Open Architecture Review

A Window on the Embedded Computing World

January 4, 2010

Interview: Rodger Hosking on Signal Processing

By: dl

Rodger Hosking, a cofounder and vice president of Pentek, has seen many generations of signal processing come and go. He discussed a variety of current topics of interest with *Open Architecture Review*, including the emergence of point-to-point serial links, the roles of VXS and VPX, mezzanine trends, the FPGA phenomenon and the uncertain future of general-purpose DSPs in military and government electronics.

OAR: The system interconnect of the past was primarily the parallel bus, and now it's the point-to-point (P2P) serial link. How has this changed signal processing?

Hosking: It has pretty much revolutionized our architectures. About three years ago, we migrated our VME designs towards VITA 41 VXS, which augments parallel VME with Gbit serial interfaces on the same board. That's very effective for solving some problems where VME used to be a bottleneck.

OAR: And VXS lets you continue to use legacy VMEbus boards, right?

Hosking: Yes, and there are going to be an awful lot of designs going forward that will still use VME and legacy boards over the next ten to fifteen years. And when the VME backplane is not sufficient to handle the bandwidth you need, you can use VXS to give you rates that are twenty times higher than what VME can do.

OAR: Doesn't having to handle all that traffic introduce some complications?

Hosking: Sure. As I said, we've pretty much revolutionized our architectures because of it. We've moved from communicating over a shared resource, a bus, to communicating via direct pathways with no contention. The architecture is fundamentally different.

OAR: Can you point to one relevant difference in board architecture?

Hosking: To support the change for one particular processor board, we incorporated a Gbit serial crossbar switch to serve as traffic cop for on-board resources, connecting a PowerPC processor, XMC mezzanines, FPGA resources and backplane. That's a major change for us.

OAR: How is the market trending in mezzanine boards? Where is there movement to XMC and serial P2P links? Where are people staying with PMC and its legacy parallel PCI bus?

Hosking: PMCs are being replaced by XMCs to add the Gbit serial links. We originally made PMC modules, then shifted to hybrid XMC modules with both PMC and XMC connectors on them, and our latest introduction has XMC connectors only. You can see the progression from parallel, to parallel plus serial, to straight serial.

OAR: Pentek chose PCI Express as its P2P link. Why PCI Express?

Hosking: PCI Express is really nice because you can reuse a lot of the software you've previously developed for PCI and PCI-X under Windows or Linux and not have to do very much to it to support the shift from parallel buses to serial P2P links. A lot of the software remains the same.

OAR: Do you support PCI Express exclusively?

Hosking: At this point we're promoting PCI Express as the standard control and data plane fabric, but we're also supporting and delivering the [Xilinx] Aurora protocol. We just announced a family of beam forming products for software radio, communications and radar. They use PCI Express for controlling and loading data on and off the modules, but they also use Aurora protocol as a high-speed interconnect between boards in kind of a daisy chain.

OAR: Doesn't PCI Express provide that capability?

Hosking: Aurora is a very low-level, lightweight protocol with very low overhead, whereas PCI Express has many layers of protocols and conventions you have to follow in order to be a true PCI Express entity. Something like Aurora is the best way to handle the high speed summation across multiple channels, which is the critical operation in beam forming.

OAR: Do you run Aurora over the VXS backplane?

Hosking: No, over cables, like people always did (and some people continue to do) with the old FPDP [Front Panel Data Port]. We can daisy chain boards down a card cage as many slots as

required, pretty much without limit, by simply attaching cables for the Gbit serial connection. At the end, you pull out the final sum and send it out across the PCI Express interface; but from the beginning to the end of the daisy chain, you don't have to rely on the motherboard or carrier at all. It really simplifies a lot of the beam forming operations people need to do.

OAR: You're supporting Aurora on a PCI Express board. If you did the same on a VXS or VPX board, would you continue to use cable for the traffic?

Hosking: That depends on what it's attached to. The cables that we provide for the PC environment are definitely not what people are looking for in a military VPX chassis. There, Aurora links would probably be done through backplane connectors, and the boards could either be joined by a switch board or by a custom backplane that's hardwired between slots to provide the daisy chaining capability.

OAR: Regarding VXS, I see very little activity in support of it these days, and most VXS suppliers seem to be plunging into VPX as rapidly as possible. Is there a shift going on here? Was VXS just a short-term, transitional standard?

Hosking: VXS was definitely a transitional move. It represents a shift from parallel only to parallel plus serial. The next shift is to serial only: that is, to VPX. The same shift is going on in mezzanines: first from PMCs to XMCs with both PMC and XMC connectors, and then to XMC-only boards.

OAR: So what are your VPX intentions?

Hosking: We will be announcing a VPX product line in Q1 2010, primarily for 3U and eventually for 6U VPX.

OAR: Pentek didn't have a small form-factor 3U line for VXS.

Hosking: There's no room for a P0-type connector on a 3U board. But VPX made provisions ahead of time for both 3U and 6U, and things are extremely well laid out and defined. The 3U VPX specification is more mature and more uniform across suppliers than 6U.

OAR: Has the stampede to VPX we've been seeing been influenced by the OpenVPX collaboration?

Hosking. Very much so. The OpenVPX initiative, the migration to VITA 65 in October, 2009, and its eventual adoption by ANSI will do a lot towards unifying the market.

OAR: When do you expect ANSI adoption?

