed did mention the reliability risks and there are also challenges with test probes. How do you probe C4 and especially how do you do that at speed? The promised land is area array with no limitations. I will expand a little on what we are trying to do but I think that eventually you would want to be able to drop a pad anywhere on the face of the die not worrying about the consequences.

If you look at the test trends, what we are basically doing today is testing and burning-in package devices. There are difficulties with that approach, especially with expensive packages, feedback into the wafer fab is significantly lengthened. The next step is that there will be more and more AC tests done at the wafer level. The best solution is performing test and burn-in operations on the wafer. Known good die will become the cost effective reality only when those operations are performed at the wafer level.

For the package and assembly today there are a dizzying multitude of packages. There is a cost issue. Package site limits the mobile revolution, and FormFactor driven products demand either packages that are the size of a chip or just bare die. The next step is a transition to flip chip packages, because there will be a transition to C-4. There is higher performance but there are reliability limitations putting it directly on FR-4 or BT resin boards. You have again a cost issue. Then the second immediate trend is chip scale packages and I see this step as a transitional phenomenon. Chip scale packages have the advantage of being a smaller form factor. Being a chip carrier, there is additional cost as compared to bare die. There is also significant infrastructure building to be done in terms of new materials. You cannot just buy those from your existing lead frame supplier. An infrastructure for huge volumes has to be developed at these new companies in the industry. We are completely inundated with requests at FormFactor to develop sockets for test and burn-in of 0.5 mm or 0.8 mm chip scale packages. There is a significant amount of work to do.

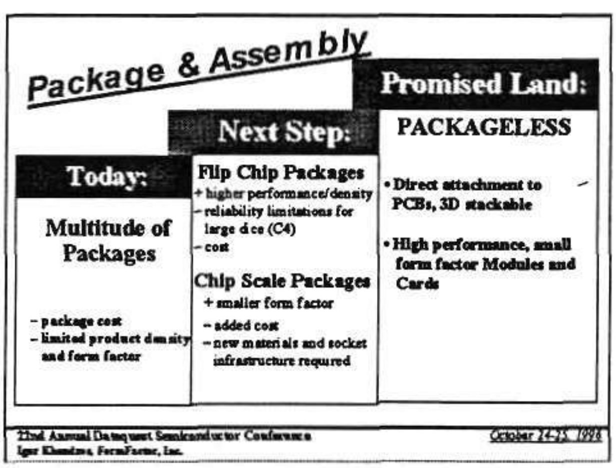

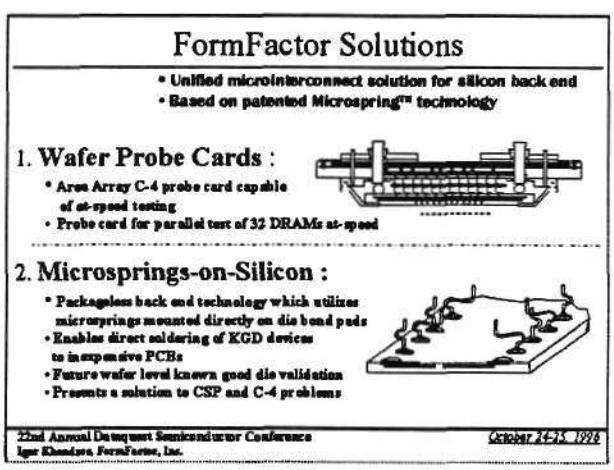

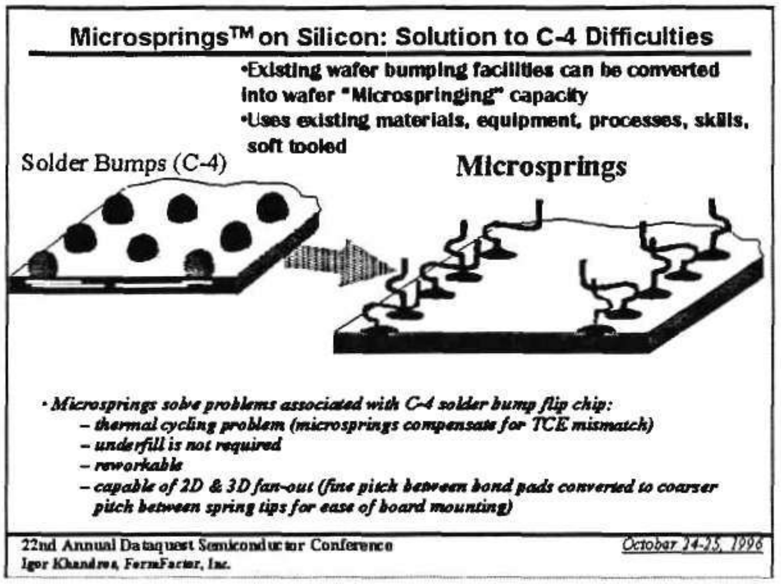

The best solution is packageless, and that will allow direct attachment to PCBs and will enable 3D stackable form factor driven products. Not all of them will be thin. Some will follow an unusual form factor as society goes mobile. What is demanded is high performance, small form factor modules and cards

FormFactor is a company that developed the miniature microspring technology. We can put small microsprings on any substrate, including directly on silicon die pads, or we could put them on a controlled impedance substrate, and make probe cards. It is a unified solution for back end. I feel that solutions that will stay will not address just the package or just test, but will be solutions that address the total pipeline. Back end is a delivery system for deliveries of silicon wafers into final products. Therefore solutions must be system solutions. New wafer probe cards are needed both for C-4 transition and packaging, and also for parallel wafer probing of multiple devices like DRAM. In the next stage in FormFactor business development, we are bringing to market microsprings on silicon, which is packageless back end technology. It is not exactly packageless, it is silicon with little legs attached to it, but we believe it does address problems associated with C-4 and chip size packages.

FormFactor probe card not only looks beautiful but it does serve a very important purpose. It is designed to contact the solder bumps on silicon, and we are now developing the ones that would test aluminum bond pads. It is a complete solution. This probe card is preplanarized when it is shipped to the customer. It provides a 50 ohm impedance substrate, with very short springs as opposed to existing tungsten needle technology, and we believe it is scaleable to large areas, possibly to wafer size. It has been analyzed to potentially be capable of operating at one gigahertz frequency, and beyond.

We would like to use our expertise in test, and move onto what we see as a unified solution for how back end flow is done. We will leverage our ability to put miniature springs directly on silicon bond pads. We are capable of doing 3D fan out by changing the shape and dimension of the springs, so a DRAM that is a five mil pitch, could conceivably fan out to 25 or 30 mil pitch. We have now wonderful reliability data on thermal cycling and it does not require underfill.

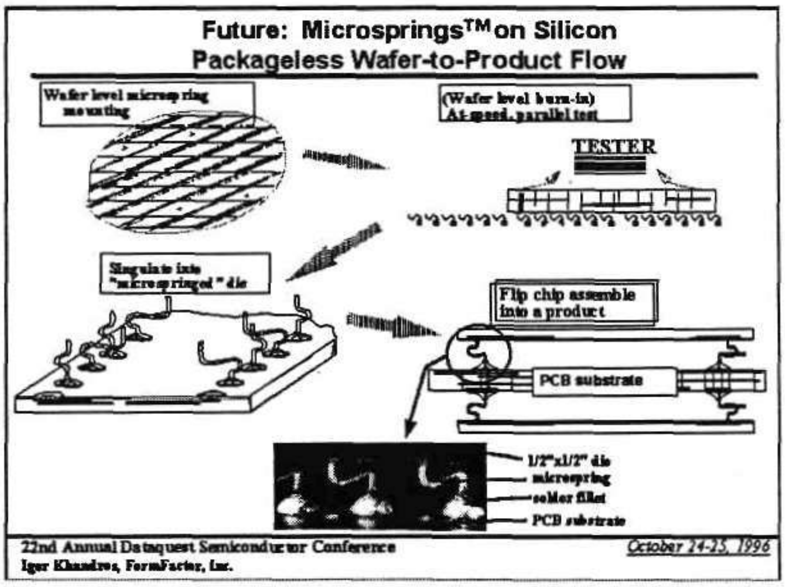

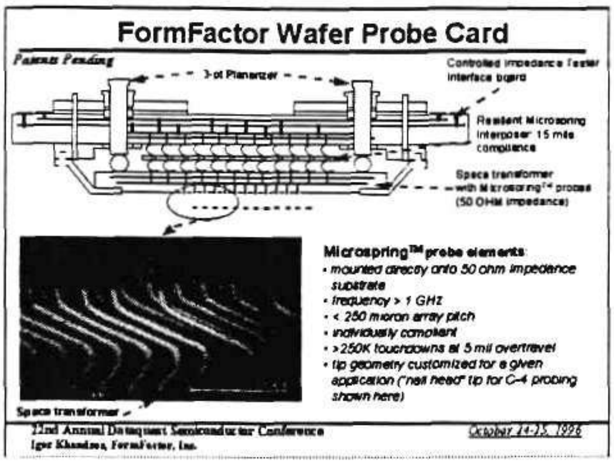

With springs being put directly on wafers, you have a wafer with its own socket sitting on it, and you can just use a substrate to test several devices, or, it can simplify wafer level testing and burn-in. Then you singulate these wafers and end up with die with little legs on it, and you can solder it on both side of a substrate. That, is a simplified solution, and a possible solution for some of the problems associated with chip size packages. Companies that are capable of doing bumping on wafers, with FormFactor providing technology to them, will be capable of, instead of bumping wafers, "springing" wafers. Once they do that, the thermal cycling problem is, potentially solved. Underfill is not required, it is reworkable, and capable of 2D and 3D fan out.

