

The Journey InsideSM: Microprocessors Background Information, Part 1

The Chip at the Heart of a Computer

The word “chip” is commonly used to refer to an integrated circuit. Microprocessors are one of many types of integrated circuits. The microprocessor is central to the functioning of a computer. It is used to process information by keeping all the parts of the computer working together to complete a given task. The microprocessor is the most complex and expensive of the chips needed to make a computer work.

Today's microprocessors may contain tens of millions of transistors and other electronic components, yet be smaller than your thumbnail. Today's microprocessors work very rapidly and are very reliable.

Microprocessor functions can be described as a simplified three-step process. The fetch step involves getting an instruction from the computer's memory. The decode step involves deciding what the instruction means. The execute step involves carrying out the instruction. A modern microprocessor can complete this three-step process many millions of times in a single second. Special areas on the microprocessor are designed to enable the computer to complete these three steps.

Each type of microprocessor has its own design and organizational scheme. However, all microprocessors must perform similar tasks. The transistors—the building blocks of the microprocessor—are arranged in circuit groups. Each circuit group is designed to perform a specific task. The following circuit groups are needed to allow the microprocessor to perform the fetch, decode, and execute operations:

Arithmetic Logic Unit: A set of circuits dedicated to numerical calculations and logical operations

Control Unit: Circuits that control the sequence of the processing done by the microprocessor and hold the instructions in the correct sequence until needed

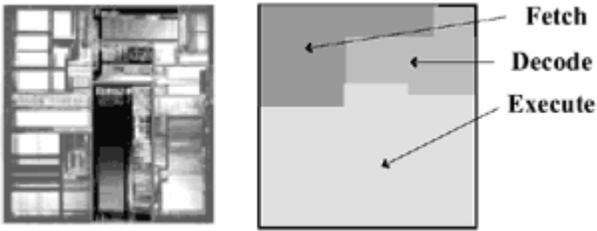
Decode Unit: Circuits used to translate the instructions into control signals and directions, and hold them in sequence until they are requested by the control unit

External Bus Unit: Pathway provided to move data to and from the microprocessor

Internal Bus Unit: Circuits that manage the flow of information from one part of the microprocessor to another

Memory Management Unit: Circuits designed to keep track of information that has been processed, reclaim the space when the instructions are done, and help organize the way instructions and data are stored

Because every microprocessor uses the fetch, decode, and execute cycle, the internal circuit groups of most microprocessors perform very similar functions. The special circuit areas work together to accomplish the three steps. The general areas used for fetch, decode, and execute on an Intel® Pentium® III microprocessor are located as indicated in the diagram.



The surface of a Pentium® chip with the special areas for fetch, decode, and execute indicated.

Each instruction that a microcomputer can execute is quite simple. For example, a typical instruction that the microprocessor can understand is "add two numbers." Another instruction is "compare two numbers to see if one number is larger than the other." A typical microprocessor has about 150 different built-in instructions with the necessary circuitry to decode and execute them.

Microprocessors neither think nor reason. They simply follow the instructions given to them by software programmers. The smallest error in such a set of instructions can lead to large errors in the final results produced by the computer.

Programmers try hard to avoid errors in the software they write. However, it is difficult to write long programs without making mistakes. In large software companies, many programmers are employed just to test the software that others write. These programmers work to find any errors in the software.

An error in a computer program or computer hardware is called a bug. The process of detecting and correcting an error in a computer program is called debugging the program. Computer programmers spend a lot of time testing and debugging the programs they write.

The Journey InsideSM: Microprocessors Background Information, Part 2

Creating Chips

There are two major phases to creating computer chips: design and fabrication. Chips are designed to perform specific tasks. When a totally new type of chip is designed, a prototype is made and extensive testing on the prototype is done. There may be several revisions of chip design before a chip is produced in quantity.

When a new chip moves from the design stage to the fabrication stage, problems may arise that were not evident earlier. A procedure that worked well for producing a small quantity of chips may need to be modified to produce thousands at a time.

Companies that design and fabricate chips continue to work toward developing more and more powerful chips. To increase the chip's ability to handle greater amounts of information at once, more and more transistors are placed into the circuits. To increase the chip's speed in moving data from one place in the chip to another, the size of the chip is decreasing. For each specific chip being created, many steps must be completed before a working chip is produced.

