1 Products

A. MEMORY

Memory and data storage take many forms, even in a single computer system. Figure 1-1 shows a simplified block diagram of a personal computer, indicating the various active sections needed to make the machine function. Several of the blocks are for storage of both the data and programs used by the computer. The CPU chip is the actual brain of the computer. But in order for this "brain" to function, it must call on data and programs stored in other parts of the computer. The CPU chip and the storage sections are shown in Figure 1-2.

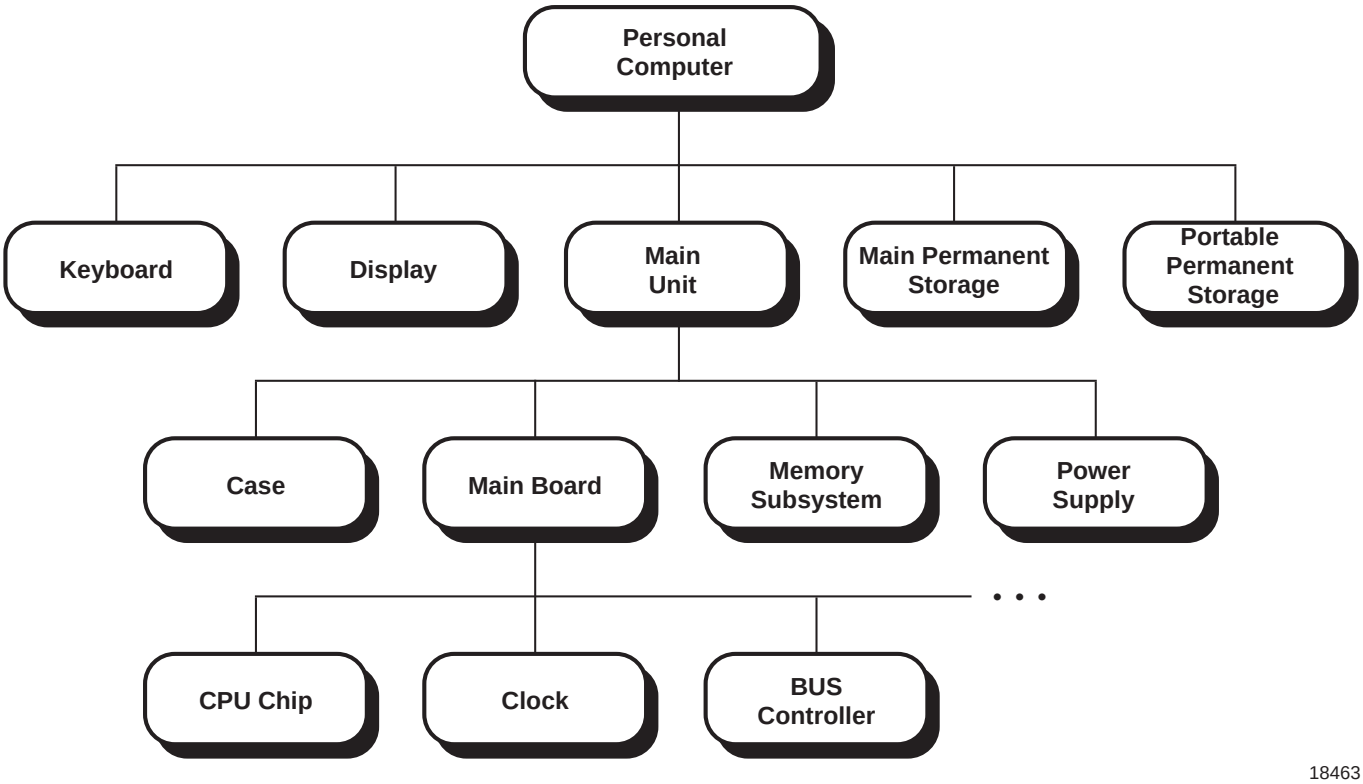

Figure 1-1. System Partition of Personal Computer

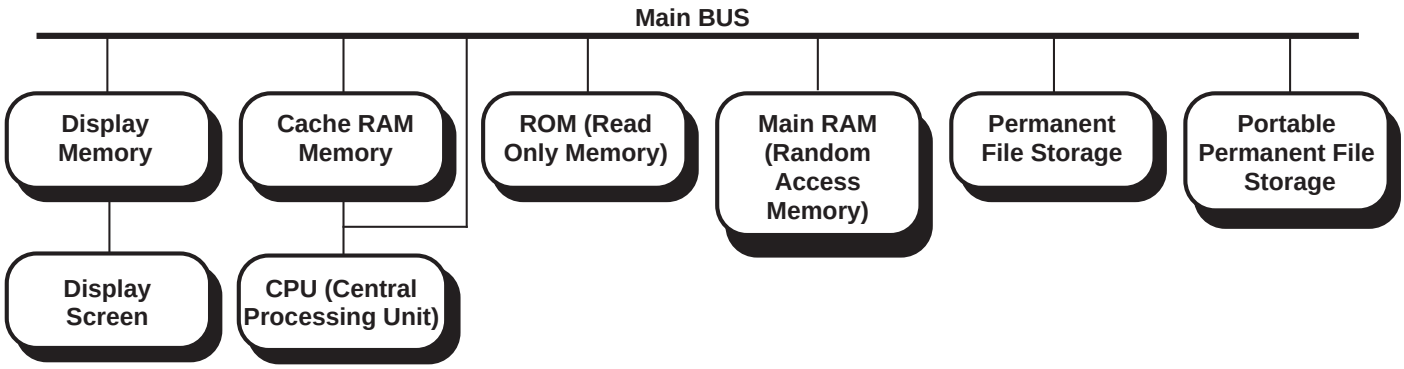

Figure 1-2. Memory/Storage in a Computer

To better understand the workings of a computer, the requirements of the blocks must be understood. Figure 1-3 describes the attributes of the various sections used for storage. Each function can be served by one or more types of semiconductor or magnetic device, as shown in Figure 1-4, which also describes the pros and cons of each alternative. As can be seen from this table, an alternative may be well suited for desktop machines (e.g., low cost, high power), but a different type is best suited for portable machines (e.g., low power). Mainframes, large minicomputers and large workstations use similar components, but these components are usually much larger in scale (and cost). Figure 1-5 shows a rough comparison of the differences in the computers described, without getting into performance and cost differences.

Circuit Block	Main Use	Typical Storage Size	Driving Attribute	Other Attributes
Main permanent file storage	Active data storage	50 MBytes or more	Low cost	Non-volatile and alterable
Portable permanent file storage	File exchange and backup	0.5 to 3 MBytes	Removable	Non-volatile and alterable
ROM	Machine startup and screen drawing routines	0.2 to 2 MBytes	Non-volatile	Reasonably fast, non-alterable
Main RAM	General temporary storage during operation	8 MBytes	Reasonably fast (<100nsec)	Low cost and alterable
Cache RAM	Very fast temporary storage	4 to 256 KBytes	Very fast (<15nsec)	Alterable
Video RAM	Stores data displayed on screen	1 to 24 bits/pixel	Fast enough for screen refresh	Low cost and dual access for fast read/write

18465

Figure 1-3. Types of Memory/Storage in a PC

Requirement	Storage Alternative	Advantages	Disadvantages
Main permanent file storage	Hard Disk	Low cost	Large size, high power, not rugged
	Flash EEPROM	Rugged, small, light	High cost
Portable permanent file storage	Floppy disks	Very low cost media	Small storage per disk
	Flash EEPROM	Rugged, small, light	High cost
	Optical & Magneto-optical	Very large capacity, low cost media	High entry cost, not a standard yet
ROM	Mask ROM	Very low cost	Un-alterable
	EPROM	User alterable at factory	Slow to program
	EEPROM	Field alterable	Very high cost
	Flash EEPROM	Field alterable	High cost
Main RAM	DRAMs	Lowest cost	High power
	SRAMs	Low power, fast	High cost
	PSRAMs	Low power	Medium cost
Cache RAM	SRAMs	Fast	Very high cost
Video RAM	DRAMs	Low cost	Single mode access
	VRAMs	Low system cost, fast	Expensive for small systems

18467

Figure 1-4. Memory/Storage Alternatives

	Size Main Memory	Size of Cache Memory	Permanent File Storage Size	Portable Permanent File Storage
Laptop	6 Meg	0	80 MByte	Floppy
PC	8 Meg	32K	200 MByte	Floppy
Workstation	32 Meg	256K	300 MByte	Floppy
Workstation as Network server	128 Meg	512K	1 GByte	Tape
Super Mini	256 Meg	512K	40 GByte	Tape
Mainframe/ Super Computer	1 GByte	512K each Processor	100 GByte	Tape

18466

Figure 1-5. Memory/Storage in Various Machines

The semiconductor devices in Figure 1-4 are described below in terms of both architecture and storage mechanism. Each type also requires slightly different manufacturing steps, leading to very different wafer costs.

1. DRAMs

DRAM (Dynamic Random Access Memory - Figure 1-6). This is the main memory used for all desktop and larger computers. Each "storage cell" consists of a single MOS transistor and storage capacitor, and can store one bit of information. The DRAM is unique in the fact that the data in each "cell" is lost over time due to leakage of the cell transistor, and must be refreshed several times each second. This refresh is done automatically and is actually not a problem in desktop computers.

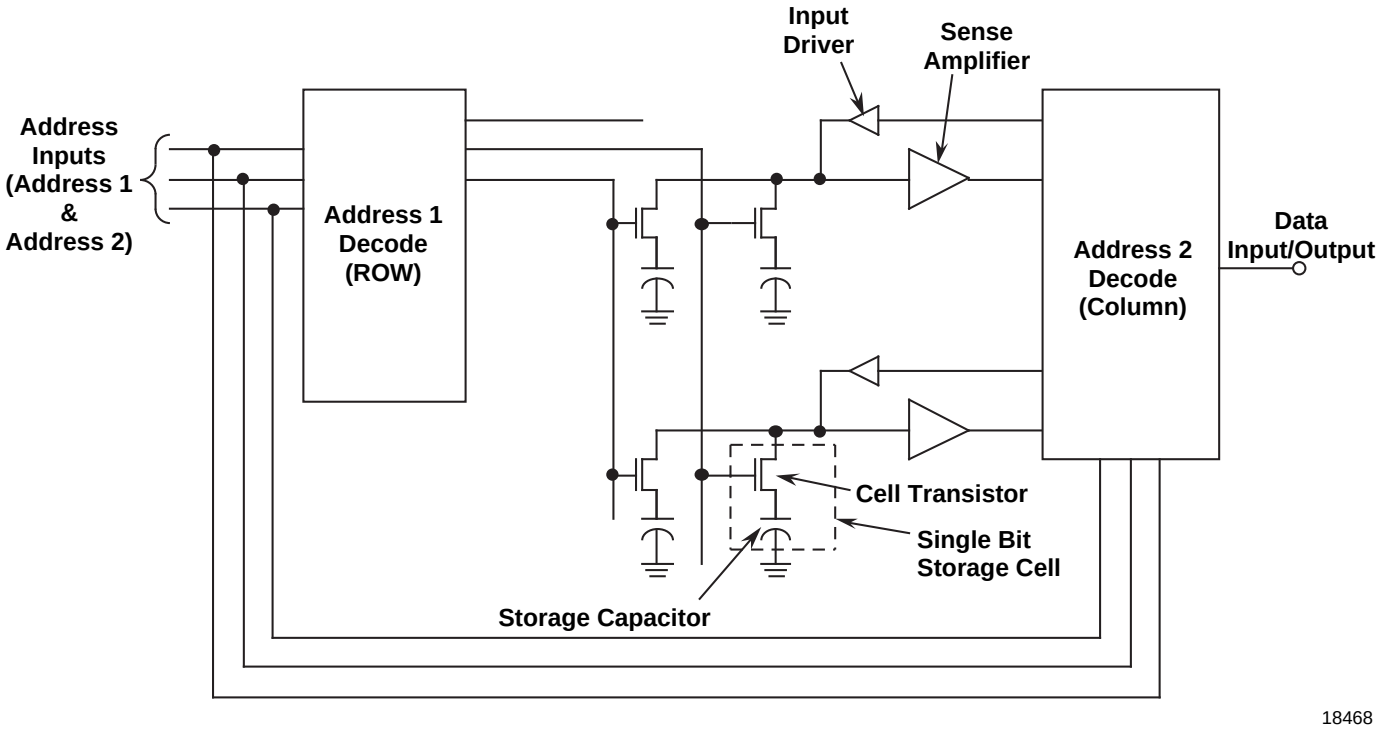

Figure 1-6. DRAM

The big advantage of the DRAM is that the amount of silicon required per bit of storage is very small, and the cost per bit of storage is the lowest of all current memories. DRAMs do suffer from a speed problem, as the circuitry on the chip required to read the data in each cell is inherently slow. As such, the DRAM speeds have not been able to improve at the same pace as the speeds of the CPUs.