The question for this panel was: Is the industry ready? Companies like FormFactor are proof that the industry leaders are ready. That is very important, because the industry looks today somewhat different. There are large players who actually have the ability to move infrastructure. There is something going on today that I did not see three or four years. The users of silicon, and the silicon are pressing companies to come up with miniaturized high performance solutions to packaging. Technologies and methodologies that will win in this transition, will address the total system. It is a system, and it will include addressing test and assembly issues, and the very important question of board routing.

Companies worry less and less about NIH. Large companies, that are interested in solutions create a team of several suppliers and ask them, even force them, to work together to solve their problems. These partnerships will accelerate the infrastructure.

We are interviewing PhDs from Berkley, Stanford, and CalTech, and MIT. Their professors tell them that the back end is really a very interesting area now to focus on. It has become an area where we attract talent.

NADER: We have heard all the proponents or the exponents about chip scale, flip chip, migration, and the issues involved with known good die, and what the definitions are, and then definitely the promised land of 'no package.' Can you comment and elaborate on a few promising chip scale and also flip chip implementations? How affordable are they, and what are the price issues involved for commercialization?

REED: There are very few actual chip scale implementations other than some things coming out now in products like Sony's most recently introduced hand held video camera. The preponderance for those is a stud bump version of a flip chip onto ceramic substrates and organic substrates. I do not really know how to address a cost right now. In our company we have not yet put a two scale package on the market.

We do have products that are shipping with flip chip. We have been supplying flip chips to some major customers for well over twenty years. They have been low lead count, so they are not addressing the issues that we are talking about here, with the high performance, and again, these are low cost. The flip chip market is segmented into two areas: one is the high performance, high lead count things that are used in high end computing, and the other is way down on the consumer range. A good example is a Citizen watch company that assembles one million flip chips a day, but they are only 15 leads, and they are very inexpensive chips. They do not worry about known good die, they do not have the issues with underfill – if the thing does not work, they throw it away. That is on the low cost end, but I do not know how to put a dollar figure on that.

RAY: Whether you use a stud bump process with a wire bonder and break the wire off, or plated bump technologies, this is clearly one of the things that the industry in general is very aggressively driving for. There are a number of industry groups that have set a target of getting wafer bumping cost down to less than $50 a wafer, to make it an attractive alternative for wide scale products. There are both industry groups and some of the joint consortia, MCNC for one, who is looking at plated technology as a way to get that bump cost down, and help build the infrastructure up to the point that it will be a cost effective process for the higher lead count devices

Q: Do you have a figure of merit? If we have the same size in pin count, and we go from the traditional package technology to flip chip, what would be the cost increment?

RAY: People that are moving into flip chip today are doing so based on it being the most cost effective alternative for that specific function, and typically it can be done in wire bond. It will be cheaper to keep it in wire bond. It will the impact on either die size, package cost, or inability to meet the functional requirements, giving you a less valuable product in the market, that would push you into flip chip. There are few cases today where it will be cheaper in flip chip. I do not know many, if any, that immediately jump to mind

REED: I can give you some examples where a power PC, for example, is designed as a flip chip. Right now we are using evaporated bump technology which is probably in the $300 per 8 inch wafer range, just to put the bumps on the wafer, and the quotes for electroplated wafers are somewhere in the $150 to $200 per wafer. The guesstimates for electroless plating are somewhere around $80 to $100 right now, but the goal is to get that below $25.

BARRY: What will drive the high end microprocessor desktop market to flip chip is this problem with power distribution, because going from the periphery of the die wire bond, to the center of the die where you need to power the transistors is simply not going to be doable. So this idea of arraying the bumps, or arraying the interconnects, most of those interconnects will be powers and ground. The economics of doing a Citizen watch chip and a Pentium Pro would be significantly different.

Q: On the test side, what do you see as the most significant technology challenges? I would like to add something about next generation devices. We keep hearing about the diversification of the silicon, and it seems it is going way beyond just the generic VLSI versus DRAM, the system on the chip. What are the barriers?

RP: Cost is always one of them. The number of pins is one of the driving factors to the cost of a tester. Once you have an architecture that allows you to expand as far as you want to, then it does not become an issue of whether you can do it, but whether you can afford to do it. As the performance gets to a point, like I mentioned 800 MHz, that is pretty high. You have a cycle time of 1.2 nanoseconds, and within that, you have to do things. That is tough, and I am not sure how that one is going to be addressed.

Q: How are we going to address that?

BARRY: The reason for high frequency testing (which we call functional test at Intel), is basically to guarantee the performance of a part at its external pins to a data sheet. In order to do that you throw in some headroom (called guardbanding) which is like error stacking. When you start doing the stack up tolerance, you run out of margin, where you cannot guarantee the performance of a product to a data sheet. So test flows in the future need to comprehend this. Test flows will certainly change.

IGOR: It was very interesting what RP talked about, to make his testers faster he actually has to come up with multi-chip modules for pin electronics. The solution is packageless, whichever way we can get the highest performance and wafer level. As much testing and validation as possible should be done on a wafer, and at some point, possibly wafer level, or total wafer, and really coming up with a way to tile silicon on an inexpensive substrate. It is the same old challenge of assembling several pieces of silicon on a common substrate. Whereas getting off one die and getting on to another die, there is no performance penalty, as compared to the case when they are on one monolithic piece of silicon. When that happens I think the back end will be on the same exponential growth.

I think it requires a very serious effort by the industry leaders, because the industry leaders are going to push this infrastructure, and infrastructure is going to go where they go. Companies possess the influence to convince providers of back end solutions, packages, testers, probe cards, to really invest in the future and it will take a unified effort.

NADER: Could you share with us Tesarah's approach and comment.

REED: Their approach has been to use a tab-like structure that is a fan-in tab with an elasmer between the polyimid film in the chip as a stress barrier. Right now, I think that the concept looks good, and a lot of people are looking at it. There is no manufacturing experience with it, and there are a lot of different approaches to chip scale. It complicates the whole thing, because you do not know which one is going to win. One of the key areas for chip scale is it has to survive through thermal cycling on a PC board for a substantial number of cycles before the joints break, and that if underfill is required for chip scale, it will kill it. Chip scale would not be a viable way of doing it. There would be no reason then to not go with a flip chip other than the fine pitches that are required on the board.

Q: I have another question, which has come up several times during the conference, and that has to do with DRAM, and DRAM testing. We know that bit lengths are getting larger. The 1 by 16, how much was there a bottleneck for the introduction of 1 by 16? The fact that, actually, the 4 meg remained at $13, the one by four, where people were actually looking for the 1 by 16, and the prices remained that high. Actually, it spurred the capital spending, and we had under-capacity and so forth. How much of that was the capability, or lack of, from the test side?

RP: When the 16 meg came out, people were building it in a 4 meg by 4 configuration, and all of a sudden, realized that what the world really wanted was a 1 meg by 16. The capacity switched over from the two, and it took a while to get it ramped up.

It does have an impact about the width. As the width gets wider and you want to test, 16 devices in parallel at probe, and they are 16 wide, you need 256 I/Os on the tester. On a memory tester, the I/O pins are the most expensive part of the tester. It is not the address clock pins, and as the widths keep going up, to 32 and 64, we cannot drive the cost of the tester down. You have to provide those pins. You can give less accuracy or worse timing, and get cost out of it, but to keep the same performance, and double and quadruple the number of I/O pins required to do it, you cannot keep the cost out of it.

Q: Will the industry have the same transition pains to the 1 by 64, or 2 by 32, for the 64 meg?

RP: The 64 meg, people are already looking and planning on what that is going to be in various widths. As such, when the 64 meg is in production, I think there will be equipment there to test it.

RAY: The point about some of the problems that arose in that transition have caused people to think about the applications that are driving those specific memory architectures, which one is likely to be more prevalent, and perhaps looking a little more carefully, so that the test strategy is planned out and ramped up at a more consistent pace with the volume.

NADER: Any comments from the panel?

IGOR: An interesting question is: What is the overall contribution of test cost to the cost of DRAM? I hear numbers like 20% now when you go to wide DRAM and it is going to go up. When we go through times when price and cost of DRAMs are almost the same, it becomes a pretty serious business. We see a huge amount of interest in trying to drive the test cost down, and I think Teradyne is seeing the same.

Q: Which companies have used or tested your microsprings? What potential problems could arise with using microsprings, and how does the cost compare with other ways of building the structures for test and bonding?

IGOR: We are starting to aggressively push microsprings and silicon technology. In the beginning we are going to focus on memory, and we have gone through internal modification. We have two programs now with memory companies, and they are going to go through modification. So we have not gone externally through that yet. The goal is to make sure that the overall cost is lower than even the commodity packaging, and then on top of that to enable form factor driven systems with this technology. We hope that it will be used, not onlybecause it is a way to solve problems with CSPs and C-4, but a way to really reduce overall back end costs.

Q: What are the potential reliability problems, and how are they actually made?

IGOR: We have reached over 2000 cycles of -55°C to +125°C. We are starting -65° to +150°C thermal shock testing. In terms of reliability, it seems to be very good, at least the initial tests we have done. We are going now through a complete battery of tests. We have patents pending on the process, but basically we use a combination of several processes, that includes micromechanical automated equipment, and the rest is processes, and simplified wafer back end processing.

Panel:
1997/1998 Capacity Status: What Will Be Tight and What Won't?

Eberhard Klasse

Vice President,
Sales and Marketing,
Wacker Siltronic Corporation

Mr. Klasse is vice president of marketing and sales for Wacker Siltronic Corporation in Portland, Oregon (part of an exchange program for key personnel to enhance Wacker's global approach to the global semiconductor industry).

Mr. Klasse began his career as a salesman for PVC in Stuttgart, Germany. Other responsibilities in his career have been sales manager for Vinyl Acetate Polymers in Munich, Germany; salesman for PVC in Stuttgart, Germany; in charge of market development project for Specialty Polymers in New York City; sales/marketing manager for the Polymer Division in Munich, Germany; marketing manager for the Silicones Division, Wacker's largest division; manager in Wacker's Corporate Planning Group; and vice president of marketing for the Semiconductor Division in Burghausen, Germany.

