

CHAPTER 2: TECHNOLOGY, PACKAGING AND MANUFACTURING

Required

1. Manufacturing references - CM

Optional

- ~~1. Packaging, power, I/C metrics - GB~~
- ~~2. Manufacturing testability of VAX, LSI-SRL - CM~~

It is customary when reviewing the history of various industries to ascribe the events therein to either market pull or technology push. The history of the auto industry contains many good examples of market pull, such as the trends toward large cars, small cars, tail fins, and hood ornaments. The history of the computer industry, on the other hand, is almost solely one of technology push. Whereas the rest of this book gives examples of the effects of technology push, this chapter explores the individual elements that have made up the technology push.

Rather than discuss all of the technologies shown in Fig. Menta, this discussion will concentrate on semiconductors and magnetic recording. Semiconductors have been the dominant factor in technology push within the computer industry, but magnetic recording density on disks and tapes has evolved rapidly too, and must be understood as a component of cost and a limit of performance. Communications links have been left out of the discussion

because during the past twenty years they have not evolved as rapidly as other technologies have. Rapid evolution is expected in the immediate future as satellites and fiber optics become more economically attractive. The packaging (cabinets and boxes), interconnection, and power aspects are discussed, principally because these do not represent pushes but rather are large, high inertia objects that have to be pushed against in the hope that they will eventually yield.

The semiconductor portion of the discussion begins by presenting a family tree of the possible technologies arranged according to the function they carry out and showing how these have evolved over the last two or three generations to affect computer engineering. The cost, density, performance, and reliability parameters are briefly reviewed, and then the application of semiconductors, using various logical design methods, is discussed with particular emphasis on how the semiconductor technology has pushed the design methods. For example, one question of interest is whether users should design ICs or design with ICs. The conclusion in this particular case should be obvious...IC design is expensive and time consuming and should be avoided for all but a few applications. The same advice could also be carried over to those who believe it is necessary to program in assembly language or to have access to microprograms. For some reason, many people believe that there is a need to have access to lower level building blocks when building many large objects. This is akin to starting a building construction project by looking for a clay formation and a sand pit from which to make bricks and mortar. Successful contractors buy their bricks and mortar.

The discussion of semiconductors is followed by a section on memory hierarchies. The memory hierarchies section describes the notion and how they operate to utilize virtually every type of memory that's either cheaper or faster than other types of memory.

The next section of this chapter presents some general observations about technology evolution: how technology is measured, why it evolves (or doesn't), cases of it being overthrown, and a general model for how it operates and is managed.

Semiconductor Logic and Memory

A section is given on packaging and interconnections. These form the physical structure by which computers are actually formed. The discussion parallels views 1 (the structural levels) and 3 (packaging levels of integration) of Chapter 1.

The final section briefly describes the manufacturing process.

The best way to observe the evolution and interaction of logic and memory technology as they effect computer design is via a time line diagram (see Fig. LMtimeline). For logic, small memories, read-only memories, and primary memories, four lines give the significant events that have affected the technology push.

A single transistor circuit performing a primitive logic function within an

integrated circuit (IC) is among the smallest, most complex of man-made objects. Alone, such a circuit is intrinsically trivial, but the fabrication process for a set of structures to form a complete integrated circuit is complex. To the digital IC users there are several relevant parameters:

1. The function an individual circuit performs within the IC, the aggregate function of the IC, and the functions a complete IC family performs.
2. The number of primitive digital switching circuits (or transistors) per IC. This density is a measure of the capability of the IC process.
3. Cost.
4. The performance of each circuit and the performance of the aggregate IC and/or set of ICs within a family as they affect system performance (as measured by the time it takes to perform its function). The semiconductor circuit configuration family (Transistor-Transistor Logic = TTL, Schottky TTL = TTL/S, Emitter Coupled Logic = ECL) usually determines this performance.
5. The number of interconnections (pins) to communicate outside the IC. This determines the complexity of the functions to be performed by a complete IC family.
6. The reliability. This parameter is a function of the circuit technology, density, number of pins, the operating temperature and use (or misuse).

7. Power consumption.

Figure ICtree shows a family tree (taxonomy) of the most common digital IC's, roughly in order of increasing complexity and evolutionary time. The secondary ordering is roughly by the regularity of the function being implemented, and whether there is memory associated with the function. The clustering of circuit types is also by technology generations (from the second to possibly the fifth). Note that large totally regular functions can be built with only memory (various registers and memories) and with no memory (adders, multipliers, multiplexors). With large scale integrated circuits it is desirable to implement regular structures to simplify understanding and to aid testing.

In the early third generation only completely unconnected logic design components were built. Collections of the basic logic primitives (AND, NAND, Exclusive OR, Adders) and sequential circuit components (individual flip flops or registers formed from groups of flip-flops) permitted logic functions that had occupied a printed circuit module of the first and second generations to occupy a single IC. This had the benefit of providing a drastic reduction in size between second and third generation designs, as shown most vividly by comparing the PDP-9 and PDP-15 (page 00), but had the drawback that modules now contained a wide variety of functions and were thus specialized.

As the densities began to improve to a hundred gates, the construction of complete arithmetic units on a single chip became possible. The earliest and most famous function, the ALU (Fig. ALU), provided up to 32 functions of two

4-bit variables. By the fourth generation, it became possible to construct very large combinational circuits, such as a complete 16 x 16-bit multiplication circuit requiring about 2^{16} gates, on a single chip.

Without well-defined functions such as adders and multipliers, semiconductor suppliers cannot provide high density, high volume products because there are few large scale, general purpose universal functions. The alternative for the users is to interconnect simple logic circuits (AND gates, flip flops), but this does not permit efficient use of the technology and the cost per function remains high (about that of the third generation) because the printed circuit board and IC packaging costs (pins) limit the attendant cost reduction.

Two methods of effectively customizing LSI are shown in Fig ICtree and discussed in greater detail later in the chapter. These are the PLA and the gate array. The PLA (Programmable Logic Array) is an array of AND-OR gates that can be interconnected to form the sum of products terms in a combinatorial design and will be described subsequently. Gate arrays are simply a large number of gates placed on the chip in fixed locations where they can be interconnected during the final metalization stages of semiconductor manufacture.

There is a special branch of the tree shown in Figure ICtree for the purely memory functions. Memory is used in the processor as conventional memory, but can also be used as an alternative to conventional logic for performing combinatorial logical functions. For example, the inputs to a combinatorial function can be used as an address, and the output obtained by reading the

contents of that address. Memory can also be used to implement sequential logic functions. For example, the memory can be used to hold state information for a microprogram. Since memories have so many uses, this branch is discussed separately in the memory section.

The remainder of the interesting logical functions include combinations of logic and memory. There are various special functions such as LPC (Linear Predictive Coding) encoding algorithms for use in real time applications and data encryption algorithms for use in communication systems. One of the most useful transducers, and the first one to use LSI, was the UART (Universal Asynchronous Receiver Transmitter).

There is a special section on bit-slice components (actually sets of 1, 2, 4, 8 and 16-bits) that can be combined together to form registers of arbitrary length. These are being used to construct most of today's high-speed digital systems, mid-range computers, and computer peripherals. Although there have been several bit slice families, the AMD 2900 series, whose register transfer diagrams are shown in Figure AMD, has become the most widely used. Note that all the primitives of this series were present in the RTM family, including the microprogrammed control unit referred to in Chapter xx as the K(PCS) (Programmed Control Sequencer). RTMs were also the archetype antecedent in the case of Fairchild's macrologic, according to the person responsible for the family, Kris Rallapalli.

The final branch of the tree is the most complex, and is used to mark the fourth (microprocessor-on-a-chip) generation of technology and the beginning

of the fifth (microcomputer-on-a-chip) generation. The fourth generation is marked by a complete processor being packaged on a single silicon die, and the fifth generation, by this measure, has already begun since a complete computer (processor with memory) now occupies a single die. The evolution in complexity during each generation simply permits larger word length processors or computers to be placed on one chip. At the beginning of the fourth generation, a 4-bit processor was the benchmark, whereas toward the end of the fourth generation, a complete 16-bit processor, such as the PDP-11, could be placed on a chip. The LSI-11 (Chapter 00) was designed at about the middle of the fourth generation and hence, four chips were required for its implementation.

The function performed on a chip is clearly dependent on the number of gate that can be placed on a chip. Thus, density in gates/chip is the single most important parameter determining chip functionality. From this measure, one can predict the functions likely to be implemented by just following the tree. It should be noted that the whole tree is relatively alive and has dense areas of new branches everywhere except at the top, where unconnected gate and register structures have been relatively static. In the growing areas, as density increases sufficiently, a new branch grows. For example, the processor-on-a-chip started out as a 4-bit processor (or rather as 2 chips for a single processor) and then progressed to have 8-bit and 16-bit processors on a single chip. Similar effects can be observed with the arithmetic logic unit and with memories.

The number of gate circuits per IC is the measure of density as seen by a user

(see Fig. Semiden). This metric is actually the product of the circuit area and the number of circuits per unit area. Circuits are photographically reduced in size to yield higher speeds and higher densities. The processing techniques to create the semiconductor materials have also been improved to have better manufacturing yields (and lower costs). Circuit and device innovation have also contributed to density and yield increases.

Semiconductor device (individual circuit) performance is also correlated with the implementation technology and density because a reduction in size also reflects a reduction in power, allowing greater speed of operation without exceeding IC package heat dissipation limits.

An Operational Model for Memory Size versus Time

The model given in Fig. Semidens is exponential and correlates with previous observations that the number of bits per chip for a MOS technology memory doubles every two years according to the relationship:

$$\text{number of bits per chip} = 2^{t-1962}$$

There are separate curves, each following this relationship, for ROMs in prototype quantities, ROMs in production quantities, RAMs in prototype quantities, and RAMs in production quantities. Thus depending upon the product and the maturity of its production process, one must scale the state of the art line appropriately by one or two years according to the following rules:

bipolar read-write memories lag by two years
bipolar read-only memories lag by one year
MOS read-only memories lead by one year
production quantity availability

This gives the availability of various sizes of semiconductor memories as shown in Fig. LMtimeline.

~~[insert figure]~~

The significance of these values is that they determine (technology push) where certain architectures and implementations can occur. The chapter critiquing the PDP-11 uses this model to show how semiconductors accomplish this push.

The most important characteristic of ICs after density is cost. The cost of ICs is probably the hardest of all the parameters to identify and predict because it is set by a complex marketplace. For circuits that have been in production for some time and for memory arrays, the price is essentially that of a commodity like eggs or bacon, and users generally consider these ICs as very similar to commodities, with the attendant benefits (cost) and problems (having a sufficient source of supply). In these cases, the prices are proportional to the die cost (which is proportional to the die area) and the manufacturing volume. For those circuits that have not yet reached commodity status, the prices are higher and depend on the strategy of the supplier - whether he is willing to encourage competition.

{Two curves are presented to reflect the price of various components/IC.

{Figure Compprice shows the price per component for an IC assuming LSI. There
{is a price band for the circuit size and circuit technology. Table GateComp
{gives some idea of how circuit density (in components) relates to actual
{functions.

The most useful data to understand past and future computer structures is the semiconductor memory cost curves (see Fig. Memprice). Here, the basic cost/bit of various sized memories is given. In 1978 the cost per bit was roughly .08 and .07 cents per bit for the 4 Kbit and 16 Kbit IC chips respectively, giving IC costs of \$3.30 and \$11.50 respectively. Whereas the chip density improves by a factor of two each (Fig. Semiden) year, the cost per bit (at the IC level) is declining at only a factor of 2 every two years. The line drawn in Fig. Memprice has the equation:

$$\text{cost/bit (in cents)} = .3 \times .72^{t-1974}$$

It is also interesting that the cost compares favorably with the price decline observed in core memory over the period since 1960-1970 for the 18-bit computers, (page 00) and for the cost declines in both PDP-11 and the PDP-8 (pages 00 and 00).

Performance

The performance for each semiconductor technology evolves at different rates depending on the cumulative learning associated with design and manufacturing process together with the marketplace pressure to have higher performance for

the particular technology. One may hypothesize that each technology can be looked at as being relatively appealing or relevant to the particular design(er) styles associated with the computer market levels (view 4, page 00). One would expect the evolution to continue along the lines shown in the following table for the period around 1980.

Table SemChar: Characteristics of Dominant (1978) Semiconductor Technologies

<u>Type</u>	<u>Evolution</u>	<u>Use</u>
TTL (Transistor-Transistor Logic)	TTL TTL/Schottky TTL/LS	logic, bus interfacing higher speed than plain TTL same speed as TTL, but low power
ECL (Emitter Coupled Logic)	MECL II, III MECL 10K, 100K	very high performance easier to work with
MOS (Metal-Oxide-Semiconductor)	p-channel n-channel	low cost greater densities, cost evolving to performance (memory)

CMOS (Complementary
MOS)

low power, higher speed,
better noise immunity

Some of the lesser used technologies such as I^2L (Integrated-Injection Logic) and SOS (Silicon on Sapphire), which both promise higher speed than TTL and lower power than MOS have been omitted. Both of these technologies have been touted as replacements for various technologies shown in the table. But, if an entrenched technology has evolved for some time and continues to evolve, it is difficult for alternative technologies to displace it because of the cumulative investment in process technology and understanding.

Semiconductors appear to be characteristic of other technologies in that usually only a single technology is used for a given problem. Each technology has a niche and there is only one winner.

The early predecessor technologies, RTL (Resistor-Transistor Logic), TRL (Transistor-Resistor Logic), and DTL (Diode-Transistor Logic) have also been omitted. They were too deficient in some aspect of cost, performance, or reliability to become standards of note. These technologies were important to the history of ICs because they were used in the first ICs, but many manufacturers, including DEC, did not use them because they did not represent a sufficient advance over the discrete transistor circuits already being used. In addition, these early circuits were packaged in flat packages rather than in the dual in-line package used today, and automated manufacture using the components was thus not possible.