Designing Chips

Designing one of today's powerful microprocessors involves the efforts of up to 600 engineers. The design process includes four distinct teams, each focused on a different aspect of chip design. After the customer and the chip manufacturing company have begun the negotiations on what the new chip needs to do, many more people get involved. The design process moves in sequence from chip architects, to logic designers, to circuit designers, and finally, to layout designers.

The **design architects** begin the process. They make decisions about what the chip must do—perhaps store data, operate a remote control for a VCR, or provide the information needed to get a computer up and running. The customer requesting the chip is consulted to make sure that the architect's decisions satisfy their needs.



The team reviews the design for the chip one more time.
Photo: Ruth Carranza Productions-taken from Silicon Run Lite [videotape].

The placement of each component is extremely critical because of the small size of the chip. The layout designers determine which portions of the chip carry circuits that must interact with each other. Circuits that work on the same information in sequence need to be placed close together. Circuits for each specific task must be able to function without interfering with any other circuit. There must be enough circuits for temporary storage so that processing is not slowed down by a wait for instructions or data.

The Journey InsideSM: Microprocessors Background Information, Part 3

Chip Design and City Planning

The engineers designing chips face many of the same decisions that a city planner does when designing city neighborhoods. How much land should be set aside for residential, commercial, and civic uses? How big can the city become and still be manageable? Where should the downtown area be located? How many schools are needed, and where should they be located? How many police and fire stations are required, and where should they be located? Each type of building needs to have water, gas, electricity, telephone, and cable. Where are each of these items to be placed, and what is the best routing to make sure that every place receives the services it needs?



A well-planned city viewed from the air.

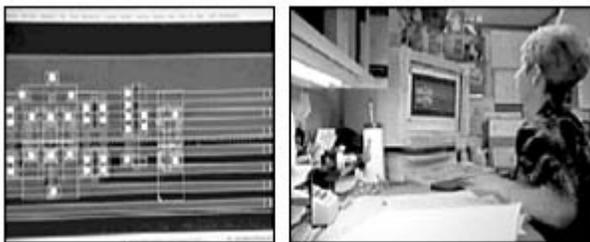
"Every succeeding generation of microprocessors represents a new, more difficult challenge," according to Sunil Shenoy, a CPU engineering manager for Intel. Consider the complexity of the Intel® Pentium® 4 processor. It has 40 million transistors squeezed into the space of a nickel. The engineering team has to design both the processor and the tools and methods to produce it, as well. It's similar to having to design and develop not just a new sophisticated jet aircraft, but also the metallurgy, computer-aided design terminals, and design software required to make it. Each generation of microprocessors requires a monumental effort. But the end result is always worth it. A modern microprocessor is a beautiful tapestry composed of millions of circuits, each doing its part to direct the electronic synapses powering the high-powered computational devices we use today.

From Design to Manufacture

Once the layout designers have a schematic drawing of the planned circuits, computer aided design (CAD) software gives them a way to test each component and the timing of pieces that work together. The software can determine if the conductors are too narrow or too close together. The software can determine if the insulation layers are thick and wide enough to prevent electricity going where it shouldn't. The software can determine if the transistors are spaced far enough apart to prevent them from influencing each other inappropriately.

The circuits are modified and tested to ensure an optimal, error-free design. Even though the software is efficient at many of the decisions to be made as the circuits are finalized, software cannot substitute for many of the skills provided by the designers. Humans are still the most efficient decision makers when deciding which part of the chip is allocated to each task so that the design is completed in the most space-efficient manner.

After the design is complete, the schematics can be translated into the actual masks or templates that are needed to arrange the components on the chip. The masks provide the necessary pattern for placing each layer of material needed to complete the circuit components.



The mask designer works to translate the circuit schematics into masks. On the left, a close-up of her computer screen. Photos: Ruth Carranza Productions-taken from Silicon Run Lite [videotape].

Before the chip design is released for full-scale manufacturing, several additional steps are needed. A prototype is created using the same steps that will be taken once the chip is produced in volume. About 20 wafers of the chip will be produced for testing.

The prototype chips are thoroughly tested. Do they work as intended? If not, how does the design get corrected? Does it run as fast as planned, or do some circuit placements need to be changed? Are the electrical signals going in and out of the chip properly? Does it work correctly at high and low temperatures? Does it work correctly at various voltages?

If any problems are identified, the chip is sent back through the design steps until the group that makes the corrections sends the chip on to the next design group. Recurring cycles of improvement are critical to putting a perfect chip design into place before full-scale fabrication begins.