The early CPUs were introduced with clock speeds of 1MHz (one million cycles per second). Today the faster CPUs in desktop PCs are 100MHz, a 100 times improvement. The early DRAMs had access cycle times (the time required for the DRAM chip to supply the data back to the CPU) of 250nsec, and the fastest units today are 50nsec or less, an improvement of more than five times.

Newer technologies are now being introduced to reduce the speed bottleneck. One innovation was to synchronize the read and write commands to the DRAM to the system clock, thereby enabling the CPU chip to receive data on a more rapid basis once the first address instruction has been received by the DRAM memory. These newer chips are called synchronous DRAMS for this reason.

2. PSRAMs

PSRAMs (Pseudo Static Random Access Memory) are actually DRAMs optimized for very low power consumption and designed with a simpler interface to the system bus. The cell structure is nearly identical to the DRAM, so the cost is low. The PSRAM sacrifices speed for the lower power consumption, so the main application is laptop computers where a small sacrifice in speed is preferable to reduced battery life. The PSRAM is a good compromise between the DRAM and the more expensive SRAM.

3. SRAMs

SRAMs (Static Random Access Memory) are designed to fill two needs. The first is the need for memory that can operate at the speed of the CPU. In this role it is called cache memory, which will be discussed later. These chips may be accessed in as little as a few nanoseconds, versus 50 to 100 nanoseconds for a DRAM. In the case of most high-end processors a large block of SRAM is integrated onto the processor chip.

The fast SRAM cell is usually composed of four transistors and two resistors. The storage cell area of an SRAM is about four times as large as the cell of the comparable generation DRAM. The data in an SRAM cell is volatile, i.e., the data is lost when the power is removed. However, the data does not “leak away” like in a DRAM. The SRAM doesn’t require the refresh cycle and the associated power consumption. Figure 1-7 shows the cell schematic of the four-transistor cell (4T cell) SRAM.

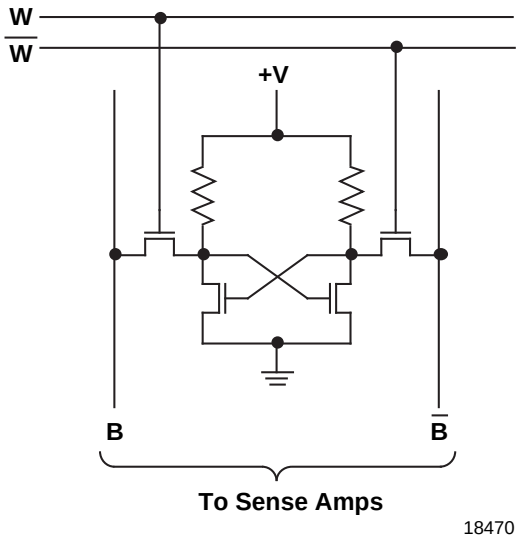

Figure 1-7. SRAM 4T (Four-Transistor) Cell

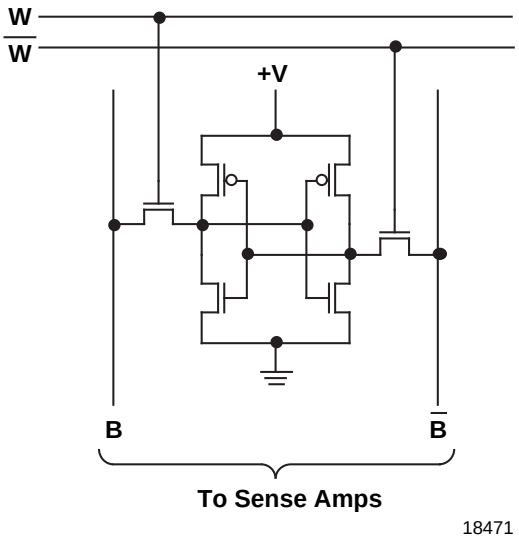

Figure 1-8. SRAM 6T (Six-Transistor) Cell

A second version of the SRAM is designed to accentuate the low-power capability while sacrificing speed and cost. These SRAMs are composed of six-transistor cells. This structure is shown in Figure 1-8. These devices are used where extremely low power consumption is mandatory, such as palm top computers operating from AA batteries.

Self-timed SRAMs are a special case for very high-performance applications. The address information is passed from the CPU to the SRAM in a short burst, instead of through the entire read or write cycle as with the standard SRAM. This allows the system designer more latitude in the design of the timing of read / write cycles, as the other timing signals are created inside the SRAM chip and presented a fixed period of time later to the CPU. Unfortunately, these features add complexity and cost to the SRAM design. It is usually only found in very high speed mainframe computers.

a. Cache Memory

To implement a cache memory requires the use of special circuits that keep track of which data is in both the main memory (DRAM) and the SRAM cache memory, and which data is only in the main memory. This function acts like a directory that tells the CPU what is or is not in cache.

The function can be designed with standard logic components, with SRAM chips for the data storage. An alternative is the use of special memory chips called Cache Tag RAMs that, in conjunction with the SRAM data memory chips, perform the entire function. Figure 1-9 shows both the Cache Tag RAM and the Cache Buffer RAM along with the main memory and the CPU (processor).

Another innovation is the CDRAM, a DRAM with a cache controller and cache memory (made with SRAM) integrated on each DRAM chip. This approach effectively offers a much simpler solution for the systems designers, and improved performance over a system with no SRAM cache memory.

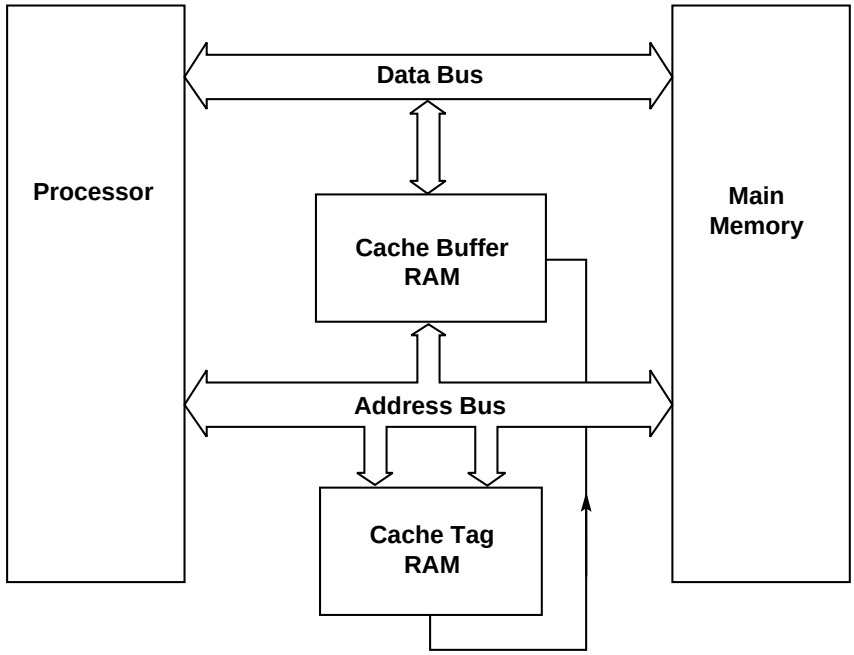

Source: TI 18472

Figure 1-9. Typical Memory System with Cache

b. FIFO Memory

A FIFO (First In, First Out) memory is a specialized memory used for temporary storage to aid the timing of non-synchronized events. A good example of this is the interface of a computer system to a Local Area Network (discussed later).

Often, a computer system will receive data along the network at some inopportune time, such as when the CPU is in the middle of calculating the cells of a spread sheet or saving a file to disk. In these cases, there is a priority conflict, because the data coming along the network must be captured immediately or it will be lost.

One approach is for the computer to "drop what its doing" and service the network request. This approach is quite inefficient, but can be done in software as part of the operating system. This is a relatively inexpensive approach.

Another alternative is to store the data received from the network until the current task is complete. This can be done with special circuitry that stores the network data in special memory chips as it is received. This memory area is operated on a "First In, First Out" basis, because that is the way the data is received.

The FIFO memory can be constructed with logic components and SRAM chips, or with the use of FIFO memory devices. The latter case is usually much more efficient use of board space and is typically found in all high-speed communications applications.

c. Multi-Port and VRAMs

Multi-port (usually two port, but sometimes four port) memories are specially designed chips using either SRAM or DRAM memory cells, but with special on-chip circuitry that allows multiple ports (paths) to access the same data, at almost the same time.

One application is where multiple CPUs are tied together by common memory. The multi-port chip acts as a bridge for temporary storage between two CPUs, so that each can operate at maximum efficiency even when data from the other is needed. Because of the speed requirements, only SRAMs can be used for this.

A special case of the two-port memory is the Video Random Access Memory (VRAM). The VRAM is used for high-performance video memory, usually for large screens (high pixel count), where the screen information storage is very large. This data storage area must be "read from" and "written to" very rapidly. In small-screen computer applications, standard DRAMs are often used.

The VRAM is usually constructed using the same fabrication process and cell structure as the DRAM. The VRAM, however, has special circuitry for multiple access to the memory data — one for writing and another for reading. Figure 1-10 shows the extra size required for this circuitry. Both chips have one million bits of storage, and are constructed with the same storage cell size.