Mr. Klasse received business administration education in Germany.

Panel:
1997/1998 Capacity Status: What Will Be Tight and What Won't?

John Luke
President,
TSMC-USA

Mr. Luke is president of TSMC-USA, a wholly owned subsidiary of Taiwan Semiconductor Manufacturing Company (TSMC). He has more than 30 years of experience in the semiconductor and electronics industries.

Before joining TSMC-USA, Mr. Luke was vice president of Sales for Monsanto Electronic Materials Inc. His previous positions include vice president of Worldwide Sales at Signetics Corporation; vice president of Marketing and International Sales for American Microsystems; vice president of Marketing and Sales, Monolithic Memories; and vice president of Sales, Fairchild Corporation. Mr. Luke also spent 10 years with Texas Instruments in design and sales management positions.

He received a B.S. in electrical engineering from Marquette University in 1960.

Panel: 1997/1998 Capacity Status: What Will Be Tight and What Won't?

Evert Wolsheimer

Vice President,
Product Technology,
Xilinx

In 1991, Dr. Wolsheimer accepted a position at Xilinx Inc., San Jose, California, as director of Technology Development. Early in 1996 he was appointed to vice president of Product Technology and is currently responsible for product engineering, wafer foundry, silicon technology, assembly, and reliability engineering for all Xilinx programmable logic products.

Dr. Wolsheimer joined Philips Research Laboratories Sunnyvale (PRLS), California, in 1982 as a device engineer in the process development organization. He became manager of the Device Physics Group at PRLS in 1985. In 1987 he moved back to the Netherlands, as a department head in the joint Philips-Siemens MEGA project, where he was responsible for device development for the 1M SRAM products. In 1989 he joined Philips Semiconductors in Sunnyvale, California, as the marketing manager for nonvolatile memory products.

Dr. Wolsheimer received a Ph.D.E.E. from Delft University of Technology, the Netherlands, in 1982.

Session # 22: 1997/1998 Capacity Status: What Will be Tight and What Won't?

Moderator:
Clark Fuhs, *Director & Principal Analyst, Semiconductor Equipment, Manufacturing, and Materials Program, Semiconductors Worldwide, Dataquest*

Panelists:
Eberhard Klasse
Vice President, Sales and Marketing, Wacker Siltronic Corporation
Arthur W. Zafiropoulo
Chairman & CEO, Ultratech Stepper Inc.
John Luke
President, TSMC-USA
Evert Wolsheimer
Vice President, Product Technology, Xilinx

CLARK FUHS: We have a distinguished group of panelists to talk about infrastructure issues and capacity issues in the semiconductor industry and we expanded it to some of the manufacturing infrastructure issues that have to develop over the next several years as well.

How does the silicon area split out, in terms of what types of devices are being manufactured in the fabs today? Power/Discrete consumes about a quarter of all silicon consumed worldwide in terms of square inches. Very low revenue per square inch, but a lot of square inches, so a lot of capacity is required in order to feed that particular need. In terms of microprocessors and leading edge technology and memory, of that memory only about 75% is DRAM. About 30% of the world's capacity is driving the leading edge technology. Everything else tends to trickle down to

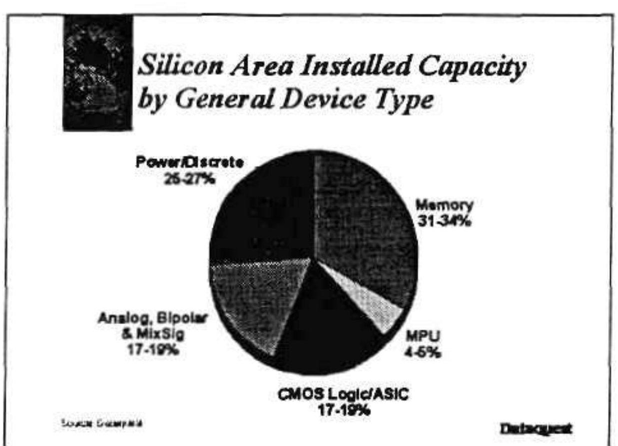

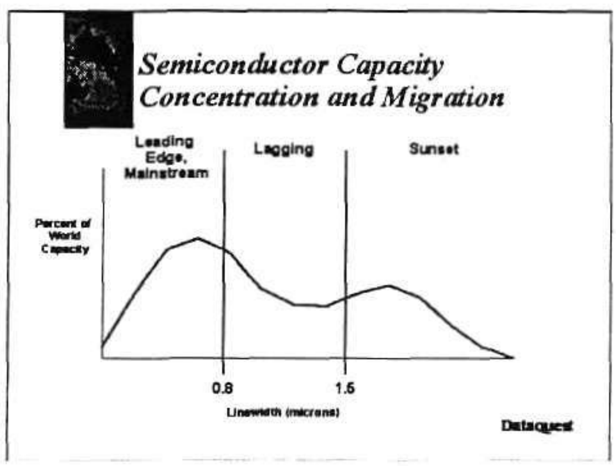

feed capacity as time goes on.

When plotted another way, the capacity in industry migrates through product channels. The majority of capital spending is in leading edge and mainstream products, and then above 0.8 micron, it would migrate and trickle down into some lagging technology areas. There is a segment of our business that is in the sunset area, above 1.5 micron. The types of products that are used in the lagging area, are microcontrollers and the telecom chips, mixed signal, smart power.

These are some of the fastest growing products areas in semiconductors today, and an infrastructure change that occurs in the middle area will increase in importance. As capacity trickles down this chain of product migration, there will be less

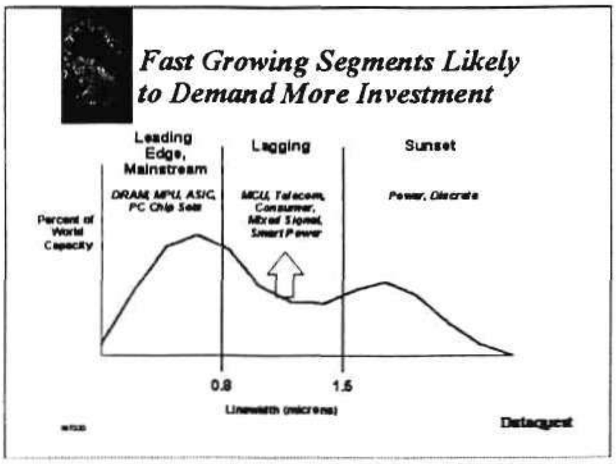

to trickle down as time goes on. More capital spending is going to be required above 0.8 micron.

When you look at the types of new fabs that are being announced today, between power/discrete and consumer telecom, there are 13%, and roughly about 15% of all new capacity in the way of new fabs, is in 0.8 and higher. This has evolved a whole new set of equipment companies that are targeting these lagging and sunset technologies, Ultratech Stepper is one of them.

Another area of infrastructure that is developing is the foundry, and the fabless model. About 9% of all silicon that is processed goes through a foundry today, and in the next five years, we expect that to grow to about the 15% or 16% level, and ultimately could be about 25% or 35% of all production is done through a foundry.

To introduce our panelists: Eberhard Klasse is VP of Sales and Marketing for Wacker Siltronic,

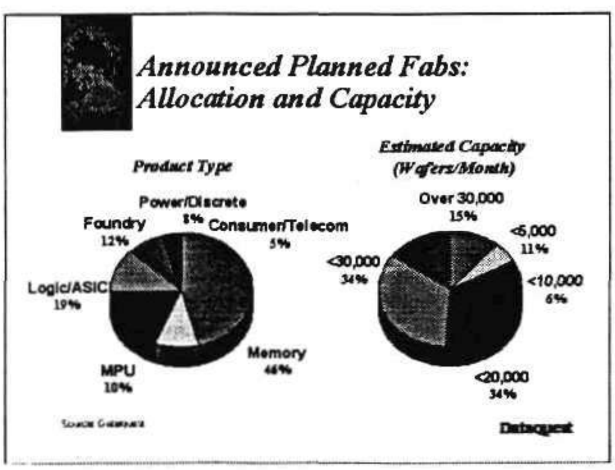

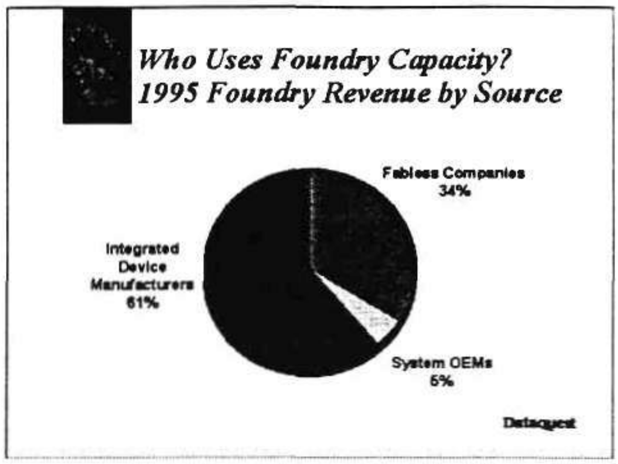

one of the leading wafer suppliers in the industry. Art Zafiropoulo is Chairman, and CEO of Ultratech Stepper, one of the five existing and major stepper companies in the world. John Luke is the President of TSMC-USA, the largest dedicated foundry in the world today, and Evert Wolsheimer is the VP of Product Technology of Xilinx. Our first speaker is Eberhard Klasse.

EBERHARD KLASSE: Wacker, my company, is one of the pioneers in polysilicon as well as in wafers.