The Journey InsideSM: Microprocessors Background Information, Part 4

Fabricating Chips

Many chips are fabricated at a time on a wafer—a thin slice of silicon crystal. Silicon is purified to 99.9999 percent, and then made into a large crystal about 8 inches in diameter and 5 feet in length. This ingot is then sliced to make the wafers.



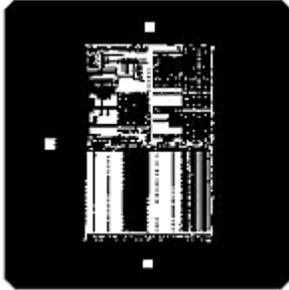
An ingot being prepared to be sliced into wafers (Photo: Ruth Carranzo Productions).

The chip-fabrication process involves many steps, building the chips layer by layer onto the silicon wafer. The processes below occur several times during the manufacturing of a chip. The order and the number of processes used to create each layer depend on the design of the particular chip being fabricated.

Oxidation is the combination of oxygen with another substance. In the production of chips, oxygen reacts with silicon to produce an oxidized silicon layer. The wafers are put into a sterile furnace with steam or pure oxygen and then heated. The temperature of the furnace and the time the wafer stays in the furnace control the growth of the oxidized layer.

Deposition is the procedure used to place thin layers of new material on the top of the wafer. These layers are known as thin films because they are less than 1/1,000,000th of a meter (1 micron) thick. The wafers are placed in a reaction chamber and chemicals are introduced which react and deposit on the wafer surface. Films used include metals, semiconductors, and insulators. The films are patterned using a process called photolithography.

Photolithography means to write with light. Photolithography is used to place a design on a layer of the wafer. When the circuit layer is designed, the design is transferred to a mask or template. There are one or more masks for each layer of the wafer.



A sample mask, one of many used in fabricating a chip

First, a thin film of liquid material called photoresist is put on the wafer. The wafer is then baked to dry and harden the photoresist. A mask is then placed over the baked wafer. When an intense light beam is focused on the mask, the photoresist material that is not covered by the mask is exposed to the light. This step is similar to light going into a camera and exposing the film. The wafer then goes into a solution that washes away the exposed photoresist, leaving a pattern identical to the pattern on the mask.

Implant is used to introduce chemical impurities into specific locations on the wafer. In chip fabrication, impurities must be added to the silicon so electricity can flow through the circuits. The impurities produce n-type silicon or p-type silicon. The desired impurity atom is accelerated to a high speed in an ion implanter and is driven into the wafer surface.

Diffusion is defined as the spreading of one substance through another. Impurities that were implanted into the wafer surface must diffuse deeper in order to make the circuit operate properly. High-temperature furnaces are used to diffuse the impurities to the proper depth. The diffusion is controlled by the length of time the wafer is heated and the temperature used.

Etching removes material not covered by the photoresist. After the photoresist is patterned in the photolithography process, the exposed portion of the layer is removed. This is done by chemical, wet etching, or by plasma etching. In wet etching, the exposed layer is dissolved when the wafer is immersed in a chemical bath, usually an acid or a solvent. In plasma etching, the wafers are placed in a plasma reactor which bombards the wafer with a high-powered beam of ions. This removes the material that was exposed during photolithography. This second method is usually preferred because it is more easily controlled and allows circuit elements to be placed extremely close together.



A worker examines the surface of a wafer.

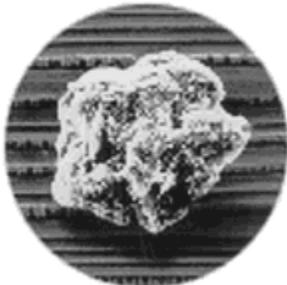
During the fabrication of the chips, the wafers move through many steps using the processes outlined above. A chip, such as the Pentium® microprocessor, may have 20 layers before it is complete. Many layers take multiple steps to produce. At each step, great care is exercised to make sure the masks are transferred precisely and the material is placed as required. Testing is done at several points in the fabrication of the chip. The smallest mistake, such as a slight misalignment of a mask, can destroy some or all of the chips on a wafer. From start to finish, it takes about two months to complete the job.

The Journey InsideSM: Microprocessors Background Information, Part 5

The Clean Room

Quality control is a critical part of the fabrication process, and testing is done frequently. When a wafer is completed, each chip is carefully put through a number of special tests. Any chips that fail these tests are discarded.

