6 GENERAL MEMORY TECHNOLOGY TRENDS

OVERVIEW

Memory and data storage take many forms, even in a single computer system. Figure 6-1 classifies the various memory categories that are discussed in the following pages. There are two main classifications of memory families. These are RAM (Random Access Memory) and ROM (Read Only Memory) devices. RAM devices are volatile, which is to say they lose their memory content when the power to the host system is turned off. ROM devices are non-volatile, meaning they retain their stored data when the power is removed.

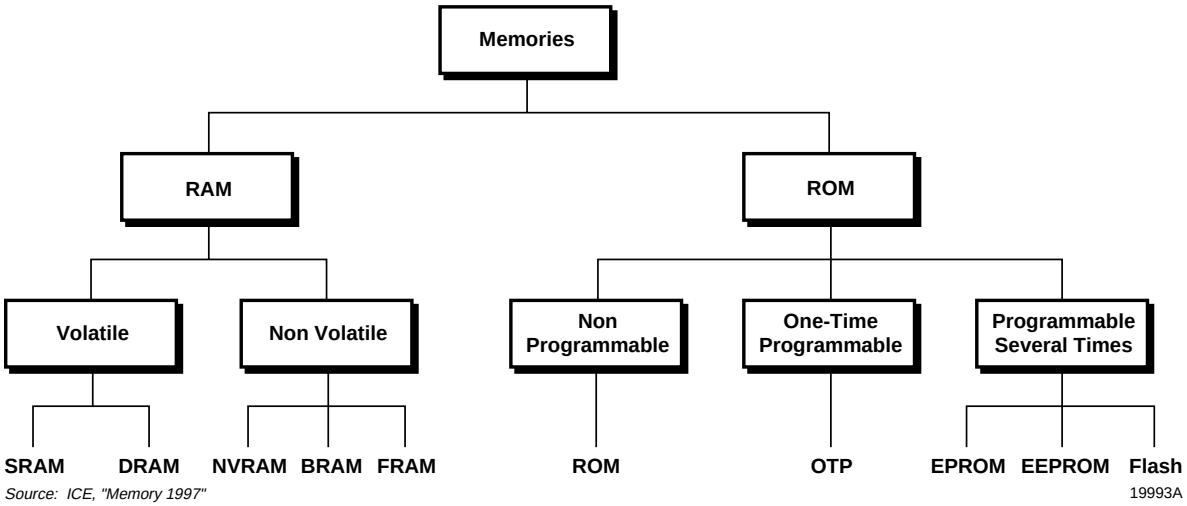

Figure 6-1. Memory Classification

RAM products are read and written to with the same electrical characteristics. They are easy to write but are volatile, meaning that when the power is turned off, the devices lose the memory content.

ROM-based devices—ROM, EPROM, EEPROM, and flash memory devices are easy to read but are more difficult to program (write) than RAM devices. EEPROM and flash devices are more difficult to program than they are to read. EPROMs need a mechanical step (UV-light) to erase the memory cells prior to re-programming the device. One-time programmable (OTP) EPROMs can

be written only one time by the user. ROM products are programmed during process manufacturing. Figure 6-2 shows the different types of memories and some of the main characteristics of the devices.

Characteristic	DRAM	SRAM	EPROM	ROM	Parallel EEPROM	Flash	NVRAM	FRAM
Cell Organization	1T + 1C	1 Flip-Flop + 2T	1T with Floating Gate	1T	1T + 1T with Floating Gate	1T with Floating Gate	1 SRAM Cell + 1 EEPROM Cell	1T + 1C
Storage Method	Charge on Capacitor	Flip-flop circuit	Charge on Floating Gate	Masked in Production	Charge on Floating Gate	Charge on Floating Gate	SRAM + Back-Up in EEPROM	Charge on Capacitor
Number of Devices in Cell	1.5	4-6	1.0	1.0	2.0	1.0	8-9	2.0
Relative Cell Size	1.5	4-6	1.5	1.0	3-4	1.5	9-10	2.0
Density	64Mbit	4Mbit	16Mbit	64Mbit	4Mbit	64Mbit	256K	256K
Overhead Cost	Refresh Logic	No	UV Erase Programmer	Mask Charges	No	No	No	No
Volatile (Power Off)	Yes	Yes	No	No	No	No	No	No
Data Retention (D.C. Power On)	4ms	∞	10 Years	∞	10 Years	10 Years	10 years	10 Years
In System Reprogrammable	Yes	Yes	No	No	Yes	Yes	Yes	Yes
Number of Reprogram Times (Endurance)	∞	∞	100	—	1,000,000	10,000	1,000,000	10^{12}
Typical Write (Reprogram) Speed	100ns	25ns	30min	—	2.5s	2.5s	—	235ns
Typical Read Speed (ns)	100	25	100	100	200	200	200	150

Source: ICE, "Memory 1997" 22610

Figure 6-2. Characteristics of MOS Memory Product Types

GENERAL TECHNOLOGY ISSUES

Memory devices have historically been considered “process drivers” as well as revenue producers. A process driver is a product that is manufactured in large wafer volumes, that pushes the state-of-the-art in processing, and has a die yield that can be used to measure the effectiveness of the process.

Memories are no longer considered the only "process drivers" in the industry. That distinction is shared among memory, MPU, and ASIC devices. Intel, in the mid-1980s, made a strategic decision to abandon the memory market for economic reasons. Intel probably had a good notion that it could apply state-of-the-art processing techniques to microprocessors, thus making these devices process drivers as well.

Feature Size

The most critical issue in the advancement of IC technology is feature size; more specifically, the progression to reduced feature sizes every few years. Today's processes require considerable process adjustment with each feature size reduction. In addition, each new generation of memory usually requires one or two additional mask layers. Figure 6-3 shows the past and projected feature sizes for DRAMs and several other experimental devices. Design rules used to build these devices has decreased from about 3.0µm in 1980 to about 0.25µm in 1997. This represents about a 14 percent decrease every year. This trend is expected to continue and feature sizes are forecast to be about 0.15µm by 2000.

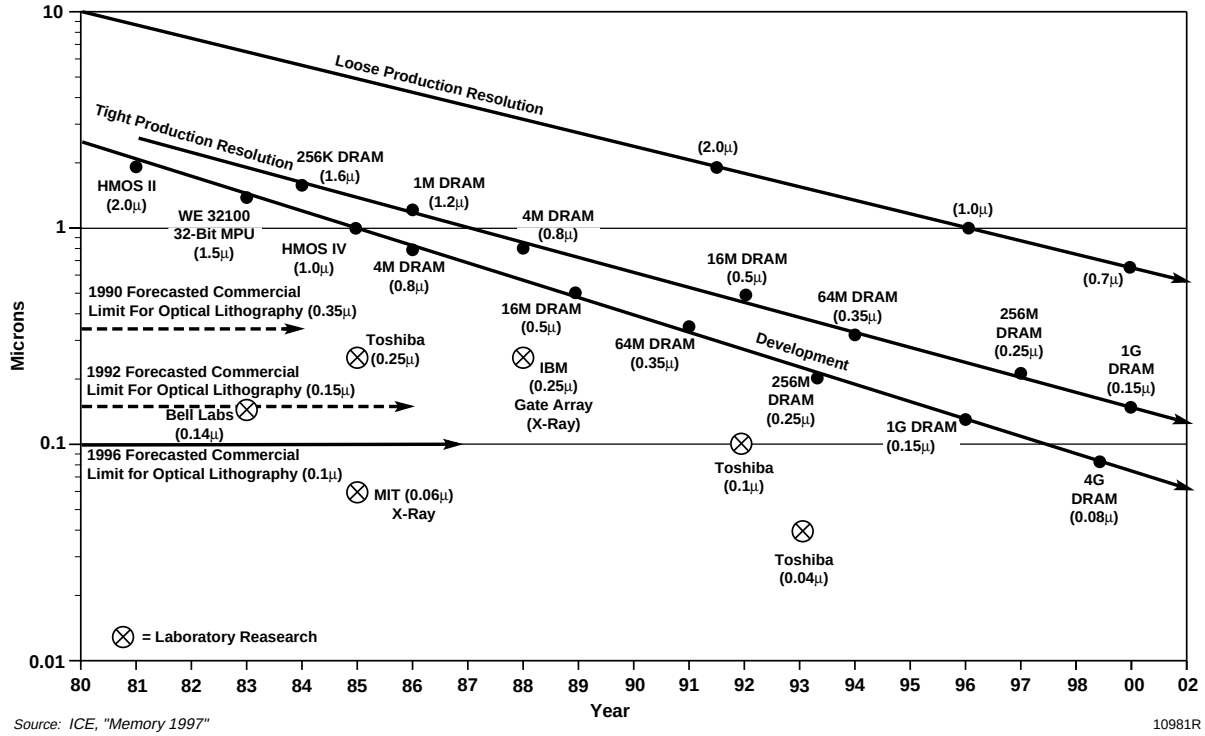

Figure 6-3. IC Feature Size Trends

Figure 6-4 shows SIA's technology roadmap through the year 2010. As shown, the integration levels are forecast to continue along historical trend lines.