DEC's use of the various IC technologies shown in the table is probably typical of most of the computer industry: TTL for mid- and high-sized

minicomputers; ECL for the larger scale machines (DECsystem 10); MOS for memories, microprocessors, and specialized high density circuits; and CMOS for special microcomputers (CMOS-8).

Table Gatedelay gives the speed-power product and gate delay which are the two most useful measures of performance for the various technologies as they have evolved with time. The speed-power product metric for a technology at a given time indicates that the user may tradeoff performance against power. There are limits to this tradeoff, as only about one watt can be dissipated by an IC package unless special cooling is used. The table was formulated by Jerry Luecke of Texas Instruments (TI) at a time when I^2L technology had just been introduced (Oct. 1975) by TI:

Table Gatedelay: Gate Delay of Various Semiconductor Technologies [Luecke, 1976]

Year	Type of Logic	tp (nsec)	P D W	Speed-Power Product (Picojoules)
1963	DTL	-	-	200
1964	RTL			180
1965	TTL	10	10	100
1967	TTL/H-series	5	20	100
1968	TTL	30	1	30
1970	TTL(Schottky)	3	20	60
1972	TTL(low power)	10	2	20
1967	ECL	2	30	60
1974	ECL	0.7	43	30
	MOS			
1970	PMOS	200	0.1	20
1973	NMOS	100	0.1	10
1973	CMOS	30	1.0	30

1974	SQS	15	0.05	7.5
	I ² L			
1975	I ² L	35	0.085	3.0
1976	I ² L	20	0.05	1.0
1980	I ² L	10	0.10	1.0

The number of pins that communicate outside the circuit indicate the packaging technology level. Low and medium scale IC's often need more pins than LSI since the latter is capable of carrying out a complete function independently--and hence can operate with less interaction with the rest of the world. Memories are especially easily adapted to living within pin limitations, since only one pin is required for every factor of two increase in size. In addition, even this modest pinning requirement can be reduced by time multiplexing the address signals.

Reliability

Over the past 15 years, the failure rate for standard ICs has been reduced by two orders of magnitude to the neighborhood of .01% per 1000 hours. This corresponds to 10^7 hours mean time to failure (MTTF) per component. Figure Reliab, from a recent survey article by Hodges [1977] shows the trend. The lower curves show the higher reliability obtained when more extensive testing and screening are employed. The improved MTTF of between 10^8 and 10^9 are obtained at a cost increase of 4 to 100 times per component.

EVOLUTION OF SEMICONDUCTOR USE

The LSI Dilemma

The economics of the LSI circuit industry make it essential that circuits with a high degree of universality be produced. Because of the learning curve of a

manufacturing process cost is inversely proportional to volume. For a design to be sold in high volume, it must be usable in a large number of applications. However, the trend in circuit complexity, which allows semiconductor manufacturers to put more transistors on a constant die area each year, lends to increasing specialization of function. The more specialized the function, the lower the potential volume and the higher the price. Figure generality shows the dilemma. The LSI product designer is therefore continually in search of universal primitives or building blocks. For a certain class of applications, such as controller applications, the microprocessor is a fine primitive and has been so exploited [Noyce, 1977]. For some applications, circuit complexity can embrace even higher functionality at the PMS level. The Intel 8XXX is an interesting example here: two processors, a 2.5(?) microsecond byte-processor and a 200 nanosecond bit-processor are combined in one LSI circuit.

Moore [1976] discusses this dilemma in a paper on the role of the microprocessor in the evolution of microelectronic technology. He points out that a similar situation existed when i.c.'s were first introduced. Users were reluctant to relinquish the design prerogative they had when they built circuits from discrete components. It was not until substantial price reductions were made that the impasse was broken. Then the cost advantages were sufficient to force users to adopt circuits that fit the technology.

The difference, however, between the i.c. switching circuit level of RT-level building blocks on the one hand and the PMS-level block on the other is that the latter uses the powerful concept of a shared-program computer. Its

function is varied by programming.

For many applications, including most computer systems, the microprocessor on a chip is not a cost-effective building block, and other solutions to the dilemma are found. For example, microprogramming is a highly general way of generating control signals for data path elements, table lookup is a highly general technique. Both methods are attractive because they use memory, an inherently low cost LSI circuit. Microprogramming, however, does have limitations. The extra level of interpretation extracts a performance penalty and some potential datapath parallelism is often given up to reduce cost. A more subtle, but practical, limitation is the development cost of microcode. Assuming the writing rate to be 700 microwords per man year for wide-word, unencoded (horizontal) micro machines, a desire to limit the effort to 20-24 manyears would limit the maximum control store size to about 16 Kwords.

Other techniques involve changing the pattern of interconnections via the top level(s) of metallization of an integrated circuit. This can be done by a factory made change or by the user, a much more flexible solution. Some of this customization is required in nearly all digital system design of the fourth and fifth generation. At the very minimum, the values for a read only memory have to be set. Table funvar contrasts the building blocks available in the fourth generation.

As a result of the increased basic circuit functionality available at each new generation, design methods have changed with the generations. This book provides an example of each, as summarized in the following table.

Table: Design Method versus Generation

<u>Design Method:\Generation:</u>	<u>First</u>	<u>Second</u>	<u>Third</u>	<u>Fourth</u>	<u>Fifth</u>	<u>Examples in this book</u>
Combinational and sequential; use of "standard" modules, IC's.	s	s	s			18-bit; PDP-8
Read only memory and PLA; microprogramming			s	m		PDP-9; PDP-11
Microprogramming with standard RT elements (high perf.) minor logical design				s	m	CMU-11
Programming using micros and logic for interfaces		p	p	s	x	LSI-11
PMS design using completely specified and pre-designed microcomputer components					s	Cm* (almost)
Customized chip design and standard logical design (high performance)			m	m	m	LSI-11

s - the standard method for most digital systems

m - done by manufacturers of basic equipment

x - also used

p - prelude to micros, also done using minis

Table funvar

<u>Building block</u>	<u>Technique for varying function</u>	<u>Degree of generality</u>	<u>Permanence of change</u>
Computer module	program	v. high	none
microprocessor	program	high	low to medium
bit slice	microprogram	medium	medium
ROM	factory mask change	v. high	irreversible
PROM	field change	v. high	irreversible
EAROM	field change	v. high	low

PLA	factory mask change	medium	irreversible
FPLA	field change	medium	irreversible
gate array	factory mask change	medium	irreversible
RAM	write	v. high	none

The design of most relatively high speed digital systems (including low- to mid-range minicomputers is carried out using standard register transfer ICs complete with data path and memory. Part IV discusses the current standard RT modules set. These bit-slice components will evolve and take on specialized functions to achieve both higher performance (e.g., a multiplier) and to be specialized to particular tasks (e.g., the interpretation of a particular computer's ISP). As an example, the series may be modified for the coding and encoding required in signal processing.

For higher performance computers, there is no alternative to using either tightly packed standard ICs or building a unique set of ICs using some form of customization. Although Seymour Cray continues to build the high speed computers (in the CDC 6600, 7600 and Cray 1) with no custom logic, he does so by using impressively dense modules with high density interconnection and freon cooling. The high performance IBM and Amdahl machines use custom ECL circuits to improve packaging.

The current spectrum of IC's and their use is summarized in Table ICspectrum.

Table ICspectrum: IC Organization and Use in Various Computers

Organization	Technology	Unique Chips	Per. (MIPS)	Cost	Examples
Microcomputer	MOS, (VLSI)	1	0.1	Lowest	Intel 8048, MOSTEK 3870
Microprocessor	MOS	1			Intel 8080, Zilog Z80, Motorola 6800
Microprocessor	MOS	2-4			DEC LSI-11, Fairchild F-8
Microprocessor	MOS	> 4			Burroughs B80, National IMP 16
Bit slice (micro-programmed)	TTL	few			DEC 11/34
Gate array	TTL	all			Raytheon RP16, IBM Series 1
MSI	TTL	few			DEC VAX 11/780, 11/70, HP 3000
Gate array	ECL	all			IBM 370/168, Amdahl 470/v6
SSI	ECL (SSI)	std.	100	highest	CDL 7600, CRAY 1

With the advent of the processor on a chip, digital system design has been or soon will be converted completely to computer design. Problems such as controlling a CRT, controlling a lathe, building a billing machine, or implementing a word processing system become just computer system design problems similar to those attacked over the first three generations. The hardware part of the design, the interface to the particular equipment, is straightforward. The major part of the design is the programming. Since the late 40's three complete generations have learned about computer design, especially programming. The first generation discovered and wrote about it. Then it was rediscovered and applied to minicomputer systems. This time, it is

being learned by everyone who must use and program the microcomputer. Each time, for each individual or organization, the story is about the same: people start off by programming (using binary, octal or hexadecimal codes) small tasks, using no structure or method of synchronizing the various multiple processes; the interrupt mechanism is learned, and the symbolic assembler is employed; and finally some more structured system--possibly an operating system is employed. Occasionally users move to higher level languages or macro assemblers, but usually not because engineers are taught to minimize product cost (as opposed to development, life cycle, maintenance, etc. costs).

In view of this cyclical history, it seems likely that current digital systems design practice, which consists of building simple hardware interfaces to relatively poorly defined busses together with programming the application, will be relatively short lived. The design method of the future will be at the PMS-level component, although at the moment it is still too difficult to be done reliably and cheaply by large numbers of engineers. The components from which the microcomputer systems will be formed will be significantly more advanced using much better packaging, clearly defined busses, standard more general interface, and base level operating systems that are embedded in hardware (i.e., placed in read only memory to give the feeling of permanency so that users are less likely to embark on the expensive, unreliable rediscovery path). Standard components will be built which can be interfaced to a wide range of external systems using parameters that are specified by a field programming method instead of using logical design and building with interconnection on modules. In this way, the complexity of individual ICs can be increased and having a standard method for interconnection, higher volume

and lower costs will result.

Before discussing the alternatives associated with IC design, it is important to characterize the various costs. Figure Design.cost shows, at a crude level, what one might expect the relative design costs to be for various inter- and intra-IC design methods. Even the design cost is highly variable depending on the project size, its goals, manufacturing volumes expected, and more importantly on the computer aided design programs.

The lowest cost designs stay completely away from modifying the ICs, except for programming read only memories. There are two elements to the cost of read only memories. One is the cost of the chips and the other is the cost of programming them. There are two kinds of read only memories, the field programmable read only memory (PROM) and the factory programmable (masked) read only memory (ROM). PROM chips have a higher initial cost than ROMs but provide some inventory advantages in a manufacturing environment because a common stock of unprogrammed parts can be divided up into various programmed parts rather than stocking a full supply of each required part. In many high volume applications, however, there is no substitute for the price advantage offered by masked ROMs. At the RT level, the standard microprogramming design method is (conservatively) only twice as expensive per instruction as conventional programming, but likely on the order of 5 to 10 times as expensive to solve the same problem that a program written for a microprocessor solves (the speeds are at least a factor of 10 more too). Given that one must design controllers with ROMs and PLAs, the cost can be still higher, but lower than with the standard sequential circuits we used in the first few

generations--particularly if the module etch is used for the {interconnection structure (see also comments about interface design for the {18-bit computers, page 00). Figure semiuse depicts the relative cost per function of various technologies.

Design of ICs (Intra-IC Design)

Despite the prospects of higher design cost with custom ICs than standard ICs, and, in some cases, higher manufacturing cost, there are numerous reasons why a designer is often forced to design ICs. In some engineering environments where there are extremely small space, low power, or extremely high reliability requirements, the engineer is forced into building ICs, without attendant cost savings. Another motivating factor for custom IC design is to obtain high performance in high volume products. As was pointed out in mentioning Seymour Cray's computers, he chose not to design custom ICs, and obtained greater design flexibility, lower design costs, and faster design turnaround at the expense of special high density packaging and freon cooling to obtain the desired performance. Both IBM and Amdahl chose to make custom ICs instead. Their computers do not run as fast, but their performance is very impressive and they build many more computers more cheaply and there are many more in operation. Therefore, the added complexity of ICs may be the only way that a high volume product at a given performance can be obtained. The use of custom ICs to reduce the number of discrete components or to reduce the total number of ICs in a machine also has a salutary effect on the reliability, since the reliability of a system is mostly a function of the man-made explicit connections, including the bond to the semiconductor die. The failure rate for the IC die interconnections is so low relative to the man-made connections that

it can be ignored. Thus reliability is simply measured by counting discrete circuit pins, IC pins module and connector pins to determine the reliability.

To summarize, IC design is along the same lines as other logical design and has even more stringent physical layout constraints. Custom ICs are built for performance and for reliability (it may be the only way that a complex design can be maintained). They are built to get some decreased cost and higher performance, permitting mid-range performance designs to be more cost effective. Finally, they are used to build very high volume, lower performance computer components (microprocessors and microcomputers) because it is the cheapest way to do a task...provided sufficient numbers are produced. The secondary reasons, to get size and cost reduction for some design, occasionally enter in too. With greater semiconductor densities for the non-microcomputer components used to build more general digital systems, the increase in density is double edged--and may force specialization. That is, with these more complex components, there is a greater risk that as they become too large they will become less useful for building a particular system (recall fig. Generalization). This is akin to building with bricks that continue to grow in size. At some point, the brick size is well beyond the size of the object being built, and the only way to get the intended function is to sculpture the brick back to a useful form.

The various design methods that might be used for various objects and densities are given in the following table.

Table: IC design method versus semiconductor density.

