

TECH —

RAM Guide Part I: DRAM and SDRAM basics

It's long been a truism that, no matter what kind of system you have, the best ...

JON STOKES - 7/18/2000, 4:16 PM

DRAM refresh: Making it work

Remember how I said that most of the challenges in using DRAM center around "making it work" when it actually doesn't. Well, this is the point where we'd normally go into all of those techniques that DRAM designers use to turn DRAM into a useful storage medium; specifically, we'd discuss DRAM refresh and then go on to Fast Page Mode, EDO, etc. I've organized this guide a little differently, however, so we're going to save the in-depth discussion of DRAM refresh and its attendant problems (and their attendant work-arounds) for a later part of the guide. Right now, I'll just give you quick rundown on what refresh cycles are, why they're necessary, and what their effects are on performance.

I've mentioned a couple of times that DRAM cells leak their charges out over time, so that charge has to be refreshed if the DRAM is actually going to be useful as a data storage device. Reading from or writing to a DRAM cell refreshes its charge, so the most common way of refreshing a DRAM is to read periodically from each cell. This isn't quite as bad as it sounds for a couple of reasons. First, you can sort of cheat by only activating each row using /RAS, which is how refreshing is normally done. Second, the DRAM controller takes care of scheduling the refreshes and making sure that they don't interfere with regular reads and writes. So to keep the data in DRAM chip from leaking away the DRAM controller periodically sweeps through all of the rows by cycling RAS repeatedly and placing a series of row addresses on the address bus.

Earlier in the article, I noted that a RAM grid is always a rectangle, and not a perfect square. With DRAMs, you want fewer rows and more columns because the fewer rows you have, the less time it takes to refresh all the rows. So DRAM makers design DRAMs with fewer rows than columns, hence the rectangular layout.


Even though the DRAM controller handles all the refreshes and tries to schedule them for maximum performance, having to go through and refresh each row every few milliseconds can seriously get in the way of reads and writes and thus impact the performance of DRAM. EDO, Fast Page, and the various other flavors of DRAM are mainly distinguished by the ways in which they try to get around this potential bottleneck. Again, we'll save the discussion of these technologies for a later edition where we'll talk in more detail about memory timing and refresh cycles. For now, you know the basics

of how DRAM and SRAM chips work, and you've been introduced to some of the major challenges and limitations of designing a DRAM-based system.

RAM Module basics

Earlier, I said that each of the cells in an SRAM or DRAM chip traps only a 1 or a 0. Also, the DRAM and SRAM example chips that I used only had one Data In and one Data Out pin apiece. Now, you know that the CPU actually sees main memory as a long row of 1-byte cells, not 1-bit cells, so you've probably been wondering how it is that a whole byte gets stored in RAM.

There are two answers to this, a simple one and a slightly less simple one, both of which are usually combined in modern systems. We'll discuss the less simple one first, because it's the most important for now. (The second one will come up later.) To store a byte (8 bits), you just stack eight, 1-bit RAM chips together, and have each chip store one bit of the final byte. This involves feeding the same address to all eight chips, and having each chip's 1-bit output go to one line of the data bus. The following diagram should help you visualize the layout. (To save space, I used a 4-bit configuration, but it should be easy to see how you can extend this to eight bits by just adding four more chips and four more data bus lines. Just pretend that the picture below is twice as wide, and that there are eight chips on the module instead of four.)




Sticking those eight chips on one printed circuit board (PCB) with a common address and data bus would make an 8-bit RAM module. To the CPU, the above, single-module configuration would look just like one big RAM chip that, instead of holding just one bit in each cell, holds four.

So let's say that in the above picture, the address bus is 22 bits wide and the data bus is 8 bits wide. This means that each single chip in the module holds 2^{22} or 4194304 bits. When the eight chips are put together on the module, with each of their 1-bit outputs connected to a single line of the 8-bit data bus, the module appears to the CPU to hold 4194304 cells of 8 bits (1 byte) each (or as a 4MB

chip). So the CPU asks the module for data in 1 byte chunks from one of the 4194304 "virtual" 8-bit locations. In RAM notation, we say that this 4MB module is a 4194304 x 8 module (or alternatively, a 4M x 8 module. Note that the "M" in "4M" is not equal to "MB," or megabyte, but to "Mb," or megabit.)

The now-obsolete SIMM on which I've based the 4M x 8 module discussed above is the TM4100GAD8, from Texas Instruments. The TM4100GAD8 is a 4MB, 30-pin SIMM that's organized as eight, 4M x 1 DRAM chips. Here's some info from the datasheet, which I scanned in. (BTW, if anyone knows where you can get online copies of datasheets for obsolete ICs like DRAMs and SRAMs, please let me know.)



So this SIMM can provide the CPU with eight bits of data at a whack...



SITREP: F-16 replacement search a signal of F-35 fail?

Footage courtesy of Dvids, Boeing, and The United States Navy.



SITREP: F-16 replacement search a signal of F-35 fail?



Sitrep: Boeing 707



Steve Burke of GamersNexus Reacts To Their Top 1000 Comments

[+ More videos](#)

← PREVIOUS STORY

NEXT STORY →

Related Stories

Today on Ars

Long COVID rates have declined, especially among the vaccinated, study finds

Apple Vision Pro's content drought improves with new 3D videos

Netflix is kicking US subscribers off its cheapest ad-free plan soon

Witness the rise of the Bene Gesserit in new *Dune: Prophecy* teaser

Space colonizers battle
ultimate killing machines
in *Alien: Romulus* trailer

Elon Musk's X may
succeed in blocking Calif.
content moderation law
on appeal

FCC closes "final
loopholes" that keep
prison phone prices
exorbitantly high

Report: Apple TV+ will
soon get a lot more
movies made by studios
other than Apple

STORE
SUBSCRIBE
ABOUT US
RSS FEEDS
VIEW MOBILE SITE

CONTACT US
STAFF
ADVERTISE WITH US
REPRINTS

NEWSLETTER SIGNUP

Join the Ars Orbital Transmission
mailing list to get weekly updates
delivered to your inbox. [Sign me up →](#)



CNMN Collection

WIRED Media Group

© 2024 Condé Nast. All rights reserved. Use of and/or registration on any portion of this site constitutes acceptance of our User Agreement (updated 1/1/20) and Privacy Policy and Ars Technica Addendum. Ars may earn compensation on sales from links on this site. Read our affiliate link policy.

Your California Privacy Rights | [✔✕](#) Your Privacy Choices

The material on this site may not be reproduced, distributed, transmitted, cached or otherwise used, except with the prior written permission of Condé Nast.

[Ad Choices](#)