Let's look at wafers first. After growing 26% in 1995, it looked like wafers would grow another 30% in 1996, until mid-year. Since then, growth came down steadily and we estimate that, by the end of this year, it will have come down to a growth rate of about 12% for all of the silicon. There is still a solid growth of over 10% for the 6 inch wafers, and of even over 50% for the 8 inch wafers. Compared to last year, 4 inch wafers and 5 inch wafers, have come down.

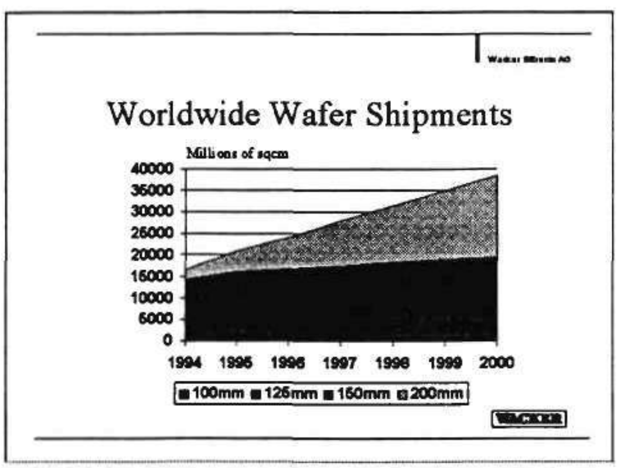

We expect this pattern to continue, with a growth of the 200 mm slowly going down as these wafers command a higher and higher share of total wafer area. Right now, their share in total area is about 30% and, by the end of this century we expect that they will be at 50%. The compounded growth rate until 2000, as we estimate it, will be about 27%. The compounded average growth rate for total silicon we see in the range of 12.5%.

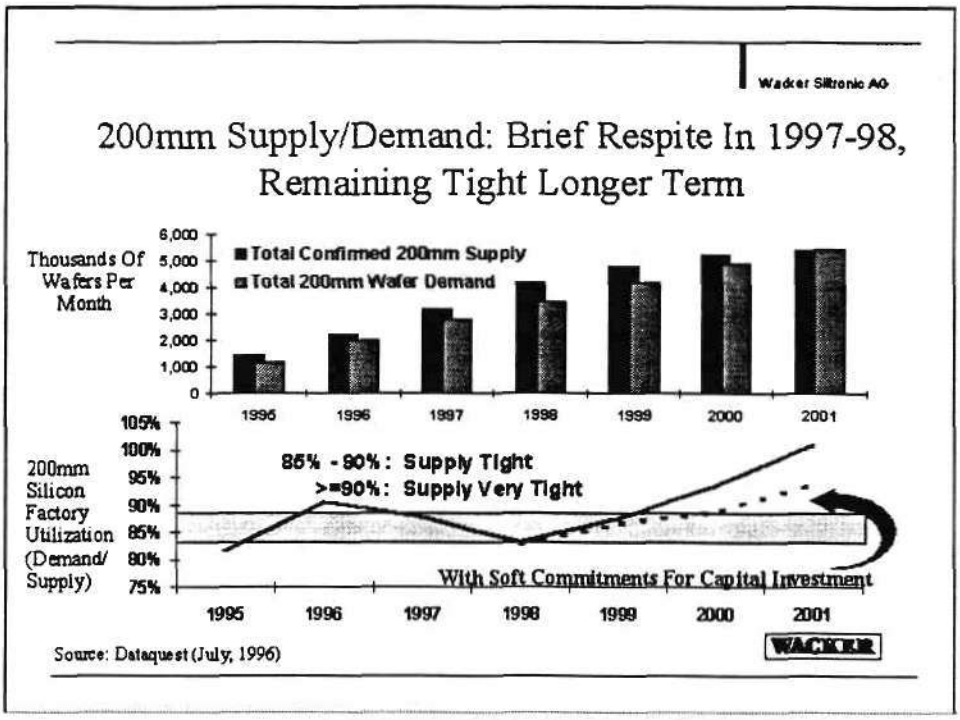

With most of the growth coming from 200 mm, I will look into that diameter a little bit more. Here you see the distribution of the regional markets in 200 mm as well as the distribution by product type, being polished and epi. As you can see, the U.S. and small Europe command the major share of the epi wafers worldwide, whereas Asia and Japan command the highest share in polished wafers. Both merged together, the picture of the polished wafers prevails, because epi is only 15% of the total.

Now a word of caution, as far as forecasting is concerned. Neither partner in this industry, I think, has had a good record, and so we focus on the year 1998. The development of our in-house forecast as well as Dataquest's forecast, and you can see very well, the peak of our expectations was in October '95. Since then we have lowered them quite a bit.

As far as the supply is concerned, our own assessment is in line with that of Dataquest. Wafer shortage should no longer prevail. Basically, the supplies for 1997 were geared at the earlier, higher expectations, so there should not be any major problem in wafer supplies for 1997 and the years to come. It is extremely difficult to make an exact prediction now for epi wafers, with the 64 meg DRAM being partially on epi. Our assumption is that with the major expansions in the industry underway, epi wafers should not be tight. Actually, they were hit most in the slump of Q4

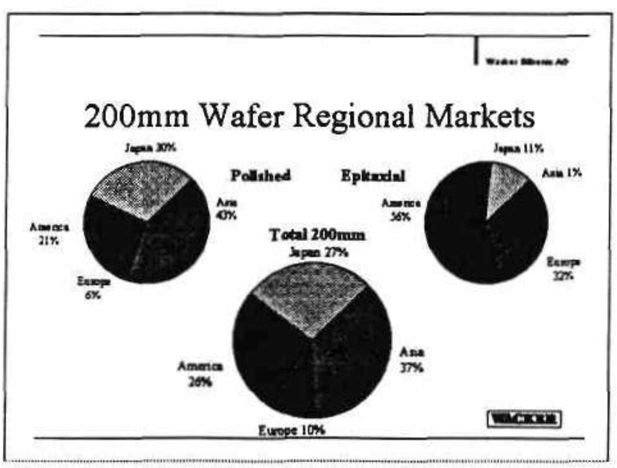

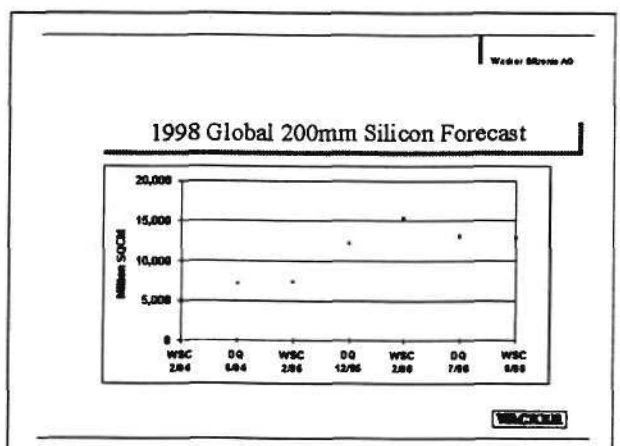

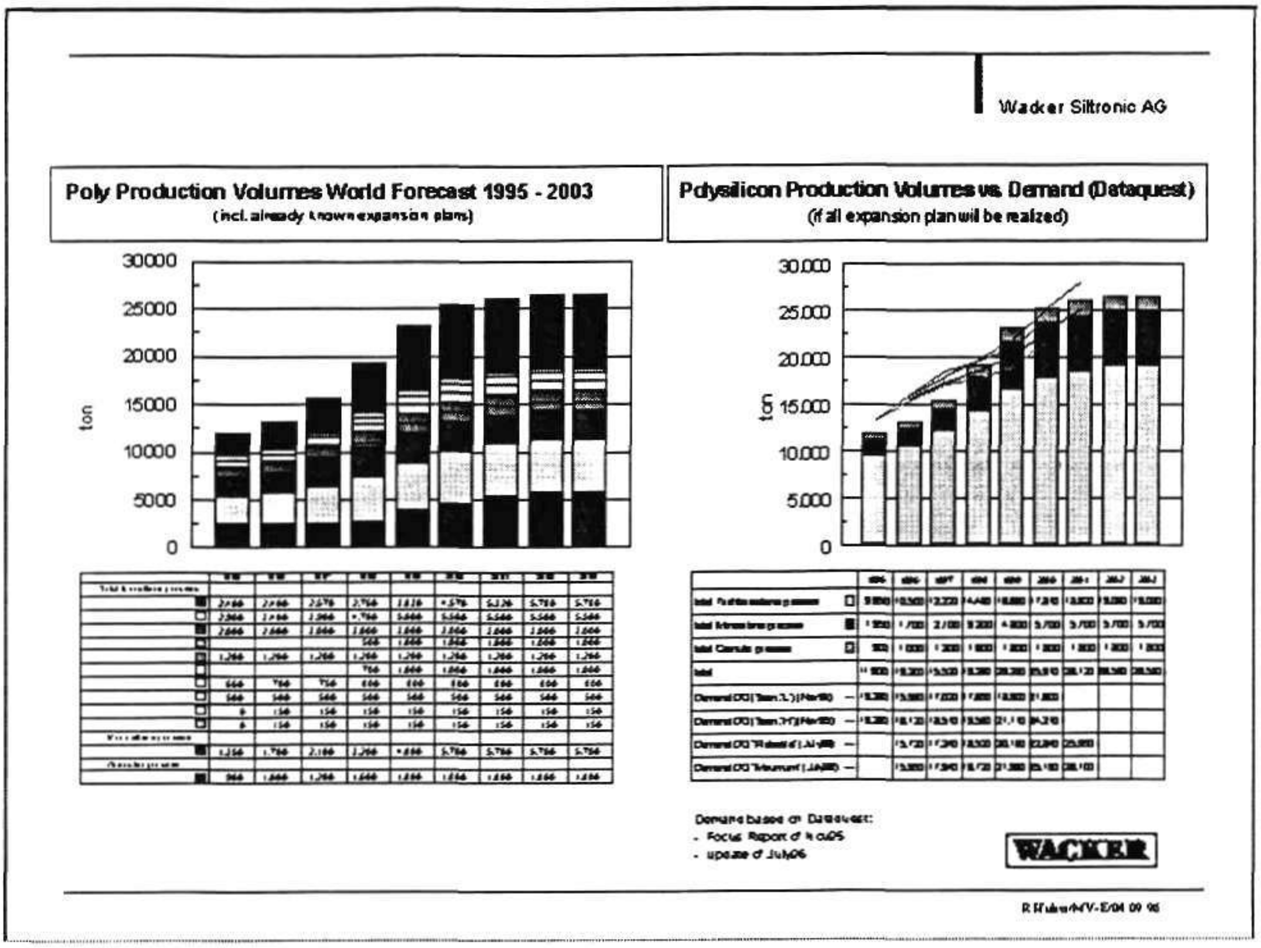

this year.

