The Things I Hear: The Engineering to Purchasing Gap in Embedded Memory Selection

Look beyond reference designs and marketing materials. Some time spent studying data sheets and exploring less well-known manufacturers and then communicating this information to your purchasing department can go a long way toward cost savings and longer life for your design project.

by Nicholas Urbano, Memphis Electronic

As the lead engineer on your latest project, you've chosen your memory, finished your design, and sent it on to the next phase of development. Then one Friday afternoon, the phone rings. It's someone in a purchasing department, either half a continent away or just down the hall. The problem? The memory you designed, be it 2 months or 2 years ago, is EOL. You scramble to find replacements, only to discover that the pin out has been discontinued, there is no recommended replacement from the original manufacturer, and despite your consumption of ten thousand units per month (or more!), there is nothing you can do but re-spin the board … immediately.

There goes your weekend.

This scenario plays itself out somewhere on the globe every week. It costs your company valuable resources and costs our industry hundreds of millions of dollars each month. Of course, it also brings stress, frustration and internal company strife while everyone tries to solve a problem, plug the hole or bridge the gap.

Me? I'm on the receiving end of the frantic calls—listening to each request, identifying the parameters of the specific application, and working to understand the concerns of each engineer. I typically advise on embedded solutions covering a wide range of applications—medical, infrastructure, aerospace, networking, automotive, civilian, military and everything in between. My job is to work with the engineer to identify the best solutions, but more importantly, to provide confidence in the sustainability of a design.

So this is about the things I hear each week. After more than a decade in this industry, I have gained a considerable amount of insight from the engineers I work with on a daily basis. Because each project is so unique, I often can't generalize the nature of questions from project-to-project. These are some of the most common inquiries and situations I've been presented with, especially in the last year. And most importantly, these are the things I have to understand to ensure that I can provide you confidence in your design.

Should I Commit to DDR2 or DDR3?

The most common question I have received from engineers in the past year is whether to commit to DDR2 or DDR3 for long-term projects. With the influx of processors supporting DDR2 and/or DDR3 interfaces, the answer isn't cut and dried and requires a bit of background knowledge.

First, it is important to understand that *memory is a commodity* in the electronics industry. Many manufacturers make their version of the same product. However, each manufacturer has its advantages and specialties. If you know the niche that each manufacturer occupies, you can leverage this knowledge for your benefit.

Secondly, consider the history of DRAM. Over the past fifteen years, we've gone from Asynchronous to Fast Page and EDO to SDRAM, and DDR through DDR2 to DDR3. While FPM and EDO DRAM have become all but extinct (there are still some suppliers that support both!), SDR still remains in heavy usage in applications all over the world. The key to its success has been that engineers working