Year of first DRAM shipment	1995	1998	2001	2004	2007	2010	Driver
Minimum feature (μm)	0.35	0.25	0.18	0.13	0.10	0.07	
Memory Density Bits/chip (DRAM/flash)	64M	256M	1G	4G	16G	64G	D
Logic Density (High volume: Microprocessor) Logic transistors/cm^2 (packed) Bits/cm^2 (cache SRAM)	 4M 2M	 7M 6M	 13M 20M	 25M 50M	 50M 100M	 90M 300M	L(μP)
Logic Density (Low volume: ASIC) Transistors/cm^2 (auto layout)	2M	4M	7M	12M	25M	40M	L(A)
Number of Chip I/Os Chip to package (pads) high performance	900	1,350	2,000	2,600	3,600	4,800	L,A
Chip frequency (MHz) On-chip clock, cost-performance On-chip clock, high-performance Chip-to-board speed, high performance	 150 300 150	 200 450 200	 300 600 250	 400 800 300	 500 1,000 375	 625 1,100 475	 μP L
Chip size (mm^2) DRAM Microprocessor ASIC	 190 250 450	 280 300 660	 420 360 750	 640 430 900	 960 520 1,100	 1,400 620 1,400	
Oxide Thickness (nm)	7-12	4-6	4-5	4-5	<4	<4	μP
Junction Depth (μm)	0.1-0.2	0.1-0.15	0.07-0.13	0.05-0.1	<0.07	<0.05	μP
Maximum number wiring levels (logic) On-chip	4–5	5	5–6	6	6–7	7–8	μP
Minimum mask count	18	20	20	22	22	24	L
Power supply voltage (V) Desktop Battery	 3.3 2.5	 2.5 1.8–2.5	 1.8 0.9–1.8	 1.5 0.9	 1.2 0.9	 0.9 0.9	 μP A
Maximum power High performance with heatsink (W) Logic without heatsink (W/cm^2) Battery	 80 5 2.5	 100 7 2.5	 120 10 3.0	 140 10 3.5	 160 10 4.0	 180 10 4.5	 μP A L

A=ASIC D=DRAM
L=Logic μP=Microprocessor

Source: SIA/ICE, "Memory 1997"

20286C

Figure 6-4. The 15-Year SIA Roadmap

Lifecycle

As with any commercial product, memories have a lifecycle. A memory will pass through the different steps of the lifecycle, from its introduction where the IC manufacturer concentrates on capital resources and R&D efforts, to the end of its life when the product becomes obsolete.

Die Size Trends

In 1975, Intel Chairman Gordon Moore predicted that engineers could shrink semiconductor device dimensions by approximately 10 percent each year, creating a new generation of chips every three years with four times as many transistors. Twenty one years later, Moore's prediction was impressively accurate. DRAM devices actually exceeded his expectations (Figure 6-5).

Figure 6-6 shows how the die area of leading-edge memory devices has increased about 13 percent per year. The trend toward larger die sizes is forecast to continue. The die sizes of the 1Gbit DRAMs described at the 1995 and 1996 ISSCC conferences ranged from 901K sq. mils to 1,451K sq. mils (Figure 6-7). As shown, the NEC 1Gbit DRAM, if square, would be about 1.2 inches on a side.

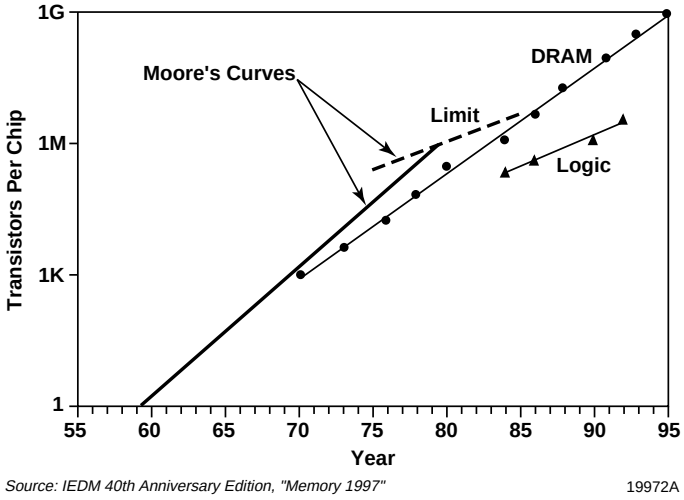

Figure 6-5. Growth in Chip Complexity Since 1959

Wafer Size

The IC industry is quickly moving to the new 300mm (12 inch) wafer standard. Figure 6-8 shows the wafer area increase that occurs each time the industry moves to the next wafer size. Currently the standard for high-volume advanced IC production is 200mm (8 inch). The area gained by moving from 200mm to 300mm wafers will be 125 percent.

Advancing from one wafer size to a new, larger size takes several years of development. In fact, development time increased substantially to transition to 200mm wafers and will be the same for the transition to 300mm wafers (Figure 6-9). Companies such as Samsung, Texas Instruments, and many other leading memory suppliers have indicated their willingness to build manufacturing facilities to support 300mm wafers. The big question is which company will be the one to take on the huge headaches and huge amount of capital needed to work out all the wrinkles associated with the transition to 300mm wafers.

