

*Address Conflict.* We use the term *address conflict* to describe a bus error that results when interfaces are misconfigured so they respond to the same address. Most bus protocols include a test for address conflicts — if two or more interfaces attempt to respond to a given request, the bus hardware detects the problem, and sets a control line to indicate that an error occurred. When it uses a bus, the processor hardware checks the control lines to determine whether an error occurred, and if so, reports the error.

*Unassigned Address.* An *unassigned address* bus error occurs if a processor attempts to access an address that has not been configured into any interface. To detect an unassigned address, most bus protocols use a *timeout* mechanism — after sending a request over the bus, the processor starts a timer. If no interface responds, the timer expires, which causes the processor hardware to report the bus error<sup>†</sup>.

## 14.18 Address Configuration And Sockets

It may seem that the easiest way to prevent most bus errors consists of installing enough memory to cover all possible addresses. Then, when the computer boots, arrange for the processor to test all memory locations. In practice, however, architects usually design a memory bus to accommodate expansion. That is, a bus typically contains enough wires to accommodate more physical memory than is installed in the computer. Second, we will learn that devices other than memory can attach to a bus, and many devices do not occupy contiguous bus addresses.

Fortunately, architects have devised a scheme that helps avoid memory configuration problems: *special sockets*. The idea is straightforward. Memory is manufactured on small printed circuits that each plug into a socket on the mother board. To avoid problems caused by misconfiguration, all memory boards are identical, and no configuration is required before a board is plugged in. Instead, circuitry and wiring is added to the mother board so that the first socket only receives requests for address 0 through  $K-1$ , the second socket only receives requests for address  $K$  through  $2K-1$ , and so on. The point is:

*To avoid memory configuration problems, architects can place memory on small circuit boards that each plug into a socket on the mother board. An owner can install memory without configuring the hardware because each socket is configured with the range of addresses to which the memory should respond.*

As an alternative, some computers contain sophisticated circuitry that allows the MMU to configure socket addresses when the computer boots. The MMU determines which sockets are populated, and assigns each a range of addresses. Although it adds cost, the extra circuitry provides additional flexibility because an owner can place memory in any socket rather than filling sockets in a particular order.

---

<sup>†</sup>The timeout mechanism also detects malfunctioning hardware (i.e., a memory that is not responding to requests).