PCIe OVER CABLE

The HIGH SPEED universal serial bus

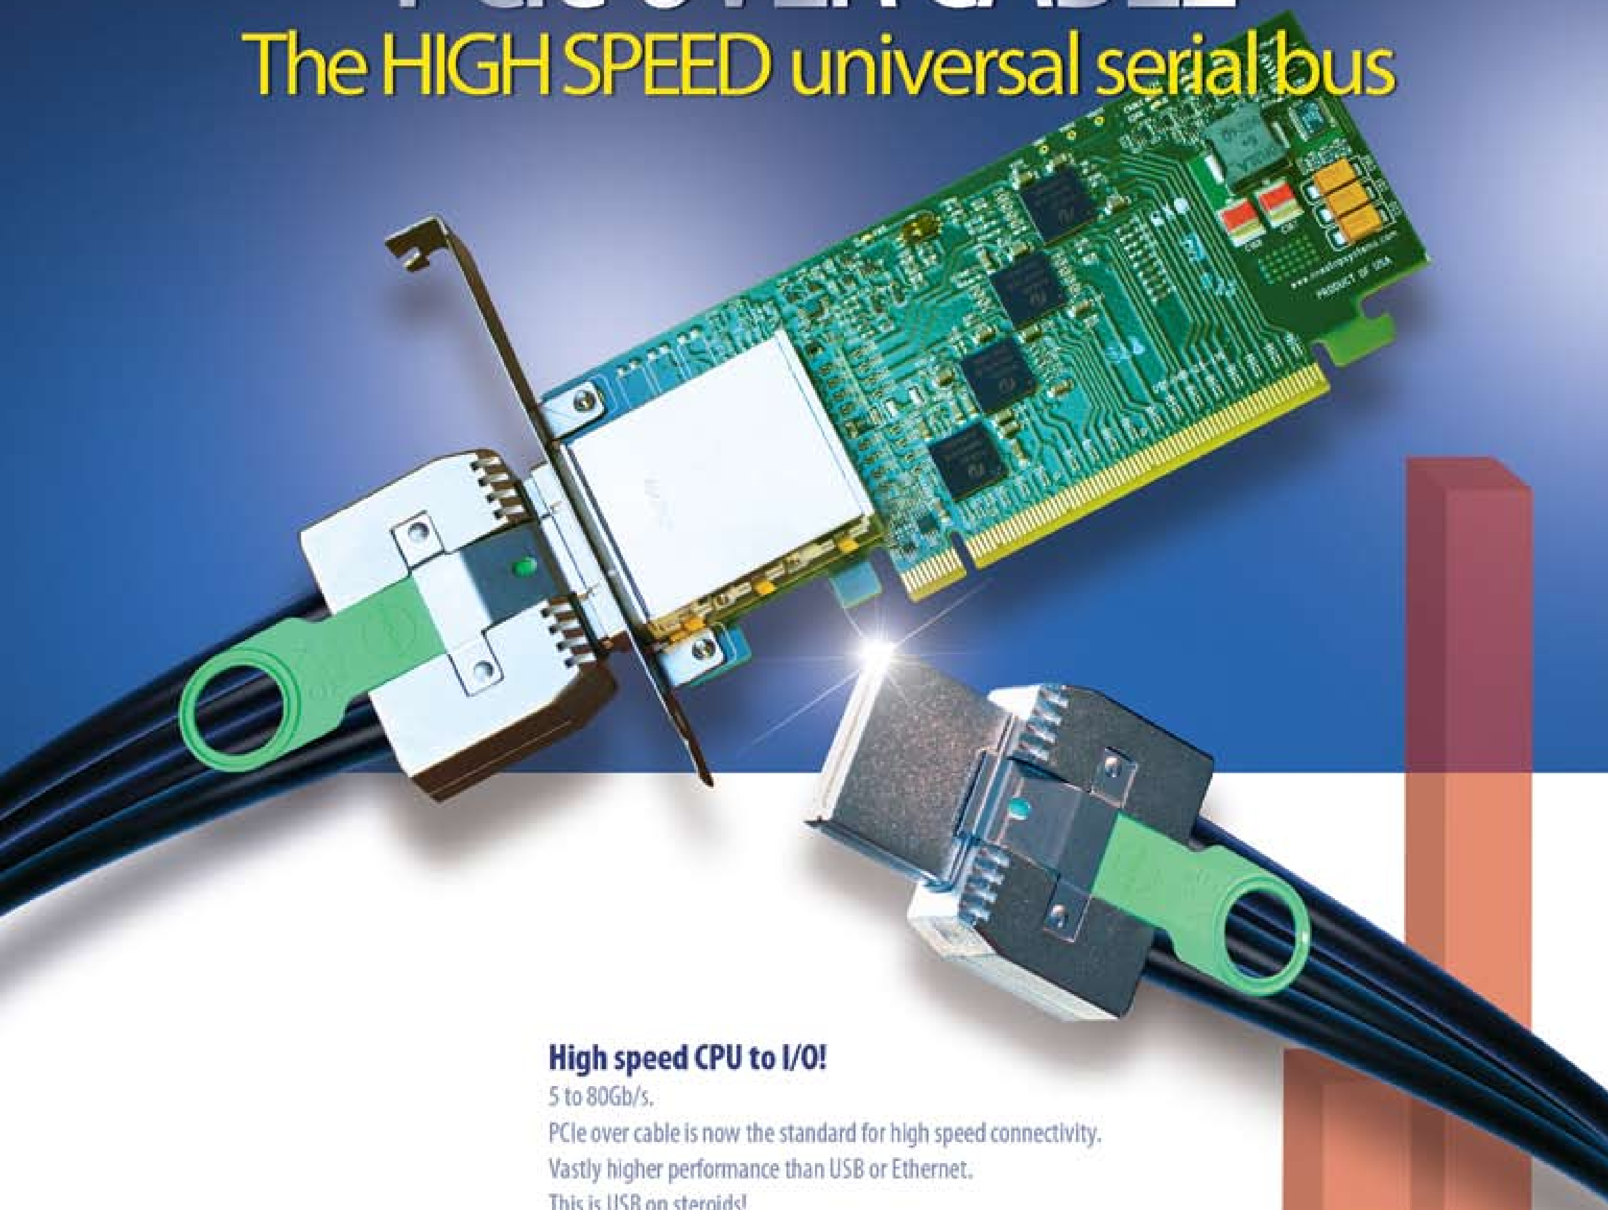

High speed CPU to I/O!

5 to 80Gb/s.
PCIe over cable is now the standard for high speed connectivity.
Vastly higher performance than USB or Ethernet.
This is USB on steroids!

Simple, Reliable, Cost Effective!

Easy installation of adapters and cables.
No software to install means no overhead and low latency.
Proven reliability across a multitude of applications.
Variety of adapters to suit price-performance needs.

CALL TODAY
(877) 438-2724
www.onestopsystems.com

ONE STOP SYSTEMS

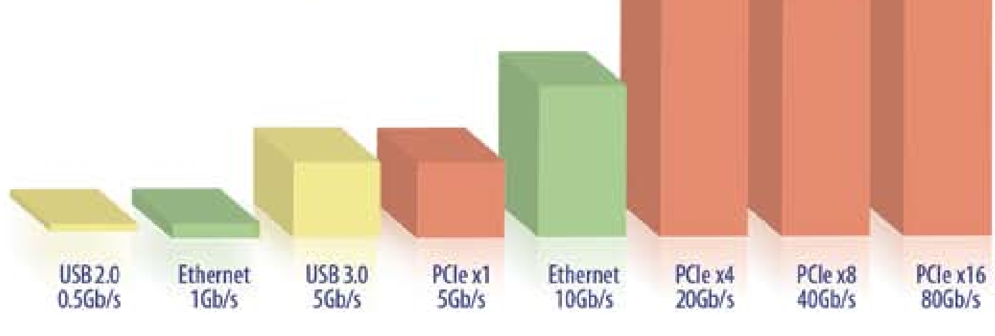

FIGURE 1

Data sheets may be many and unexciting, but they contain important details, which, when carefully compared to the requirements of your design, can result in significant savings. These details must be clearly communicated to purchasing.

on applications utilizing SDR allowed for increased densities. In fact, only now are we starting to see concerns about the continued support of lower-density 64 Mbit SDR in the long term.

Fast forward to 2011; the industry has moved from DDR2 to DDR3 as the primary DRAM device in PCs, with the bulk of consumer applications soon to follow. However, we all know too well that embedded designs don't always follow at the pace of the consumer. So what to do and what can we learn from SDR?

Positively, most major memory categories introduced since SDR in the last decade are still active and in production in common configurations and densities, and that trend will continue. As long as you choose a common configuration and are flexible with your choice of manufacturer, *both DDR2 and DDR3 are safe* for long-term embedded designs. That leaves density as the most important variable in your selection.