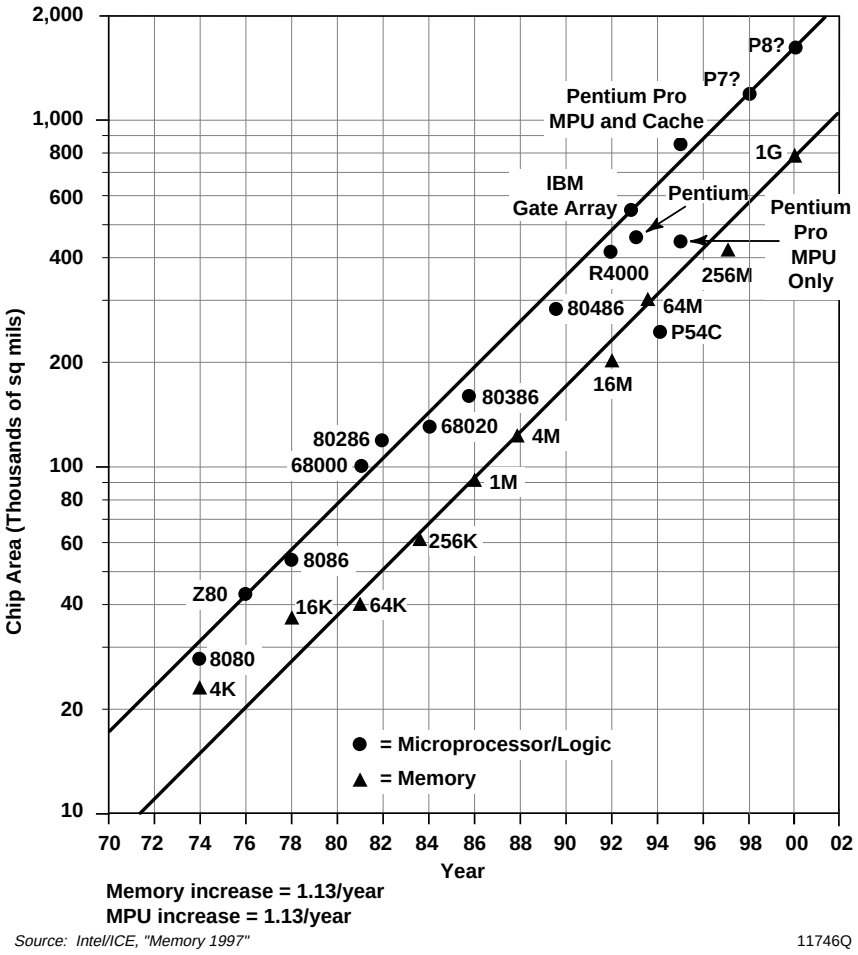

Figure 6-6. IC Die Size Trends

Company	Density	Feature Size (μm)	Cell Size (μm²)	Chip Size (K sq. mils)	Access Time (ns)	Organization	Conference
Hyundai	256M	0.3	—	561	36	32M x 8	ISSCC '95
Matsushita	256M	0.25	0.72	638	—	16M x 16	ISSCC '94
Mitsubishi[1]	256M	0.25	0.72	472	34	32M x 8	ISSCC '94
Oki[2]	256M	0.25	0.72	530	—	32M x 8	ISSCC '94
Mitsubishi[3]	1G	0.14	0.29	901	32	—	ISSCC '96
Samsung[4]	1G	0.16	—	1,010	—	—	ISSCC '96
Hitachi	1G	0.16	0.29	1,108	33	64M x 16	ISSCC '95
NEC	1G	0.25	0.54	1,451	—	—	ISSCC '95

[1] Produced using KrF excimer-laser lithography.
[2] Packaged in a 64-pin 600-mil TSOP, produced using e-beam lithography.
[3] SDRAM produced using synchrotron-generated x-ray lithography.
[4] SDRAM produced using KrF excimer-laser lithography.

Source: ICE, "Memory 1997" 20289A

Figure 6-7. ISSCC Advanced DRAMs

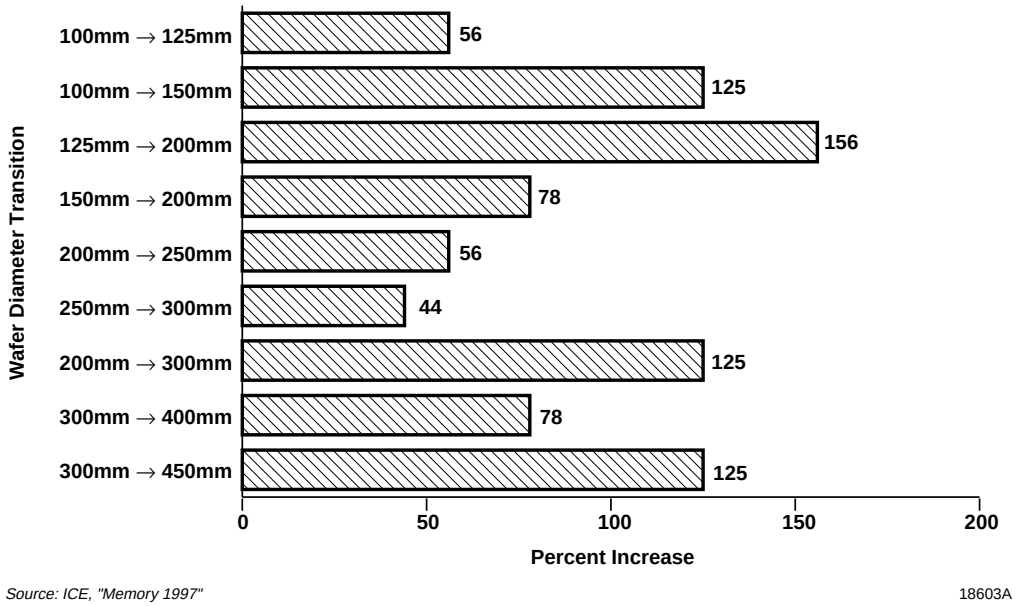

Figure 6-8. Wafer Area Increases (Percent)

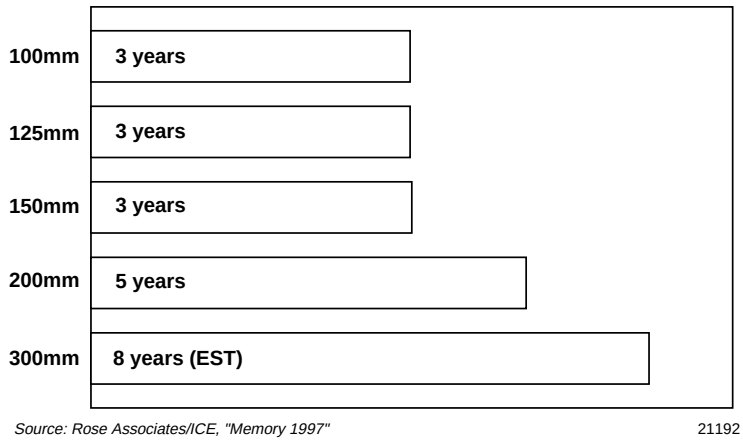

Figure 6-9. Wafer Development Time Requirements (Time to Reach 100 MSI Production Rate)

Defect Density

The cost to manufacture ICs is governed by many factors. One of the most important is the number of good dice per wafer started into the wafer fab. This number is dependent on the number of potential dice per wafer, and the number of defective dice.

Figure 6-10 shows the potential number of dice for various wafer sizes. The number of potentially good dice per wafer is dependent on the wafer size and the die size. The number of defective dice on a wafer are the result of random killer defects, usually expressed in defects per square centimeter, and the number of defective dice caused by parametric defects. In a world class fab, the number of dice lost because of parametric defects is very low.

As shown in Figure 6-11, random killer defects have a dramatic effect on the number of good dice per wafer, especially as the die size increases. Figure 6-12 shows the yield effect of various defect densities versus die size, where the yield is defined as the number of good dice divided by the number of potential dice.

Defect density control is extremely important for memory fabs since the cell array area in a memory IC contains extremely compact circuitry. Each generation (256Kbit, 1Mbit, 4Mbit, etc.) of memory chip is about 50 percent larger than the previous with four times the number of bits of storage.

As feature sizes become smaller, ICs are susceptible to smaller and smaller particles causing random killer defects. This means that if everything else remains constant, the killer defect density will increase as smaller particulates become killer defects. Yield also decreases with increasing die sizes. Therefore, to maintain acceptable yields, extreme care must be taken to reduce particulates in the ambient air, the equipment, the process gases, and the process liquids.

Fab processes themselves generate particulates that can cause defects. In modern processes, each process step is given a defect "budget" for the number of defects per square centimeter added for the step. As the number of process steps increases, the job of reducing the defect density becomes more and more difficult.

Redundancy

As mentioned earlier, the effect of defect density on memory chips is a much larger problem than with most other products due to the circuit density in the storage cell array. Nearly any particulate-caused defect in this area is a killer defect. To enhance memory chip yields, manufacturers use spare rows and columns (redundant rows and columns) that can replace defective rows or columns. During 100 percent wafer probe, defective rows and columns are "replaced" by these spares through the use of laser blown fuses that alter the decode mechanism. When external signals try to access a defective row or column, the decode circuitry selects a spare one instead. Figure 6-13 shows a simplified logic of redundancy programming. A normal decoder contains half as many decoding transistors as a redundant decoder. If redundancy is not required—that is, the chip is perfect—spare decoders will be deselected regardless of the input address.