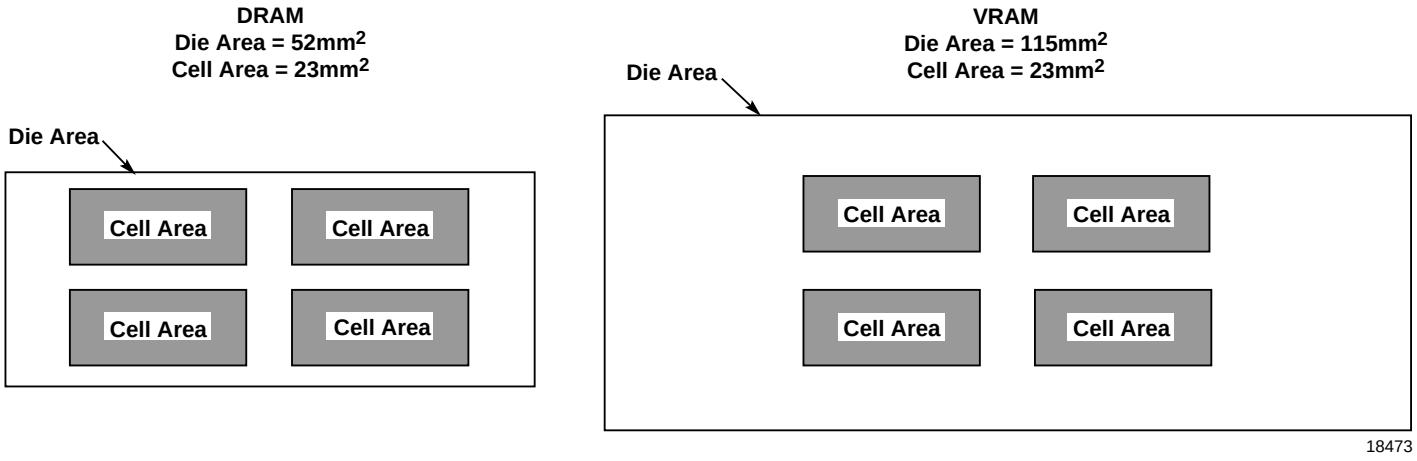

Figure 1-10. DRAM Versus VRAM

4. EPROM

EPROM (Electrically Programmable Read Only Memory) is a special case of semiconductor memory that is truly non-volatile, i.e., the memory retains the stored information when the power is removed. The device is programmed by forcing an electrical charge on a small piece of poly-silicon material (called a floating storage gate) located in the memory cell. When this charge is present, the cell is "programmed." The programming (write) cycle takes several hundred milliseconds per byte. The read time is comparable to that of fast DRAMs. Figure 1-11 shows the cell used in a typical EPROM. The floating gate is where the electrical charge is stored.

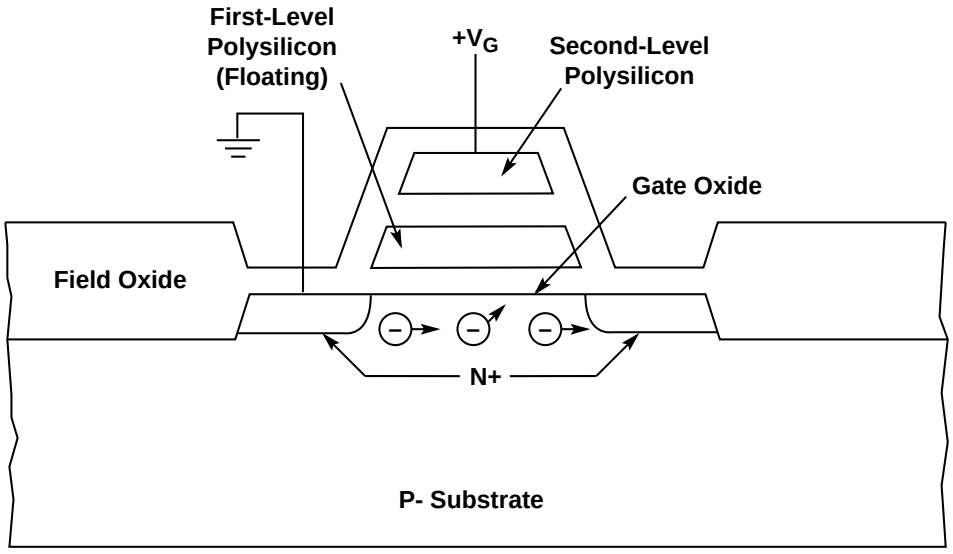

Source: Intel 18474

Figure 1-11. Double-Poly Structure (EPROM/Flash Memory Cell)

Some EPROMs are packaged with glass windows over the chips. If a strong ultraviolet light is shined on the chip for several minutes, the entire chip can be erased. Another alternative is the One-Time Programmable (OTP) version packaged in standard black opaque plastic. In this case, the chip can only be programmed once, hence the name.

The erasure capability is the EPROM's main advantage over the Mask Programmed ROM. It allows the user to buy mass-produced devices off the shelf and program each device for a specific need. The Mask Programmed ROM, however, is the lowest cost ROM available.

5. MROMs

MROMs (Mask Programmable ROMs) are used in certain applications where the identical "program" is repeated in each of hundreds of thousands of sockets. Examples of this are electronic games, the special screen graphics routines in the Macintosh computer, and the very small startup circuit in PCs (required to read the disk drive before the operating system is loaded), called BIOS. MROMs come from the IC factory already programmed and cannot be changed. This limits their applications to very high volume and stable products.

6. EEPROMs

The EEPROM (Electrically Erasable Programmable Read Only Memory) is actually incorrectly named, as it is really not a "Read Only Memory" at all, but can be electrically written as well as read. The cell structure is similar to the EPROM, but the storage transistors are manufactured so that, when the high voltages used for programming are applied in the reverse direction, the program will be erased. The programming and erasure mechanisms are both relatively slow, tens of milliseconds. The read times are as fast as DRAMs. Figure 1-12 shows the cross section of an EEPROM storage transistor.

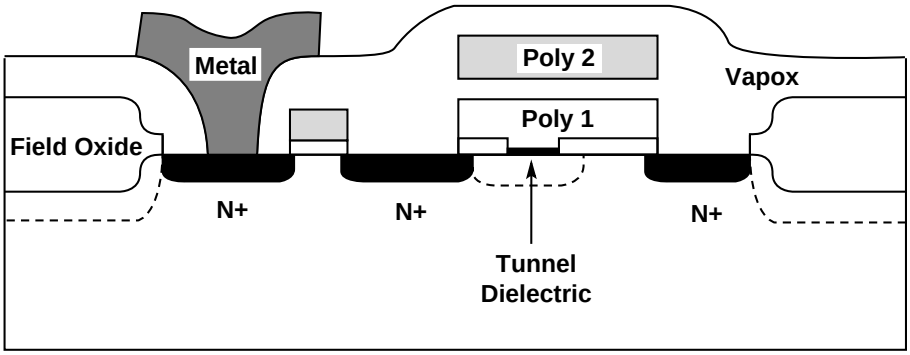

Source: Seeq Technology, Inc. 18476

Figure 1-12. EEPROM Cell Cross Section

The advantage of this device is the in-circuit erasure capability. The disadvantage is the very large cell size required to allow both the high positive and negative voltages (for writing and erasing, respectively) to be directed to each individual cell. This is shown in Figure 1-13.

Another disadvantage of EEPROMs is the limitation on the number of erase cycles allowed. This limitation is due to the slight damage created in the oxide film each time the device is erased. Typically, the limit is 100,000 erase cycles per cell.

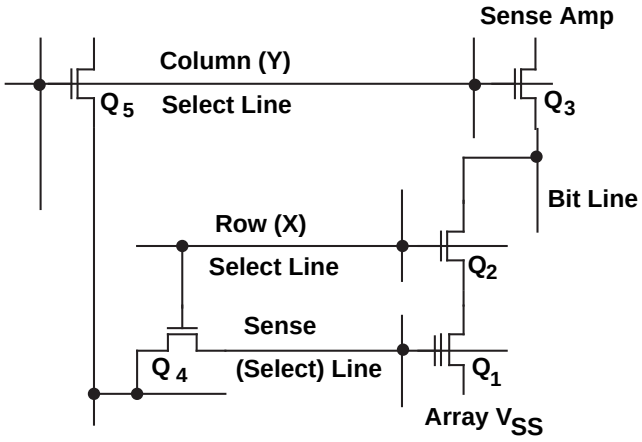

Figure 1-13. Generic EEPROM Memory Cell

7. Flash

A modified version of the EEPROM called a Flash EPROM (or Flash EEPROM) is a compromise between the EPROM and the EEPROM. These flash devices are programmed like the EPROM and EEPROM, and can be erased electrically. The difference between the flash device and the EEPROM is that individual cells of the flash device cannot be erased. Rather, large groups of cells are erased as a block. This block size ranges from hundreds to thousands of cells. Figure 1-14 shows the cell organization of the flash device.

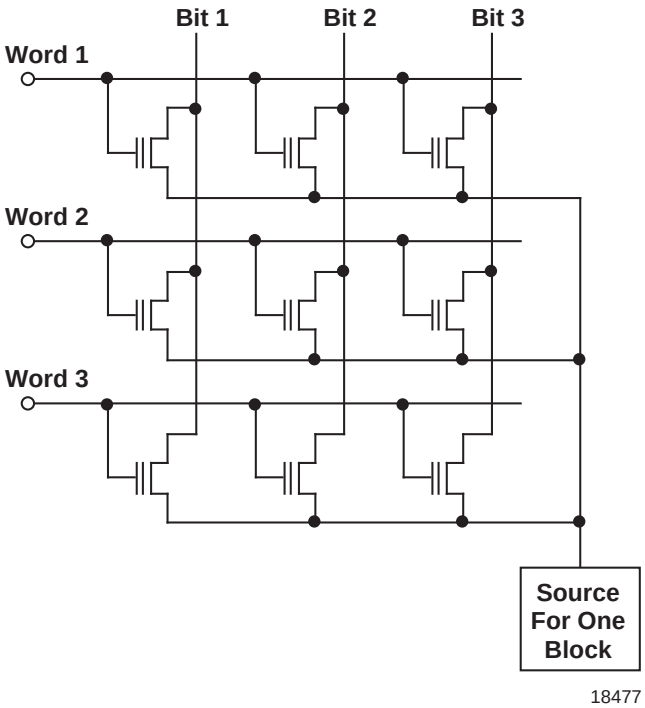

Figure 1-14. Flash EPROM Cell Connections

The advantage of the flash device is that the potential cost per bit of storage approaches that of the EPROM, whereas that of the regular EEPROM will always be considerably higher due to the large cell size. The disadvantage of the flash device is that individual cells cannot be erased. Also, the flash structure suffers from the same limit to the number of erase cycles as the EEPROM. In the case of a the flash device, however, the limitation applies to an entire block rather than a single cell. This is due to the fact that the entire block must be erased when a single cell in the block is erased.

The early flash devices used a NOR structure. A structure that has been offered as a lower cost alternative is the NAND structure. The basic difference is in the layout of the individual cell. The NAND has an advantage in that it uses a smaller cell by stringing several EEPROM transistors together, without read / write contacts between each one. The disadvantage is that the peripheral circuitry required to read this type of transistor organization is complex, slow, and consumes a great deal of silicon. The slower (serial) readout NAND approach is sufficiently fast to be used as a replacement for rotating media, which is one of the major markets for flash devices.

The NOR connection used for many flash devices is also the standard connection used for all Mask ROMs, EPROMs, and EEPROMs. This configuration allows a fast read cycle, faster than most DRAMs. The terms NAND and NOR originate in the wiring of the cell transistors, which are connected together with common source and drain for the NOR, and in series, with the source of one cell transistor connected to the drain of the next, for the NAND. Figure 1-15 shows the cell arrangements.