IC Design Method	Density			
	SSI	MSI	LSI	VLSI
Minor variations in standard ckts. (high performance computers)	busses	RT components	Special interfaces - (e.g., UART)	
Gate arrays	-	Set useful for integrating a large system (e.g., a computer)		possible alt, to custom design
Standard cells	-	Small system relatively high volume system		"
Custom design (for high volume parts)	bus interface, signal convers.	high perf.	memories, microcomputers including peripherals	

The most straightforward intra-IC design method is to modify an existing design. This has been used extensively to get the components for implementing computers. If this approach cannot be used, the next most straightforward method is to use arrays of gates and interconnect them to form the desired function. Design with gate arrays occurs in a completely defined environment because there is only one circuit from which the gate is formed and the gate can be completely parameterized and defined. Furthermore, interconnection is a well understood aspect of logic design and is used to form the more complex macro structures (various flip flop types, adders) and then to form the higher levels of design by using arrays of gate arrays.

It should be noted that gate array design is not a new idea brought about by the need for a simple method of customizing LSI. Rather, it was one of the design philosophies advocated in the first few generations. The concept then was to have a single module containing a set of gates, and all subsequent logical design would be done in terms of that module. For example, flip flops would be constructed by interconnecting the gates. Note that a design predicated on a single module type simplifies the spare stocking and servicing aspects immensely, and it is possible to troubleshoot a problem by simply replacing circuits according to a pattern.

Gate array design uses a single well-understood, well-defined component and the fabrication of all but the last few semiconductor processing steps is identical for all designs. The interconnection of the gates by metal is carried out last. Gate array design methods do not permit the high density possible with the more custom methods such as the standard cell design methods. Standard cell design is identical to the logical design of the first few generations (e.g., PDP-1) since there is a well-defined set of components (AND gates, flip flops) in which the design is carried out. The advantage of the standard cell design methods is that higher densities than those possible with gate arrays can be achieved. However, each cell occupies a different space and hence the improvement in packing density may not be as substantial as direct comparisons between gate array technology and standard cell technology might at first indicate. In addition, there are a large number of circuit types and the set and their interconnection rules may not be characterized well enough to be reliable.

A representative gate array is a Raytheon RA-116. It has 300 Schottky TTL gates, of three configurations:

120 internal driver gates (3-input NAND)

60 external driver gates (4-input NAND)

120 internal expansion gates (7-input NAND

or

2-input OR expanders)

The gates have a typical propagation delay of 5-6 nanoseconds and dissipate 5.5-6 milliwatts per driver and 1 milliwatt per OR expander. Two metal layers are used for tailoring, and the resulting circuitry can be connected to the outside world by means of 64 external pins.

Since the designer can arbitrarily interconnect, he constructs flip flops, adders, decoders, etc. Because the use of IC gate arrays is recent, data on package count reduction is scarce, but one informal study for the Raytheon RP-16 aerospace computer measured a factor-of-three improvement over a system constructed with standard components.

For higher speed applications, an ECL gate array is used. This device exploits the inherent properties of current mode logic to obtain a particularly flexible element [Gaskill et al., 1976].

The most common form of intra-IC design is custom design. It is in some ways a variant of the standard cell, since a designers typically have a set of

favorite circuits which they interconnect to create designs for specified applications. With custom design the designer can (theoretically) specify a circuit for each use within a particular logical design. For example, upon observing that a particular gate or flip flop only drives a certain load, the designer can modify that gate or flip-flop to provide only the appropriate driving capability. Therefore, with custom design, the whole IC can theoretically be an optimum, since each part is no larger than it need be. The advantages are clearly size, cost, and speed. The design costs are high because each part can, in principle, be customized. The quality of the circuit design is totally dependent on the single designer who must analyze each circuit geometry in terms of his expectation of performance, operating margins, etc. To the extent that this analysis is carried out, the circuit is clearly optimal. Design of this type, where both circuit and logical design are varied for the context, rarely occurred in the first few generations.

Also on the graph is a hypothetical line for universal logic arrays. For at least 15 years, academicians have studied the possibility of designing a single array of logical design elements, or a collection of such arrays, that could be interconnected on a custom basis to carry out a given function. The gate array can be looked at as the simplest example of this type of design. While many are skeptical that such a device exists, a line representing it is placed on the graph as a target for those who search for the one truly universal logic array.

Both Read Only Memory (ROM) and Programmable Read Only Memory (PROM) are commonly used, but trivial, forms of the truly universal arrays, because they

can be used in a table look up fashion to create several functions of a number of input variables. For example, a 1,024 word ROM arranged in a 256 x 4-bit fashion can generate 4 independent functions of 8 variables. This is a distinct alternative for using a conventional gate structure to carry out combinational functions. A disadvantage of this method is that the required ROM size doubles for each additional input variable.

The PLA is a combinational circuit which remedies the disadvantages of the ROM implementation of combinatorial functions by allowing the use of product terms rather than completely decoding the input variables. Fig. PLA shows a typical circuit, which consists of separate AND and OR arrays. Inputs are connected to the AND array, and outputs are drawn from the OR array. Each row in the PLA can implement an AND function of selected inputs or their complements, thus forming a boolean product term, and the OR array can combine the product terms to implement any boolean function.

A simple application is operation-code decoding. For the PDP-11, the 16-bit Instruction Register could be directly connected to a PLA and the output thereof used to specify the address of the microprogram that executed that instruction. Three different types of operation-code decoding are customarily applied to PDP-11 instructions: source mode decoding, destination mode decoding, and instruction decoding. With a PLA implementation, a PLA could be used for each of these decoding operations, and only three chips would be required. A ROM implementation, on the other hand, would require 128x8 for address mode decoding and 64Kx8 for instruction decoding. Using 2Kx8 ROM's, 33 chips would be required. For this reason modern minicomputers, such as the

PDP-11/34 use PLAs rather than ROMs or combinatorial logic for instruction decoding. The technique is also extended downward into microcomputers such as the LSI-11, where PLAs are used to conserve the die area used by their control units.

The PLA becomes an even more useful building block when it is made field programmable -- the FPLA. The programmable connectors shown in Figure PLA are fusible nichrome links that are burned out when the unit is programmed.

When a register is added to the outputs of the PLA and incorporated in the same integrated circuit, a simple sequential machine is obtained in one package. Since register circuits are pin intensive, adding registers to PLAs or ROMs permits about a factor-of-two package count reduction in typical applications and has been undertaken by IC manufacturers (Signetics 82Sxxx) as soon as i.c. density improvements permit.

The first PLAs had propagation times of the order of 150 nanosec [check with Spencer] and were thus suitable building blocks for slow, low-cost computers. Propagation times of 45 nanoseconds are quite common today, and the PLA is more widely used. A straightforward application with these higher speed components is the replacement of the SSI and MSI packages used to implement the control logic for Unibus arbitration. Alternatively, such a structure may be more cost-effective when implemented in gate arrays.

A more complex application than instruction decoding has been documented in (Logue et al., 1975). An IBM 7441 Buffered Terminal Control Unit was

implemented using PLAs and compared with an SSI/MSI version. The PLA design included two sets of registers fed by the OR array (PLA outputs): one set fed back to the AND array (PLA inputs), and the other set held the PLA outputs. A factor-of-two reduction in printed circuit board count was obtained with the PLA version. The seven PLA's used in the design replaced 85% of the circuits in the SSI/MSI version. Of these circuits 48% were combinational logic, and 52% were sequential logic.

MEMORY TECHNOLOGY (AND SEMICONDUCTORS USED AS MEMORIES)

The previous section discussed the use of memory for microprogramming and table look up in logical design, but that is not the principal use of memory in the computer industry. Rather, the more typical use is within a hierarchy to store information both on a short term basis while a program runs and on a longer term basis as permanent files. This section will present the various parameters and discussion relevant to this use. While the principal focus will be on core and semiconductor memories, slower speed electromechanical memories (drums, disks, and tapes) will be considered superficially, as their performance and price improvements have pushed the computer evolution. Since the typical uses for memory usually require read and write capabilities, write once or read only memory such as video disks will be excluded from the discussion.

It should be possible to discuss memory using a minimal number of measurement parameters because memory is the simplest of components. Many of the parameters are time variants, particularly price, which has declined at a rate

of 30% per year compound, (which amounts to over 50% in two years). The price is expressed only as price/bit, but it is important to know the price (or size) of the total memory system for which that price applies because of the economy of scale. In order to get the lowest price per bit, a user may be forced to a large system. (Is there economy of scale? What is the smallest or largest memory that can be built with the technology?)

Performance for electromechanical memories is expressed in two parameters: the number of bits that can be accessed per second after a transfer begins; and time to access the start of a block of memory.

The operational environmental parameters of power consumption, temperature, space and weight effect the utility of memories in various applications, and the reliability measures are needed in order to see how much redundancy must be placed in the memory to operate at a given level of availability and data integrity.

In summary, the relevant parameters for a given memory are:

1. state of development of the technology at the time the measurements are taken relative to the likely life span of the technology
2. price per bit
3. total memory size or total memory price

4. performance
 - a. the access time to the first word of the block
 - b. the time to transfer each word (data-rate) in the block

5. operational power, temperature, space, weight

6. reliability and repairability

A good example of a technology that is young relative to its total lifetime is semiconductor memory. Figure Memprice gives past prices and expected future prices of semiconductor memory. These memories have declined in price every two years by a factor of two, and that rate of decline is expected to continue well into the 80's because of continued increases in semiconductor densities. Figure Memsizperf, a graph by Dean Toome, Vice President of Engineering for Texas Instruments, shows memory size and performance improvements with time and includes Charge Coupled Devices (CCDs) and magnetic bubbles. These latter devices show slower performance figures than the other semiconductor memories because they are cyclically accessed in a fashion similar to disks.

Core and Semiconductor Memory Technology for Primary Memory

The core memory was developed early in the first generation for Whirlwind (1953) and remained the dominant primary memory component for computers until it began to be superseded by semiconductor technology. The advent of the 1 Kbit memory chip in 1972 started the demise, and the crossover point occurred for most memory designs with the availability of the 4 Kbit semiconductor chip in 1974.

The description and operation of the core memory is given briefly in Chapter 0 on the PDP-8 (page 0), and several of the significant circuit and production innovations are given in Chapter 0 on the 18-bit computers (page 0).

Over the period since the early 60s, the price of core memory declined roughly at a rate of 19% per year. This decline can be seen in the 12-bit (page 0), 18-bit (page 0) and IBM 360/370 memory prices (since 1964). The price of DECsystem 10 memory has declined at 30% per year (page 0), although it is unclear why. We can conjecture that the modular memory structure had a high overhead, and that with subsequent implementations the memory module size was increased, thereby giving an effective decrease in overhead electronics cost and a greater decrease in the cost per bit.

The cost of various memories was projected by several technology marketing groups in the period 1972-1974. Each study attempted to analyze and determine the core/semiconductor memory crossover point.

Three such studies are plotted in Fig. Core.wst along with Turn's [1974] memory price data and Noyce's [1977] semiconductor memory cost (less overhead electronics) projection. Note most crossover points were projected to be in 1974, whereas one study showed a 77 crossover. Even though all studies were done at about the same time, the variation in the studies shows the problem of getting consistent data for technology forecasting.

Figure Memprice shows the cost of semiconductor memory decreasing at a rate of 28% per year. Therefore, future memory costs may decline at an even more rapid

rate than in the past.

While these graphs of core and semiconductor prices and performance permit an understanding of trends in the principal use areas for these devices, processors, primary memory, cache, and small paging memories, additional information is needed for disk and tape memory in order to complete the memory hierarchy.

Disk Memories

Disk memories are a significant part of most systems costs in the middle-range minicomputer systems and for larger systems, dominate the costs. As a model of the technology improvement, we include figures supplied by S. H. Pufhuff of Memorex Corporation, which gives historical data and projects the costs of various devices till 1985.

Although access time is determined by the rotational delays and the moving head arm speed, the single metric that is most often used is simply memory capacity and the resultant cost/bit. In the subsequent section on memory hierarchies, it will be shown that performance parameters are less important because more higher speed memory can be traded off to gain the same system level performance.

Memory capacity is measured in disk surface areal density (i.e., the number of bits per square inch) and is the product of the number of bits recorded along a track and the number of tracks of the disk. Two figures, Areal.Digital and

Areal.analog give the areal recording densities using digital and analog recording methods. Figure Large.disk.price shows the price of the state-of-the-art line for large, multiple platter, moving head disks. Note that the price decline is a factor of 10 in 9 years, for a price of decline of x% per year.

Figure Memtrends shows the performance plotted against the price per bit for the technology in 1975 and 1980.

Magnetic Tape Units

Figure Tapechar shows the performance characteristics that are relevant for magnetic tape units. The data is for several IBM tape drives between 1952 and 1973. It shows that the first tape units started out at 75 inches per second and achieve a speed of 200 inches per second (nearly a factor of 3) in 1973. While this amounts to only a 5% improvement per year in speed over the 21 year period, since mass is moving, this is a rather impressive gain. It is akin to an improvement in automobile speed of a factor of three.

The bit density (in bits per linear inch) has improved from 100 to 6250 in the same period for a factor of 62.5, or 23% per year. Since both speed and density have improved, the tape data rate has improved by a factor of 167, or 29% per year.

Tape unit prices (see Fig. x) are based on the various design styles. Slow (minitapes) tape units are built for lowest cost. The most cost effective seem

to be around 75 inches per second (the initial design) -- when we consider only the tape. High performance units, though disproportionately expensive, provide the best system cost effectiveness.

Memory Hierarchies

{Information processing within a system at a given short time interval appears to be associated with relatively
{local memory use. The location of the computation changes as processing occurs
{but still the amount of memory involved in the computation usually stays constant.

{This phenomenon has been observed to apply to a number of activities within a
{computer system, and in each case the principle of operation appears to be
{the same. Thus, armed with the observation, a new memory--providing
{proportionally more memory for less, should always be useful if it is:

{cheaper and slower than an existing memory; or
{it is faster and more expensive (implying that it is smaller).

{Of course, any memory that is better in all dimensions of speed, size and
{cost automatically replaces an existing memory.