At this point, DDR2 has identified its "sweet spot"—1 Gbit IC production will remain the largest and most common density of DDR2. 2 Gbit Monolithic DDR2 ICs are offered as somewhat of a niche product (but are readily available!), while 4 Gbit DDR2 is available only as stacked die in a single package (DDP). On the low end of the spectrum, if you're working with a processor that recommends 256 Mbit, forget it. Unless your processor cannot read more than a 512 Mbit DDR2 IC, it's best to work with the 1 Gbit. The highest cost savings and safest long-term sustainability come with the largest segment of production, which is the 1 Gbit DDR2. Even if you do want to go up in density with a 2 Gbit or 4 Gbit, the pin out is the same as the 1 Gbit (or the 512 Mbit for that matter). While there might be a bit of a cost savings in a 512 Mbit DDR IC, the price does not directly correlate to density; half the density (512 Mbit vs. 1 Gbit) is not half the cost. The most commonly produced density and configuration remains 1 Gbit DDR2 in x8 and x16.

So where does that leave DDR3? If space constraints, processor selection, or application demands push you over the limits of what DDR2 1 Gbit ICs can easily do for your application, then use DDR3. Presently, 2 Gbit DDR3 production is the "sweet spot." By the end of 2011, most major manufacturers will have brought out a 4 Gbit Monolithic (some already offer a DDP 4 Gbit), and still others have even announced 8 Gbit DDR3 ICs. So long as your configurations match, then the best solution is to work with the "sweet spot" as much as possible and plan and allow for increased densities in the future.

It is also important to understand that in designing with DDR3 now, your embedded applications will not be in full consumption for several quarters. DDR3 will continue to be dominated by consumer applications. Stick with the most common densities, but don't limit yourself to specialty CL callouts and configurations. 2 Gbit and 4 Gbit DDR3 in x8s and x16s will be safe to design with for the long term.

It's Not in the Reference Design!

Manufacturers devote a great amount of attention to sample production and sales support in an effort to get their memory ICs included in reference designs. As we all know, these reference designs drive purchasing patterns for a very long time. However, engineers will often see only the major manufacturers for each new design that is debuted.

Here's a real-world example. I had a client call for support on a 1 Gbit DDR2 IC in a x16 configuration, CAS Latency (CL) 5 at 2.5ns. The 64Mx16 configuration in DDR2 is incredibly common in embedded memory applications. For some manufacturers, a CL5 at 2.5ns carries a specific part number add-on, and with it, a price premium. When my client wanted to match the original spec completely, I asked if it was really necessary to have CL5 at 2.5ns. After some deliberation and discussion with engineering, it turned out that the application did not specify nor require CL5 at 2.5ns. In fact, many of our customers' feedback involved running 200 MHz at 5ns clock cycle times. The reality therefore is that nearly all DDR2 components can operate at CL5, CL4, or CL3 when clocking at such low speeds as 5ns. However, since the component specification was originally on the reference design and added to the Approved Vendor List (AVL; sometimes also called a BOM or AML), purchasing was forced to unnecessarily seek out a costly IC. In this case (as in many others), specific nomenclature is included in the part number that adds premium to your supply chain requirements.

By thoroughly studying the simple differences on a datasheet, you can provide your purchasing department with multiple support options for future buys instead of just one. Additionally, you can prevent price bloat with add-ons that aren't actually necessary for your specific application. Without review, cost savings and supply chain easing options might never make it to your AVL (Figure 1).

Ask your suppliers and reps which speeds are most commonly manufactured. We've seen lots of EOLs on major brands for SDRAM this year. The most common speed previously was 133 MHz. However, most new manufacturers are only producing at 143 MHz. Some clients have balked at the speed increase, when in reality it rarely makes a difference in your application when you put it to test.

Additionally, it is imperative to look for other manufacturers that make your chosen memory IC. Some of the major memory manufacturers continue to market SDR, DDR1 and DDR2 products in industrial and commercial temperature ranges. There are also several mid-tier manufacturers and fabless memory companies that offer competing products as legacy vendors. Their business models and long-term outlook are solid and can often provide deep cost savings.

When dealing with mid-tier fabless companies, the key is to find local distribution or representation for those manufacturers. Once you have the support of a representative or distributor, then you have a direct channel to the lesser known manufacturers that may not have the marketing budget to comprehensively support reference designs. These representatives and distributors can then support your applications with sampling, technical support, cost saving and long-term design sustainability.