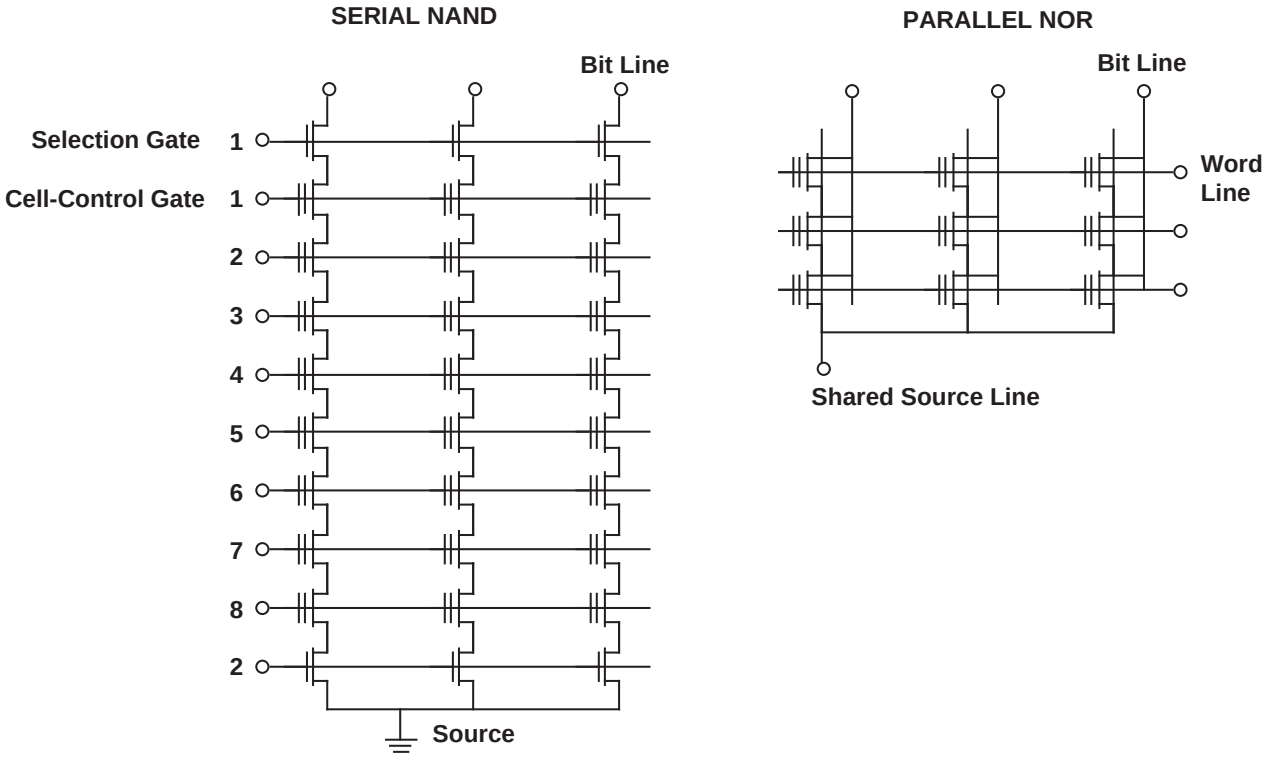

Source: Intel 18478

Figure 1-15. The NAND Versus NOR Contest

Any of these ROMs (MROMs, EPROMs, EEPROMs, or Flash) are logical places to store small computer programs. Because they are programs and not data, they are executed by the CPU, but never altered. A high-volume use of this is in microcontroller applications. Here the ROM can be either on the same chip as the CPU or a separate chip. Typical uses are appliance control and engine control in automobiles

The font cartridges in printers contain ROMs. In this case the ROMs contain the data needed to print particular fonts, giving the printer greater flexibility. Here, the data is stored but never changed.

8. Shadow and Battery-backed SRAMs (NOVRAMs)

Certain applications require the speed of normal memory, but have the additional requirement that the data must retained when the power is shut off, i.e., that they are non-volatile. Normal EEPROMs are non-volatile, but the write and erase times are far too slow for normal system operation.

One alternative that allows the use of SRAMs in these applications is combining SRAM and EEPROM on the same chip. In normal operation, the CPU will read and write data to the SRAM. This will take place at normal memory speeds. If the chip detects there is a power failure about to occur, the special circuits on the chip quickly (in a few milliseconds) copy the data from the SRAM section to the EEPROM section of the chip, thus preserving the data. The EEPROM "shadows" the data that is in the SRAM in normal operation. When power is restored, the data is copied from the EEPROM to the SRAM and operations can continue as if there were no interruption. Figure 1-16 shows the cell schematic of one of these devices.

Another alternative is the use of a small battery in the package of the SRAM chip. SRAMs can be designed to go to a sleep mode where the power consumed is very low but the data is retained. The combination of the SRAM and the small battery can be very cost effective, with retention times of five years. Note that notebook and laptop computers have this "sleep" feature, but here the regular system battery supplies the power.

B. STANDARD LOGIC

Standard logic was the first application of integrated circuits replacing discrete transistors. All early families used bipolar transistors and the names referred to the circuit design, such as Resistor-Transistor Logic (RTL), Diode-Transistor Logic (DTL), Emitter-Coupled Logic (ECL), and Transistor-Transistor Logic (TTL). These are shown in Figure 1-17.