Transparency is another consideration about the utility of a given physical memory within a hierarchy that effects use. At the language level, nearly all

memory technologies are transparent, and the programmer does not have to consider the characteristics of the memory when writing programs. There are some exceptions, however. For example, the original FORTRAN had explicit statements to read and write tapes, and the modern counterpart has operations to access files in different access arrangements; hence, these operational improvements are not quite transparent. However, whether a processor has zero, one, or sixteen is irrelevant to the FORTRAN programmer.

The following table, in order of memory speed, gives the memory hierarchy as it currently exists. There is a continuum based on need together with memory technology size, cost, and performance parameters.

Table: Computer System Memory Component and Technology

Part	Transparency (to machine language programs)	Based-on
Microprogram memory	yes ¹	very fast
Processor-state	no	very small, very fast register set (e.g., 16 words)
Alternative processor- state context	yes	same (so speed up processor context swaps)
Cache memory	yes	fast. used in larger machines for speed
Program mapping and segmentation	yes	small number of association, or large map
Primary (program) memory	no	relatively fast, and large ¹ depending on

		Pc speed
Paging memory	yes	can be electromechanical, e.g., drum, fixed head disk, or moving head disk. Can be CCD or bubbles.
Local file memory	no	usually moving head disk, relatively slow, low cost
Archival files memory	yes (preferably)	very slow, very cheap to permit information to be kept forever

Nearly every part of the hierarchy can be observed in the computers in this book. Chapter 0, 0, and 0 describe PDP-11 implementations than use microprogrammed memories. These memories are transparent to the user, except in machines such as the PDP-11/60 which provide user microprogramming via a writeable control store. Mudge (chapter 0) describes the writeable control storage user aspects associated with the 11/60 and the user microprogramming. Chapter 0 describes similar possibilities in the LSI-11, although the writeable control store option was not available at the time the article was written.

{In principle, it is possible to have a cache to hold user microprograms, {hence, there could be another level to the hierarchy. Without memory of this {size and speed it is not possible to build microprogrammed machines that are {cost effective with hardwired systems. Mudge also discusses this aspect and {the technology tradeoff. In retrospect, the small, read only memory for {PDP-9 in 1967 was adequate and could have been enhanced to be used in later {machines. This would have permitted earlier entry of more complex ISPs at {lower costs. (Similarly this ROM could have been used to implement lower {cost PDP-10s earlier.)

To the machine language program, the number of registers in the processor state is the most non-transparent part of the architecture. This number is solely dictated by the availability of fast access, low cost registers. It is also occasionally the means of segmenting architectures (e.g., single accumulator based, general register based and stack based).

In 1964, even though registers were not available in single IC packages, the PDP-6 adopted the general register structure because the register was only a small part of the system cost (see chapter 00). In the chapter on the DECsystem 10 there is a discussion of whether the architecture should use general registers in an explicit (transparent) fashion, or whether the stack architecture should be used (also requiring a number of registers), but the use is transparent to the programmer. We adopted the general register structure to give better program control of a small number of local variables (i.e., locality) and permit the performance advantages. The change in register use and their number can be observed between the 12-bit and 18-bit designs and the later DECsystem 10 and PDP-11 designs.

As the number of registers increased, and technology improved, the processor state storage was increased to provide multiple sets of registers to improve process context switching time. In this way several multiple programs could be active at a time and selected on the basis of interrupt urgency, thus providing better real time and multiprogramming response.

In the late 60s, the cache memory was introduced for large scale computers.

This concept was applied to the KL10 and 11/70 in 1975 when the relatively large (1 Kbit), relatively fast (factor of 5), memory chip was introduced. The cache is described and discussed extensively in chapters 00, 00 and 00, so the reader may want to peruse page 0 if not familiar with the concept. It derives much power by the fact that it is an automatic mechanism and hence transparent to

{the user. It is the best example of the use of the principle of memory {locality. For example, a well designed cache of 32 Kbytes can hold enough local computational memory so that, independent of program size, 90% of the accesses are via the cache.

A similar memory circuit is required to manage (map) multiprogrammed systems by providing relocation and protection among various user programs. The requirements are similar to the cache and may be incorporated in the caching structure. The KI 10 used an associative memory for this mapping function.

The Atlas computer (Kilburn et al 1962) was designed to have a single, one level large memory and paging drum. This structure ultimately evolved so that multiple users could each have a large virtual address and virtual machine. However, the concept of paging mechanism works because there is not equally random access to each page, but rather only local access to various parts of a program by the processor at a given time. Denning pointed out the clustering of pages for a given program at a given time and called this phenomenon the working set [19]. For most programs the number of pages accessed locally is small compared with the total program size. Initially a magnetic drum was used to implement the paging memory, but as disk technology began to dominate

the drum, both fixed head and moving head disks (backed up with larger primary memories) were used as the paging memories. In the next few years, the relatively faster and cheaper CCD semiconductor memories and bubble memories are clearly the candidates for paging memories.

For medium sized to large scale systems there is no alternative to disks, with archival files on magnetic tapes. These permit files to be stored cheaply on an indefinite basis. For smaller systems there are usually fewer memory technologies used than in larger systems, because the smaller systems cannot afford the overhead costs (disk drives, tape drives, etc.) associated with the various technologies. At most, two levels of storage would probably exist as separate entities.

Alternatively, one might expect a combination of floppy disk, low cost tape, and magnetic bubbles to be used to reduce the primary memory size and at the same time provide file and archival memory. Currently the floppy disk operates as a single level memory. Here one can see two alternatives for technology tradeoff using the hierarchy: a tape or floppy disk can be used to provide removability, and archivability, whereas bubbles provide performance.

Terminals

Figure Terminalprices shows the price of the basic alphanumeric CRT versus time since 1970. Note that of all the computer components, the CRT has improved the least.

The particular curve is a least squared fit to the data points shown near the curve.

The lack of progress may have several explanations. First, it is constructed mostly of components (CRT, keyboard, plastic cabinetry, power supply) have been learned about and do not decline further in price.

Only a few semiconductors are employed in a CRT and these constitute the main component of improvement. Alternatively, there are many CRT manufacturers, all of which have insufficient demand to permit truly high volume manufacturing; as such, all have inordinately high costs.

MEASURING (AND CREATING) TECHNOLOGY PROGRESS

The previous sections have presented technology in terms of exponentially decreasing prices and/or exponentially increasing performance. This section presents a basis for this constant change. The metric of technology, $T(t)$ at some time, t , has been classically observed to be just:

$$T(t) = K \times e^{ct}$$

This can be converted to a yearly improvement rate, r , by changing the base of the exponential to:

$$T(t) = T \times r^{t-t_0}$$

where T = the base technology at t_0

and r = yearly increase (or decrease) in the technology metric

This is the same form used for declining (or increasing) cost from base c

$$C = c \times r^{t-t_0}$$

The questions that may arise upon studying the previous graphs (e.g., semiconductor prices) are:

1. Under what conditions does cost decrease exponentially?

2. Under what conditions does technology improve (i.e., in performance) exponentially?

3. When does a technology reach a limit of improvement?

Clearly there are manufactured goods that neither improve nor decrease in price exponentially, although many presumably could with the proper design and manufacturing tooling investments. The notion of price decline is completely tied to the cumulative learning curves of a) people building a product for a long time, b) process improvement based on learning to build it better, and c) design improvement by engineers based on learning from the history of design. Production learning per se is inadequate to drive cost and prices down because after an extremely long time in production, a few more units contribute little to learning. With inflation in labor costs, the costs actually rise when the learning is flat. In order to provide a base for predicting the inflationary effect, the consumer price index has been plotted Figs. CPI and CPI.log.

Learning curves don't appear to be understood beyond intuition. They are (empirical) observations that the amount of human energy, E_n , required to produce the n^{th} item is:

$$E_n = K \times n^d$$

where K and d are "learning constants". Thus, by producing more items, the repetitive nature of a task causes learning, and hence the time (and perhaps cost) to produce an item decreases with the number produced and not with the

calendar time an object is produced. Fusfeld [] conjectures that the technology of the i^{th} unit produced also improves exponentially with the number produced just as in the case of the learning curves. Thus, using the technology measure:

$$T_i = a \times i^b$$

he found the following technology progress constants:

Item	Measure, T_i	Quantity produced (i)	Technology progress (b)	Change observed in study	Total change
light bulbs	lumens/bulb	10^{10}	.04;.19	33	80
automobiles	vehicle h.p.	$3 \times 10^7; 10^8$.11;.74	10	6;13
titanium	p.s.i./\$/16	3×10^8	.3;1;1.04	10	350
aircraft	max.speed	2×10^5	.33-1.2	6	56
turbojet engines	fuel consumed, weight	1.6×10^4	1.06	2	2.9×10^4
computers	mem.size x rate	10^5	2.51	10^9	3.5×10^{12}

Computer technology has evolved substantially more than any other technology, and recent advances suggest this will continue, but perhaps to a slightly lesser degree. This rapid improvement has occurred in part because in computer technology there are so many materials to store, transmit, and process information and nearly all of which are electronic based. In the above technologies, most are mechanically oriented with associated physical limits. In essence the table is comparing systems constrained by Newton's Laws with those determined by Maxwell's Equations.

Fusfeld also showed that if the number of items produced increases

exponentially (with time):

$$i = e^{c/b \times t},$$

then calendar time and units produced can be used interchangeably.

There has been exponential growth of computer systems, so calendar time can be used instead of units. Time is an easier and more accurate method to measure than learning curves based on units.

The question of why the cost declines exponentially can be conjectured by using Fوسفeld's observation that it is because of learning curves and the exponential increase in quantity produced. Furthermore, this exponential increase raises a fourth question;

4. Why is the demand exponential?

The demand or quantity, q , sold per unit time is completely elastic, (exponential), according to the expression:

$$q = k/\text{price}$$

This creates a positive feedback market system whereby decreasing prices increase demand exponentially causing decreased costs (through learning) which support the decreasing prices as shown in Fig. Mkt.cycle (a variant of Fig. 1.).

There has been no attempt to answer question 2 of why technology improves exponentially nor is the answer for why cost declines exponentially at all satisfactory. Simple production learning does not account for the rapid technology changes in the integrated circuit, for example, where totally different production processes have been evolved to support the greater technology. It appears best to simply observe that the three situations have been true, and can be extrapolated to hold over the next few years because one can see ways by which each limit can be overcome.

Management scientists studying technology evolution (von Hippel, 1977, Fusfeld 1973) have made a number of observations about technology progress:

1. Quantity produced is simply a good dummy variable to measure improvement. (With exponential increasing sales, time can be used as the dummy variable.)
2. Technical problem solving is correlated with business activity. Inventors tend to be stimulated by sales and slacken efforts when sales are low. (This is a counter-initiative observation.)
3. Production alone does not stimulate innovation. A lesser number of inventions are stimulated by production needs. Of these, the same user-supplier relationship is the best framework. The users of equipment (the producer for the end product) stimulate the production equipment suppliers.

4. Major innovations come from use (sales). This is opposed to innovation arising from a technical possibility for which use is subsequently discovered. In the case of semiconductor technology, the computer designer (user) is responsible for many of the design innovations.

Computers also evolve rapidly under this and the levels of integration model because the suppliers are also significant users. The problem of market coupling is diminished significantly. Alternatively with computers, the users write (develop) the applications programs; hence again the user and developer are one in the same--this stimulates use!

The gains in packaging (i.e., printed circuit boards, backplanes, and cables) have been driven by production, versus product functional needs. Testing technology of all types has been motivated solely for production needs. These change least.

The cost of computing is the sum of costs which correspond to the various levels of integration described in Chapter 1 plus the operational costs. The levels were integrated circuits, boards, boxes, cabinets, operating systems, standard languages, special languages, applications components, and applications. In actual practice additional level-of-integration is often looked at as overhead. Using standard accounting practice, the basic hardware cost, at the lowest level, is then multiplied by an overhead factor at each subsequent outer level. While an overhead-based model may work operationally for a stable set of technologies, such a model will not adequately allow for

rapidly evolving technologies or the elimination of levels, for example. By examining each level, as this technology section and the packaging section attempt to do, observations can be made about the use and substitution of technology. More importantly, conclusions can be drawn about how structures are likely to evolve.

In preceding chapters and sections, semiconductor technology has been singled out as the main determinant of a computer's cost, performance, reliability and memory size. Magnetic storage technology is of equal importance because disk and tape memory densities evolve rapidly, even though the cost of a given physical unit does not. These units have become an increasingly major component of the price of computer systems as the costs of processors and primary memory have dropped, but their performance/price improvements have been notable, nonetheless.

Cost, Performance and Size (Economy of Scale) Relationships Design

{We often simplify the understanding of technology to just cost as the only
{dependent performance parameter. This simplification omits the most
{significant parameter, calendar time. In a similar way, we often fail to
{understand whether there is any associated economy of scale when considering
{the performance (utility) of a given set of devices. That is, do larger
{units perform significantly better than a set of small units when taken
{either independently or as a collection? Therefore, a more correct
{assessment, taking into account economy of scale, and performance, would be:
{

{ performance = $k \times \text{cost}^s \times r^t$

{

{ where k = base case performance

{ s = economy of scale coefficient

{ r = rate of improvement of technology

{ t = calendar time

{

{For many technologies a second parameter, size, is important because it is a
{measure of usefulness of the product that uses that technology. Size is
{related to performance in some situations, but is often independent. For
{example, a certain memory size would permit a certain number of accesses.
{For some applications, we would ask whether the size were adequate and for
{other applications we ask whether there is adequate time to access the data.
{Thus the systems performance in such an example is bottlenecked by either
{size or access and is the minimum of one of the resources. That is:

{

{ system performance = $k \times \min(\text{memory-size}, \text{memory performance})$

Figure Perfsize shows the various options of the amount of economy of scale:

1. Economy of scale holds. A particular object can be implemented at any price, and the performance varies exponentially with price.

$$\text{performance} = k \times \text{price}^s; s > 1$$

2. Linear price performance relationship.

- a. performance = $k \times \text{price}$

-
- b. performance = base + K x price; base overhead (recall this case is similar to 1, as described on page 0)
3. Constant performance, price independent
- $$\text{performance} = k$$
4. & 5. Only a particular device has been implemented. The performance (or size) is a linear sum of such devices.
- $$\text{performance} = n \times (k \times \text{price})$$

Sometimes, economy of scale effects are observed in situations where their applicability would not normally be expected. For example, assume a performance improvement feature exists that costs the same whether it is added to a large computer or added to a small computer. Adding that feature to a product that is already high priced will have a modest effect (say 5%) on the cost but a substantial effect (say 100%) on the performance. Adding the same constant cost feature to a lower cost product will have a substantial effect (say 200%) on the cost, but a performance effect similar (100%) to that obtained with the higher cost system (or perhaps even less). This condition is especially true in disks and computer systems. Use of a particular recording method employing costly logic for encoding/decoding, or addition of a cache memory is often employed to the high priced systems first. With time, and learning, the technique can then be applied to lower cost systems. For example, cache, a nearly perfect example of the constant cost add-on, first appeared in such large machines as the IBM 360/85 in 1968 and later migrated down to large minicomputers such as the 11/70 in 1975. This is very clear in

the case of the cache-based PDP-8, page 0.