DIE SIZE** √AREA		DIE AREA		CANDIDATE NUMBER OF WHOLE DICE					
mils	mm	mil^2 x1000	mm^2	3-INCH	100mm	125mm	150mm	200mm	300mm
100	2.5	10.0	6.5	508	936	1,528	2,276	4,168	9,700
110	2.8	12.1	7.8	432	780	1,272	1,892	3,444	8,040
120	3.0	14.4	9.3	356	656	1,060	1,576	2,892	6,772
130	3.3	16.9	10.9	300	548	904	1,348	2,480	5,784
140	3.6	19.6	12.6	256	468	780	1,152	2,136	4,992
150	3.8	22.5	14.5	216	416	688	1,012	1,844	4,344
160	4.1	25.6	16.5	188	360	600	880	1,640	3,828
170	4.3	28.9	18.6	164	316	524	780	1,444	3,388
180	4.6	32.4	20.9	148	284	460	688	1,280	3,024
190	4.8	36.1	23.3	132	256	416	616	1,148	2,692
200	5.1	40.0	25.8	120	224	376	556	1,036	2,448
210	5.3	44.1	28.5	112	208	332	500	936	2,212
220	5.6	48.4	31.2	96	188	308	456	864	2,000
230	5.8	52.9	34.1	88	164	284	424	780	1,836
240	6.1	57.6	37.2	80	148	256	384	716	1,688
250	6.3	62.5	40.3	76	148	240	356	656	1,560
260	6.6	67.6	43.6	68	120	208	324	608	1,428
270	6.9	72.9	47.0	60	120	200	300	556	1,328
280	7.1	78.4	50.6	52	112	188	276	524	1,232
290	7.4	84.1	54.3	52	96	164	256	476	1,148
300	7.6	90.0	58.1	52	96	156	240	448	1,060
310	7.9	96.1	62.0	44	88	148	224	424	1,012
320	8.1	102.4	66.1	44	80	140	208	392	936
330	8.4	108.9	70.3	40	76	120	192	376	880
340	8.6	115.6	74.6	32	76	120	188	340	820
350	8.9	122.5	79.0	32	68	112	164	332	780
360	9.1	129.6	83.6	32	60	112	156	308	732
370	9.4	136.9	88.3	32	52	96	148	292	688
380	9.7	144.4	93.2	32	52	96	148	268	656
390	9.9	152.1	98.1	24	52	88	140	256	616
400	10.2	160.0	103.2	24	52	80	120	248	600
410	10.4	168.1	108.5	24	52	80	120	240	556
420	10.7	176.4	113.8	24	44	76	112	216	540
430	10.9	184.9	119.3	24	44	76	112	208	508
440	11.2	193.6	124.9	16	40	68	104	208	492
450	11.4	202.5	130.6	16	32	68	96	188	460
460	11.7	211.6	136.5	16	32	60	96	188	440
470	11.9	220.9	142.5	16	32	52	88	180	432
480	12.2	230.4	148.6	16	32	52	88	164	392
490	12.4	240.1	154.9	12	32	52	80	156	392
500	12.7	250.0	161.3	12	32	52	76	148	376
510	13.0	260.1	167.8	12	32	52	76	148	356
520	13.2	270.4	174.5	12	24	52	76	148	340
530	13.5	280.9	181.2	12	24	44	68	132	332
540	13.7	291.6	188.1	12	24	44	68	120	316
550	14.0	302.5	195.2	12	24	40	68	120	308
560	14.2	313.6	202.3	12	24	40	60	120	292
570	14.5	324.9	209.6	12	24	32	52	112	284
580	14.7	336.4	217.0	12	24	32	52	112	268
590	15.0	348.1	224.6	12	16	32	52	112	256
600	15.2	360.0	232.3	12	16	32	52	104	256
610	15.5	372.1	240.1	12	16	32	52	96	248
620	15.7	384.4	248.0	4	16	32	52	96	240
630	16.0	396.9	256.1	4	16	32	52	88	224
640	16.3	409.6	264.3	4	16	32	44	88	216
650	16.5	422.5	272.6	4	16	24	44	88	208
660	16.8	435.6	281.0	4	12	24	44	80	208
670	17.0	448.9	289.6	4	12	24	40	80	200
680	17.3	462.4	298.3	4	12	24	40	76	188
690	17.5	476.1	307.2	4	12	24	32	76	188
700	17.8	490.0	316.1	4	12	24	32	76	180
710	18.0	504.1	325.2	4	12	24	32	68	180
720	18.3	518.4	334.5	4	12	24	32	68	164
730	18.5	532.9	343.8	4	12	24	32	68	164
740	18.8	547.6	353.3	4	12	16	32	68	156
750	19.0	562.5	362.9	4	12	16	32	60	156
760	19.3	577.6	372.6	4	12	16	32	60	148
770	19.6	592.9	382.5	4	12	16	32	52	148
780	19.8	608.4	392.5	4	12	16	32	52	148
790	20.1	624.1	402.6	4	12	16	24	52	148
800	20.3	640.0	412.9	4	12	16	24	52	132

*** Corner of die at center of full radius wafer.**
3mm band around edge of wafer not used.
**** Size after die separation, 3 mil saw kerf.**

Source: ICE, "Memory 1997"

8317D

Figure 6-10. Die Size Versus Die Count*

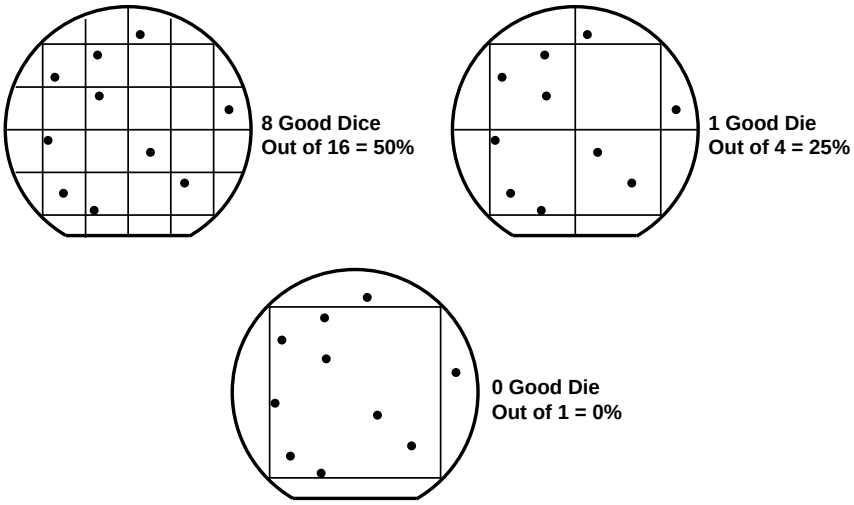

Figure 6-11. Effect of Die Size on Yield

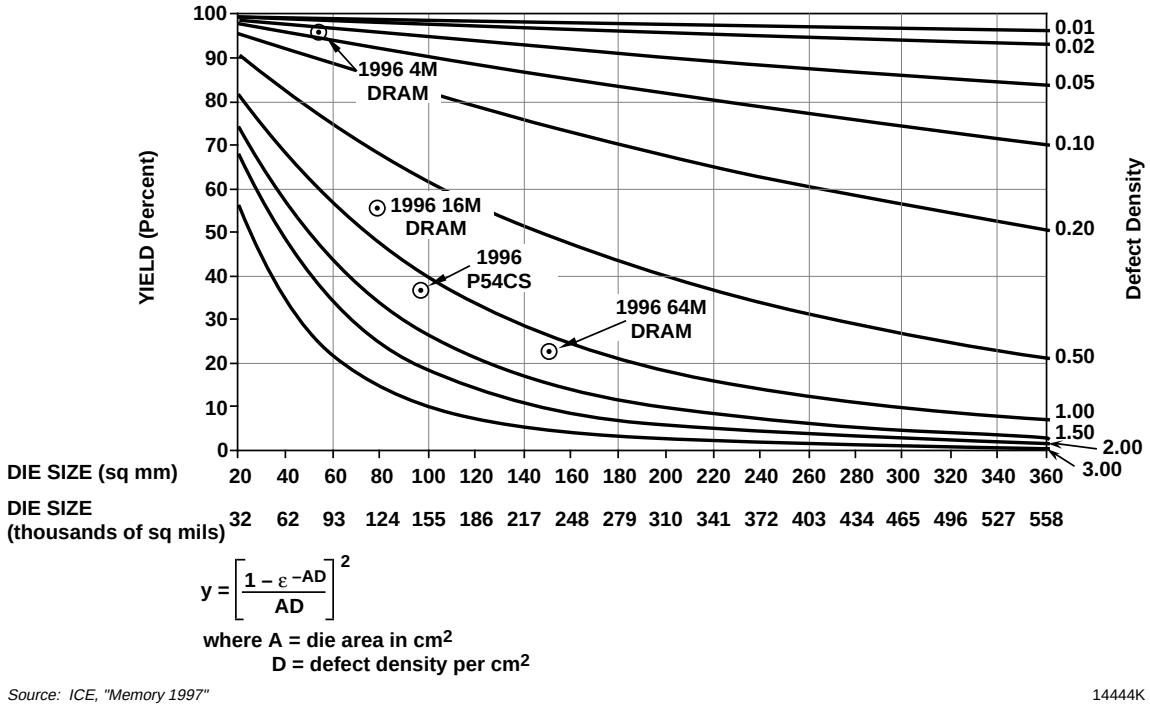

Figure 6-12. Murphy's Probe Yield Model (As a Function of Defects per sq cm)