Hosking: The first quarter of 2010 or so.

OAR: What motivated OpenVPX was the real-world interchangeability difficulties users have been having with VPX boards. Does the new specification solve the problem?

Hosking: There are still quite a few government agencies and organizations that are concerned that VPX is not yet stable enough. They want to know that the VPX system they buy today will accept future boards from different vendors, and that everything will work together. I expect their fears will fade over the next year or so.

OAR: What has been causing the incompatibilities?

Hosking: It's just based on the individual choices that designers made when they put together the first VPX boards and systems.

OAR: Examples?

Hosking: Clock distribution and management, chassis configuration and control, naming the various Gbit lane bonding conventions, and defining slot profiles for different types of VPX cards. Many of these issues were addressed in the OpenVPX effort.

OAR: Did the same type of problem arise with VXS?

Hosking: Because VXS is simpler with only two 4X VXS ports on any card, there were fewer design choices. Even so, there are several flavors of both switched and switchless VXS backplanes. With VPX, though, each card can have 20 or more different ports ranging from 1x to 16X, all needing to go to specific destinations. VPX has the capability for doing a lot of very, very high-performance embedded computing, but that freedom comes with a penalty in how all those interconnects are actually implemented.

OAR: So what does the near term look like for VPX?

Hosking: We see a lot of vendors lining up with available products. One complication is that VPX will require a diverse range of backplanes, which is quite different from VME. There will probably be four or five popular backplane configurations, but major programs will probably require custom backplanes with specific interconnections. That wasn't the case for VME.

OAR: Will VXS just wither away and eventually die?

Hosking: Someday, but VME will definitely help prolong VXS. For us, VXS was a very positive thing and an important stepping stone. It really helped enormously to get VPX started because it got us building hardware and software to support interboard Gbit serial links. Once we developed software for VXS boards, we were immediately able to take advantage of the interfaces and drivers for VPX, as well.

OAR: Have the P2P serial links in any sense been a leveler in terms of performance differences between motherboards and backplanes?

Hosking: You can do a tremendous amount of processing and achieve a very high interconnect bandwidth in the lower cost motherboard environment, but bandwidth isn't everything. The trouble is that PCs are not very good in unfriendly environments. Also, as you get into higher

density products, thermal management and adequate air flow become a problem because the PC is not a very good environment for cooling. VME or VPX or CompactPCI are much better at handling the heat than a PC.

OAR: Any other caveats?

Hosking: Yes. The PCI Express slot connectors that are found on motherboards often don't have enough amperage on the pins to power up the devices that are being plugged into them. So, for the PMC carriers we've done for PCI Express, we've had to implement separate power connectors that go directly to the power supply.

OAR: In the past, I've seen FPGAs here and there on various board-level products, but these days they're being widely used, especially in DSP environments. Why and why now?

Hosking: FPGAs have a traditional role of providing multiple hardware engines that perform a lot of DSP operations in parallel. What has happened more recently is that Xilinx and Altera incorporated Gbit serial engines as dedicated blocks right inside their FPGAs, making it easy for people to incorporate those serial interfaces.

Also, the high-speed devices we are typically involved with — high-speed A/Ds, D/As, network interfaces, serial FPDP, etc. — are all supported beautifully with the configurable I/O that FPGAs provide. So, you can comply with a lot of different logic levels and signaling rates up in the hundreds of MHz, and tune the I/O to match the characteristics of the peripherals you're trying to connect to.

OAR: That's a pretty powerful story.

Hosking: It gets better. You can now also get memory controllers, either from the FPGA vendors themselves or third parties that install inside the FPGA to allow you to interface to virtually any kind of external memory available. And the icing on the cake is that the devices have become faster and denser with each new generation. The latest version of Virtex-6 from Xilinx has over 2000 DSP engines on board, each capable of performing multiply/accumulate/addition operations in parallel — that's tough competition for a general-purpose DSP!

OAR: So will general-purpose DSPs wither away and die?

Hosking: They've definitely waned in popularity in our markets. That's because over the past ten years, the leader, Texas Instruments, has focused on the telecom market. As a result, the DSP chips are mostly fixed-point devices and they tend to have dedicated telecom peripheral interfaces on board. Our markets in military/government electronics like to use floating point because it's easier to have an algorithm work properly without having to worry about scaling and dynamic range.

OAR: So Pentek doesn't see any general-purpose DSPs in its future?

Hosking: That depends. If, for example, the telecom market appears to be getting saturated and TI gets competition for its DSPs from overseas knockoffs, they might start seeing our embedded markets as an opportunity and perhaps extend their offerings. We'd look closely at anything they might do in getting back to floating-point DSPs. We'd definitely not just give up FPGAs because they do so many different things, but we'd look at a good credible, competitive DSP.

OAR: Competitive in what sense?

Hosking: There are two things a DSP could compete on: power consumption and cost. FPGAs are expensive and they draw a lot of power. If TI could come in and handle the algorithm number crunching we're using FPGAs for, but at a lower cost and lower power, that might be interesting.

OAR: Do you see any other trends in signal processing?

Hosking: We definitely see more and more people using Windows and Linux, as opposed to the more traditional real-time embedded operating systems like VxWorks. Customers don't like the run-time licensing and the development tool costs. VxWorks is an excellent product, but we're seeing that when people can get away from it, they often do.

http://www.oareview.com/?page_id=1698