A lot of concern in the recent years was directed to polysilicon. If there should be any bottleneck at all in 1997, it could be polysilicon. Our expectation is that, for 1997, there will be enough poly, if the wafer growth stays at the 12% that we assume. Should it approach 15% then poly could become a bottleneck again. In any way, poly will be tight, but are quite a few expansions coming on stream, including our own, in 1997. There will be a major increment of capacity coming on stream in 1997. From 1998 on, poly should not be a problem any longer.

The whole situation is rather volatile, because not all wafer fab capacity in the IC arena is going to be utilized. There is idle capacity that could quickly be used in case of an upswing which could lead to constraints in supply again. With cautious forecasting, that should not happen.

ARTHUR ZAFIROPOULO: 1997 may be a question mark on the issue of poly-silicon, but it is not a question mark of equipment; it will be available.

Looking at some of the semi numbers that were generated regarding book to bill ratio in equipment, we have seen a dramatic change in the past half a year. This change did not just begin when we peaked out in February but began the end of last year. In February the book to bill ratio was the highest in the past number of quarters, about 1.4 to 1. In each quarter, this industry was increasing the sales, so this book to bill number is enormous. The demise of memory (DRAMs, SRAMs, and Flash), has really impacted the rest of the industry.

The demand issues are still in place. We see strong demand for the end product in all areas. There is some weakness at Motorola, and analog communications, but there is strength in digital. In general, I think the demand is very strong, but capacity has far exceeded demand. On the equipment side, book to bill ratios have suffered significantly, and are now reaching 0.75. The numbers will probably settle out around this 0.75 as we go forward.

The worst downturn this industry has had was in 1974. We had two drivers: calculators and watches – that was it. The next bad one was 1984-85. The present downturn is going to have some depth, and some breadth, and I believe that the, next two to three quarters, extremely difficult. We are getting close to the bottom. The best part of this whole thing is that most equipment companies that are solid today, the leaders, have great balance sheets. We have never had balance sheets like this ever before. Applied Materials, KLA, my company, have virtually no debt, and strong cash positions.

For the equipment book to bill, there are three areas, which are front end, test, and assembly. Looking at the U.S.-based semiconductor equipment, in the test equipment area, there is a great separation between the shipments and the bookings. We are taking backlog down and that backlog will continue to be eroded until those lines cross. Some companies report a backlog of one year and some, six months. We report six month backlogs once a year, and we give trends quarterly on book to bill ratios. A one-year backlog can be very misleading, so that you can ship things out two quarters and still be within your window of one year. It is hard to ship things out two quarters when you have a six-month backlog. We still see cancellations, push-outs, in the new factories, and we see some impact in the older factories.

Generally, again, we will not see this industry recovering until these two lines cross. It is my feeling this will occur in second to third quarter next year and the upswing will begin again third to fourth quarter next year. Some of the comments made earlier about 1998, and how bad that is going to be, are overly pessimistic. I think we will see things recovering by the end of next year, gradually increasing through '98, and then strongly increasing in '99. This should be a great three-year run in our industry.

In the test equipment there is a divergence between bookings and billings, again indicating people are eating the backlog. We have yet to see the real true bottom of some equipment companies. This will come in Q1 or Q2.

The assembly equipment has done a cross-over in book to bill above one. If you look at the monthly shipments, it is about 30 million per month. This is U.S. So the Japanese have a major role in the area of assembly. This reflects things such as yield increases, and the slow rate of die bonders and probers. I think that this recovery will be short-lived. Then we will see a book to bill again below 1 to 1, until the full recovery occurs.

Looking at the front end, which is the major section of our entire industry, we see rates dropping from a maximum of about 1.1 billion per month, to about 700 million in bookings. We will see companies look at a sharp decrease in bookings going forward, and I am projecting that these numbers will be about half of what they were at their peak. We are getting closer to that point, so downside is far less than what the upside potential will be going forward.

How should you form your corporation? How should you drive it? What strategies do you put in place for these cyclical markets? One is to develop advanced technologies, so there is a leading edge to your company. You feed the R&D and preproduction area, and that area continues to grow, even through a downturn. The second thing is market share. If you have large market share in a down market, that is not a good place to be. If you have small market share in a downturn, that is a good place to be. You can grow market share in a downturn.

If you are diversified, with a core technology, it is extremely important. As an example, we are in three businesses. The disk drive industry, it is very robust right now. It is the same as 1991. This industry had a small downturn in '91. The drive industry had a rapid increase in '91. The same thing holds true today. Last year our sales in the drive industry were about 29%. This year they

Graphic materials for this speaker were not available at the time of publication

will be just under 40%.

Having a counter-cyclical strategy and diversification are very important, and the other part is retrofitting older factories. People still spend money in a downturn. They look at their older factories more carefully, in a downturn. So by having a strategy in a business of counter-cyclicality, it will prevent serious cyclicality within your business and you can manage your people more effectively. The other thing is the low-cost producer. In bad times, people still spend money, but they look more carefully how they spend it. If you live in New York City, and you take a vacation, and you have had a great year, you will take your family to Hawaii. If you have a bad year, you will take your family to Cape Cod. You will still take your vacation. The same thing applies with equipment. You look more carefully at the ASP in a downturn than you do when things are well.

Another thing is financial packages. One should be creative in these kinds of environments. We began the first equipment company forming a leasing corporation, to form a true partnership with our customers. With our strong balance sheet, with lots of cash, we wanted to share that with our customers, and have them own our equipment, that we pay for. It is operating expense, we maintain it for five years, and then we take it back. They have a true partnership with us, a long-term relationship, and we invest our money on their behalf. These kinds of programs, I think, will be more exciting, in the years to come. As we grow our industry, and our businesses, we have to be more creative on our product strategy, on our pricing, and our financial packages.

JOHN LUKE: I am going to talk about the fabless semiconductor model, and the role of the foundry in that model.

The data we are showing here says that, out through 1998 there will be a $240 billion industry growth, or of sales in '98 growing at a compounded annual rate of 20%, and the foundry numbers will grow at around 25%.

What fuels our business is the number of fabless semiconductor people coming into it, or fabless semiconductor organizations coming on, as well as vertically integrated people coming to foundries for additional incremental capacity. What are some of the issues that customers face in growing their businesses? Shortening product life cycles – we do a lot of business with graphics chip suppliers. They run through two or three designs a year, so you can see how short the life cycle is there, continuously changing product lines. We are going from 2D to 3D in graphics, people putting more and more on a chip in mixed signal. We are talking about supplying things like 4 meg embedded SRAM, with DSP in it. The market drives more and more complex devices.

One of the other issues for customers is how to focus resources on the lifeblood of the company, which is a continuous stream of new products. You like to put your energy in capital and money, into design and marketing, as opposed to distractions of production, and of course building a fab is a billion dollar investment. Our typical fabs at TSMC are 35 to 40 thousand 8 inch wafers a month out, 0.35, 0.25, and you get into things like $1.2 billion, when it is fully facilitated. Then if one does have his own fab, you run into the situation of insufficient capacity, or capacity problems, too much, not enough, or the wrong kind.

How does a foundry solve those problems? If you get a foundry that can offer you the advanced technologies, and a path to migrate down the technology curve on feature size, you can ramp up pretty quickly for new designs, and hit the market windows as they open. We have cases where people have gone from a tape out, to a market position in three to six months, particularly the fabless community. That is quite an achievement and with all the design tools, most of the stuff works the first time through. One of our customers told us that the importance of getting it done right the first time, whether it is the customer's design, or our fabrication, he said, "Every spin of the device costs a quarter of market share." That is pretty dramatic when you look at the return on

Graphic materials for this speaker were not available at the time of publication

investment.

Another thing a state-of-the-art foundry can do is provide technology that is not resident in a company, particularly if one is a vertically-integrated company. They may be experts in DRAMs or logic and not have mixed signal capability. More and more device functions are incorporating analog functions on the chip itself. That implies a whole different kind of a technology in fab.

It also allows a customer, by going to a foundry, to focuses energies on new designs. There is certainly a preservation of capital, and we hear numbers going to 1.5 and 2 billion, as you get down to 0.18. Then, for a vertically-integrated company, it is easy to augment existing capacity by going on to a foundry.

What drove a semi, or fabless, model? Back in the '70s, the typical model was a vertically-integrated company, with its own process, application, design, etc. It designed and put out products that a systems company used. That IC user was at the mercy of the supplier – some companies had the attitude: If we do not make it, you do not need it. The devices were driven by the semiconductor manufacturer, not by the needs of the consumer. We were talking about 10 or 20 thousand transistors per chip.