{In Fig. Costvsstime, we see what typically happens as the cost of the lowest
{price unit is kept to a minimum and decreasing whereas the mid range product
{continues to increase, and the highest performance product increases the
{most, because it can afford the overhead costs. If we look at the basic
{technology metric, then there are really three curves too as shown in Fig.
{Techvstime. The first curve is applied to get the greatest improvement and
{be applied to the large price unit. With time the technology evolves and is
{reapplied to the first level copy in the middle range products (to most
{likely provide the best cost performance) and finally, several years later
{the technique becomes commonplace and is applied on low cost products. The
{resultant cost/performance ratios are shown in Fig. Cost/tech.

In most industries, the management of technology by applying it to products in various price and performance ranges in a more or less ordered fashion, has not occurred to the extent that it has in the computer industry. This is probably because no other industries have evolved in the same rapid and broad fashion as have the computer and semiconductor industries. The computer industry is fundamentally driven by the semiconductors technology push on the one hand, and by IBM on the other. IBM follows the strategy of applying technology on an economy of scale basis. This permits the technology to be first tested at the high performance, high price, lower volume systems before being introduced in higher volume production. The following examples (from IBM) show this at work. In printing, the high price/low volume to low/price high/volume introduction cycle was followed in the use of dot matrix printing,

chain printing, ink jet printing, and computer printing on flexible disks as a precursor to use as systems products using xerography. In magnetic storage, the cycle saw the basic technology for large disks as a precursor to the use of similar technology on smaller disks.

Technology Substitution

Since each constituent technology evolves at its own rate, the cost and performance of a system are roughly the additive and multiplication functions, respectively, of the parts. Usually when one component begins to dominate (e.g., packaging), then pressure occurs to more rapidly change and improve the technology to avoid the cost or performance bottleneck. Sometimes a slowly evolving technology is just eliminated as a substitute is found.

Some of the substitutions that have occurred:

1. Semiconductor memories are now used in place of core memories. Since the latter has evolved more slowly in terms of price decline, semiconductors are now used to the exclusion of cores. (This has not occurred where information must be retained in the memory during periods of time without power.)
2. Read-only semiconductor memories are now substituted for semiconductor logic elements.
3. In a similar way PLAs can be potentially substituted for ROMs and true

content addressable memories can replace various read write and ROM memories.

4. By using the principle of locality smaller, faster, more expensive memory can be introduced to replace intermediate speed, and moderately priced memory as in the cache.
5. The judicious use of CCD memory can cause drastic reduction (and quite possibly the elimination) in MOS random access, for primary memory. Note the fixed head disk is eliminated at the same time.
6. For small systems the main operational memories could be completely non-electromechanical; electromechanical memories (e.g., tape cassettes and floppies) would be used for loading files into the system and for archives. For yet lower cost systems, semiconductors ROMs could replace cassettes and floppies for program storage as in the programmable calculators.

After a while those components of computer system cost which are decreasing less rapidly than other components, or are remaining static, or are rising (like the packaging and power) may become a significant fraction of the total cost. Costs are additive and hence exponential improvements have disproportionate effects causing pressure for structural change.

For instance, although the PDP-8 is normally considered to be the first minicomputer, it post dates the CDC 160 (1960) and DEC's PDP-5 (1963).

However, the PDP-8 was unique in its use of technology to:

1. It eliminated the full frame cabinets used by other systems. This also presented a new computer style such that users could embed the computer in their own cabinets. A separated small box held the processor, memory and many options.
2. Automatic wire wrap technology was used to reduce printed circuit board interconnection cost. This also eliminated errors and reduced check out time.
3. Printed circuit board costs were reduced by using machine insertion of components.
4. The Teletype ASR33 (also used on PDP-5) was connected as the peripheral. It had a combined printer, keyboard, and paper tape i/o device (for program loading). It eliminated the paper tape reader and punch.

{Tradeoff Between Size and Performance

{A memory system might permit an encoding to be used such that memory accesses
{could be traded for memory size. This is about the only example where such a
{condition might exist. This condition is shown in Fig. Tradeoff for 3
{points. Here, the effective memory size can be increased or reduced by
{either using memory system accesses for computation or for storing
{precomputed values. This tradeoff permits the operating region for the
{system to be greatly extended to either include the memory (by giving up 1/2

{of the processor or 2.5 times the amount of processing/by giving up 1/2 the
{memory).

The Effect of the Research, Advanced Development Process of the State of the
Art Line

The complete development process is pipelined as shown along the lines of Fig.
1 in Chapter 1, page 0 with the with the stages: research, applied research,
advanced development (product breadboard), development, test, sell/build, and
use. In this model, ideas flow through the various organizations in a
process-like fashion, culminating in a product. Each product type has a
different set of delays associated with the parts of the pipeline and, at the
end, the

{education of use (also a delay) culminates in market demand. For well
{defined, commodity-like products such as disks and primary memory, the
{education of use delay is zero as each user "knows" the product. For a new
{language (e.g., PL/1) the delay is slow to build requiring user education, and
a new product may even
{eliminate one that potentially would create a demand.

A new product idea such as a virtual memory requires somewhat less time.

In order to understand the pipeline model of the development process, consider
the disk supply process. The technology (as measured by the number of bits
per areal inch) doubles about every two years (i.e., the density improves 41%
per year). IBM is estimated to invest about 100 million dollars per year in
the pipeline and associated pipeline for the manufacturing process, not

counting any improvements in semiconductor technology (semiconductors also influence the manufacture of disks). In this way, the IBM disks essentially establish the state of the art line in a structure that is typified by Fig. Techvstime. Note there are two other lines, each delayed about two to three years. Using the pipeline development process, development of competitive disks by other companies would lie somewhere about four to six years behind the state of the art line. This can be seen by looking at the development process and taking into account the delays through each stage. In order to be more competitive, the disk industry short circuits the various delays as it engages in reverse engineering; this results in only two year lags. In reverse engineering, the tools are micrometers and reverse molds. At time of the first ship of a new product by the technology leader, the product is purchased by competitors and basically copied on a function per function basis. The more successful designs use pin for pin compatibility so as to take maximum advantage of existing IBM design decisions.

From the process, it is also easy to see how by merely copying competitive products, one has built a process that guarantees products that will be two to four years behind leadership products and most likely sub state of the art. This structure is most likely the result of having a strong market function which operates to define products based on existing product use. If the design and manufacturing process is quite rapid, then such a strategy can be effective. This process can be very effective for software products because while there are no delays associated with manufacture, the time to learn about the products provide a time window in which any competitive pressure is likely to be slow to build. However, software engineers are the most creative and

most likely to ignore all previous work in an area.

A high technology, exponentially increasing (volume) product is denoted by:

1. Exponential yearly cost improvement (price decline) rates through product technology improvements as measured by price decline of $> 20\%$ (e.g., disks = .8, cpu = .79, memory = .7).
2. Short product life < 4 years.
3. Various types of learning curves. Some products have negligible {manufacturing learning and some denote forgetting through turnover, {increased training to meet demand and inflation.

The Product Problem (Behind the State-of-the-Art)

Typical product situations, including competitive "problems" can be seen in Fig. Product.cost.

When a product is introduced to the market, it has a relationship to the state-of-the-art line. There are five possible situations:

1. ideal (on the state-of-the-art line)
2. advanced (moves below the line)
3. late (slip in time to the right)
4. expensive (more than expected in cost, straight above the line)

5. late and expensive (to the right and above the line)

Situations 3, 4, and 5 are product problems because they are behind the state-of-the-art line and hence less competitive. This implies increased sales costs, lower margins, loss of sales, etc.

Note that a late product could be OK if somehow the cost were lower.

Similarly an expensive product is OK if it appears earlier in time.

Time is Money (and vice versa)

Thus product problems can be solved by either:

1. movement in time (left) to get on the line; or
2. movement in cost (straight down) to get on the line

since

$c = \text{cost at time, } t \text{ (in years)}$

$b = \text{base cost}$

$r = \text{rate of price decline}$

therefore, with exponential price declines a family of products over a long time will follow a cost curve, c .

t

$$c = b \times r$$

now

dc = change in cost above (or below) to get back to the state-of-the-art

dt = delay (or advance) in time to get back to the state-of-the-art line

let

$f = dc/c =$ fraction (%) of cost away from line

dt

$f = 1 - r$ poor cost, expressed as
project slip

and

$dt = \ln(1 - f) / \ln(r)$ poor timing,
expressed as poor
cost

These formulas permit the interchange of time and money (cost).

For example, in disks or cpu's where $r = .8$ and $\ln .8 = .22$

dt

1. $f = 1 - .8$

2. $dt = - 4.45 \times \ln (1 - f)$

(using 1.) A 1 year slip is equal to a 20% cost problem

(using 2.) A 10% cost increase is equal to a .47 year slip

Who does what, and their effect

By and large engineering, by establishing the product direction, has the greatest effect on the product. However, since most product problems may have multiple components, it's worth looking at each.

1. Timing

a. Engineering schedule slips translate into a competitive cost problem as a sub state-of-the-art, late product.

b. Manufacturing - ramping up the learning curve quickly by risk taking has a high payoff when considering the apparent cost or delay.

2. Cost

A number of components and organizations contribute to the total product cost, as shown in Fig. Net.

- a. Engineering is perhaps the major determinant of cost by the product design - number of parts, ease of assembly, etc. The most common cost problems occur by continued product enhancement during the design stage to provide increased functionality (called "one plussing the design").
- b. Manufacturing - direct labor and overhead really count. Making major {product moves with decreased learning (increased forgetting) all {serve to stretch the product time out and put pressure on getting {newer products. One curve shows the direct costs associated with {manufacturing assembly and then some learning should take place as {long as product volumes increase exponentially. New technology {materials show the greatest cost improvement for computers--here we {assume semiconductors and other electronic materials improve with {time. Note that by capital equipment investment (tooling), there can {be stepwise cost decreasing in materials.
- c. Inflation - while not a direct cost function, it combines with labor cost to negate the benefits that were obtained from learning effects.
- d. Note the costs are taken altogether. In terms of a sub state of the art product, the costs are compound. A one year late, 10% overcost product has the effect of being about 1-1/2 years late or about 30% too expensive.

3. Manufacturing learning

Learning curves and forgetting curves really matter. Left alone, a typical product may go down three alternative paths (see Fig.

Product.price):

t

1. $c = b \times .95$ decrease as little as 5%/year.
2. $c = b =$ stay constant with little attention.

t

3. $c = b \times 1.06 =$ increase with inflation.

Mid Life Kicker for Product Rejuvenation

By enhancing an existing product (the so called mid-life kicker) one can improve the cost/performance metric of a given product. This is non-trivial, and for certain products must be inherent (i.e., designed in). Under these conditions, improvements in cost go immediately to get the product back on the curve.

For example a factor of 2 in performance halves cost/performance. The effect, of doubling the density of a disk is to move the product back to the state-of-the-art line by a time shift. Plugging the factor of two improvement into the formulas:

$$dt = 4.45 \times \ln(0.5) = 3.1 \text{ years!}$$

This situation is shown in Fig. Product.improve and is compared with a 5%/year learning curve.

PACKAGING

Figure 1 suggests that packaging is a completely recursive process where boxes (packages) are placed within boxes (packages) on an indefinite basis.

Packaging is iterative, not recursive, because for each level (IC, module,...cabinet, room) of packaging, a completely different design problem exists. However, for a computer system consisting of a set of components, each component can be thought of as a system itself which in turn is decomposed into lower level components. The computer systems described herein can be nicely decomposed into three to five distinct levels.

Packaging the electronic and electromechanical components that constitute a system at a given level consists of: interconnecting the set of the components via signal carrying links and holding (packaging) them within some basically mechanical structure such that as a unit system they carry out the objective function. Figure 2 introduces a few of the packaging problem subtleties. There are some visual effects required in addition to the function for which the system was designed. The visual effects may, in fact, be the dominant reason (or deciding factor) in whether the object is purchased. While thusfar only basic electronic signal processing has been discussed, there are two important side effects to that signal processing: power is consumed and there is electromagnetic interference (EMI). Power not only needs to be delivered to

the object, it also needs to be taken away, and cooling or heat dissipation to some outside environment by convection, radiation or conduction is required so that the system will ultimately not burn up...or malfunction when stored or operated outside a limited temperature range. The EMI interference is bi-directional: signals emanate from the package and must be controlled in order to avoid interference with the radio frequency spectrum; and other equipment broadcasts radio frequency interfering (RFI) signals which must not interfere with the package's function. In addition to cooling (or cooling), other environmental constraints must be considered. These include humidity and air quality (dust, corrosiveness, salinity, etc.).