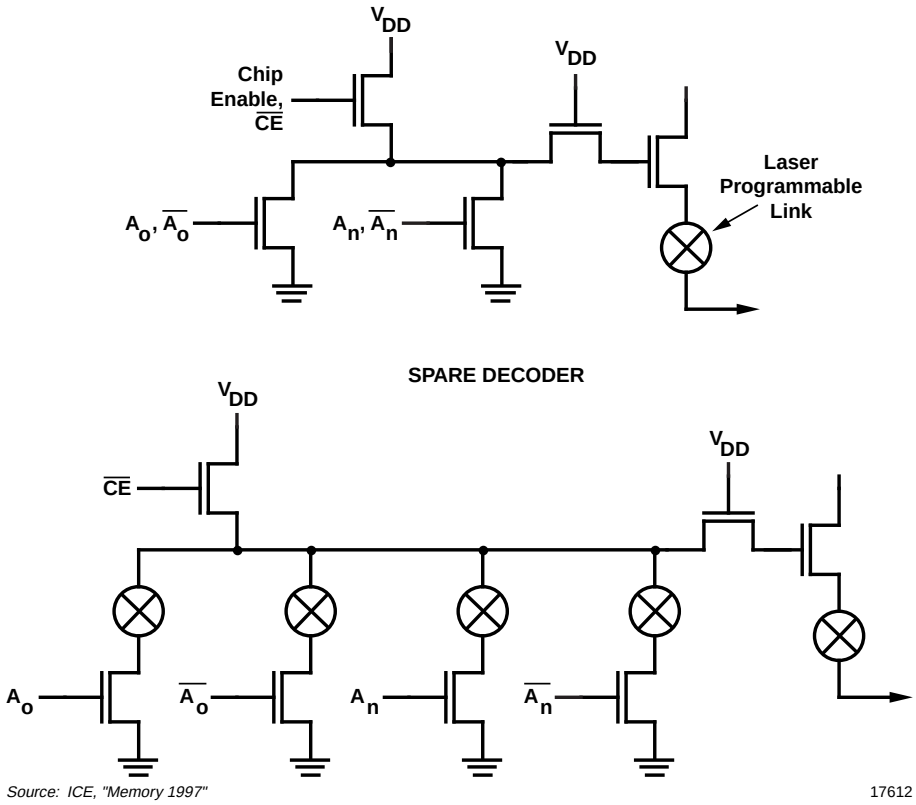

Figure 6-13. Row Decode for Redundancy

Process Complexity

Process complexity is usually measured by the number of critical mask layers, plus any special processes that may have a detrimental effect on yield. As feature sizes have become smaller, structures under the surface of the silicon wafer have become more shallow. The structures above the wafer surface have remained relatively thick. The thicker interconnect layers are the result of the need to separate the conductive layers as much as possible, thus reducing the capacitive loading of the long thin conductors (polysilicon and aluminum) running the length and width of the chip. If the deposited dielectric films are reduced in thickness, the parasitic capacitive load will increase, and the device will be considerably slower.

Several manufacturers have lowered the effect of parasitic capacitance through the use of materials with lower dielectric constants. Polyimide is one such material. The dielectric constant of polyimide is below 3.0, versus 3.9 for deposited silicon dioxide. The improvement in reduced parasitic capacitance is a linear function with respect to dielectric constant.

Although the number varies somewhat from vendor to vendor, over twenty critical mask steps are required to produce today's most advanced memory chips. Each new generation requires one or two additional mask layers. Each mask layer is composed of ten to twenty actual process operations, some major and some minor.

With today's processes composed of hundreds of steps, the requirement for tight process control on each step is critical. When processes are planned, budgets are established for process variations at each process step. These budgets are modeled and used to predict manufacturability. Uniformity across the die, across the wafer, from wafer to wafer, and from wafer lot to wafer lot can be predicted.

The size of memory chips is obviously very dependent on the size of a single memory cell. Reducing the feature size reduces the cell (storage area for one bit) somewhat, but there are portions of the cell that do not reduce with feature size. To continue the reduction of cell area, designs are now utilizing three-dimensional structures. These will be described later in the various product descriptions, but in all cases, the addition of these structures leads to additional process steps.

POWER SUPPLY

Two factors are driving the reduction in memory operating voltages: power dissipation and process geometries. The market for low-voltage devices is growing fast, and IC manufacturers are concentrating significant R&D efforts to develop low-voltage versions of their memory ICs.

Power Dissipation

The demand for low power dissipation comes from two main factors. First, system and package power constraints demand lower operating voltage due to higher board integration and the use of more integrated chips.

Systems with dual power supplies (5V and 3.3V) consume system space and require careful interface design between the two supplies. For this reason, system manufacturers push ICs suppliers to produce low voltage devices so that they can design single low-voltage supply systems. Figure 6-14 shows the market transition from 5V to 3V systems.

The second reason for the rise in low-power ICs comes from growth in the portable electronics market including portable computers and digital cellular phones. This equipment requires low power and low voltage to increase battery life. Dissipation drops 36 percent when going from 5.5V to 3.3V, for example. Using 2.7V in a typical digital cellular phone design could extend talk time by 25 percent over 3.3V operation. Figure 6-15 shows 5V to 3.3V power savings in a typical notebook computer.

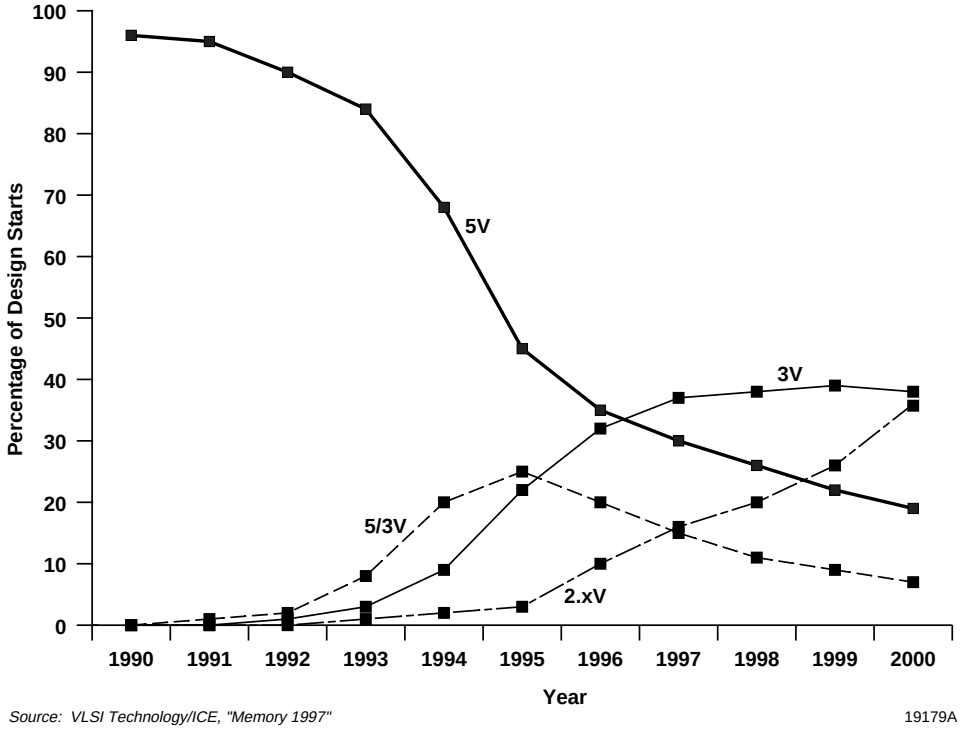

Source: VLSI Technology/ICE, "Memory 1997" 19179A

Figure 6-14. Transition from 5V to 3V Systems

Notebook Computer System Components	Power (W) at 5V		Typical 5V Power (W)	Power at 3.3V Logic (W)	Expected 3.3V Power (W)
	Logic	Other Functions			
CPU Plus Core Logic	2.20	0.00	2.20	0.95	0.95
System Memory	0.50	0.00	0.50	0.22	0.22
Display Controller Subsystem	1.50	0.00	1.50	0.65	0.65
LCD Panel Plus Backlight	0.30	3.20	3.50	0.13	3.33
Hard-Disk Drive*	0.20	0.30	0.50	0.09	0.39
Miscellaneous Circuits	0.50	0.00	0.50	0.21	0.21
DC/DC Conversion	0.00	1.70	1.70	0.00	1.20
Total System Power			10.40		6.95

***Hard-disk drive power estimates reflect a mix of active and idle time.**

Source: Cirrus Logic/Electronic Products/ICE, "Memory 1997" 19217

Figure 6-15. 3.3V Logic Power Savings in a Typical Notebook Computer

An extra incentive for designing 3.3V ICs is the U.S. government's EPA Energy Star Program. According to the EPA, computer systems account for five percent of commercial electricity consumption in the United States. Energy Star mandated energy reduction in any PC the federal government purchased.