In the 1980s, we saw a phenomenon called foundries, either dedicated foundries or people selling excess capacity as a foundry. That helped spawn a fabless semiconductor community. Not only the foundries, but a decentralization I think, of the industry. We find ourselves now in a industry that is horizontal, as opposed to vertically-integrated semiconductor companies. Now the systems level user, or systems customer, IC user, had the option of going to a vertically-integrated company, that gave birth to certain devices, and they has this fabless semiconductor community which was defining devices, we think, more in line with what the customer needed. We were assembling up to 30,000 to 50,000 transistors per chip.

A company that had its own silicon production, its own design tools, its own applications group, circuit designers, made their own processing equipment, made their own test equipment, made their own steppers, had their indigenous packaging operations, was a self-contained company. There was an evolvement of an industry going horizontal.

vertically-integrated company, at one time owned and operated functions internally. Almost every one of those functions spawned an industry, or a subgroup industry – silicon wafer suppliers, a la Wacker, stepper suppliers like Art's company, and not quite a complete decentralization, and would now let people come into this fabless semiconductor type of phenomenon, much more readily. This type of phenomenon helped spawn the fabless semiconductor community, which is going to fuel some of the growth for people like ourselves.

Another thing that has happened in the last six months or a year, is the advent of third party design houses. People like Cadence and Mentor are getting into the business, not only supplying design tools but offering design services. The semiconductor user can now go to one of these third party designers, who we support with our latest SPICE models, design rules, process flows and PCMs. This chip user can now take a net list to one of these design houses, have that particular design house a database tape, for example, and that system house can go directly to a foundry, with his own custom-designed devices. We think that this is one of the phenomenons that are just starting. In the future, we see a distinct shift in our customer base, where right now probably 90% of our customers are other semiconductor suppliers, and we will see an advent of system companies coming directly into a foundry because of this further decentralization of the industry.

EVERT WOLSHEIMER: I would like to talk to you about taking the fabless concept one step further, and so discuss the manufacturingless company.

The manufacturingless company does their own product design, marketing, and sales. All the other activities are done elsewhere. For example, mask making, wafer manufacturing, electrical test, wafer sort, assembly, final test, and even shipping to customers is done by subcontractors. Some of this has developed over the last couple of years, starting with the concept of fabless semiconductor companies, of which Xilinx is an example. However, this concept can be taken a lot further, if the right infrastructure is put in place.

There is a lot of information flow between the company that does the design and the subcontractor or subcontractors. What makes it even more complicated is that the company that builds the masks, the company that builds the wafers, and the company that does the assembly are often not the same company. While dealing with this whole flow, a company like Xilinx often has to deal with five, or six, or even ten different subcontractors. What we would like to see, and what we would like to actually stimulate is the concept of a turnkey subcontracting company.

To give you an example of the type of data that goes across the boundaries between the companies, I have listed a few of them here. During product development, design rules are exchanged, SPICE models, package designs. During wafer manufacturing, expected yields, speed distributions, wafer forecasts, etc. At wafer sort, it is the same thing, wafer sort program has to be installed at the subcontractors. Problems do occur, every now and then, with wafer lots, and there has to be somebody in place to make a decision on the wafer lot, and what happens to it.

None of the infrastructure to take care of this in an integrated fashion is in place today. There are great foundry companies, like, for example, TSMC, UMC, and Chartered Semiconductor, but none of these companies offer the full flow that I have just shown you.

One of the other complications is that the electronic data interchange system is not in place. If you are a small start-up company, you do a design, you run a few wafers. Then it is okay to do all the data exchange by fax, by telephone, or even by e-mail. However, when you become larger, like, for example, Xilinx, and we start more than 100,000 wafers per year, we would like to see everything go through one point, and have a real sophisticated EDI system in place, where we have direct and immediate access to the status of the wafers, the status of the wafer sort yields, etc, the status of die bank, etc. None of this infrastructure is currently in place, or has not been integrated to the right level.

My conclusion is that true turnkey subcontractors do not exist today. They will develop over the next couple of years, and companies like Xilinx, and other fabless manufacturingless companies, will work with the subcontractors to set up all the right EDI systems.

During this current over-capacity situation, it is an excellent opportunity for some subcontractors to differentiate themselves from other companies, by offering great cycle times, great yields, great service, and great pricing. They can start taking over managing of die banks for fabless companies. They can enhance their engineering support level. What is also key, of course, is the reverse situation, which happened last year, when there was a great under-capacity situation. Some of the smaller fabless companies were in deep trouble during that time. It is key that the foundry suppli-

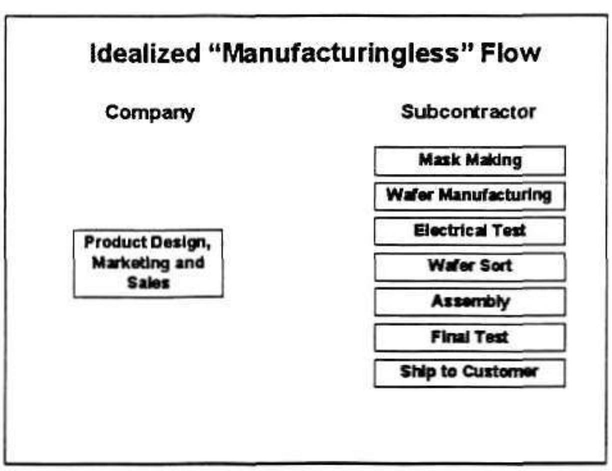

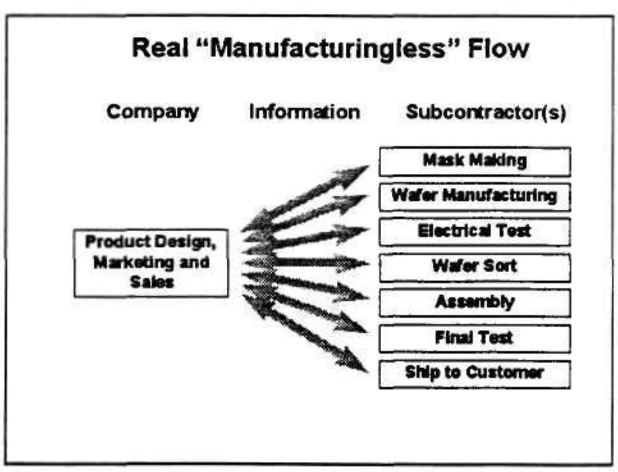

ers realize that the way they treat the companies during the tough times will have an impact on how they are looked at in the over-capacity situations.

Q: It appears that Ultratech has pioneered leasing equipment as an alternative to purchasing. Do you see this becoming a trend in the equipment sector? What is the next step after that?

ART: Let's deal with leasing first. I began studying leasing about two years ago, to see what value it had to the end user. We talk about partnerships with our customer, and there has been a lot of animosity between customers and suppliers in the past. When Intel had gone out there, and for years was, sort of hitting the industry, the equipment people, about their making too much money. So there was a lot of uncomfortable feelings between the parties. There has to be a better way to have a relationship, a partnership, a true partnership, where each company does not reach for their wallet and hold it, when people talk about partnerships.

So I thought this leasing program would be kind of interesting from the equipment side, the equipment for semiconductor manufacturing. If we look at Xerox Corporation, 98% of what Xerox sells, they lease. In your homes, each of you have a home, I hope, and you have a mortgage, probably, on your home. The bank takes that mortgage, and so you own it with the bank. So it is sort of a lease, with a buy-out at the end, over a number of years, 20 years, 15 or 30. So my thought was: How do we really create a bond between semiconductor manufacturers, people with fabs, and people that make the equipment? My feeling was that today, we have a very strong balance sheet. Many companies have that balance sheet: KLA, Applied Materials. A number of others have, over the past three years, gone to the public market and raised a great deal of money.

My company has just under 160 million in cash, and last year's sales were 158 million. This year, we hope to reach around 200 million. We see that there is a real value in having cash. Now, how do we use that constructively with our customers? If we lease the equipment to them, it becomes an operating expense, not a capital expense. Effectively, they do not pay that lump of money, in terms of cash flow, and they pay it only as a period of operation monthly. It is like paying your mortgage every month.

If it is a five-year period and it is my machine, I am going to support that machine, as I own it, not as you own it. Therefore, I am going to make sure after five years, that piece of equipment that is coming back to me, has value to me, and I know the best residual value for that piece of equipment. Also, for me it eliminates third parties coming in and servicing my machine. I get that revenue, in service, that I would not get normally, as a partnership.

During this five-year period, if technology changes and the semiconductor manufacturer decides to send it back, he has a right to do that, and there is a vehicle for him to do that, so the risk is reduced. I think it is truly a true partnership, over a length of time, more than the warranty period. I can put my service people into a location for five years, not one year, where I have to relocate them at huge expenses. It is a savings to me, it is a cash flow opportunity, and a tax opportunity for the semiconductor companies, and it is a true partnership.

I truly believe this is just the beginning of this industry, those strong financial companies moving in that direction. We began that in the U.S. We are now moving to Japan, where cash is really important issue. We believe we can help, selectively, in certain products, globally, in this enterprise of leasing. Now our business is making machines, but our true business is making our customer happy, and we believe this is the vehicle to get us there.

Q: What is the availability of six inch wafers over the next number of years? Do you see the demand for that increasing or decreasing over the coming years?

EBERHARD: We see the demand for 6 inch wafers going to increase at a rate, slightly below 10%. A lot of that increase will come from the conversion of 5 or 4 inch lines to 6 inch. It will be to the largest extent, in our view, at the expense of 4 and 5 inch. The total segment of 4, 5, and 6 inch will only grow very slightly. It may even decrease if too much unutilized 200 mm capacity is unused. At Wacker Siltronic we will, in parallel with our customers, migrate the smaller wafers to 6 inch diameter, and we will also exploit further opportunities for debottlenecking, and increasing 6 inch wafer available capacity.

Q: Last year there was a tight supply of four and five inch wafers. Are we going to assume that four and five inch wafers are okay to get right now as well?

