

Memory 3 – The "RAMs"

- by Stan Kaplan, WB9RQR
105 Martin Drive
Port Washington, WI 53074-9654
(414) 284-9346
skaplan@mcw.edu

Last month we looked at the physical memory "sticks", the ones you add to your motherboard to increase Random Access Memory (RAM). This month we'll finish up the topic by exploring the various types of RAM. In the past, all one had to worry about was RAM speed. Almost everyone knew that if your machine required 70ns SIMMs, it would not do to put in 80ns sticks when adding memory. Further, practical experience taught many of us that it was best to use sticks that all came from the same manufacturer, and even from the same lot! Today, there is much more to consider.

The RAM you add as sticks is really DRAM or Dynamic RAM. DRAM consists of many, many little capacitors, each associated with its own transistor. If the capacitor is charged, this indicates a 1; if uncharged a 0. The transistor associated with each capacitor is used to read the state of the capacitor, in other words, whether it holds a 0 or a 1. Simple enough, except that the capacitors are constantly discharging by themselves. Therefore, Dynamic RAM needs to be refreshed about every 2 msec or it will lose its data.

An interesting aside is that DRAM chips are currently being produced that contain 256 million transistor-capacitor pairs! Compare that with a 386DX CPU (275,000 transistors), a 486DX CPU (1.2 million transistors) or the first Pentium (3.1 million transistors). The difference is that in CPUs the transistors are arranged in very complex logic circuits, while in DRAMs, the arrangement is simple, square columns and rows.

Think of the capacitor-transistor pairs in DRAM as being laid out in a square array, just like the columns and rows on a spreadsheet page. Indeed, a given pair is accessed by selecting a column and row address before querying the transistor at that site concerning the state of its capacitor. Logically enough, this process is called paging. It works well, but takes time to accomplish.

Fast Page Mode DRAM uses a slightly different method to reduce the time taken by this process. When accessing the pairs, the row address is kept the same and only the column is changed, which simply takes less time. Another method to speed access in FPM DRAM is to simply read the data held by the next three capacitors at the same time. After all, we know what row we are working with so why not just go ahead and read the next three columns as well? It actually takes a bit less time to do so without first going back to select a new column, which speeds things up considerably. This ploy is called a bursting technique, and DRAMs designed to employ this are called Burst Mode Access DRAMS, a subset of FPM DRAMS. Yet another subset uses an interleaving technique to speed things up. Two separate banks of memory (two or four sticks) were used alternately, so that one could be charged while the other was being accessed for data. Again, this made access faster, but the method was dropped with Pentiums (64 bit machines) because it would require installing memory 128 bits at a time (four 72 pin SIMMs or two 168 pin DIMMs at a time).

During the mid 1990s, Extended Data Out (EDO) memory was invented as yet another subset of Fast Page Mode. About 10ns per cycle were saved by not turning off the output drivers on the memory chip before the next memory access was started. This saved nearly as much time as interleaving, mentioned above, and didn't require installing sticks in pairs. Further, it costs no

more to produce EDO sticks, though your Pentium motherboard must support this type of memory if you are to use it.

SDRAM (Synchronous DRAM) is likely in your machine today if it was built after about 1997. These DIMM modules are designed to run in synchrony with the memory bus of your system. Using a clocked interface, these chips deliver information to the memory bus in high-speed bursts of information that overcome the limitations of earlier memory designs. As an example, a task that took 1.4 seconds for Fast Page Mode memory and 1.1 seconds for Extended Data Out RAM would take only 0.8 seconds for an SDRAM module. SDRAM is also currently capable of supporting a 100 MHz bus speed, and this speed may well double in the future. If you see modules advertised that are PC/100 certified, these are sticks that adhere to the Intel standard that sets criteria needed for them to work reliably in a 100 MHz system (which means they can technically operate in a 125 MHz environment).

What is next? Look for RIMMs later this year (Rambus Inline Memory Modules). These are similar in size and form to DIMMs, but they are of a radical new design. The data bus in the chip is doubled to 16 bits, and the clock speed is up to 800 MHz! You will purchase these amazing sticks in 256-Mb sizes initially, but they will eventually be available in 1 Gb sizes or even larger. Imagine yourself saying to a friend, "Yeah, I need to get another Gigastick of memory for my computer." It may well happen, sooner than you think! Another development is DDR (Double Data Rate) SDRAM. This stick does much the same as standard SDRAM, but twice as fast by sending two bursts of data for each clock cycle. These are likely to be less expensive than RIMMs, since the technology is license-free while the company that invented RIMMs (Rambus) must license them.

Finally, a prediction. Eventually, we will all be using SRAM (Static RAM) for standard computer memory. Even today's SRAM designs can handle processors running up to 500 MHz. There are no capacitors holding data in SRAM, just a cluster of six transistors for each bit. Since there is no capacitor leaking its charge, no refresh cycle is needed, and the chips will hold their data so long as power is applied. Today, these chips are up to 30 times bigger than DRAM chips (all those transistors take up space) and are also up to 30 times more expensive to produce, so it is impractical to even think about using them in place of DRAM. However, I believe that future advances in design and production will eventually make SRAM practical for standard memory modules. When that happens, battery backup will keep them "alive" even when you turn off the power. That means, when you boot your machine in the morning, the letter you were working on will still be on the screen, just as you left it the night before. There will be no need to load it from the hard drive, since it is still up in memory!

Don't forget to make arrangements to attend Field Day, and have fun there! Happy computing!