Seymour Cray*, the designer of the CDC 6600 (12/64) and 7600 (12/69) computers and most recently the Cray 1 (12/75) computer, described packaging as the most difficult part of the computers designer's job. The two major problems are heat removal and the thickness of the mat of wires which cover the backpanel. His rule of thumb indicates that with every generation of large computer, the size decreases by roughly a factor of five--and each generation takes roughly five years.

While it is easy to understand why one of Cray's super computers is so dominated by packaging, it is more interesting to examine the effect of packaging on small computers. At one extreme, the first hand-held scientific calculator, the HP35, was simply a new package for a common object, the calculator--which has been around about 100 years.

Although semiconductor density had been improving for several years, it was

*Lawrence Livermore Laboratory lecture in December 1974.

was not until densities were high enough to permit implementation of a calculator in a few chips, and not until those chips could be re-packaged in a particular fashion, that the hand-held calculator came into existence. Currently this embodiment is synonymous with the calculator name--in the future, the calculator might take on some other form (e.g., watch, pencil, voice actuated, hearing aid, notebook).

Packaging also seems to be the dominant reason for the PDP-8 and minicomputer phenomenon--although marketing, the coining of the name, and the ease of manufacture (also part of packaging) are alternative explanations. The packaging advantages of the PDP-8 over predecessor machines are listed in the previous section.

The Packaging Design Problem

Packaging is the complete design activity of interconnecting a set of components via a mechanical structure in order to carry out a given function. In order to package a large structure such as a computer, the problem is further broken into a series of levels each with components that carry out a given function. Figure 3 shows the hierarchy of levels that have evolved these last twenty years for the DEC computers. There are eight levels which describe the component hierarchy resulting in a computer system.

For each packaging level there is a set of interrelated design activities as shown in Fig. 4. The activities are almost independent of the level at which they are carried out, and some design activities are carried out

across several levels. For some levels a design activity doesn't exist. For example, since integrated circuits operate at low voltages, there is no safety design problem associated with IC chip design. However, since there are many circuits on a single IC and each can be operated at a special voltage, there are both power supply design and power dissipation (cooling) problems associated with IC design.

While the initial design activities indicated in Figure 4 are each aimed at solving a particular problem, the solving of one problem in computer engineering usually creates other problems as side effects. For example, the integrated circuits and other equipment that do information processing require power to operate. The power creates a safety hazard and is provided by power supplies that operate at less than 100% efficiency. These side effects create a need for designing insulators and providing methods of carrying the heat away from the power supply and the components being powered. In this way, cooling problems are created. Sometimes cooling is carried out using conduction to an outside surface so that it may be carried away by the air in a room. Most cooling problems are solved by convection with a fan in the cabinet which carries air into the room so that the room air conditioning system is left with the problem of carrying the heat away. In this process, the fans create acoustical noise pollution in the room, making it more difficult for humans to work in the room. If the computer is used in an unusually harsh environment, a special heat exchanger is required in order to avoid contamination of the components within the computer by the pollutants present in the cooling air flow.

Finally, a particular package exhibits mechanical characteristics such as weight and size. These parameters directly affect manufacturing and shipment costs. They determine whether a system can be built, and whether it can be shipped in a certain size (e.g., Boeing 707 airplane) or distribution channel (e.g., parcel post or United Parcel). The dynamic characteristics determine the type of vehicle (special air ride van for electronic equipment) in which equipment must be shipped.

It is also necessary to examine the particular design parameter in order to determine whether it is a constant (meets the German VDE standard x), a goal (cheap as possible), or part of a more complex objective function (measured in bits/sec/\$ or part of a system benchmark--jobs/sec/\$). Thus a system may be determined by a government standard (e.g., safety and EMI), a market consideration, a distribution channel (shipped via parcel post), a method of sales (catalog and customer installable), a competitor characteristic (e.g., no louder than x, faster than y, or cheaper than z), or a physical operating environment (dusty, hot and with high RFI).

The following table lists the various kinds of design activities and whether they are likely to goals, constraints, or part of more complex objective functions. The table also gives the dimensions of various metrics (e.g., cost, weight) available to measure the designs and many of these metrics are used in subsequent comparisons.

Table Design Activities, Metrics and Environment Setting Goals and Constraints

<u>Design Activity</u>	<u>Determining Environment and [Metrics]</u>
Primary function and performance (e.g., memory)	Market, next highest level (i.e. the consumer) of system [memory size (bits), operation-rate (bits/sec.)]
Human engineering	Human factors criteria, competitive market
Visual/Aesthetics	Market, other similar objects, the environment in which the object is to exist--usually only important at outer-most level
Acoustic noise	Government standards, operating environment, market [decibels]
Mechanical	Shippability (e.g., air cargo container size, truck vibration), handling, assembly/disassembly time [weight, floor area, volume, expandability, acceleration, frequency response]
<u>EMI</u> radiation input	Government standards, must operate within intended environment (e.g., high noise) [power vs. frequency] [standard no. 000]
Power	Operating environment, market [watts, voltage supply range]
Cooling and environment	Market, intended storage and operating environment, government standards [heat dissipation, temperature range, air flow, humidity range, salinity, dust particle, hazardous gas]
Safety	Government standards [standard no. 000]
Cost cost/metric ratios	[cost/performance (its function)--cost/bit and cost/bit/sec., cost/weight, cost/area, cost/volume, cost/watt]
density metrics	[weight/volume, watts/volume, operation-rate/volume]
power metrics	[operation-rate/watt; efficiency = power out/power in]
reliability	[reliability--failure rate (Mean Time Between Failures, Mean Time To Repair (MTTR))]

Given the basic design activities, one may now examine their interaction with the hierarchy of levels (i.e., the systems) being designed (see Table 00). This is done by looking at each level and examining the interaction of the design activities for that level with other design activities (e.g., function requires power, power requires cooling, cooling requires fans, fans create noise and noise requires noise suppression).

Table Interrelationship of Hierarchy of Levels and Design Activities

Activity \ Level	Level						
	Chip	IC	Module	Backplane	Box	Cabinet	Computer Systems
Functional	logic----->			spatial\		what	selection of right components
	electrical					config-uration options	
	circuit design		physical layout	physical layout	will fit and operate	boxes, what config-urations will operate	user is designer
	physical layout						
Human interface						location of console, size for use	placement for use
Visual					somewhat seen bought for integration	deter-mines system appearance	set of cabs, attractive place to be
Acoustic				Airflow vibration			quiet for operators and users
Mechanical	buildable	signal transmission	shippable	serviceable			floor load room size
EMI	noise coupling and rejection of ext.RFI		inter/intra module noise		coupling, RFI	RFI containment, external RFI shield	away from RFI input (outside operating range)
			containment				

			and shielding			
Power	special on-chip	dist. and regulation	dist. and regulat- ion	control, dist. and regulat- ion	inter- connect- with computer system	by user. special power supplies for high avail- ability
Cooling and other environ- ment	chip to IC cooling module special envir. special envir.	IC to cooling	module	cooling & covering	source	interbox coupling to room air envir.
Safety		power-----> for various systems		deter- mines safety if used at this level	deter- mines user safety	
Dominant design activities	circuit logic logic	----->		mech- anical, power, cooling, EMI, acoustic	config- tion visual, shipping EMI, safety	user config- ation design

Note: The box and backplane levels can be considered as a single level.
 (Alternatively, the box level may be eliminated in large systems)

Computer Systems Level

The topmost level in the preceding table is the "computer system", which for the larger mini and DECsystem 10 computers consists of a set of components within cabinets, housed in a room, and interconnected by cables. The functional design activity is the selection and interconnection of the cabinets, with a basic computer cabinet holding processor, memory, and interfaces to peripheral units. Disks, magnetic tape units, printers, and terminals occupy free standing cabinets. The functional design is usually carried out by the user and consists of selecting the right

components to meet cost, speed, number of users, data-base size, language (programming), reliability and interface constraints. Aside from the functional design problem, cooling and power design are significant for larger computers. For smaller computers, accessibility, acoustic noise, and visual considerations are significant because these machines become part of a local environment--and must "fit in".

Cabinet Level

The purchaser or user of a "computer system" takes the collection of computer system components (the cabinets) and connects them to form a system as described above. For the computer hardware designer, the component associated with the cabinet is often his largest system. For example, the central processor designer is often responsible for primary memory and interfaces to peripheral equipment. This functional design insures that the various components (i.e., boxes) that make up a cabinet level system will operate correctly when interconnected. Since the cabinet is the lowest level component that most users interface to and observe, the physical design, visual appearance, and human factors engineering are also dominant design activities. Safety and EMI characteristics are also important because the cabinet also serves as the outermost place that shielding can be installed. Cooling and power distribution can be a problem within cabinets, since a number of different boxes may be mounted within the same cabinet. Thus there may be thermal and EMI noise pollution problems when box components interface poorly. The mechanical structure of a cabinet is responsible for maintaining its physical integrity when shipped.

Box Level

Box level functional design consists of taking one or more backplanes, the power supplies for the box, and any user interface such as an operator's console and interconnecting them mechanically. The PDP-8 is the original integral box level design (see Fig. 00, page 00). For systems that are not sold at the box level, no separate box is required, and backplanes are mounted directly in a cabinet; hence box and backplane design are merged as one. In a similar fashion, when only a single backplane is used, the power, the cooling, the backplane, and the box to cover and support the other components become a single design.

If systems are sold at the box level, then the visual characteristics may be important, otherwise, the design is basically mechanical and consists of cooling, power distribution, and control of acoustic noise. The structure must be basically quite sound in order to protect the unit during shipment.

Backplane Level

This level of design is the final level of interconnection for the computer components that are designed to stand alone, such as a basic computer disk, or a terminal. Backplane design is part of the computer's logical design. Figure 00, page 00 shows the Omnibus backplane for PDP-8 and how it is used to interconnect the basic components. Secondary design activities include holding, powering, and cooling the modules so they will operate correctly. Since the signals are transmitted on the backplane, there is an EMI design problem. For industrial control systems whose function is to switch power (see PDP-14 on page), additional safety problems are created.

Module Level

In the third and fourth generations, module level design is one of the physical layout design problems associated with logical design. In the second generation, module level design was a circuit design activity taking discrete circuits and interconnecting them to provide a given logic function. Chapter 00 describes the electronic circuit designs of the second generation. Now, this interface between circuit and logical design is within a chip level design. This shift in roles (function) will be described below on the section discussing the interrelationship between the technology generations and packaging. The integrated circuits that perform the functions are assigned to different positions on the module. Note in a recent implementation of the LSI-11, a processor with 32 Kbytes of memory is packaged on a single double sized module (see Fig. 00, page 00). Since module level design is basically electronic, power, cooling and EMI considerations dominate.

Integrated Circuit Package and Chip Level

[need brief description of . ceramic, plastic DIP . TO-5, flatpack, alternatives to DIP]

Most integrated circuits used in the computer industry today are sold in a package configuration that has two rows of pins, and is hence called a Dual In-line Package, abbreviated DIP (Fig. 5, page 00). Because of the popularity of this packaging style, the terms IC and DIP are often used interchangeably. This is not strictly correct, as the integrated circuit is actually that part of the semiconductor wafer called a chip (Fig. 00, page 00). Until multiple chips are packaged within a single DIP, creating

different design problems at the chip and DIP levels, the chip and DIP can be discussed as a single level. In some applications the multiple chip in one DIP is already here, as Figure 5 shows a single DIP with two chips as used by the LSI-11.

Logical design is a part of the functional design problem at the chip level. The task is to supply a terminal behavior that can be used at the next (module level) and, depending on the technology generation, this function can vary from a simple gate (in the early third generation) to a full computer (late in the fourth generation). A completely correlated design activity is the circuit design with transistors; as such, the logician can interconnect the various transistors in accordance with a well established set of rules so that the circuit will carry out the required logical function. The physical assignment of circuits (functions) to geometry forms the remaining part of the functional design.

Other design activities at this level include generic to electrical signal processing: power, heat dissipation and EMI. Since some ICs are designed to operate in hostile environments, there is a considerable mechanical design activity associated with packaging, interconnection and manufacturing.

The Packaging Evolution

Figure 6 shows the relation of packaging and the computer classes for the various computer generations. For each new generation, there is a short, evolutionary transition phase, but ultimately the new technology is

repackaged such that a complete information storage or processing component (bit, register, processor) occupies a small fraction of the space and costs a small fraction of the amount it did in the former generation. Quite discrete events mark packaging characteristics of each generation, starting from 1 bit per vacuum tube chassis in the first generation and evolving to a complete computer on a single integrated circuit chip in the fifth generation. Not only the size of the packaging changed, but also the mounting methods. In the first generation, logical units were permanently mounted in racks, whereas in the later generations they were made removable for ease in servicing.

Although it is difficult to do precisely, the second packaging time line indicates the transition in package size for a basic processor and primary memory. For the minicomputer, WHIRLWIND is used as an example (perhaps overzealously) which was roughly the size and power of the LSI-11 (a four chip, 16-bit processor). Whirlwind occupied a complete building, and the processor portion occupied a large room.

In the case of the computer classes, several computers are given to mark the various technologies and sizes. Here, it is clear that the package plays a significant event in the determination of the new classes. Note the super, mini, micro, calculator, and terminal lines.

It is difficult to be too precise about the class distinctions, as a given computer class can neither be precisely defined in price nor performance terms. A workable method is to say that combination of the price, packaging

 hierarchy, and market distribution structure determines the computer class.

The following table shows the packaging hierarchy with time for what might roughly be considered to be a minicomputer.

Table Physical Structure Packaging Hierarchy versus Technology Generation for Roughly Constant Performance Computers

<u>building</u>	<u>room</u>	<u>room cabinets</u>	<u>cabinet box(es) backplane</u>	<u>cabinet ? module(s)</u>
rooms	cabinets	backplanes	modules	integ.ckts.
bays	?	module	integ.ckts.	chip
chassis	chassis	disc.ckts.	chip	
disc.ckt.	disc.ckts.			
early first (e.g. Whirlwind)	late first ?	late second (PDP-1)	late third (PDP-8/E)	early fourth (CMOS-8)

Note: The line shows what is required for the processor part only. Only one each (e.g., 1 room...1 integrated circuit) of next lowest level component is required per processor.