One reason chips fail is the damage caused by tiny contaminants in the clean room. There are many possible sources of contamination. Dust in the air, small pieces of hair, or flakes of skin from workers' bodies are all possible sources of contamination. When compared to the components on a chip, which are less than a micron in size, even tiny specks of pollen or dust seem like huge boulders.



A flake of dust ruins the surface of a chip.

The extremely clean area where chips are made is called the clean room. Clean rooms are designed to protect against such contamination.

Within a clean room, fresh air takes on a new meaning. The air we breathe is filled with contaminants. One cubic foot of air normally contains about 15 million dust specks that can potentially harm a wafer during the fabrication process. In the clean room, almost all such dust particles are filtered out.

The air pressure inside the clean room is kept slightly higher than the pressure outside. This difference in pressure prevents dirty air from entering, because clean room air tends to push out through any openings. Air enters the clean room through ceiling filters and is removed through small holes in the floor. This air flow keeps unwanted particles beneath the working space. The air in a typical processing area is extremely clean, since it is replaced and filtered at least seven times every minute.

Researchers have also investigated the best lighting conditions for chip fabrication. Because some of the fabrication steps are done by controlling the exposure of the chip surface to a special light, the choice of light is critical. The photolithography area uses a soft yellow light rather than the fluorescent lighting used in most manufacturing sites. Research has shown that this type of light does not interfere with the steps needed to place the circuits onto the chips.

No one is allowed to simply walk into a clean room. Each worker completes approximately 50 steps before entering. Even a freshly washed set of street clothes is not clean enough. The workers wear special clothing known as bunny suits over their street clothes. The fabric has been specially designed to be lightweight, yet prevent any particles from moving from the worker's body into the room. No part of a worker's clothes or body is left exposed.



Clean room workers dressed in their bunny suits.

The bunny suit even includes an air filter system. Exhaled air is drawn through a hose and filtered through a small unit attached to the bunny suit belt. This system eliminates the possibility of workers releasing impurities into the air when they exhale. In a short conversation between two people, about 300 drops of saliva are released into the air. Even though these particles are too small to be seen by the naked eye, they are large enough to cause serious damage to the surface of a chip.

Resources

The first five resources are annotated and are appropriate for both teachers and students:

Bodanis, D. (1995, April). "It's in the Air: Skin, Stardust, Radio Waves, Vitamins, Spider Legs." *Smithsonian*.

This article describes in detail the cubic foot of air immediately in front of your face and provides an excellent perspective on why clean rooms are needed for successful chip manufacture.

Carranza, R. (Producer). (1997). *Silicon Run Lite* [videotape]. (Ruth Carranza Productions, PO Box 391025 Mountain View, CA 94039)

This videotape provides an excellent presentation of the fabrication steps used in turning ordinary sand into computer chips—suitable for teachers and older students.

Hassig, L. (Ed.). (1990). *The Chipmakers*, (rev. ed.) Richmond, VA: Time-Life Books.

An excellent collection of graphics and text provide an understandable explanation of the complex ideas to be considered in the design and manufacture of chips.

Malone, M. (1995). *The Microprocessor: A Biography*. Santa Clara, CA: Springer-Verlag.

A readable history of the invention of the microprocessor which also provides a rare glimpse of the people instrumental in making the microprocessor business so successful— teachers and older students.

Wyant, G. & Hammerstrom, T. (1994). *How Microprocessors Work*. Emeryville, CA: Ziff-Davis Press.

This book provides an excellent presentation of the microprocessor: its evolution, function, architecture, manufacture, and possible future are all addressed.

Carranza, R. (Producer). (1996). *Silicon Run I* (2nd ed.) [videotape]. (Ruth Carranza Productions, PO Box 391025 Mountain View, CA 94039)

Carranza, R. (Producer). (1993). *Silicon Run II: The Sequel* [videotape]. (Ruth Carranza Productions, PO Box 391025 Mountain View, CA 94039)

Evans, A. (1996). *Basic Digital Electronics*. Richardson, TX: Master Publishing, Inc.

Flaherty, T. (Ed.). (1993). *Memory and Storage*. Richmond, VA: Time-Life Books.

Hassig, L. (Ed.). (1990). *Input/Output* (rev. ed.). Richmond, VA: Time-Life Books.

Malone, M. (1995). *The Microprocessor: A Biography*. Santa Clara, CA: Springer-Verlag
Studio Interactive (Producer). (1996). Digital Lab [CD-ROM, workbench and components].
CyberCrafts: Hands on Learning. Philips Media.