Process Geometries

For memory devices to become increasingly dense, feature sizes must shrink. Both gate oxides and gate lengths need lower voltage to avoid an increase in electrical field (Figures 6-16 and 6-17). For a similar process on a given generation, 3.3V parts have longer access times than 5V parts. Figure 6-18 gives the supply voltage decrease versus the technology and the memory density.

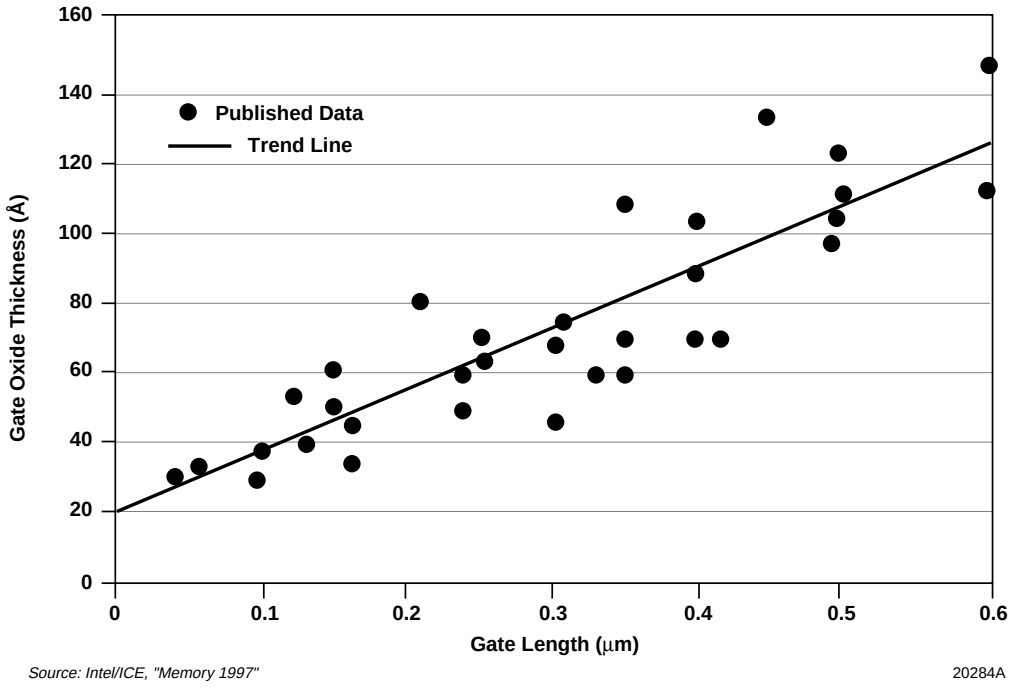

Figure 6-16. Gate Oxide Versus Gate Length

Unfortunately, it is more difficult to yield high-speed devices for low voltage ICs than it is for high voltage parts. Figure 6-19 shows a typical schmoo plot of the access time versus power supply.

DRAMs

During 1994, there was a significant ramping-up of low-voltage DRAMs (Figure 6-20). ICE forecasts that in the year 2002, more than 95 percent of all DRAMs sold will be classified as low-voltage (3.3V or lower).

Some PC manufacturers have delayed shifting to pure 3.3V motherboard designs because price premiums for 3.3V DRAMs would negatively impact the cost of their system. However, during 1996, all major DRAM suppliers were making the transition to 3.3V at the 16Mbit DRAM level.

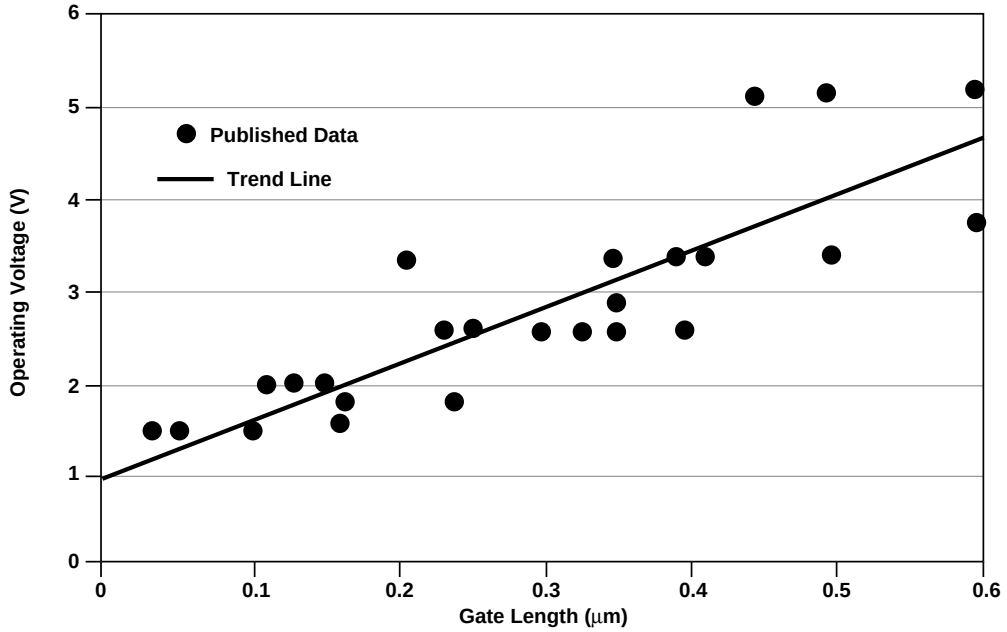

Source: Intel/ICE, "Memory 1997" 20285A

Figure 6-17. Gate Length Versus Operating Voltage

Price parity between 5V and 3.3V devices should quickly ensue. 64Mbit DRAMs are proposed only at the 3.3V level. The upcoming 256Mbit DRAM will operate internally at around 2.5V but will interface to 3.3V.

Certain drawbacks exist for DRAMs operating at lower internal voltages. For instance, it is more difficult to read the charge on a cell capacitor when a lower voltage is stored on the capacitor. The cell is also more susceptible to errors due to ground bounce (electrical noise) caused by the simultaneous reading of several hundred cells at one time.

SRAMs

Advanced SRAMs are shifting to 3.3V technology even though there is a speed penalty versus a comparable 5V part. In terms of data sensitivity (i.e., noise immunity, alpha particles), the 6T CMOS cell is the best design. The benefits of a TFT (Thin Film Transistor) cell are described in Section 8. One additional benefit of the TFT is reduced power consumption. However, due to process complexity, most SRAM manufacturers stay away from this design.

ROMs

ROM technology is not necessarily state-of-the-art, but these products are moving to lower voltage operation like the other memory technologies. Several manufacturers have proposed low-voltage ROMs. Sharp offers 16Mbit and 32Mbit ROMs with power supplies ranging from 2.7V to 3.6V. Other ROM makers such as AMI, Hitachi, and Ricoh also offer low-voltage ROMs.