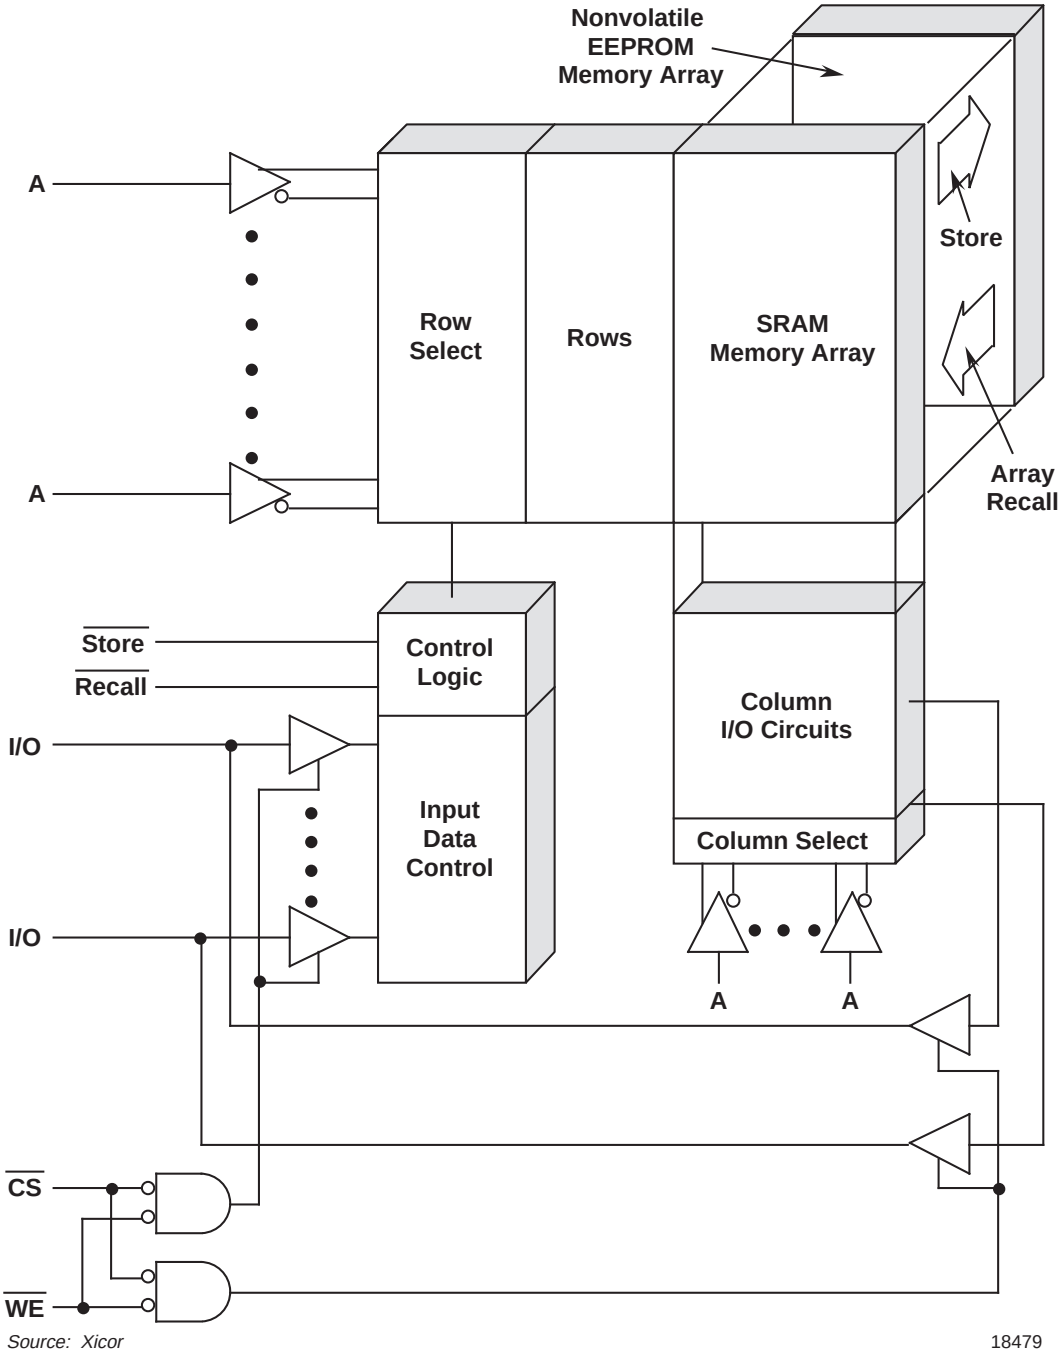

Figure 1-16. Block Diagram of the Xicor NOVRAM Family

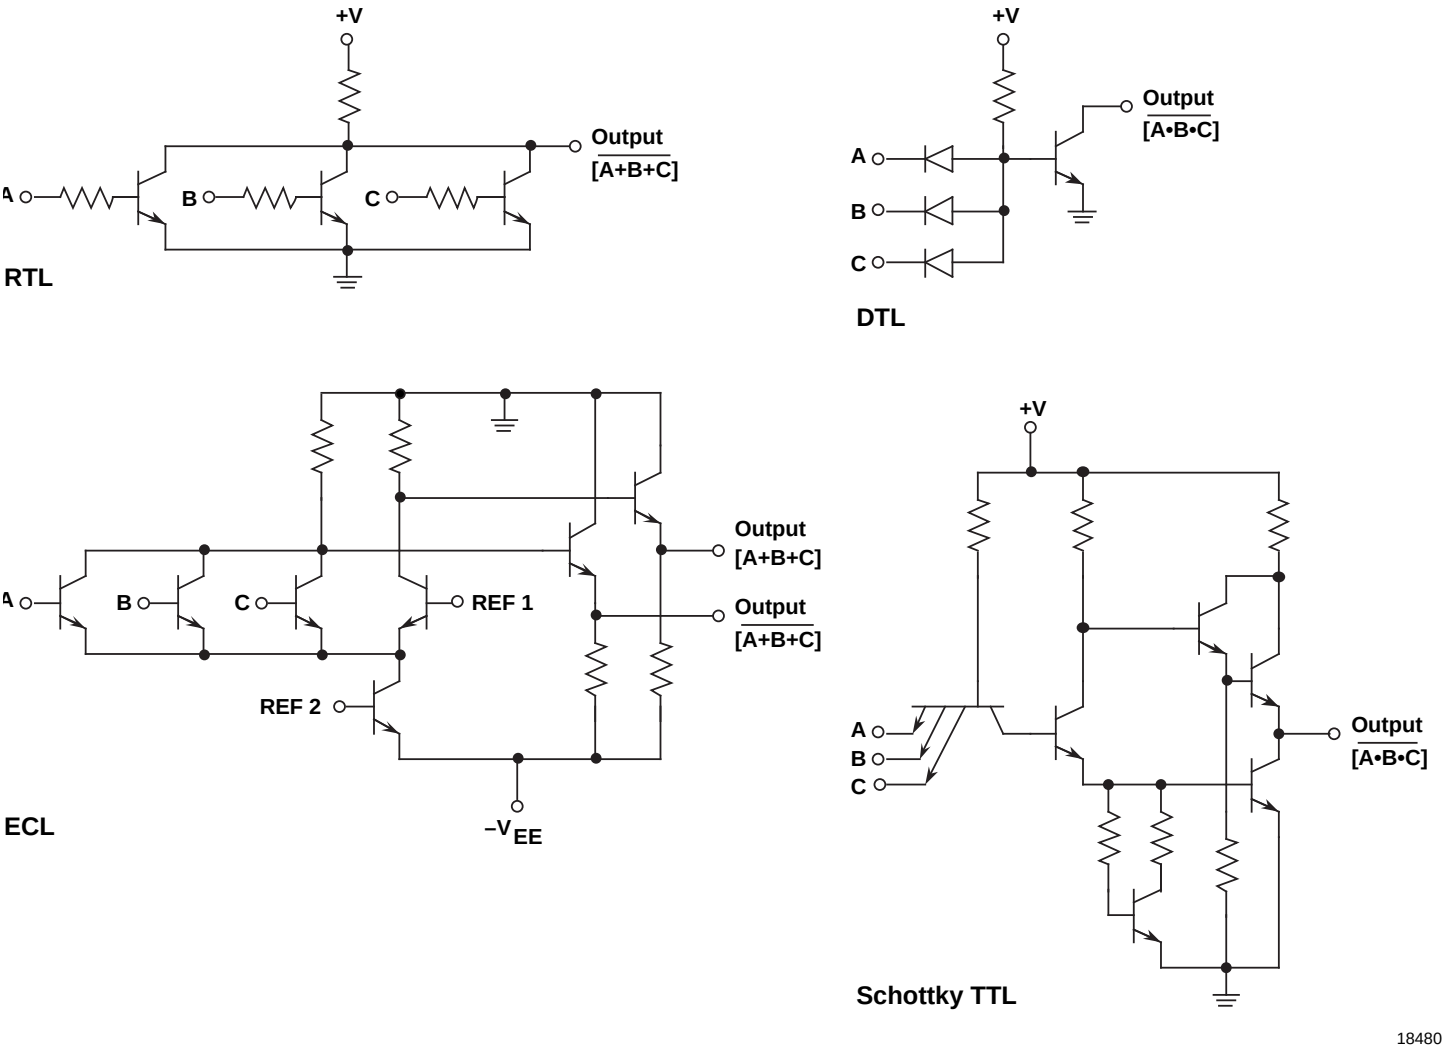

Figure 1-17. Various Standard Logic Circuit Designs

The most popular family was TTL, introduced by TRW and Sylvania in the early 1960's. The design was popularized by TI as the 54/74 series in about 1966. As the years passed, improvements in transistor performance allowed new families to be introduced that offered better speed or lower power, while retaining compatibility with the older system designs. Most manufacturers even used the same suffix numbers, as these were familiar to systems designers. Some of the families are manufactured using CMOS or BiCMOS processes, as well as the more familiar bipolar, offering improved performance or considerably greater functionality.

Figure 1-18 shows the various acronyms that name the families and the relative merits of each family in terms of speed and power.

Family Name (Common Reference Name)		Technology	Relative Speed	Relative Power Consumption/Gate
Schottky TTL	(Schottky)	Bipolar	4	7 (Highest)
Low Power STTL	(LS)	Bipolar	6	5
Advanced Schottky TTL	(AS)	Bipolar	1 (Fastest)	6
Advanced Low Power TTL	(ALS)	Bipolar	5	4
High Speed CMOS TTL	(HCT)	CMOS	7 (Slowest)	3
Advanced CMOS TTL	(ACT)	CMOS	3	1 (Lowest)
BiCMOS TTL	(BCT)	BiCMOS	2	2

18481

Figure 1-18. Logic Families

Virtually all the parts in these families are general purpose logic, i.e., they are suitable for virtually any application that needs that particular logic function. Today, Very Large Scale Integration (VLSI) CMOS is designed to perform the function of hundreds of these chips, and for very specific applications.

C. MICROPROCESSORS AND MICROCONTROLLERS

1. Microcontrollers

The terms microcontroller, microcomputer, embedded controller, microprocessor, and CPU are used in the trade press without precise definition. ICE defines a microcontroller (MCU) as a device that contains all of the necessary functions of a computer on a single piece of silicon or other semiconductor material. These necessary functions include the central processing unit, Random Access Memory (RAM) used for read-write memory, nonvolatile program memory (ROM or ROM variations such as EPROM, EEPROM or flash), and at least one input/output port. The main thing that separates MCUs from microprocessors (MPUs) is that MPUs do not have on-chip ROM.

MCUs are sold almost exclusively as embedded controllers. Therefore, the statement "all MCUs are embedded controllers" is almost 100 percent correct. An embedded controller is a processor with its support circuitry whose function is to control the operation of some piece of equipment. However, all embedded controllers are not MCUs, as in many large systems regular MPUs are used as embedded controllers.

2. RISC vs. CISC Designs

The terms RISC (Reduced Instruction Set Computer) and CISC (Complex Instruction Set Computer) refer to the instruction vocabulary available to the programmer that will use the processor. The RISC concept was proposed by Patterson and Ditzel back in 1980 at IBM. The computer created was technically very different from the typical IBM computer of that timeframe.

The idea was to create a computer that was much simpler in design, and therefore could have other features that would allow the machine to operate considerably faster. As an example, the area on the processor chip taken up by the complex decoding required for the large repertoire of instructions in the CISC could be used for a large number of on-chip registers for the RISC, making the RISC very fast.

Today, the limitations on transistor count per IC chip no longer cause many significant design compromises. Many "good ideas" from the RISC world have been implemented in modern CISC designs. Likewise, some of the early RISC ideas have been found to be impractical and have been dropped.

Modern RISC chips will undoubtedly remain leaders in performance for two reasons. First, most modern CISC chips are designed to be compatible with the original designs first introduced in the 1970's, while the RISC chips are based on 1980's knowhow. Second, RISC chips can be applied to parallel applications much more easily, and this is a definite trend for all high-performance applications.

D. ASIC

The level of integration possible with Very Large Scale Integration (VLSI) caused a major change in the design of ICs. The capability of putting thousands of transistors on a single chip meant that more than one standard logic function could be placed on each chip. This, in turn, meant the system design would determine the circuitry of each IC chip. The IC could then be customized for a very specific application or product.

Application Specific Integrated Circuits (ASICs) are those ICs designed by or for the customer to fill a very specific function. Another type of IC designed for a specific application is the Application Specific Standard Product (ASSP). The typical differentiation between the two is that an ASSP is purchased by more than one customer, and it will typically have a datasheet describing the performance of the chip.

1. Custom

Several methods are used to convert a logic level design (a design composed of logic gates such as AND, OR, NAND, NOR, etc.) into the design of the IC chip. Full custom design refers to those chips designed at the transistor level, i.e., where each transistor is sized to fill the exact requirements. It is used for very high volume applications because the design cost is very high.

2. Standard Cell

Standard cell is the most common design approach used for VLSI ICs. The concept consists of two phases. First, "cells" or "macros" are designed to perform logic functions. These macros may be as simple as a two-transistor inverter or as complex as a several thousand transistor Arithmetic Logic Unit. The macros are characterized in terms of performance, power supply connections, input/output wiring requirements, etc. Figure 1-19 shows a "2:1 multiplexer" macro containing eight transistors. This would be considered a small macro.

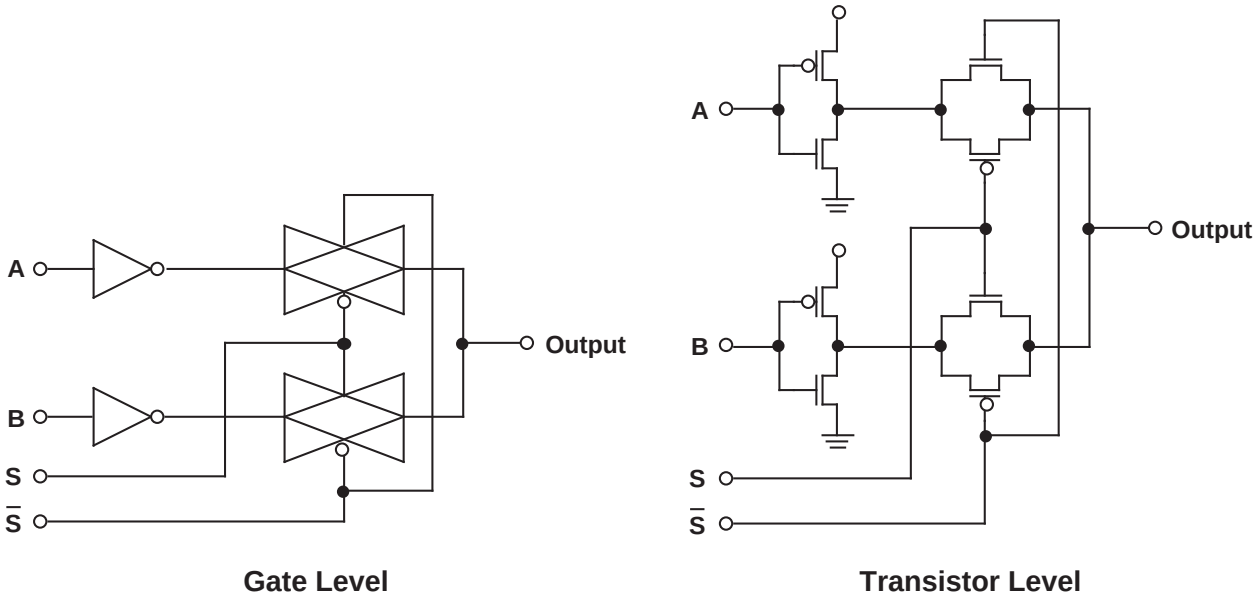

Figure 1-19. Two-Input Multiplexer

The second phase of a standard cell design is the actual design of the total chip. The designer uses the macros to create the total chip function. The computer then "places" the cells in an IC layout and connects them to achieve maximum performance. The standard cell approach is not as efficient as the full custom approach in terms of the amount of silicon used, i.e., the IC chip will be larger than with full custom. However, the standard cell approach requires considerably less engineering effort to design, as the macros can be re-used for different designs.

3. Gate Array

Another ASIC approach is the gate array. In this technique, only one basic cell is created, and is repeated on the chip thousands of times. The basic cell usually consists of four transistors as shown in Figure 1-20. These cells are not complete, requiring several more connections to become complete logic gates.