EBERHARD: I think that the four and five inch wafer surge, was a temporary revival, because four inch has been on a slowdown slide for years. Five inch has been stagnating, but when there was so much semiconductor business, any available line has been used to get ICs out. This is why suddenly old four inch lines, that had been thought to be closed, suddenly had big capacity increases. So I think 1995 an abnormal situation. I do not think it will reoccur.

Q: We have a couple of questions here about 300 mm wafers.

EBERHARD: It will be there earlier. There are still so many problems to overcome in the whole area of equipment and development that this would be absolutely premature to do within the next years, when all the funds are needed to expand the 200 mm capacity to satisfy all the fabs that have been built, or are in construction. So, in our company, we are actively researching on this subject. We do have a special group that is dealing with it. We are able to sell dummy wafers, or equipment wafers, and we assume that we will have a pilot line of up to 20K by the year 2000.

Q: There is two major efforts going on in 300 mm, I 300 I, sponsored by SemiTech, about $26 million of funding, and Silite a collection of Japanese companies with about 14 times that amount of funding, of which half of that budget, as I understand it, is for wafers. Does this mean that the Japanese are taking the lead in 300 mm? What is going on there?

ART: I do not think so. I think the Japanese are concerned about what has happened in the past several years in losing market share, both in semiconductor devices, and in equipment business. In 1983, the Intel Corporation moved to six inch wafers. That transition to six inch was very painful, and came very close to hurting Intel severely.

Then IBM moved to eight inch wafers, and, to some extent, these players are not that actively involved in the 12 inch program, although we have four companies in the world that are talking about doing some preproduction in the area of 12 inch. It will be driven by memory or microprocessors, because that is the biggest bang for the buck. The six inch wafer was painful; the eight was less painful; this 12 inch program is going to be extremely painful. Painful for everyone, and I agree, I do not think we will see much movement until at least the year 2000, and in production, in any reasonable quantity toward the middle of 2000.

From the equipment side, I am on the board of SEMI. SEMI is a worldwide organization that deals with standards, trade shows, and education. We have a situation in Japan where we are looking at standards there. We need one world standard, for 300 mm, and 400 mm later. It is very important that the world have only one standard. Otherwise, it is going to be chaos, for not just the wafer manufacturer, but more importantly, the equipment manufacturers. I think differentiating countries and equipment, on standards, is the wrong way to go. I believe that there will be some understanding. This is going to be very slow moving, and very expensive. For instance, as an equipment manufacturer we have these marathon tests. We have to run 10,000 wafers and these wafers run in excess of $1,000 each. Who is going to pay for that? How do you pay for that? So, these are major issues that have to be addressed by the industry, both the manufacturers of devices, and the manufacturers of systems. I think 300 mm is going to go forward much slower than people believe. There will be preproduction facilities in place by the year 2000, but I do not believe anything of any significant capacity.

JOHN: All of our capital expansion right now that is planned is in eight inch. We have three six inch fabs, then we have fab three, fab four, fab five, and a joint venture, all eight inch fabs, and we are progressing with them. From our perspective we are really not proactive, we are more reactive. When the world is there, we will have no choice but to go, but it will be expensive. I think the transition from five to six was not that dramatic, some of the handling equipment and throughput. I think it got worse at six to eight, but I think going from eight to 12 runs into single wafer processing. That is a whole new phenomenon.

ART: Back in the mid '70s, we saw the very strong drive from batch processing to single wafer processing, dry plasma etching. We have seen that evolution occurring as we continue through the scaling up of wafers. I truly believe that the 12 inch wafer will begin to really decrease the use of batch processing. It will still be used, but technologies like ion implantation will have to go to single wafer processing. Thermal issues will have to be addressed, and lower temperatures will have to be formed because of warpage. As you go to 0.25 micron, and 0.18 micron, the warpage, and depth of focus, becomes a major issue. Thermal issues, as well as single wafer processing, will be highlighted, especially as we move to 12 and larger wafers in the future.

Q: We are currently experiencing an over-capacity in all facets of the business and finally it is hitting the foundry industry. Prices have come down recently. How long do you expect this trend in downward foundry prices to last?

JOHN: We had a ride from September of '92 through September of '96. We were literally at over-demand. We started seeing softness coming in October, in our perspective going back to July. We had some excess capacity coming on because of our new fabs coming on faster than anticipated, because the equipment became available. So our ramp rates became quicker, or faster-paced. Then some of the customers slowed down. There was inventory problems, in almost every facet of the industry.

We think that the bleeding may stop in the fourth quarter. We saw a lot of intense price pressures in July and August for fourth quarter delivery. We think most of the needs have been bought out by now, and we see a little flattening. The pricing was particularly aggressive at maybe today and yesterday's technologies, which is 0.6 and 0.5, and certainly 0.8. As we are seeing a transition of our customers going down to 0.35 (we will be in production first quarter next year), we think the ball game changes. I do not think all 0.35 are going to be created equal, based on processing. We do not know how many players are going to have the sustaining power, now as you get into 0.35, and going on to 0.25 a stepper costs what, $6 million bucks? And a 0.35 stepper is $4 million bucks. Well, you have a $40,000/month Fab, and you have 30, 35 steppers in there. It gets into some money. Was it Senator Dirkson years ago, that said, "you know, a billion here, a billion there, pretty soon you talk about real money," and I think we are going to see that. Our anticipation is that hopefully we have seen a bottoming out of the pressures.

The pressures came from newcomers or players that were not in the market here historically. Particularly some that had 4 meg DRAM fabs that were sitting idle, and that was basically 0.5 type technology. As the migration happens down to 0.35 and 0.25, some of the pressures will go away. In our case, we will ship about 2 million wafers next year, on a 6 inch equivalent basis. Something over 60%, and almost two thirds of that will be 0.5 or less. That is how fast we think the industry's migrated down the feature size curve.

Q: Evert, being the buyer of these wafers, and the customers do you agree with this view? With all this new capacity from new players, from Korea, and other places coming on?

EVERT: I agree mostly with what John said. There is one extra factor there that a lot of people do not realize, and that is something that we have seen at Xilinx also. When you switch from 6 inch wafers to 8 inch wafers, with designs that are similar, you get an extra yield boost by the improved quality of the equipment. Typically people count on a factor of two, between 6 inch wafers and 8 inch wafers. They assume that the capacity has doubled if you go from a 6 to an 8 inch wafer. Because of the better yields that you see a lot of times on these brand new 8 inch fabs, it could easily be a factor of three. The over-capacity situation that we are seeing could actually be a little bit worse than anticipated by some people. From a company point of view, from a customer point of view, I hope this current situation will continue for a while.

JOHN: I would like to add one more thing as we see this migration of 0.35, I think down to 0.25, I am going to call it a CMOS recipe. When you get down to 0.35, designers get into more potential problems, and we, as suppliers, are going to have to give them more help and expertise in the process. By that I mean, if you get down to 0.35 and 0.25, the transistor parameters no longer become the driving force in the performance of the circuit. You get into the interconnectivity problems. You have four and five layers of metal and all these neat little capacitors. Right now, most of the foundry suppliers have not comprehended that, in their spice models and their design rules, and I think the ones that have the ability to do that are going to win out.

ART: If I can add to that 0.35, and 0.25. In lithography, we are seeing a severe elevation in the cost of this equipment – $4 million, $4.5, and going to $6 million for deep UV. As we move to 193 nanometers some estimates are $8 to $10 million per system. So when you have a fab of 35 and 40 steppers, this cost is really exceptionally high. You do not need those steppers for every level, so this technology was involved in the past of mix and match. It is my feeling that as this industry increases the cost ASP, of the lithography tools, the industry is forced to go to mix and match, to use non-critical steppers for non-critical layers. When we talk about 0.25 microns, not every layer is 0.25. About half the layers are, so for those layers, 0.25 to 0.5, use those tools that cost $6 and $8 and $10 million. But for those layers that are not, I think the industry will be forced to go into less costly tools for those non-critical steps.

Q: One question on the new players from the customer's perspective. You see all these new players coming on, they have capacity, and they want to off load it to the foundry customers whenever possible. We saw this in the last two downturns. Is the situation different today? Is the customer less likely to want to go to a new player? Is it going to be harder for these people?

EVERT: In the last six months the situation has been very severe. There has not been a week where I did not get a call from somebody offering either 6 inch or 8 inch capacity with very advanced technology at very good pricing. But it is very difficult to just pick up a design, pick up a product line, and take it to another fab because of $100 discount in wafer prices. You just do not do that. The cost of retooling, the cost of requalification, the whole relationship that has been built up, is not worth $100 a wafer. Which does not mean that I will stop complaining about wafer prices. In general it is very difficult to just pick up your business and go somewhere else for a discount on a wafer. It is just not the way it works.

Q: Coming back to the issue about migration of 0.35 and 0.25 micron. When is EPI going to become a more major part in the capacity situation, and the process capability of that?

EBERHARD: This is extremely difficult to predict right now. Some of the DRAM manufacturers have their 64 megabit design on an EPI substrate, on an EPI wafer, but they are not sure yet, how that will make their 256 megabit DRAM. The likelihood is that even some of those that are doing it on EPI now have already indicated strong intentions to go back to polished. We look at the DRAM as one of the major drivers. It is a 50/50 likelihood that there will be more EPI. It is definite that the big Koreans, with the 64 meg and 256 meg end will stay on polished.

Q: Is the same sort of driver toward EPI becoming evident in the logic side of the house? Do you see more demand for EPI in your customers, when you go to 0.35?

JOHN: I do not think we know yet, because we are customer-driven. The customer comes to us and says, "here is my design, here is my data base tape." If he wants EPI, he or she calls out EPI. We have seen in the last 18 months, the percentage of the wafers we supply using EPI going down. That is coupled with the migration from 0.8 to 0.6 to 0.5, and now to 0.35. Quite frankly, I think a lot of customers have learned how to design, and not worry about some of the things they worried about where they needed EPI before.

