



The New Memory Performance Figure of Merit: Address Rate

Networking system designers have a staggering degree of freedom when it comes to solving performance problems while designing network switching and routing equipment. As a result, it is impossible to say with certainty exactly how the performance roadmap for one particular component type in the system will impact ultimate system performance or even architectural choices. However, one figure of merit for memory performance can be interpreted in a fairly straightforward way in regards to some of the challenges networking system architects face going forward: address rate.

It has been suggested that the Internet is just a memory. The aphorism is more true than glib. While disk drives may hold the content that goes onto the web at one end and may catch that content on the other end, everything in between is either copper, optical fiber or some form of volatile memory—either DRAM or SRAM or embedded memory or registers—and the logic contrived to read and write them. Granted this view is somewhat self aggrandizing for a memory vendor, given all the protocols and algorithms at the heart of the grand design that makes it all work, but at a hardware level it all comes down to how data gets moved and manipulated in very short spans of time that makes or breaks the performance of a networking box. And, paradoxically, it is managing small chunks of data that turns out to be the most difficult.

The small chunks of data of interest here are minimum packets—the smallest chunk of data the equipment in question is designed to manage as a indivisible, routable piece of information. Although years ago a 40 byte packet was considered the minimum, the pervasive role of Ethernet has boosted that number (with windage) to 672 bits. In this case we are talking about a 64 byte frame and a 20 byte interframe gap. If we take a deep breath and stipulate that point for a moment, the next issue is Line Rate. Recent history and the immediate future suggest it is useful to look at 1Gb/s, 10Gb/s, 20Gb/s, 40Gb/s and 100Gb/s line rates. If we know how fast the data is coming in (and we hope, going out as well) and we know the minimum packet size, we can calculate how short a time a packet may last, in the worst case. The information about where a packet has come from and where it is going and what is in it, and on and on, is generally carried at the front end, in the earliest arriving part of the packet, the header. Once that has arrived and that information is in hand, decisions can be made and actions can be taken to dispose of that packet appropriately. But the crux of the matter is the making of those decisions and taking those actions, both of which take a finite amount of time and require some number of memory accesses to accomplish. Some memory accesses may be read-modify-write cycles that increment counters that keep track of billing or quality of service issues, there may be table walks that sort out routing information, there may be duplications or re-reads to deal with multiple destinations. The list may be endless. But in the end, it is the rate at which these sorts of memory operations can be initiated and successfully retired that sets the pace for everything that happens in the box.

So which memory performance metric gives the best feel for the suitability of a particular memory technology for the highly random, transaction rate driven jobs in a networking system? Random Address Rate. Random Address Rate—The measure of how often a new fully random access can be successfully executed in a memory device. It is the most fundamental cost-controlling performance characteristic of any memory device. Although data bandwidth is all the rage as the popular figure of merit for memory performance in compute systems from PCs to supercomputers, in networking, address rate is a fundamental driver for certain necessary functions. Granted, there are many jobs memory devices do in networking systems and not all of them require high transaction rate memory. But some jobs do, and if high transaction rate memory cannot be secured to do those jobs at the pace required, the entire system performance must be throttled back accordingly.

So it remains true that memory devices like commodity DRAMs with low random address rates but high data bandwidth are important in networking. But Random Address Rate is the root, the nub, the atomic level issue that drives the most difficult, performance critical tasks in network system architectures. Yes, there are other important issues besides Random Address Rate—power, cost per bit of storage, available memory capacity per unit area, physical space, reliability, long term source of supply—they all matter, but Random Address Rate is the deal maker or the deal breaker for high performance networking systems.

So why does nobody talk about Random Address Rate? One can suppose it is for the same reason people once only talked about processor clock rates despite the fact that clock rate was and still is poorly correlated to overall system performance. Or because gearheads talk about horsepower when it is actually torque that makes the difference in acceleration of a car. The good news is that we are all getting smarter. We are buying for actual computer performance rather than processor megahertz, seeing that electric

cars can outrun super sports cars costing hundreds of thousands of dollars more, and now even seeing that when it comes to networking memory, how often the system logic can ask for something new and get it, is much more important than how fast a memory device can deliver what the system does not need.

The following table shows how things looks in just those terms. It shows how often each of the memory types shown can successfully execute a fully random access in a minimum packet time at each of the line rates discussed. Obviously, at lower line rates the minimum packet time is pretty long (in electronic terms) and more reads or writes can be done in the time available. So it is no surprise that lower end, lower cost memory products are well suited for slower speed networking equipment. But the faster the line rate, the bigger the challenge. This paper makes no attempt to specifically identify the things that may be accomplished with the memory reads and writes that can be executed in the available minimum packet time, but it is clear that a designer working on a 100Gb/s system can do more with the 6 transactions provided by the SigmaQuad-IIIe Burst of 2 parts than with the 3 that are possible with the QDR-II+ Burst of 4 SRAMs.

One of the things that jumps out of this table right away is that megahertz is an even worse indicator of RAM performance in a networking application than it is in the compute world. The highest frequency memory products on the market, DDR3 DRAMs, lie at the bottom of the performance heap when Random Access per Minimum Packet Time is considered. Even the QDR-II+ family of SRAMs has long fallen victim to the conventional trap of putting megahertz over Random Address Rate, only recently adding a Burst of 2 version that even now only improves transaction rate 33% compared to the performance of its 5-year-old predecessor, the QDR-II B2 SRAM. The B2 SigmaQuad-IIIe SRAMs strike at the heart of the issue, delivering 1 GHz Random Address Rate performance, outperforming any other discrete memory solution on the market by 50% and outrunning the most popular networking SRAMs by 100%. Nothing else comes close.

If memory vendors fail to press forward on Random Address Rate, the challenges for networking system architects could become overwhelming. Fortunately, that is not the plan. We are moving forward and the new architectures we are bringing to the market show that we are serious. Good thing, too. Because the demand for network performance is showing no signs of slowing down.

Line Rate	1Gb/s	10Gb/s	20Gb/s	40Gb/s	100Gb/s
Bits/Second	1,000,000,000	10,000,000,000	20,000,000,000	40,000,000,000	100,000,000,000
Seconds per 64 byte frame + 20 byte interframe gap	672 ns	67.2 ns	33.6 ns	16.8 ns	6.72 ns
Number of Random Accesses per Min Packet Time					
SigmaQuad-IIIe B2 at 500 MHz	672	67.2	33.6	16.8	6.7
QDR-II+ B2 at 333 MHz	447.6	44.8	22.4	11.2	4.5
SigmaQuad-IIIe B4 at 625 MHz	420.0	42.0	21.0	10.5	4.2
QDR-II+ at 550 MHz	370.9	37.1	18.5	9.3	3.7
RLDRAM-II at 533 MHz, 15 ns tRC	44.8	4.5	2.2	1.1	0.4
DDR3-1600 at 800 MHz, 45 ns tRC	14.9	1.5	0.7	0.4	0.1