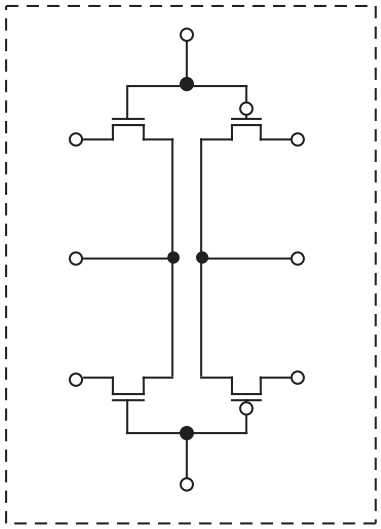

Source: Fujitsu 18483

Figure 1-20. The Basic Gate Array Cell

The final metal layer(s) of the IC is used to complete these connections, using the transistors of the basic cell to create different kinds of logic gates. This is shown with the heavy connection lines in Figure 1-21. The gate array is less efficient in silicon usage than the standard cell, but has several advantages. One advantage is faster turnaround time in the IC manufacturing fab, as wafers can be partially processed and stored, awaiting only the interconnect metal layer(s). Another advantage is that new complex cells can be designed quite easily from the basic cell building blocks.

The chip designer uses almost the same design tools whether designing for a standard cell or gate array. The libraries for either methodology will be slightly different. The software that creates the actual chip is very different, but this is transparent to the designer.

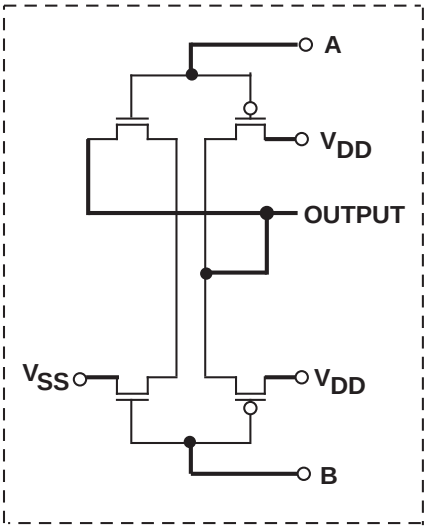

Source: Fujitsu 18484

Figure 1-21. The Basic Gate Array Cell as a 2-Input NAND Gate

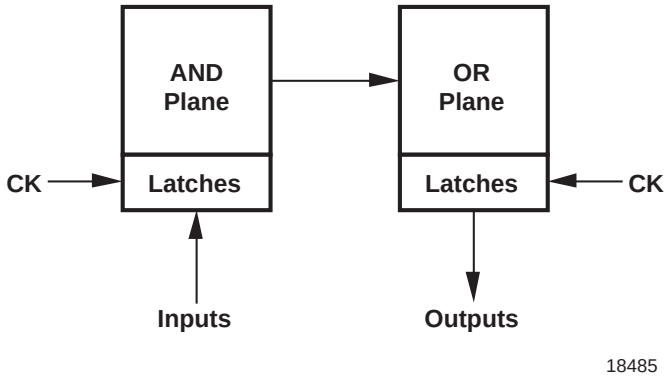

Figure 1-22. AND-OR PLA

4. PLDs

a. PLAs, PALs

The use of Programmable Logic Devices (PLDs) is another ASIC implementation approach. PLDs cover a broad range of IC chip categories. The earliest PLDs were implemented in a modified version of the general purpose Programmable Logic Array (PLA) shown in Figure 1-22, and named PAL by Monolithic Memories. These parts differed from the PLA in that the OR plane is fixed rather than infinitely variable. An example of the PAL architecture is shown in Figure 1-23.

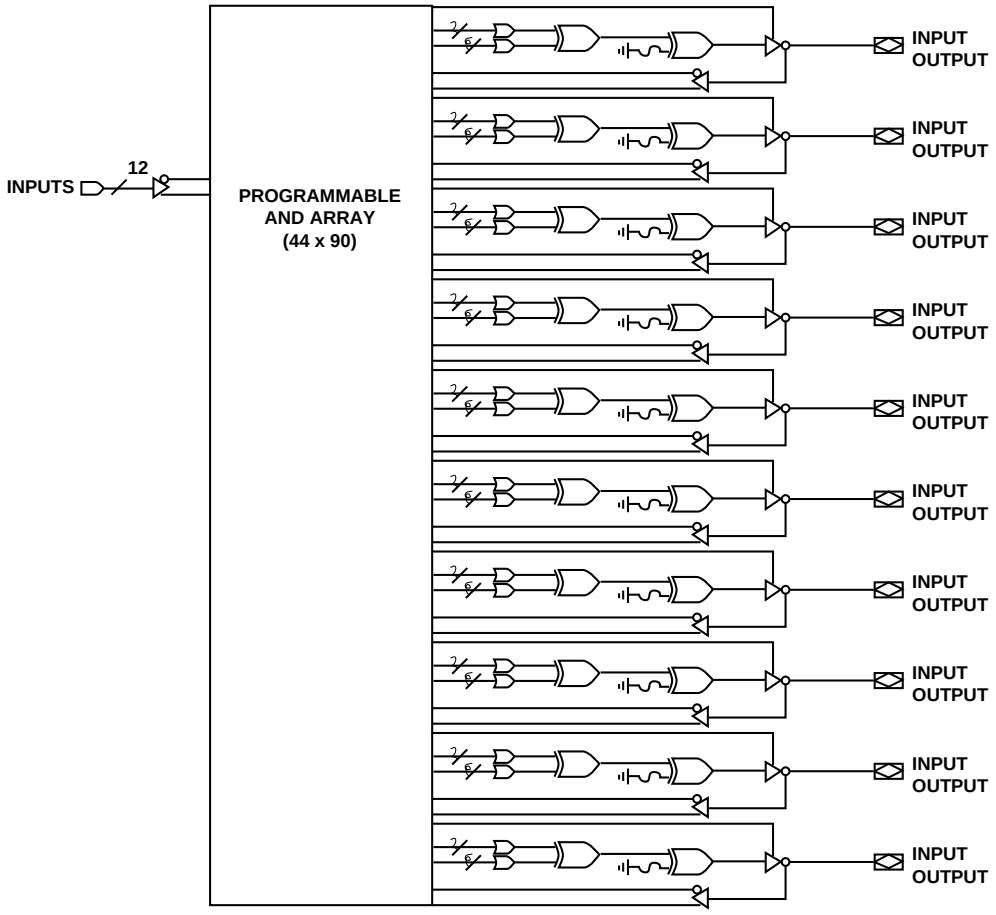

Figure 1-23. Block Diagram of an AmPAL22XP10

These early PALs were introduced in bipolar TTL logic, with metallic fuses that were "blown" to program the part in the desired logic configuration. Today, many of the new PAL devices are CMOS, using EEPROM transistors instead of the metallic fuses. In addition, the usefulness of the device has been enhanced by the addition of complex input and output "macros" (not to be confused with the macros or cells in the standard cell methodology). These macros are programmed to form a variety of logic functions. Figure 1-24 shows a common output macro for a widely used PAL device.

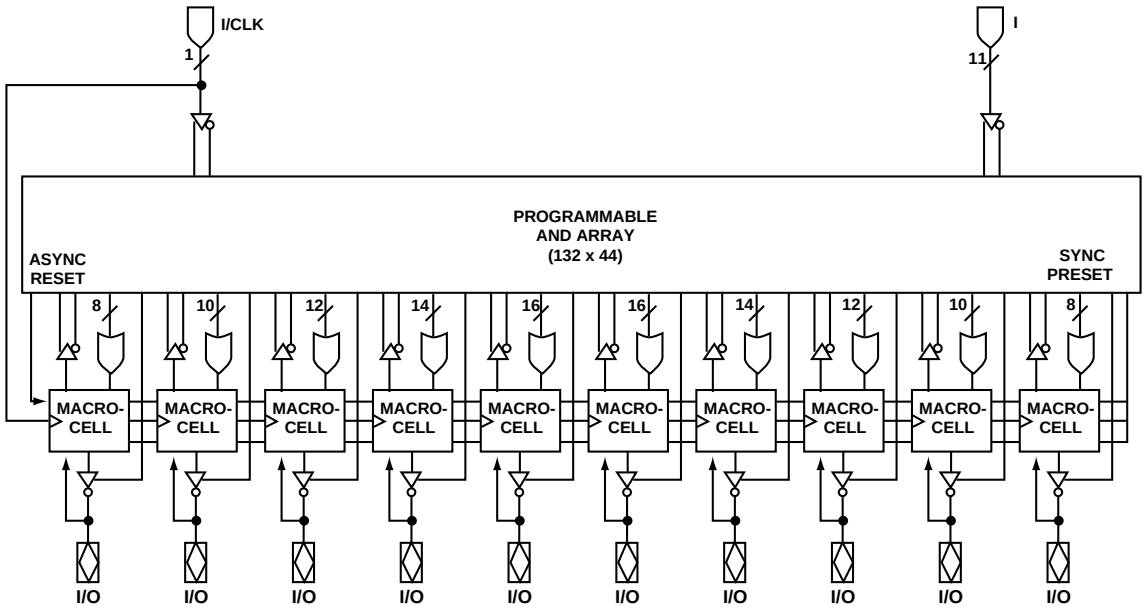

Source: AMD

18487

Figure 1-24. PAL Block Diagram

b. PGAs

In the late 1980's, new devices called Programmable Gate Arrays (PGAs) were introduced. These structures use proprietary internal architectures, and can be thought of as PALs with very small input AND and OR arrays, but with a very large quantity of output and input macros. An example of this is shown in Figure 1-25.

Although the parts don't actually resemble an array of gates as the name implies, the user designs the circuit as if it is. The "macros" are programmed by the vendor's software to implement the user's requirements. The complex part of the system is the software, as it must optimize the use of the macros very carefully to fully utilize the IC chip. The metallization on the chip also tends to be very complex. It must not only allow the macros to be connected to one another with as many paths as possible, but must also supply paths for the programming logic.