Q: Does this change go into 12 inch wafers? Are we going to need EPI as we go to 12 inch?

EBERHARD: I think it is likely that 12 inch will generally be more EPI. It also has to do with EPI costs, I think the cost of EPI deposition in principle, is falling with the size of the wafer. With the big 300 mm wafer, the premium for the EPI will be smaller as compared to 200 or 150 mm, and therefore, it could be that just for security reasons, the industry may go all EPI on 300 mm, but that is still an open issue.

Q: Evert, you mentioned in your presentation that the fabless foundry model needed to be reintegrated. John presented a story that it is disintegrating, and you want to reintegrate it. There is a bit of debate but, right now, who takes ownership of the test and packaging? Is that being transferred? You mentioned that the data is required. Do you think that has to be all in one company to make it work?

EVERT: For a company like Xilinx, that would be the best solution from our point of view. We really want to focus all our resources on what we do best, which is designing new products, writing the software that goes with it, marketing it, and selling it. From our point of view, yes, we would like to see fully integrated manufacturing, but that is not today's situation. For Xilinx, we have separate flows, we run wafers at a couple of companies, and we run assembly and tests at different companies. There are a lot of operations in between, that Xilinx owns, shipping the wafers to an assembly plant, taking the parts from the assembly plant to the test facility, etc. We prefer not to spend our resources on it, but, let that be done by people who handle that best.

Q: In order to keep the costs down, we need an integrated back end, and front end approach. Now go to the foundry manufacturer. What do you think about that? Is that necessary?

JOHN: I think it is an ideal model to a semiconductor user. We have been successful in providing turn-key services by using contract manufacturing. That is, a customer can come to us for 1,000 wafers a month, or 100 wafers a month, and he wants a turnkey approach, so he gives us test vectors, and/or a program. We can do the wafer fab, we can do the probe. If he wants us to do the assembly we will contract that out to a contract assembler, either of the customer's choosing, or of our choosing that is been approved within our process. Assemble and symbolize it to the customer's specs, final test it, ship to the customer, or drop ship to his customers. We have found that model works pretty well. It has worked well for us, because we have been able to take our capital resources, and dump it into processes and fabs, as opposed to back ends. I think there is an argument both ways.

EVERT: The only problem with that model is, who owns the data? How can a company like Xilinx keep the visibility of where the product is? How do we know what product is in diebank, and how long does it take to flow through the assembly process?

Q: If you had the information systems that we have today, if all the companies were utilizing that, wouldn't that make it more effective? I think we have the tool today, that maybe everyone is not using those tools.

EVERT: I agree with you. I think all the tools are available. They are just not integrated, not tied together. At least one or two suppliers that we work with have set up a web page, on the Internet, where, of course with a password we can get direct access to wafer status.

Q: So, essentially, it could satisfy both needs?

EVERT: Yes. I think the tools are there but it would take somebody, or an organization to put it together.

Q: So who owns that process? Do you want to own that process, or do you think the foundry should?

EVERT: I think, since the greatest benefit will probably be to companies like Xilinx, probably companies like Xilinx will have to drive it.

Q: In the foundry industry, SMIF is very widely used. The standard machine interface – the SMIF boxes – for isolation of the hemisphere. Is SMIF good? Is SMIF used, and why is it used in a foundry? What is its benefit, and where is SMIF not very effective?

JOHN: TSMC was the basic pioneer. I think we were the first full-fledged fab to go 100% SMIF back in '91, '92 when we did our fab 2A, 2B. Now these are both 6 inch fabs, about 45,000 a month out each. What we found out going to SMIF was, first of all, we saved about 30% of cost in construction for air handling. We got a fringe benefit where the operators do not have to wear bunny suits, so it is a user or an operator-friendly environment. We get a built-in quality control system because there is an electronic traveller with each cassette. If the cassette comes to the Station #12 and it has not done 11, it says, "You cannot come here. You have to go back and do 11."

We do not think we could have come down the defect density curve as quickly as we have on each of our processes, as we have come down the feature size curve, because the wafers are always in a zero-D environment. We do not think we could achieve those kind of defect densities that quickly in a ballroom. It offers us the flexibility of running a multitude of processes, and a multitude of feature sizes, literally down the same line. We could be running SRAMs down this module, and we could be running mixed signal down this module, and we can be putting in Module 3, while Module 1 and 2 are operating, so it gives us the flexibility of running. We probably run four to five hundred different mask sets a month in production. So, from our standpoint, we were 110% sold on the mini-environment.

ART: I think that flexible manufacturing probably is the principle issue but it is my belief that it has not been adopted. This began in 1980 with Hewlett Packard. This technology has not been adopted in the Japanese market, and generally some of the Korean markets. I believe that there will be a different driver for SMIF. My belief is that in the year 2000 SMIF will be used because of technology. That is, you cannot process wafers in an air environment, that the oxidation from air will create a barrier to technology. So you will be processing them in nitrogen. It is very hard to fill a ballroom full of nitrogen, with people just walking around in bunny suits. I believe that the driver will be the environment, not so much particulates. We see yields at 90% in ballrooms, that are running 20,000 wafer starts for DRAMs, or microprocessors. I think that if you want flexibility, SMIF is a great program. But I believe the major driver going forward in the year 2000 beyond, will be the environment, and oxidation of the substrate.

Q: 2000 and beyond puts us into the 300 mm realm, with the larger wafers. Does that mean that we have to completely rethink how we design a fab for 300 mm?

ART: I think so. The Japanese have some interesting programs. I think we have the same approach we used back in the '60s and '70s. A new design is going to come forward. Maybe integrating some of the back end with this area, where we see lithography now being used in the polyimide processing in back end. Maybe there will be some commonalities. The SMIF approach will maybe drive part of that. There has to be some revolutionary concepts in the architecture of how a facility is manufactured. Probably 12 inch wafers is an interesting place to start, but again, we are adding risk. I think the preproduction of 12 inch will be as it is today, so people will use ballrooms for 12 inch when they start. It may occur beyond 2004, and 5, when they understand processing 12 inch, understand the contamination issues, and then maybe they will move to a 15 or 20 thousand wafer start fab in 12 inch, with a different architecture.

Q: Both Eberhard and Art have mentioned that the cycles do not appear to be going away. Are they going to get worse? Are they going to get better? What is the general feeling?

ART: They are all painful. I think that is all relative. In 1974 we had two drivers: calculators and watches. We had more drivers in '84/'85. We have more drivers today, but we are growing our businesses, so they appear to all be difficult. I know we will survive this downturn very well.

Now I am much more comfortable, still aggressive, but more comfortable due to our balance sheet. That is going to help us go forward, and take some risk. I think that the next cycle, in 2000, 2001 will be just like this one here, based on the acceleration of capacity, and the demand issue as we go forward.

EBERHARD: The cycles will not go away, due to the difference in the income and spending pattern between the wafer manufacturer and the semiconductor manufacturer. The semiconductor manufacturer, if he runs at full load, it takes him less than a year to create the revenue for his investment. The wafer manufacturer needs almost a year, first to ramp up a factory to nameplate capacity. It takes him more than one year just to generate the revenue, not the income, to equal his investment.

Q: Do you think that the wafer manufacturer has a higher capital intensity, or a lower capital intensity than the semiconductor manufacturer? In other words, depreciation percent of revenue, is it higher or lower than the semiconductor manufacturer?

EBERHARD: I think the depreciation of course is higher for the semiconductor manufacturer, but not as a percentage of sales. That is the problem. If he loads his capacity fully, then with his higher margin, the depreciation as a percentage of sales, is lower, of course.

Q: Going back to the tests and assembly, who currently owns the test coverage? In other words, who is responsible for the testing part of the device right now generally, the fabless company or the foundry, and is that going to change?

EVERT: I think now, and in the future, that the company that owns the product will own the test coverage, because that is the only way you can guarantee the functionality of the part. I do not think that responsibility can ever be transferred. What could be transferred, and what should be transferred, I think, is the execution of the test. I think that the company that develops the product knows how it should be tested. After that the test tape, and a tester could be transferred to a subcontractor.

JOHN: I agree with that, because we get a GDS2 tape from a customer, we have no idea what that thing is, or how it plays. So for us to do any probe, wend, or final test, we need to have, from the designer, either test vectors, and/or test tape.

Q: There is a movement of putting probers in front end manufacturing, and there is a need for multichip modules, and for known good die sorts of strategies that will go to full test probers in the front end of the line before dicing. Is that something that you see a trend in, and do you agree with that trend?

JOHN: Let me ask you my way, are you saying that the trend is to supply wafers probed, in a wafer form?

Q: Fully test probed, yes.

JOHN: Including AC?

Q: For known good die.

EVERT: That is the problem, the AC testing.

JOHN: That is right.

EVERT: What we have implemented in all the fabs that we work with is what we call a wafer sort test, where we do a fully functional test, except the AC test. Immediately after the wafers come out of the fab, the fab has available feedback on yield issues. I think that is essential, there should not be a four week period in between the moment that the wafer leaves the fab, and the moment that the wafer gets tested. That should be done immediately, so I agree that part should be integrated. That is a relatively simple step.

JOHN: We are doing that basically right now, it is a continuous flow. Wafers come out the last wafer fab station, somebody transports the cassette over our test line, which is adjacent to the fab, so that is a continuous process. I thought you were driving at doing in-line probes. We have some limitations of physics there, because you cannot probe unless you get some metal somewhere.

Index

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

R

S

T

U

V

W

X

Y

TO ORDER...

- **Additional copies**

- **Bulk reprints**

CALL (800) 457-8233

Transcript Services:
DATAQUEST
26524 Golden Valley Road, Suite #401
Santa Clarita, CA 91350 USA

Produced and edited by Market Makers