I never said the SHARC was better than the PowerPC

But FPGAs enable hybrid DSP architectures best

Q & A with BittWare's Jeff Milrod

Editor's note: *Jeff Milrod of BittWare is one of my favorite industry pundits. A cofounder and chief "brainiac" of Ixthos DSP, which was ultimately acquired by Dy 4 Systems, Jeff is as outspoken as he is brilliant. In this written interview, he smacks me around a bit concerning his vision of DSP's future and at my attempts to paint him into a corner (just for fun). – Chris A. Ciufo*

VME: ***Before we begin, can you briefly describe BittWare's DSP lineup?***

MILROD: Historically, BittWare has focused on providing multiprocessing, floating point signal processing boards with the highest possible performance, necessarily using a bunch of Analog Devices SHARC and TigerSHARC processors.

We started adding FPGAs to our DSP clusters just to provide reconfigurable interfacing. But as FPGA technology has improved, we've found that they provide complementary signal processing resources to DSPs. At the risk of oversimplifying, FPGAs are very good at solving well-defined, high-speed, repetitive problems, while DSPs are generally better at implementing highly complex algorithms and applications that require floating point, some sort of decision-making, adaptive processing, or algorithms that may change frequently. Real-world embedded signal processing applications often require both of these types of processing.

Therefore, we now believe the optimal solution for general purpose signal processing boards should have both FPGAs and DSPs – a "hybrid" signal processing architecture – providing more cost- and power-efficient solutions than either alone.

> "At the risk of oversimplifying, FPGAs are very good at solving well-defined, high-speed, repetitive problems, while DSPs are generally better at implementing highly complex algorithms …"

After a few missteps, we chose Altera's Stratix II, Stratix II GX, and Stratix III families of FPGAs to combine with the TigerSHARCs for this new approach. We now have boards that combine DSP and FPGA building blocks and allow them to be symmetrically or asymmetrically scaled, overcoming challenges such as I/O distribution, communication between compute elements, and host command and control.

VME: ***Some years ago you wrote a series of articles highlighting the superiority of the Analog Devices SHARC over the PowerPC architecture. Is that still true today?***

MILROD: For the record, I never claimed SHARCs were superior to PowerPCs. I did say that SHARCs are much better at real-time signal processing – especially for complex applications requiring multiple processors. I also said there are many other applications for which the PowerPC architecture is far superior to the SHARC's. These are very different processor architectures developed to solve very different problems. While many of the detailed processor implementations have changed – especially clock rates, power consumption, I/O bandwidths, memory management, and latencies – I still believe these statements are as true today as they were four years ago.

Extending the question to FPGAs is more complicated because FPGAs actually have embedded DSPs and PowerPCs! However, even with the latest advances in FPGAs, they will never truly be DSPs or PowerPCs and they weren't intended as such. They were intended to be FPGAs, which are simultaneously inferior and superior to DSPs and PowerPCs.

VME: ***What about the new Cell BE processor?***

MILROD: The Cell is a pretty awesome processor, but it's unclear what problem this technology was designed to solve, other than video gaming and high-end graphics. [Editor's note: the Cell BE was designed by IBM/Sony/Toshiba primarily for use in the PlayStation 3.] Using this processor for signal processing seems like trying to shove a round peg into a square hole. Furthermore, all the processing in the world doesn't do anyone any good unless it's useable, and by all accounts, this processor is very, very hard to use effectively. My guess is that the total available market for the embedded/COTS space is less than one week's worth of PlayStation 3 production, so I doubt IBM will be willing to invest in developing general purpose tools, training, or support for the embedded markets. [Editor's note: Full disclosure … IBM is working with OpenSystems Publishing, publisher of *VME and Critical Systems* magazine, to help spread the word about the Cell BE.]

Another major concern I have about this processor is the thermal problem. I've gone on record stating that I believe the future of COTS/embedded systems lies in small formats with switch fabrics. It's hard to see how a Cell processor could physically support that type of embedding when a whole board needs to be well under 50 W. Other multicore processors are sensitive to the requirements of portable devices, such as laptops, so they have much lower-power versions available. But even while these multicore processors may be valuable in embedded systems, it won't be for signal processing.

VME: ***Which markets do you serve, and what are the key trends in each?***

MILROD: High-end signal processing is a unique technology with many commonalities across applications. BittWare has

> "I guess I believe that what we've called 'VME' is technically mature, and is entering legacy mode."

primarily been a "horizontal enabling" company or one that can provide a broad range of markets with signal processing technology to enable their applications.

However, we are now finding more significant differentiation between markets and are crafting targeted signal processing solutions for the military/defense, telecom/datacom, and medical/instrumentation markets. While these markets are unique and typically use different system formats and interfacing, they share several key trends:

- The need to use different technology to address different parts of the problem efficiently, such as DSPs, FPGAs, and GPPs. I think most system designers now realize that a bigger hammer (i.e., more MIPs) isn't usually the best solution for a given problem.
- Desire to have high-speed serial, or SERDES, switch fabrics for all system-level communications. Different markets use different protocols, but they all want SERDES.
- Size, power, and life-cycle management are becoming more and more critical.

VME: ***What are some of the unique technical problems/challenges facing your target customers these days?***

MILROD: Board format chaos (VME/VITA 41/VITA 46/PMC/XMC/EMC/AdvancedTCA/MicroTCA/AdvancedMC …). Switch fabric protocol chaos (PCI Express/Ethernet/Serial RapidIO). SERDES signal integrity, especially over extended temperatures. Thermal/cooling. Moving single-threaded algorithms/concepts to hybrid architectures, and partitioning between DSPs and FPGAs.

VME: ***What does the future hold for VME?***

MILROD: I've been intimately involved in VME for 20 years, and it has been the foundation of my career. But quite frankly, I no longer know what VME is. I guess I believe that what we've called "VME" is technically mature, and is entering legacy mode. I think it's been that way for a long time – I can't remember the last VME board I saw that wasn't actually a PCI board with a VME bridge bolted on.

That's not to say I think VME is dead – far from it. I think the VME community has tremendous cohesiveness, tenacity, and a huge installed base and wealth of existing products. The knowledge, legacy, and success of VME is not something a sane person would abandon. Maybe what we call "VME" has more to do with mechanics and packaging than with buses and signaling.

All the new VITA specs admit that there has been a disruptive technology change in that they all embrace high-speed serial or SERDES switched fabrics as the primary mechanism for system-level data flow and command and control, rather than the outdated multiple master, arbitrated parallel bus. However, these new specs also recognize the value of the old adage: "If it ain't broke, don't fix it," by leveraging the mechanicals, building on the environmentals, and providing backward integration that attempts to bridge the technology disruption. I predict some core "VME-ness" will be a part of many leading-edge embedded systems for a very long time to come, and that most of those will be small formats with switched fabrics. **CS**

Jeff Milrod *is BittWare's president and CEO. Before joining BittWare in 1998, he gained extensive design experience at NASA and business experience at Booz Allen Hamilton. He then merged his technical expertise with his improvisational skills and started Ixthos in 1991, one of the first companies (along with BittWare) dedicated to COTS DSP. He then ran Ixthos until it was acquired by Dy 4 Systems in 1997. Jeff holds a bachelor's degree in Physics from the University of Maryland and an MSEE from Johns Hopkins University.*

To learn more, contact Doug at:

BittWare, Inc.
9 Hills Avenue • Concord, NH 03301 • 603-226-0404
Jeff_Milrod@bittware.com • www.bittware.com