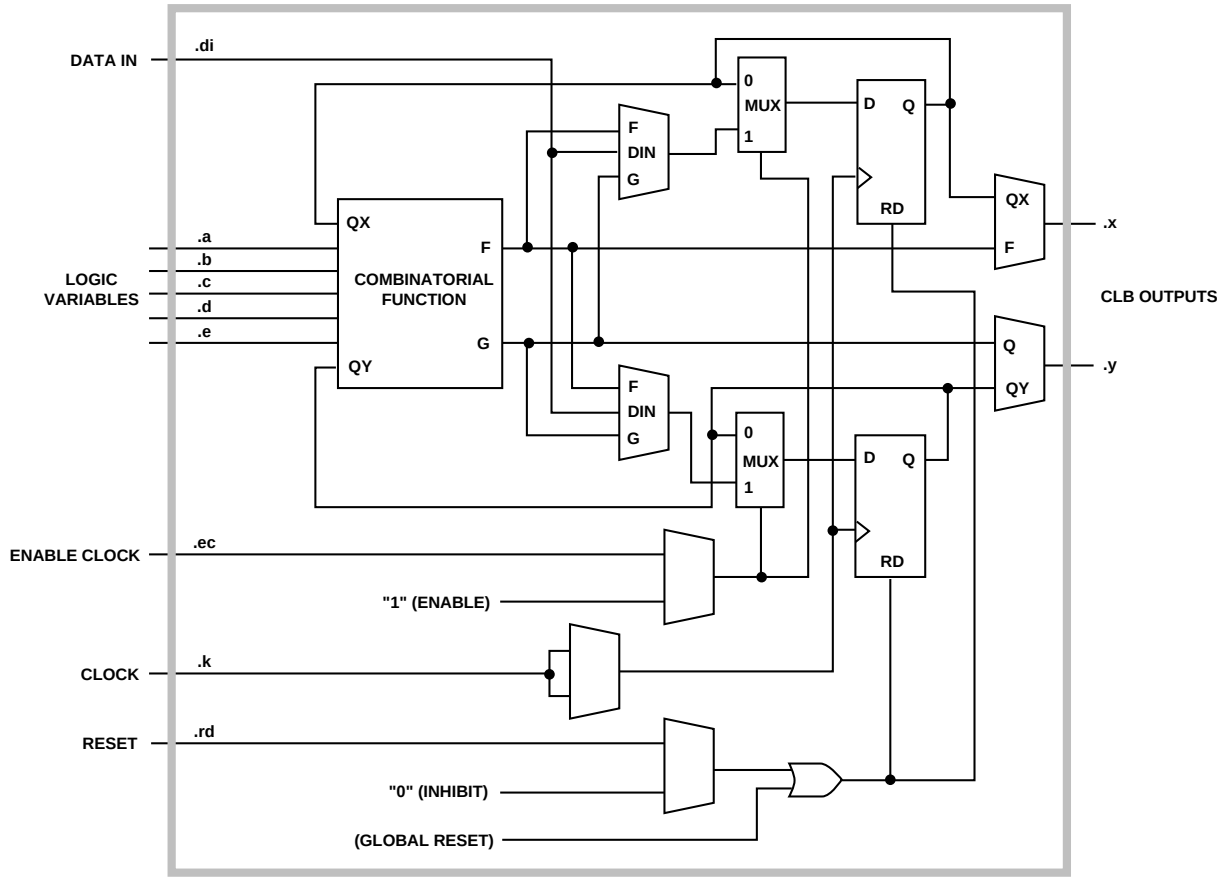

Source: Xilinx

18488

Figure 1-25. Configurable Logic Block

The storage of the programming for each macro can be stored in a number of different cells. Flash EEPROM, EPROM, SRAM, and anti-fuse cells are all used for storage of the program configuration. The SRAM is the only storage method requiring the program to be temporarily stored in another chip or on disk, as the SRAM cells are volatile and lose their configuration when the power is removed.

Anti-fuses are special layers or films that change electrical properties when pulsed with some electrical stimuli, such as a voltage pulse. In most cases the original state of the film is a non-conductor. After the electrical stimulation, the film becomes an electrical conductor. Since this is the opposite effect to that of a metallic fuse, the term anti-fuse is used to describe the structure. In nearly all cases, the structure acts as a resistor.

E. ANALOG

Analog (or linear — the names tend to be used interchangeably) is a general category of ICs that includes all ICs that do not fit into one of the digital IC categories. In most cases, the ICs are designed to operate at higher voltages than the standard digital familes, and in some cases they do perform analog functions, as described below.

1. Amplifiers

Amplifiers are circuits that are a combination of active devices (such as transistors and diodes) and passive devices (capacitors, resistors and inductors) that are designed to increase electrical signals from a low level to a high level. The circuits in a radio or TV that increase the signal level from that transmitted over the airwaves to a signal detectable by the human ear are audio amplifiers, i.e., they amplify audible electrical signals. Figure 1-26 shows an amplifier with input and output signals, both with low and high amplification settings.

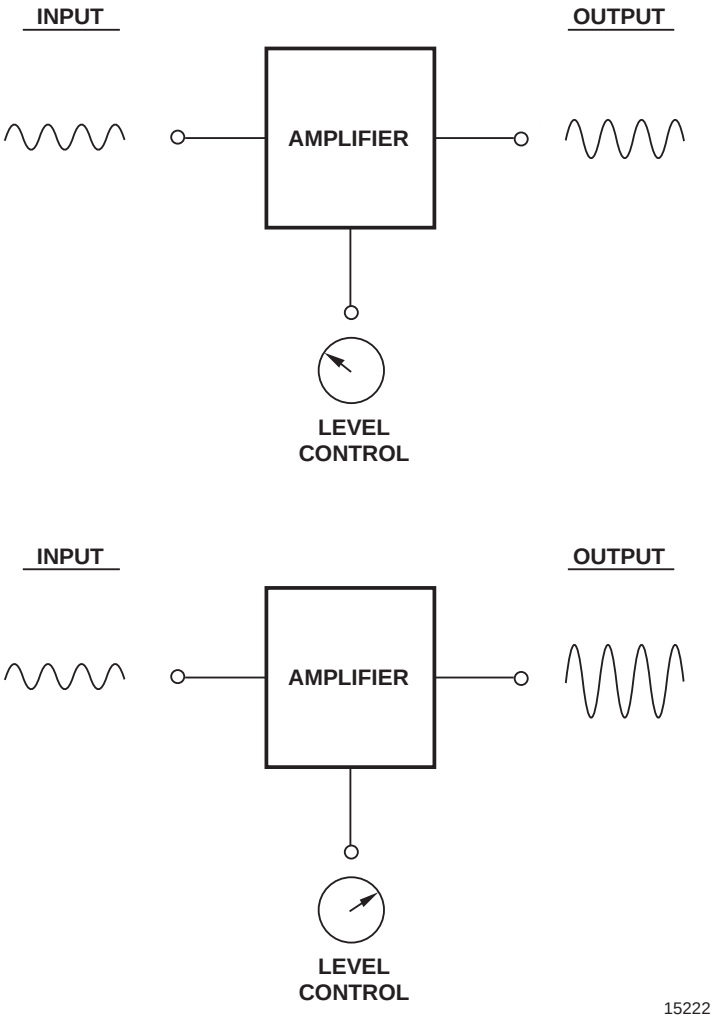

Figure 1-26. Analog Example

Many industrial applications use general purpose amplifier chips that are called operational amplifiers, a term used to denote a general purpose amplifier with infinite gain (the amount of amplification available) and very high input resistance (causing little or no effect on the signal connected to the amplifier input). Figure 1-27 shows a small amplifier schematic.

Although a simple amplifier could be built with three of four transistors, the design would not be very practical. Semiconductor devices change dramatically with changes in temperature, and these changes must be nullified with compensation circuits. This accounts for the large number of devices in the figure.

2. Voltage Regulators

Voltage regulators are circuits designed to deliver very precise electrical voltages to other ICs in a system. The input voltage to a voltage regulator may vary over 25 percent, while the output is controlled to less than one percent. There are thousands of different designs used for the various types of equipment in use. The key to high-quality voltage regulators is the ability to tolerate large fluctuations on the input, while maintaining tight control on the output, and to do so over a wide temperature range.

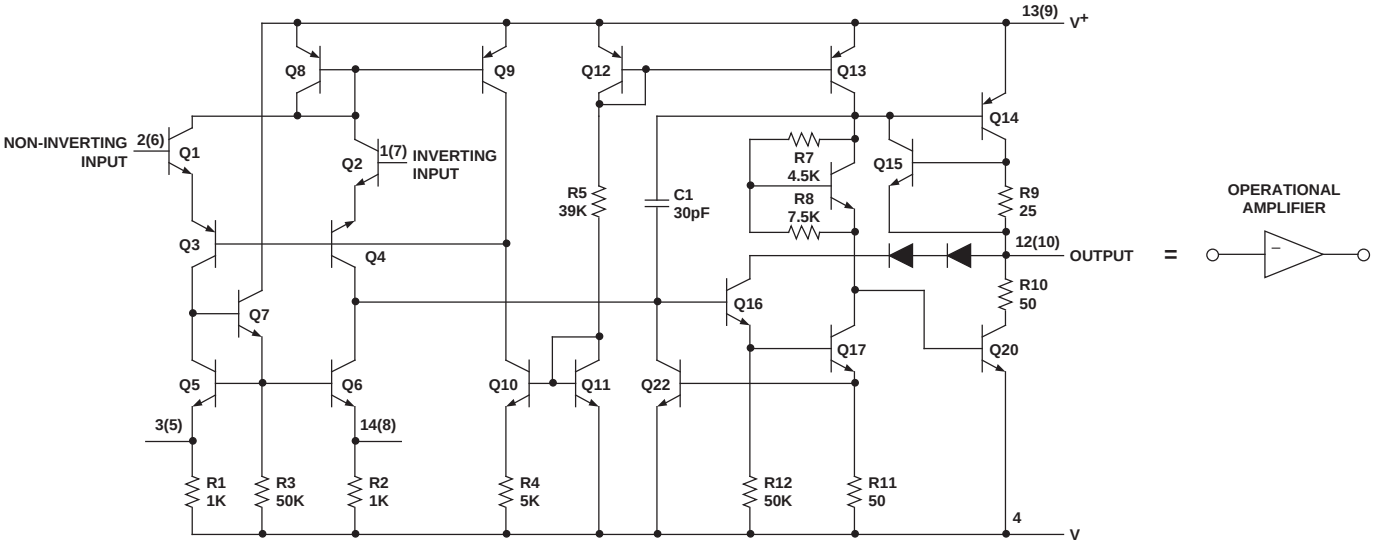

Note: Numbers in parentheses are pin number for amplifier B, DIP only.

GAIN	200,000
INPUT DRIFT	1 X 10^{-6} VOLTS/°C
MAXIMUM VOLTAGE	±22 VOLTS
OUTPUT = INPUT ±0.01%	

15224

Figure 1-27. Precision – Linear

3. Data Conversion

Data conversion circuits are those that translate from analog signals to digital signals, and back again. An example of this is the typical telephone call between cities. The voice (an analog signal) is converted to a digital signal at the local telephone office. The digital signals are sent from the local office to the office near the person called. The digital signals are then converted back to analog signals for the person called to hear.

Figure 1-28 shows a typical voice or music waveshape. This signal is converted to digital data by sampling the waveshape at specific intervals. These sample points and values are shown as bars and numbers on the figure. This "digitization" function is performed by an analog-to-digital (A / D) converter.

The digital values are transmitted, in order, along the transmission media, wire, optical fiber, or radio link, to the receiving station. At the receiving station the digital data is sent through a digital-to-analog (D / A or DAC) converter to recreate the analog signal. Figure 1-29 shows the relative functions of these devices.