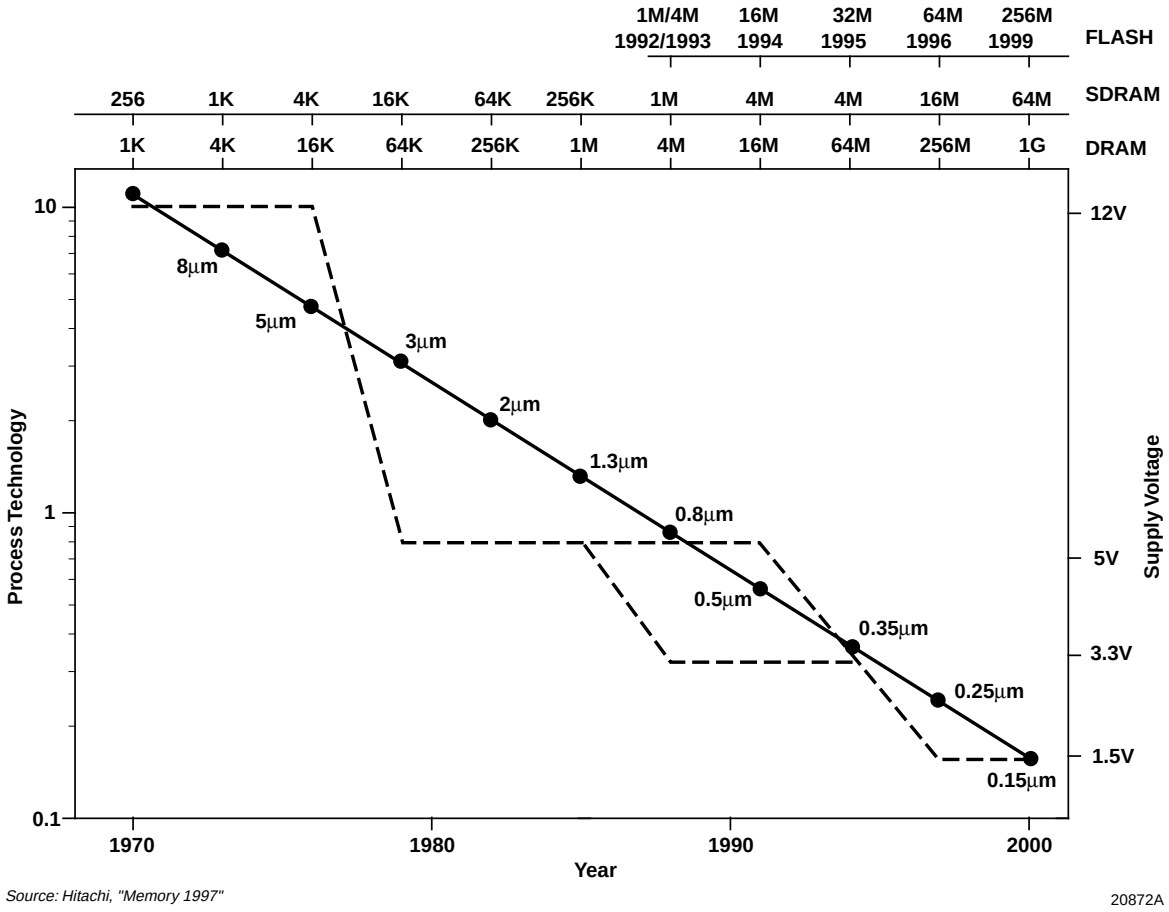

Figure 6-18. Technology Roadmap/Wafer Process

EPROMs

The development of high-density EPROMs has slowed due to the evolution of flash memories. However, like other technologies, low-voltage parts are available. The low-voltage supply is only used for read operations. High voltages are required for the write operations (Figure 6-21).

EEPROMs

EEPROMs must internally generate high voltage for write/erase operations. The V_{CC} power supply affects the internal high voltage supply (V_{PP}). Low-voltage parts are much more complex. There is, however, a need for low-power parts in applications such as digital cellular phones. 2.7V to 3.6V power supply EEPROM devices are available from several vendors.

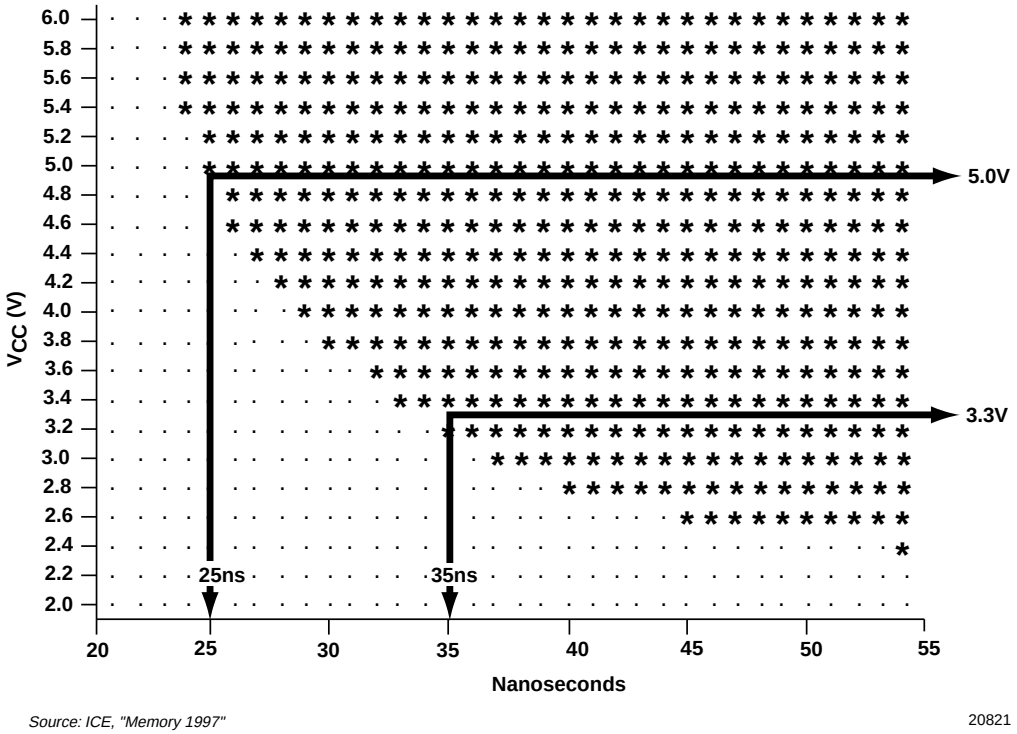

Figure 6-19. Typical Access Time Versus Power Supply

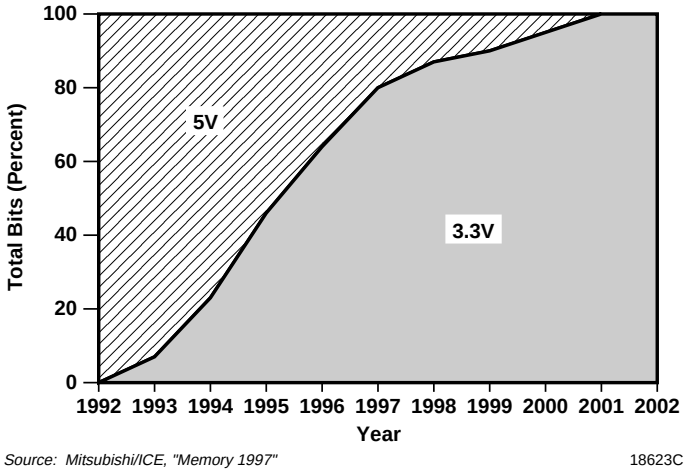

Figure 6-20. Trend of Low-Voltage DRAM

Flash Memories

Flash memories require high voltage for write and erase cycles. The first step to low voltage is to move parts from two power supplies (12V/5V) to one power supply (5V). In this case, the high voltage is internally generated.

V_{CC} (Read Mode)	3.3V
V_{pp} (Read Mode)	V_{CC}
V_{CC} (Write Mode)	6.5V
V_{pp} (Write Mode)	13V

Source: ICE, "Memory 1997" 19964

Figure 6-21. Low-Voltage EPROM

AMD has emphasized its low-voltage, single-voltage flash products. AMD introduced an 8Mbit flash device at 3.0V only (but with a range from 2.7V to 3.6V). The part has the 2.7V to 3.6V voltage range for both read and write operations.

Intel promotes its Smart 3 family belonging to the SmartVoltage concept that allows voltage flexibility. The memory can be read and written at multiple voltages. The Smart3 family includes 4Mbit, 8Mbit, and 16Mbit devices with 2.7V-3.6V for read operation and program/erase plus 12V for fast programming capability.

