Dynamic RAM

Larger microcomputer systems use Dynamic RAM (DRAM) rather than Static RAM (SRAM) because of its lower cost per bit. DRAMs require more complex interface circuitry because of their multiplexed address bus and because of the need to refresh each memory cell periodically. DRAMs are covered in section 7.2 of the text.

DRAM Structure

As with SRAMs, the DRAM memory array is implemented as a square matrix of rows and columns. Half of the address bits select a row and half select a column.

In a dynamic RAM each memory cell is formed from a capacitor and an associated transistor that is used to connect the capacitor onto a line shared by all the memory cells in that column. A dynamic RAM memory cell is therefore much smaller than an SRAM cell which needs at least two gates to implement a flip-flop.

Address Multiplexing

A feature of DRAM chips is that they multiplex the row and column addresses onto the same set of pins in order to reduce the number of pins required and thus the size of the chip.

Two strobes, RAS* (row address strobe) and CAS* (column address strobe) are used to tell the chip which part of the address is currently on the address pins. To further reduce the chip count the CAS* signal also acts as an output enable when R/W* line is high (read).

Exercise: How many address pins would be found on a 16 Mbit DRAM?

Refresh

Since the DRAM storage capacitor discharges over time it must be refreshed periodically. The DRAM's circuitry recharges the memory cell every time it is read. Since all cells in a row are read or written at the same time, it is only necessary to periodically cycle through all of the *row* addresses to refresh all of the bits in the memory array.

The simplest technique to provide DRAM refresh is to include a device (such as a DMA controller or video display circuit) that accesses the RAM in such a way that all of the rows are accessed at least once during the minimum refresh time (typically every few tens of milliseconds).

Another technique is to add a circuit that periodically forces the DRAM to refresh itself. For most DRAMs asserting CAS* before RAS* causes an internal refresh operation using an on-chip circuit that cycles through the row addresses.

Exercise: Assume a microprocessor with a 200 ns memory cycle time is using a 1 MByte DRAM with a maximum refresh time of 10 ms. How many row addresses will have to be refreshed every 10ms? What is the approximate time between each refresh cycle? How many memory cycles are there per refresh cycle? What percentage of the memory accesses are "wasted" on refresh cycles?

Comparing DRAM and SRAM

Since the SRAM devices require more circuitry per memory element they are more expensive (per bit) to produce and have lower density per chip. The typical memory density ratio between DRAM and SRAM for the same size chip is about 4 to 1.

The disadvantages of DRAMs are that they require additional control circuits to multiplex address lines and to handle refresh. If DRAMs are organized as bit-wide devices it is necessary to use a number of devices that is a multiple of the the data bus width (8, 16 or 32) in a system.

The use of large DRAM arrays in which the CPU address and data buses must drive many chips usually requires buffers for the address and data lines. Because of these reasons DRAMs are mainly used in systems that require large memories and SRAMs are mainly used in smaller systems.

The fastest RAM designs are static, so SRAMs are often used for high-speed memories such as cache or address translation tables.

CMOS SRAMs consume very little power when not being accessed so they are often used in low-power designs. With the use of battery (or a large capacitor) backup they can retain their contents for months. On the other hand, since DRAMs must be continuously refreshed, their power consumption cannot be reduced to very low values.

Due to the larger number of bits per chip and wider bus sizes, DRAMs are now being offered in nybble (4-bit) and larger organizations. The number of extra pins required to provide the additional data bits is less of a concern with modern high density packages such as QFP and SOIC.

Exercise: Consider a system using 16 Mbit X1 memories to design a memory array for a microprocessor system with a 32-bit data bus. What is the minimum amount of memory that could be provided using these devices?

DRAM Timing

The sequence of operations required to read from a DRAM is as follows:

- set R/W* high
- place the row address (the MS half of the address) on the address pins
- bring RAS* low
- place the column address on the address pins
- bring CAS* low
- read the data from the data pin

The sequence of operations required to write from a DRAM is as follows:

- set R/W* low
- place the row address (the MS half of the address) on the address pins
- bring RAS* low
- place the data on the data pin
- place the column address on the address pins
- bring CAS* low

Figure 7.11 shows the simplified read timing requirements of a DRAM.

RAS* must be asserted at least 25 ns before CAS*. Data becomes valid at least 25 ns after CAS* and 100 ns after RAS*. This last parameter is the often-quoted access time. Other parameters give the times for the tri-state outputs to switch between output and high-impedance modes.