As for the real-world example, we remedied the situation by qualifying three part numbers from the existing manufacturer, and added options from two additional manufacturers as well, one of them from a smaller brand. This has bridged the client through several revision changes by having the overlapping qualifications already done without seeing gaps in their supply.

What About Flash?

Many applications these days are now using some form of NAND flash memory devices. It is helpful to think of the category in a similar manner to DRAM. It is mostly consumer driven, and revisions will change based upon varying consumer demand and opportunities for the manufacturers to drive down costs with die shrinks and more efficient manufacturing processes.

Typically, NAND flash devices have been limited to the big three memory manufacturers: Samsung, Hynix and Micron. However, other NAND manufacturers are beginning to come onto the scene for smaller densities.

For embedded applications, your choices are two-fold: either onboard flash memory ICs or peripheral flash devices. Peripheral flash devices include such products as disk on modules (DOM), compact flash cards (CF), or secure digital card (SD & SDHC).

You can start by standardizing your flash controller if it's not already built into your processor. Most problems I see with qualification are not the memory chips, but the flash controllers themselves. If the reference design calls for a specific manufacturer, then actively test other NAND flash devices with your controller to ensure it works consistently. Lastly, remember that for rugged and high-vibration applications, onboard flash is the best choice.

If you do choose to design with NAND flash ICs on the board, then keep in mind that the bulk of production on NAND ICs is presently in TSOP and at 3.3V. We have seen an increasing number of requirements for 1.8V BGA NAND ICs. But as production output is not as high as 3.3V TSOP, you may run into supply problems with BGA 1.8V NAND product.

If your goal is to allow flexible densities for a product in an application, it is easiest to qualify a peripheral flash device. However this can be tricky in embedded design due to the wealth of manufacturers presently offering flash products in the market. The biggest problem customers create with peripheral flash devices is that engineers limit themselves to one manufacturer. Try to qualify several manufacturers who are willing to support with samples, sell industrial grade products and offer fixed BOMs.

Once you qualify an industrial flash device, be sure to specify the differences between standard products and industrial grade products on the AVL. A simple note of “Industrial Temperature/Grade Only” in the AVL or in the part description can make a difference, and will prevent supply chain confusion with retail or consumer grade products.

It is also important to be mindful that flash controllers across different manufacturers change. Once you find a flash device that works well with your application, make sure you’re on a fixed BOM with the flash device itself and stick with it. Without a fixed BOM, flash controllers will change rapidly, often times with negative results for embedded applications.

Lastly, be open to solutions for different flash devices. Understand that revisions change frequently in NAND flash. In a peripheral flash device, a chip revision or a flash controller change will affect your complete device. Be open to suggestions if there are problems that arise due to these changes.

I have one last crucial and simple tip on design. Since you’re going to be looking at more than one manufacturer to include on your AVL after reading this article, try to find the manufacturer with the oddest and biggest package to design around. I have seen multiple applications where real estate on the board is at a huge premium because the original design used the smallest package available at the time.

Packages, pin outs and FBGA layouts are standardized, *but package sizes are not*. I have had multiple clients reject alternative manufacturers because their package size was off by a millimeter or less in one direction or another. If you plan your pad space around the largest available package, you will prevent a world of headaches down the line. This is especially true for large companies that use a central library of internal reference designs and qualifications to share between projects.

I have had so many conversations with engineers and purchasing personnel through the years and I wish I could cover everything. My best overall advice in embedded memory selection is simply to have an open mind to new and different manufacturers, expect comprehensive and knowledgeable support from your representation, and don’t get too creative on the configurations. Try to find out what is happening in the memory market. We are all part of a truly global business, and events half a world away affect our business daily. Consumer trends affect us in the embedded sector more than we sometimes want to admit. Accept that roadmaps will change despite our best planning and problems will arise. All we can do is try our best to be prepared. And one last personal request…when that next EOL notice does come down the pipe, try not to kill the messenger. We really are here to help! ◢

Memphis Electronic
Houston, TX.
(713) 600-6080.
[www.memphis.ag].