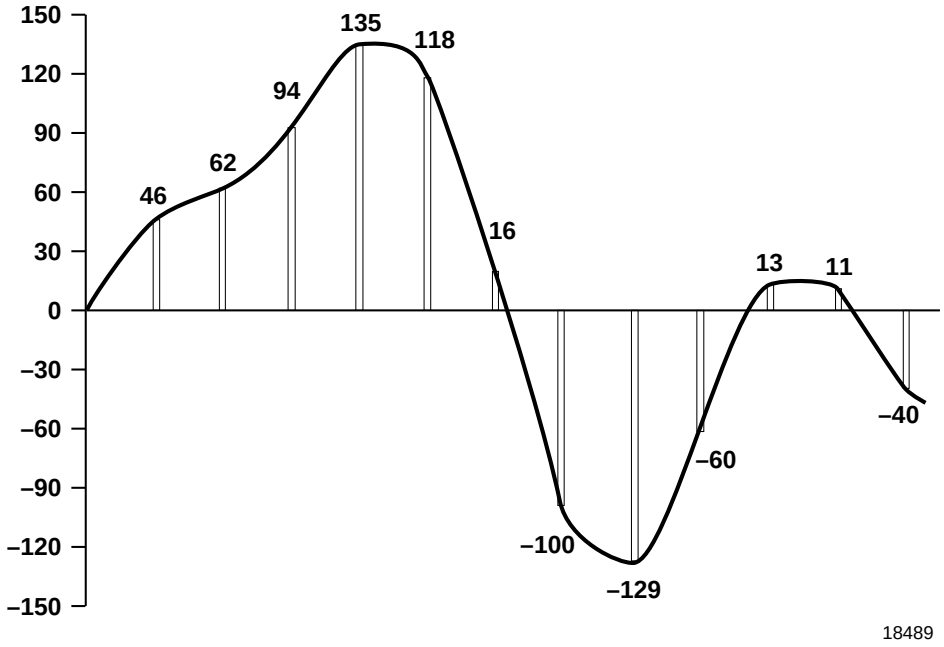

Figure 1-28. Sampled Analog Waveshape

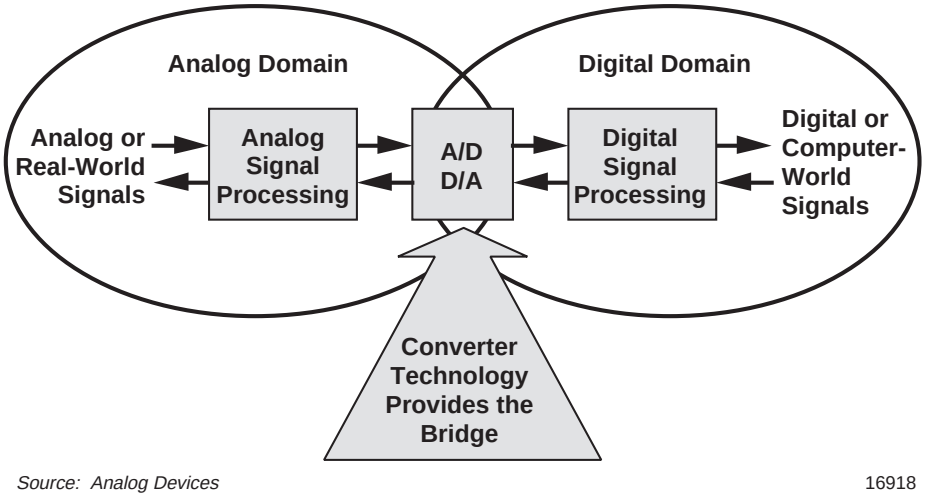

Figure 1-29. Real-World Signal Processing

A special version of these converters is the CODEC found in telephone systems. These ICs contain both A/D and D/A converters, so they support bidirectional communications. The devices follow very special rules in conversion so various brands of equipment are able to communicate.

The advantage of the digital transmission method for long distance (versus the 1950's vintage "all analog" transmission) is the quality of the signal at the receiving end. Analog signals degrade (become lower) over distance, but the noise levels tend to stay constant. Therefore at the receiving end the signal becomes mostly noise. When digital data is transmitted over long distances, the power of the signal degrades, but the "ones" remain distinguishable from the "zeros," so a noise-free version of the starting signal can be reconstructed at the receiving end.

4. DSP

Digital Signal Processor (DSP) ICs are special microprocessors designed for very high speed mathematical computations. The applications are those where the inputs and outputs are both analog, but where the processing can be digital or analog. Historically, the processing was performed by analog components because of cost. However, as the cost of digital processing has dropped, many applications began using DSP chips because of the improved performance.

In these applications, the "real world" analog signals are converted to digital data by an A/D converter. The digital information is acted on by the DSP chip and converted back to analog by a D/A converter. Figure 1-30 shows the DSP chip with the two converters and the program and data memory.

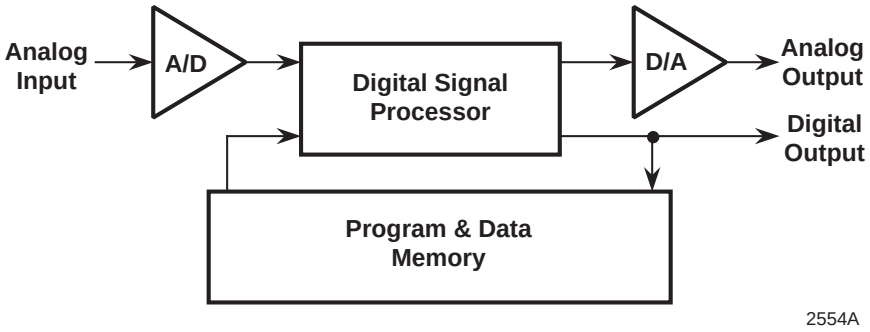

Figure 1-30. Digital Signal Processing

5. Interface Circuits

Interface circuits are a very general category of ICs that are fully digital, but act as translators between two or more signal levels. For example, most computers switch between 0 volts and +5 volts internally, logic zero and logic one, respectively. To communicate with serial printers using RS-232C interface, the signals must be changed to a signal varying between +12 volts and -12 volts. A single IC accomplishes this translation.

Subscriber Line Interface Circuits (SLICs) are used in telephone switches (at the phone company or local PBX) to connect to the telephone lines (actual wires) from the telephone sets. These devices perform the interface functions such as recognizing when the telephone user has taken the phone off-hook and wants a dial tone, passing the dial tone to the user until the first digit is dialed, passing the ringing or busy signal to the user, etc.

These ICs must be able to withstand a significant amount of "abuse" in the form of electrical transient signals, poor wiring, etc., where signal levels are in the 50-volt range, and also communicate on a logic level with the electronic circuits in the telephone switch, usually between 0 volts and +5 volts.

6. Analog Arrays

Analog arrays are an analog version of digital gate arrays, or as close as the idea can get. The ICs are groups of discrete transistors, resistors and capacitors that can be interconnected with metallization to form some type of analog or mixed-mode circuit. These circuits require the customer (or whoever designs the final circuit) to understand linear design. The advantage over standard analog is that special custom requirements can be designed into the circuit, without requiring a full custom design. The disadvantage is that the array approach uses fixed component sizes, therefore limiting the range of performance compared to a full custom approach. Figure 1-31 shows an analog array chip. The list in Figure 1-32 shows the various components on the chip available to the designer.

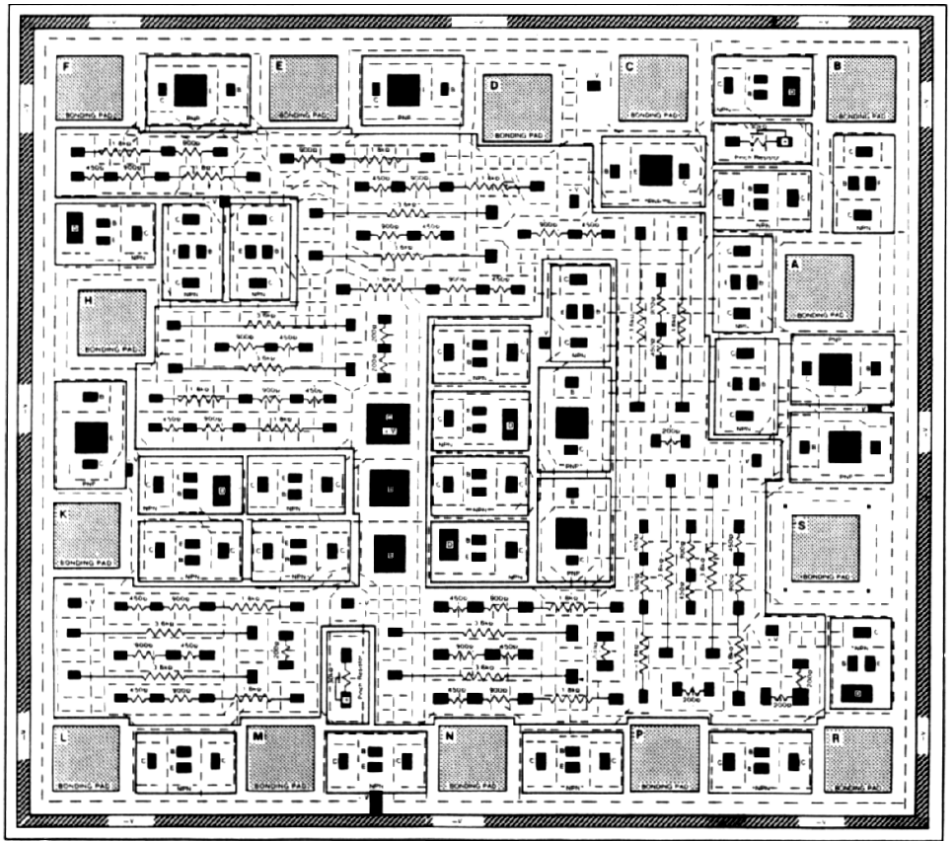

Source: Exar 18490

Figure 1-31. Analog Array

Total Component	110
Bonding Pads	14
Operating Voltage	20V
NPN Transistors	23
PNP Transistors	8
Schottky Diodes	6
Resistors:	
200Ω	8
450Ω	18
1,800Ω	20
3,600Ω	12
30,000Ω	2

18491

Figure 1-32. Analog Array Component List

7. Mixed-Mode

Mixed-mode ICs are special cases of full custom or standard cell ASIC products that include both digital and linear circuitry on the same chip. These chips are becoming popular as VLSI technology moves towards single-chip solutions.

The process technology for mixed-mode ICs requires certain compromises. The performance of a single mixed-mode chip is usually less than the equivalent function performed by an analog chip and a digital chip working together. However, the cost of a single chip is usually much less than the two-chip approach.

F. SYSTEM TERMS

1. LANs and Ethernet

Local Area Networks (LANs) are used to connect computers together. The typical wiring for a LAN differs from the old method of connecting each terminal or small computer to a host (usually a mainframe or minicomputer). The LAN allows any individual device connected to the LAN to communicate with any other device, without going through a central computer. A good example of the power of a LAN is the sharing of expensive high-speed printers.

Ethernet is a specific protocol, or set of rules, for both the software and hardware that is used to create a LAN. These rules were established in the mid-1970's by Xerox, DEC, and Intel. The rules formed a standard, enabling different companies to supply equipment that would be able to send data back and forth.

2. Displays - Active Matrix LCDs

Liquid Crystal Devices (LCDs) are visual display devices that operate by blocking the transmission of polarized light. The liquid crystal material itself is a very thin liquid film sandwiched between two glass plates. When an electric field is applied to the liquid, the liquid changes from an orderly pattern that twists the polarized light to a different orderly pattern that causes the light to be absorbed.

These displays are used in simple applications such as calculators, watches, instrument displays, etc. In these applications, each segment (usually seven segments are used for the ten numbers) is connected directly to the IC chip that applies the full battery voltage or zero volts to activate the display.

Other applications, such as replacements for Cathode Ray Tubes (CRTs) in TV sets or computer displays, are more difficult for LCDs. In these cases the voltage used to drive the LCD on or off is shared with a portion of the cells in the same rows and columns as the cell being addressed.

To make this address mechanism more effective, a transistor is needed for each pixel of the display, and it must be located on the glass of the display. These displays are referred to as Active Matrix LCDs. Silicon is deposited on the glass and thin-film transistors are manufactured in this material. The contrast of these displays is considerably better than regular LCDs, but the manufacturing cost is considerably higher.