Unified Memory Architecture (UMA)

The Video Electronics Standard Association (VESA) proposed a new architecture for PCs called Unified Memory Architecture (UMA). The goal of this architecture is to lower the price of the PCs by eliminating some memory.

In the traditional PC memory architecture, a frame-buffer memory is used for the graphics subsystem and a DRAM array is used as main memory. The new PC bus configuration proposal suggests that both main memory and the graphics frame-buffer use a common DRAM array. A graphics frame-buffer uses 1 to 2Mbytes of memory. This would be included in the main DRAM array. By eliminating the separate frame-buffer memory chip in the PC, a cost savings of between $30 and $80 could be realized.

The UMA design, however, reduces the performance of the PC. Speed degradation is estimated at anywhere between 5 and 20 percent. For this reason, UMA targets the low-end applications where price is more important than performance. Figure 6-22 shows the traditional architecture versus the VESA UMA architecture. Companies that have developed chip sets for UMA based systems are shown in Figure 6-23.

Among the difficulties UMA faces is lack of support from some major companies. For example, Microsoft says that the performance of an 8Mbyte UMA system using Windows 95 is decreased too much, as available memory is reduced by 1Mbyte.

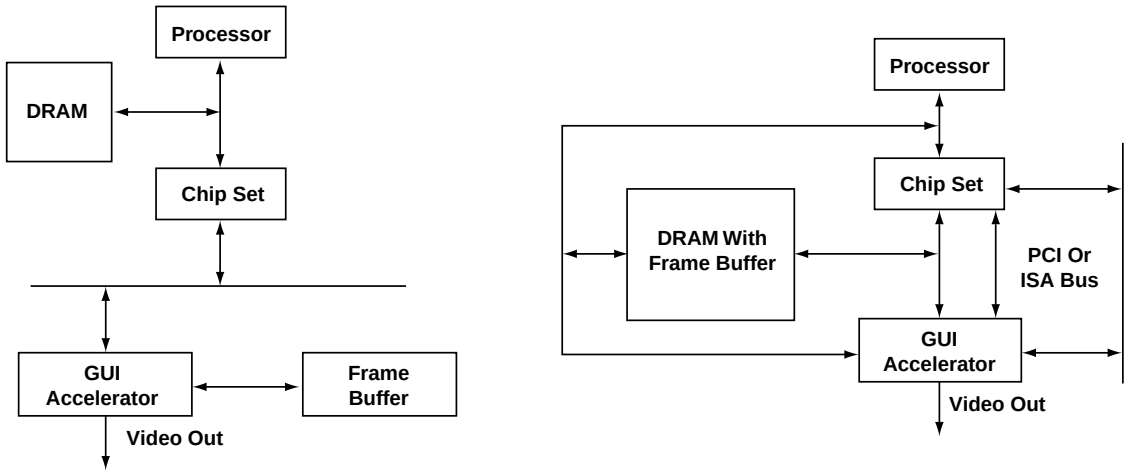

Source: ICE, "Memory 1997"

20871

Figure 6-22. Traditional Memory Architecture Versus Unified Memory Architecture

Vendor	Model	Target Systems	Clock Rate (MHz)	VESA-UMA	Intel SMBA	Remarks
Acer Labs	Aladdin III	Pentium	75-100	Yes	No	—
Cirrus Logic	CL-GD54UM36	Pentium	75-100	Yes	No	Graphs Accelerator
Opti	Viper-UMA	Pentium	75-100	Yes	No	—
S3	Trio 64UV+	Low-end Pentium	75-100	Yes	Yes	—
Trident Micro	TGUI9682	Pentium	75-100	Yes	Yes	Graphics Accelerator
Via Technology	Apollo VP-1	Pentium	75-100	Yes	No	—
VLSI Technology	In Development	P54C, P6, Power PC	100 and up	Yes	No	—

Source: Electronic Business Today, "Memory 1997"

20873A

Figure 6-23. Comparing UMA Products

Intel, also, does not support UMA. In fact, it developed another architecture called Shared Memory Buffer Architecture (SMBA). SMBA minimizes the impact on processor performance from sharing partitioned system memory with the graphics controller by using smart-buffering programmable timers with a strict latency policy.

Latch-up

A major limitation to bulk CMOS technology is the inherent parasitic bipolar transistors that occur as a natural part of the CMOS structures. The CMOS structure forms parasitic bipolar PNP and NPN transistors that are the electrical equivalent of a silicon controlled rectifier (SCR). The parasitic device is electrically connected across the V_{DD} and V_{SS} power supply and input or output

pads. Thus, if this parasitic device conducts, it establishes a very low resistance path to the power supply and can cause very large amounts of current to flow through the structure. This phenomenon is shown in Figure 6-24. Since there are many PNPN devices on a CMOS chip, it is possible to trigger any one device into conduction and cause the latch-up problem.

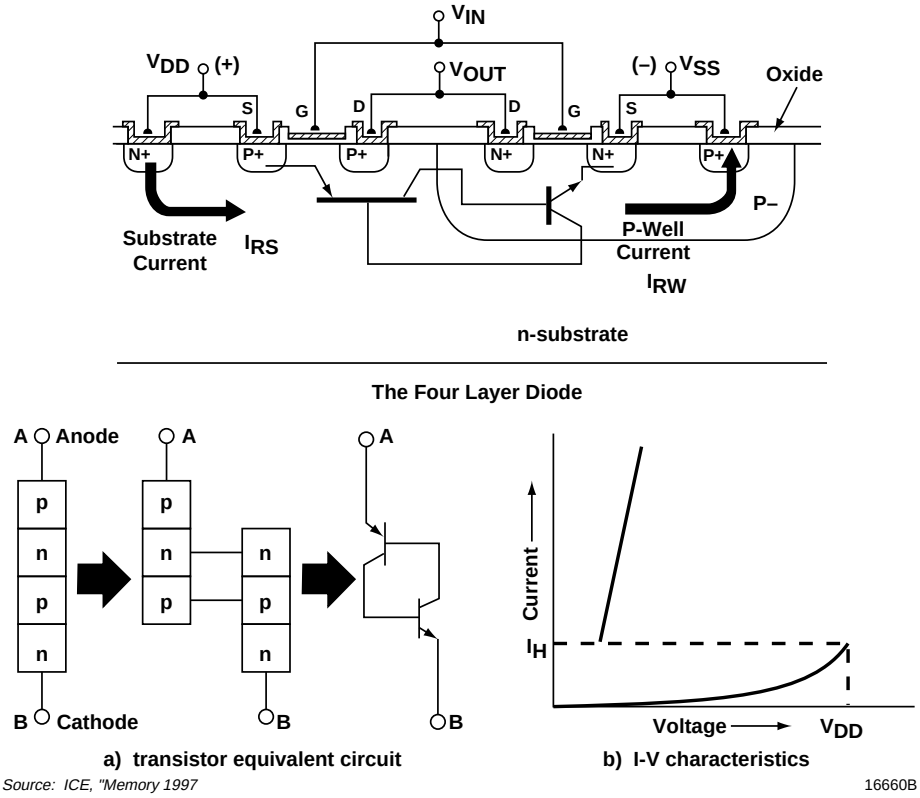

Figure 6-24. Bulk CMOS Inverter With pnpn Latch-Up Path

The phenomenon of latch-up is well understood by memory designers and many approaches have been taken to eliminate it. However, the latch-up sensitivity of a CMOS circuit increases in severity as device dimensions continue to shrink. Therefore, new latch-up reduction techniques will be needed in the future as the process geometries become smaller.

Conditions that trigger latch-up can come from several different places in an electronic system including over-voltage stress, voltage transients, radiation-induced photocurrents, or the input protection circuit being over driven.

Some of the methods used to control latch-up include guard rings, epitaxy on a heavy doped substrate, epitaxy / buried layer CMOS structures, Schottky clamps, physical barriers to lateral currents, trench isolation, total dielectric isolation, and well controlled design rules.

The Electronic Industries Association (EIA) under the direction of a JEDEC committee published JEDEC standard no. 17, “Latch-Up in CMOS Integrated Circuits.” This is a standard used for testing CMOS circuits for latch-up problems.