In a similar fashion, one can observe how a particular computer component, now called the Universal Asynchronous Receiver-Transmitter (UART) has evolved with time in the following table.

Table ___ Packaging Hierarchy Evolution for Universal Asynchronous Receiver Transmitter (UART) Telegraph Line Controller

backplane

modules	2 modules	module		
discrete	discrete	IC	IC	
circuit	circuit	chip	chip	chip area
early second	late second	early third	late thrid	late fourth

This logic carries out the function of interfacing to a communications line carrying serial data which is encoded by a DC level to denote a 1 or 0 and transforming the data to parallel on a character by character basis for entry into the rest of the computer system. The UART has three basic components: the serial/parallel conversion and buffering; the interfaces to both the computer and to the communications lines; and the sequential controller for the circuit. The preceding table shows the evolution of the UART packaging hierarchy with time.

The UART is probably the first fourth generation part computer component, since it is somewhat less complex than a processor, yet large enough to be identifiable with a clean, standard interface.

Of historical note, DEC played a significant part in the development of the UART technology. With the PDP-1, the first UART function was designed using 500 Khz systems modules (see Fig. 8), and was used in a message switching application as described in chapter 00. The interface was called a line unit, and was subsequently repackaged in the late second generation to be on two extended systems modules (see Fig. 9). The UART function was also built into the PDP-8/i using two modules, but substantially smaller

ones than those for the PDP-1. In the 680/i, a PDP-8/i driven message switch, the UART function was accomplished by programmed bit sampling (see page 00). Late in the third generation (or at the beginning of the fourth generation), some designers from Solid State Data Systems, a small company on Long Island, worked with Vince Bastiani at DEC and developed a UART that occupied a single chip. This subsequently evolved to become a standard IC, and has been used throughout the industry. Figure 10 shows one its uses at DEC, the PDP-11 DL-11 communications line interface module.

The DEC Computer Packaging Generations

With this general background on packaging, one can examine the DEC packaging evolution more specifically. Figure 3 gives the general archetype of the packaging structure at DEC, but specific schemes have been used at different generations. Figure 12 shows how the hierarchies have changed with the technology generations. The figure is segmented into the different product groupings. A product is identified as being at a unique level if it is sold at the particular packaging level. Note that with time, lower level components are available for use within other systems and different packaging hierarchies. The first computers (i.e., PDP-1 to PDP-6) were sold at the cabinet level as complete hardware systems. The PDP-8 was significantly smaller and sold at the mechanical box level so that it could be incorporated into other hardware systems and within the packaging hierarchies of OEM customers. The PDP-8 was also available at the cabinet level for a complete system.

Subsequently the systems evolved to be available down to the backplane

 level with the LSI-11, and currently to the module level for the CMOS-8.

The following tables expand on Fig. The following table shows the original packaging hierarchy for most of DEC's second generation computers, which used a relatively common packaging scheme based on the PDP-1. The most significant change occurred in the late second generation when Flip Chip modules (see Fig. 12) were introduced so that backplanes could be wire wrapped automatically, enabling lower cost and greater scale production.

Figure 12 shows:

Table Physical Structure Hierarchy for Second Generation Cabinet-Held Computer Systems (1, 4, 5, 6, 7, 9, KA10, LINC, 12)

cables	computer cabinets*	floor/room
int.cables	backplane(s)	cabinet frame
soldered wires**	module***	backplane, connect
PCB etch	discrete ckt.	PCB
(interconnect)	(component)	(holding structure)

*single cabinet pair for PDP-5

**changed to wire wrap after PDP-6

***originally called systems modules, and changed to Flip Chip modules after PDP-6

In the case of DECsystem 10 and the 18-bit family, these computers were not available as boxed components. The unavailability of "boxed" 18-bit computers may have caused their demise. Also of significance was the packaging of later third generation computers (past the 8/I and 8/L) to the M-series Flip Chip modules to reflect the change in logical design with ICs. In all recent computers the module is treated as part of a single monolithic logical design, and partitioning (assignment of particular functions to physical components across multiple levels) is across the IC,

module and backplane levels. This design activity is also shown in Table 00.

Table Physical Structure Hierarchy for Third Generation Box-Held Computers
 (8/I, L, E, F, M, A, 11/04-11/70)

	box ----->	int.cables	cabinet boxes	cab.frame
wire wrap bp.	module	box, etc.		
PCB etch	ICs	PCB		
Wire board	chip	DIP		
(interconnect)	(component)	(holding structure)		

Computer-in-a-box

Computer-system-in-a-cabinet

The change to wire wrap technology also enabled the box level production of computers as shown in the following table for PDP-8, LINC-8 and the 8/S. Here, the changes to wire wrap and two level (box and cabinet level) products is clear.

Table Physical Structure Hierarchy for Second-Generation Box-Held Computers
 (8, LINC-8, 8/S)

	box ----->	int.cables	cabinet box	cab.frame
int.cables	backplanes	box		
wire wrap	module	backplane + connect		
PCB etch	discrete ckt.	PCB		
(interconnect)	(component)	(holding structure)		

Computer-in-a-box

Computer-system-in-a-cabinet

The change to IC packaging in the third generation is shown in the preceding table for computers that are sold at the box level. Note that this set of computers includes all the PDP-8 and PDP-11 computers. In this

 way, the minicomputer tradition continues to be available at the box level of the hierarchy.

The LSI-11 hierarchy is represented in Fig. 12 at three levels. Although components are sold as separate modules (e.g., communications line interfaces, additional primary memory), a complete system requires a backplane, thus the lowest level for the product is the backplane. For larger systems, a power supply is combined and placed in a metal box (the 11V03), and finally a complete system such as the 11V03 is available with terminal, and mass storage in a single cabinet.

The CMOS-8 hierarchy is presented in the following table.

Table Physical Structure Hierarchy for Fourth Generation Module-Held Computer System (CMOS-8)

			terminal-->cables	system in desk
		-----> cables	modules term.	term.,mass store, desk
PCB etch	module	ICs	PCB	
wire board	chip	(compon-	DIP	
(inter-	ent)	(holding structure)		
connect)	Computer Module		Terminal	System

The CMOS-8 module, described more completely in Chapter xx, is a single-board complete computer with processor, 16 Kword memory, and all the optional controllers to directly interface up to five peripheral options. For a computer packaged in this fashion, a backplane is not used, and

doesn't exist at a level. The computer system within a CRT terminal exists at somewhere between a box and cabinet configuration. The third level of the hierarchy is within 1 or 2 cabinets (i.e., a desk) complete with mass storage (i.e., two floppy disks) and a separate stand alone printer (cabinet) for letter quality printing.

Specific Cabinet and Box-level Designs

Cabinet and box-level design is perhaps the most difficult part of computer design, yet it is perceived (incorrectly) to be trivial and left until last. It is also misunderstood because there is not a single discipline or complete set of well understood laws that govern the design. This deception arises from the fact that, on first glance, the only physical law is that not more than one thing can occupy the same space at a given time. Also, everyone has some (usually strong) feeling about what an attractive package is. However, a number of disciplines must be understood, including acoustical engineering, heat transfer, aesthetics, human engineering, RF engineering, and the understanding of package functionality. In addition to this understanding, each box type requires a separate specialized manufacturing process. Obviously no one person fully understands all the disciplines necessary to correctly design a package. Herein lies the problem: Packaging is the integration of a number of separate disciplines. To be specific one needs to look at the basic packaging problem by first enumerating the possibilities for placing modules within a particular box package. By being quite restrictive, one can build a grammar with six production rules that describe the possible DEC boxes. Even with this simple grammar, about three million

possibilities can be generated. The large (combinatorial) number of possibilities arise from the basic size and way a box is held, the console mounting, cooling, power supply location and the way modules are mounted within the box. The packaging possibilities can then be explored by simply writing statements that express the alternatives for the separate decision dimensions.

A Grammar and Syntax to Generate Three Million Box/Cabinet Alternatives

1) Size and Mounting

The box is (1/2 or full) cabinet depth by (3-3/4", 5-1/4", 10-1/2", 15-3/4" or 21") high, mounted with (fixed slides, right hand hinge, or left hand hinge).

2) Console

The console is (non-existent, simple power on/off, maintenance only, full feature maintenance/programming).

3) Cooling

The cooling is carried out by (box-level or cabinet-level)(fan or fans) producing (plenum or forced) air flow that is (normal or parallel) to the module mounting axis, which has eight possible orientations, or else the cooling is carried out by natural convection with air entering at the

(top, bottom, right side, left side, front, or rear), and exhausting through the (top, bottom, right side, left side, front, or rear).

4) Power Supply Location

The external power supply is mounted (separately, attached behind the box, attached to the back of the cabinet, or attached on the front of the cabinet), or an internal power supply is mounted using a mounting scheme that is the (same as or different than) that of the modules, and is located at the (top, bottom, right side, left side, front, or rear) of the cabinet.

5) Module Mounting

There are (no, one, or many) backplanes for mounting modules.

6) Module Pins

The module pin axis is (top, bottom, right side, left side, front, or rear) mounted, oriented in the (horizontal or vertical) plane with the IC side facing (up, down, right, left, front, or rear).

One can state the six sentences (1 point in the decision space) that describes the PDP-8/A:

1. The box is 1/2 cabinet depth by 10-1/2" high with fixed mounting.

-
2. The console is for programming and maintenance.

 3. Cooling is carried out by box-level fans giving forced air cooling flow parallel to the module mounting axis.

 4. The box power supply uses a different mounting scheme and is located at the bottom of the box.

 5. There is one backplane for mounting modules.

 6. Modules are rear (wall) mounted, oriented in the horizontal plane with IC side facing up.

Alternatively, one can give the conventional three views of a PDP-8/A to show the location and configuration of the various box components (see Fig. 13). In this case a top view is adequate to show the configuration and air flow--although a side or front view is necessary to show the number of modules and their length.

Obviously there are more than the 3.9 million possibilities that can be generated from the above grammar--although not all of these are feasible. It is important to have some way to relatively methodically examine the more common possibilities. In fact, the above grammar doesn't even include the original PDP-8 package which consists of a box mounted on slides with two backplanes that can be hinged apart so as to provide access. Thus, it was not purely a box mounted on slides or a book-like structure suspended

with hinges.

Of all the dimensions to consider in the design, perhaps the most important is how the box (or module mounting structure) is placed in a cabinet. This placement effects air flow, shippability, configurability, cable placement, and serviceability, and is a classical case of design tradeoffs. The scheme that provides the best metrics such as packaging density and weight, may have the poorest access for service, and the most undesirable cable connection characteristics. These characteristics are given in the following table:

Table ___ Fixed, Drawer and Hinged Box/Cabinet Mounting

	<u>SERVICE ACCESS</u>	<u>CABLING</u>	<u>DENSITY</u>	<u>COOLING</u>	<u>APPLICABILITY</u>
Fixed	Good for either backplane or module, but not both unless a thin cabinet is used	Best (i.e., shortest, and right)	Good for thin or rear cabinet PS mounting	Best (known)	Box not needed; box can be used
Drawer	One side access	Long + moves	Very high	Can be	High density, self contained
Drawer (with tilt) for service	Good	Longer + much more moves	Very high	Cooled*	
Drawer vertical mounting modules	Very good	Longer + moves	High		
Hinged (module backplane)	Very good	Short	Medium	Good (if fans are fixed to	Separate box is awkward

cage)

*Density restricts cabinet airflow

Figure 14 shows the various boxes (top view) as they fit within a cabinet profile. In some cases there is no two level cabinet/box hierarchy because these structures have only been available at the cabinet level as shown in Fig. 12. Packaging is largely a matter of designer preference because the only clear goal is manufacturing cost. Marketing considerations, especially for OEM use, drives toward getting the highest density to minimize volume, floor space and rack mounting height. Despite this, many designers prefer packaging schemes that are not the highest density, but rather are fixed so that the power, cooling, and EMI conditions can be well-understood and so that the cabling can be done more rigidly. This has been particularly true for the larger computers.

In the following table, metrics are given for the various packages in terms of their pay load, which is printed circuit board area. This measurement produces load ratios of board area per cubic foot of box, and board area per square foot of floor.

Power Supplies

while logical functions can be performed using small quantities of electrons and thus accommodated in very small physical structures, the power to move those electrons at useful speeds comes from power supplies which do not scale down in size as readily as the logical functions they support. The principal reason that power supply technology has not provided the impressive increases in capability per dollar or capability

per cubic foot that semiconductor technology has is that power supplies involve such materials properties as voltage breakdown limits, dielectric constants, and magnetic permeability. Most of the equations involving these properties include a term for the "goodness" of the material and a term reflecting a physical dimension. If progress in materials science limits the "goodness" of the material to a certain value, increased capabilities in terms of voltage breakdown rating, capacitance, or inductance is gained by making the component physically larger. Progress in improving the "goodness" of the materials used in power supplies has not kept pace with the progress in semiconductor technology.

This section gives the cost and performance characteristics of various DEC power supplies as shown in Table x. It lists a set of supplies, their (relative) cost, MTBF, volume/year, power, technology, size, regulation, voltage, efficiency, and the ratios cost/watt, cost/cu. in., and the year introduced.

Modules

{It isn't clear what should be said here. Phister has quite a bit. It {will depend on what happens in the Modules chapter. We should clearly {reference board area. We must bring in the backplane and get at the pay {load of board area, square inches of panel area and cu. inches, then talk {of efficiencies. The pin should also be a metric. Line width is the {standard metric...but I'd just as soon not give it.

AN OVERVIEW OF MANUFACTURING

The result of a design project is an entity which is manufactured, but very little is written about manufacturing in computer engineering literature. Such literature generally discusses algorithms, logic design, and circuit technology. Yet for a computer to be commercially successful, it must be manufacturable, economically operable, and serviceable. Moreover, for most of the computer engineering discussed in this book, because the designs are intended for production, engineering costs are small (1-10%) compared with other product and lifecycle costs. The product cost is determined by the price of the components and the manufacturing process. The lifecycle cost includes the purchase price, the operational costs, and service costs. For production, machines must be easy to assemble and test, repair must be rapid, engineering changes must be introduced smoothly, and the production line cannot be held up because of shortages of components -- all parts of traditional manufacturing understanding.

A detailed discussion of manufacturing is clearly beyond the scope of this text, but is discussed extensively in [Grant and Leavenworth, 19xx]. Information on test equipment is to be found mainly in manufacturers manuals, but a survey is given in []. A reference on quality control is [Juran, 1962] and a discussion of learning curves is in [].

The Life Cycle of a Product

Figure Lifecycle shows a simplistic process flow for the major phases and milestones in the life of a product. In reality, planning and designs for

many of the phases go on concurrently. The early research, advanced development, and definitional phases are not shown. Often times, products may proceed from the idea to engineering breadboard and are terminated because they may not meet original goals or because better alternative ideas arise.

The engineering breadboard usually is built with wirewrapped boards rather than printed circuit boards in order to accommodate changes. At this stage schedules are made for manufacturing start up. Other organizations formulate and execute plans: systems' engineering - for product test/verification; software engineering - special software and verification; marketing - for promotion and product distribution; sales - for training; field service - training and parts logistics; and software support.

The engineering prototype proves the design using the printed circuit modules that will be used in the manufacturing. All peripherals in the planned systems configurations are tested on the prototype. Usually a number (10-100-depending on the complexity, cost and volume) of prototypes are constructed. The complete system must meet the product specification.

Limited release (LR) to manufacturing is a major milestone. The product is placed under formal engineering change control; specifications and documentation are available for the product and manufacturing process. For the integrated circuits, second sources of supply and testing procedures are in place. Process control tapes for the numerically controlled machine

tools, e.g., component insertion, backplane wiring, and printed-circuit board drilling, are ready. Any special tooling for the mechanical packaging has been obtained. Finally, the engineering change order (ECO) procedure for the hardware and microcode is set. Testing at all levels has been specified; test programs for computer-controlled testers have been written, special test equipment has been built, and diagnostic programs are ready.

Design maturity testing (DMT) with a number of engineering prototypes verifies the design and justifies the risk of a pilot run. Tests for reliability and functionality are conducted. Environmental testing (shock, temperature, humidity, static discharge, radiation, power interrupt, safety) is conducted at this stage.

The pilot run shakes down and verifies the actual manufacturing process by building a small number of units at the manufacturing plant using the product and process documentation.

Product announcement usually occurs during the DMT period but can be as early as limited release or as late as first customer shipment -- depending on the marketing strategy. This strategy is clearly a function of the volume, novelty, and competitive needs.

Process maturity testing (PMT) verifies that the product is being manufactured with the desired cost, quality, and product rate. The steady-state phase of manufacturing continues, with possible perturbations

due to the introduction of product enhancements or process changes to lower product costs, until the product is phased out.

Manufacturing Process Flows

An overview of a manufacturing process is given in Figure Manufover which shows how a product moves through the various factories. There are often different plants for boards, peripherals, memories, and central processors that form the box-level and independent hardware (for example, memory modules) components. Integration from the other stages and stack storage occurs at the stage called final assembly and test (FA&T). Here, the software system operations manuals and other documentation that is to be run is also integrated and tested.

Figure process60 gives the complete flow for a typical volume manufacturing line -- the 11/60 central processor facility in Aguadilla, Puerto Rico. Integrated circuits are manufactured outside this factory and are sent to be inserted, soldered in, and tested. The photographs in Figures x through y show several test points in the process.

Integrated circuit failure rate versus time follows a bathtub-shaped curve, high at either end of the life cycle, as shown in Figure Bath. The manufacturing process includes extensive thermal cycling to ensure that the "infant mortality" cases are discovered early during manufacturing, because it is more expensive to find defective components at the larger, more integrated systems level. The temperature/humidity environmental chambers which house n CPU's each are shown in Fig. TChamber. Figures 2224 and

lvstation show small heated enclosures used to induce failure during the test and repair of modules.

Since testing occurs at each stage in the manufacturing process, dedicated logic must be added to the design, to provide physical access "probes" for the test equipment. To test a particular function, it must be:

specifiable, invokable, and observable. For example, the function of an adder can be clearly specified, but it cannot be easily invoked or observed if its inputs and outputs are etch runs on a printed circuit board.

Several testing strategies are used: add signal lines from the adder to the backplane where there are adequate probe access points; probe directly onto the module etch or IC pins; and subsume the adder in a function whose inputs and outputs can be more easily controlled and observed. The general problem can be modelled by the simple circuits shown in Fig. obs. Each gate could be part of a computer on a chip where there is no possibility for repair. The problems of observation and control exist at all levels of integration. Examples of observation points at each level are given in Table Obslevels for the 11/60 computer.

In Fig. QVstation, the function being tested is a complete CPU which is packaged on several printed circuit modules. Behavior is being induced by PDP-11 programs and observed by inspecting of the results in memory using a diagnostics program or manually. However, a lower level of behavior can be observed (on the special display panel at the right) and controlled (by varying the clock rate of the CPU). The lights on the display panel are driven from backplane signals and show the contents of certain registers,

e.g., micro-instruction register, program status register, and the ALU output register.

Table: Obslevels

Level in computer hierarchy	Observation point	Stage in manufacture of computer	Example
electrical circuit	transistor contacts on metallization layer	semiconductor fabrication	wafer test with micro-probe
switching circuit	leads on I.C. package	incoming inspection of ICs	
register transfer	etch run	module	probe on PC board (module-specific test) GR
register transfer	backplane	module	2224 memory exerciser for cache
PMS (Pc)	Unibus	CPU	
PMS (Pc)	contents of memory	CPU	diagnostic programs at subsystem level, e.g., memory-management unit, or processor instruction set tests
PMS (C)	contents of memory	System integration	peripheral diagnostic programs
PMS (C)	Unibus	System integration	DECX-11 bus exerciser

Figure APT shows the final acceptance of a unit (CPU, memory, power supply, and console) prior to final assembly and test of the system using a system

called Automatic P? Testor. In both Figs. OVstation and APT, the entire instruction set is tested; a central computer system loads the diagnostic programs and monitors their execution. Several hundred lines emanate from this system to all of the computers under test throughout the production line. The Unibus of the processor under test serves as the umbilical cord [Chapter 00, p. 00] to accept the diagnostic programs and to be monitored by making memory observations.

Modules are tested using several methods. One method (see step j of Fig. Process60), uses the GR tester. Here, every signal input and test point on the module is probed using a fixed "bed of nails" test probe. The testor then selects the desired input and test point. In another method, (see step k, Fig. Process60) the module under test is placed in a processor in which all of the other modules are known to be good to be verified. This later test is necessary because the first test usually doesn't run at the operational speed of the computer, nor is such a test guaranteed to cover 100% of the logic.

Following the inter-plant transfer, systems integration begins. Figure FA&T shows the process flow for the 11/60 at the Westminster Plant. For some computer systems, FA&T is carried out by having an individual technician follow a system through each of the three phases: incoming test of the CPU cabinet assembly, integration of peripherals, and final acceptance. Thus one or more have individual responsibility for fabricating a single system in what is a traditional, hand-crafted, job-shop fashion.

Alternatively, FA&T is carried out in what is fundamentally a continuous, production line environment called the Common Systems Integration (CSI) line. Here, workstations exist at each stage of the production line, and are interconnected by conveyors.

The floor plan of the Westminster Common Systems Integration (CSI) area is shown in Fig. CSI. Figure spurs shows four of the workstations in the integration phase. A turntable in the center of the spurs routes the systems to the workstations.

"Statistical Quality Control", Grant and Leavenworth, Fourth Edition, McGraw-Hill, 19xx.

"Quality Control Handbook", Juran, J. M., McGraw-Hill 1951; Second Edition, 1962.

Braeckelmann, W., Fritzsche, H., Kroos, F-K., Trinke, W., and Wilhelm, W.

"A Masterslice LSI for Subnanosecond Random Logic",
Digest of Technical Papers, 1977 Ent. Solid-State
Circuits Conference, pp. 108-109.

Fusfeld, A. R.

The Technological Progress Function
Technology Review, Feb. 1973, pp. 29-38

Luecke, J.

Overview of Semiconductor Technology Trends;
IEEE Computer Society International Conference,
Fall, 1976, pp. 52-55.

Masaki, A., Harada, Y., Chiba, T.

"200 Gate ECL Masterslice LSI",
ISSCC Digest of Technical Papers, pp. 62-63, Feb.
1974.

vonHippel, E.

The Dominant Role of the User in Semiconductor and
Electronic Subassembly Process Innovation;
IEEE Transactions on Engineering Management, vol.
EM-24, No. 2, May 1977. pp. 60-71.

Chapter 2 Figures and Tables (that are to be made as figures)

- Memtree — Family tree of memory/logic technology (courtesy of Memorex Corporation, Puthuff, 1977)
- LMtimeline — Logic and memory technology evolution time line
- ICtree — Family tree of digital Integrated Circuit functions
- ALU — Arithmetic-Logic Unit logic diagram
- AMD — AMD 2900 4-bit microprocessor slice block diagram (registers and data path)
- Semiden — Component (transistor) per single Integrated Circuit die versus time
- Compprice — Price per gate versus time
- Gatecomp — Circuit components (transistors) required for various logic functions in various technologies -- show also rom, ram, pla, ccd, flip flops, gates 2900, LSI-11, CMOS-8; includes N, P, IIL, CMOS, ECL< TTL, TTL/S
- Memprice — Cost per bit of Integrated Circuit memory versus time (Noyce-Sci.-Am.)

Generality	Generality (use) and price of an Integrated Circuit versus its complexity
Reliab	Failure rate of silicon Integrated Circuits
Design.cost	Current design cost (or time) versus circuit density using various design methods
Semiuse	Manufacturing cost versus LSI circuit density for various design techniques
PLA	Signetics field programmable logic array (FPLA)
Memsizepert	Memory size versus access time for various memories and yearly availability (Toome's data, TI)
Core.cost	Cost per bit of core memory estimated by various market surveys and future predications
Areal.digital	Areal density of various digital magnetic recording media (courtesy of Memorex Corp., 1978)
Areal.analog	Areal density of various analog magnetic recording media (courtesy of Memorex Corp., 1978)
Large.disk.price	Price per bit of large, moving head disks and semiconductor

	memories (courtesy of Memorex Corp., 1978)
Mem.trends	Memory trend 1975-1980 (courtesy of Memorex Corp., 1978)
Tapechar	Characteristics of various IBM magnetic tape units versus time
Tapeprice	Price of magnetic tape units (1978)
CRTprice	Price of various cathode ray tube terminals versus time
CPI	Consumer Price Index using 1967 as base
CPI.log	Consumer Price Index (plotted on semi-logarithm graph) using 1967 as base
Mktcycle	
Perf.size	Performance or size vs price at some time, t
Cost vs time	Cost of
Techvstime	
Cost/tech	

Tradeoff	System performance constrained by either memory size or performance
Product.cost	Product cost versus time on the state-of-the-art line showing various situations
Net	Cost of various components that contribute to product cost
Product.price	Product cost versus time with manufacturing learning
Product.improve	Product cost improvement by enhancement of cost/function

Note there is a reference to photos elsewhere in the book to illustrate the various packaging concepts.

Fig. 1. Packaging consists of placing boxes within smaller boxes...within boxes on an indefinite (not recursive) basis.

Fig. 2. A packaged system provides some function, but in addition gives a visual impression, usually requires cooling of some sort, has certain mechanical characteristics and an Electromagnetic Interface (EMI).

Fig. 3. Eight level packaging hierarchy for DEC second to fourth generation computer systems.

Fig. 4. Packaging is a set of closely interrelated design problems.

Fig. 5. A DIP with two semiconductors per chip.

Packaging Photos (need somewhere)

Fig. 6. Time line evolution of computer generations, packaging and classes (with examples).

Fig. 7. None

Fig. 8. —
(NOT IN) PDP-1 control (called Line Units) for interfacing eight teletype lines using what is now called the Universal Asynchronous Receiver-Transmitter, (UART).

Fig. 9. —
(NOT IN) 4706 and 4707 receiver and transmitter line units of the late second generation.

Fig. 10. —
(NOT IN) DL11 line unit based on early fourth generation IC.

Fig. 11. _(NOT IN) UART chip and block diagram.

Fig. 12. DEC physical structure (packaging) hierarchies vs technology generations.

Fig. 13. Three views of PDP-8/A box.

Figure

- Lifecycle A simplified process flow for the major phases and milestones in the life of a product.
- *Breadboard
- Manufover Overview of manufacturing flow.
- process60 The process flow for the 11/60 manufacturing plant in Aguadilla, P.R.
- Bath Chip failure rate plotted against time shows the three behavior regions. Thermal cycling is used to ensure that "infant mortality" cases in the first region are discovered early in the manufacturing flow.
- *Tchamber Thermal chambers. (photo)
- *2224 2224 memory exerciser used for testing cache. (photo)
- *QVstation (photo)
- *GRtester (photo)
- Obs Placement of test points for observation and control.
- *APT Final acceptance prior to inter-plant transfer. (photo)
- CSI The common systems integration area at the Westminster Plant. (photo) -- From article in Materials Handling Engineering, July 1977
- *spurs (photo) [need new photo giving better view of system being built]
- *Waferprobe Probing test points on a wafer prior to dicing. (photo)
- *These are photos -- copies of Polaroid photos are attached for allbut Breadboard and Waferprobe. CM is getting B/W glossies retaken in Aguadilla and Breadboard from the mill.