FPGAs: Are they really worth it?

By Rodger Hosking

Driven by impressive claims that they are the ultimate silicon solution, FPGAs have gained considerable attention over the past decade. Virtually every major new advance in FPGA technology has targeted the specific needs of critical computing. Now, after developing and deploying a tremendous range of FPGA-based products, embedded board vendors and system integrators have gained a wealth of valuable experience. While reaping the promised benefits, valuable insights emerge in navigating the minefields unique to FPGA hardware and software development. Better understanding of these shifting insights can be useful as the attractions of FPGAs grow more irresistible each day.

With hundreds of powerful DSP slices clocking upwards to 500 MHz, today's FPGAs can outperform DSP and RISC processors by factors of 100 to 1,000 times. No other solution better matches the digital signal processing demands of increasingly complex communication waveforms and high-resolution radar systems. Data transfer paths are challenged by wider signal bandwidths, plus more communication channels or radar antenna elements per system. Here again, FPGAs offer strategic solutions with on-chip gigabit serial interfaces that support new switched fabric standards now emerging in open architecture embedded systems as well as desktop PCs, including PCI Express, Serial RapidIO, GbE, and InfiniBand.

"... FPGAs dominate the design by nicely assimilating all of the disparate functions of complex devices ... However, the numerous benefits of FPGAs often lure inexperienced board or system designers into dangerous waters."

Compared with ASIC or discrete logic alternatives, FPGAs offer embedded board designers an unbeatable combination of on-chip capabilities, including flexible digital I/O interfaces, state machines, DRAM memory controllers, Ethernet MACs, block RAM, FIFOs, shift registers, glue logic, timing generators, bus interfaces, and status/control registers. As a result, FPGAs appear somewhere in almost all new high-performance products. In many of them, FPGAs dominate the design by nicely assimilating all of the disparate functions of complex devices like the software radio transceiver XMC module shown in Figure 1. However, the numerous benefits of FPGAs often lure inexperienced board or system designers into dangerous waters.

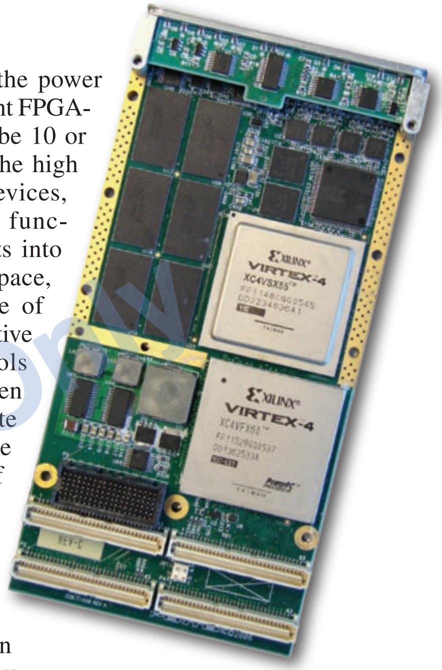

Figure 1

Compared to an ASIC, the power dissipation of an equivalent FPGA-based design can easily be 10 or 20 times higher. Given the high density of the latest devices, designers absorb the functionality of external parts into FPGAs to save board space, or to boost performance of algorithms for a competitive product advantage. Tools for predicting power often fail to deliver an accurate composite model of the normally diverse range of functional blocks within a single FPGA, each operating in response to often unrelated system events. The result can be product failures in the field, angry customers, and costly redesign. Such problems can only be avoided by conservative design, extensive modeling, and testing of the finished product using several representative high-density applications.

One project manager has likened maintaining schedules for FPGA development to the herding of cats. Often, a critical algorithm performs flawlessly on a workstation, and passes all FPGA modeling tests, but refuses to function within the FPGA device. In many cases, the problem is not with the core algorithm itself but with the interface to the surrounding FPGA circuitry. Teasing out the solution often requires the problem-solving skills of a seasoned hardware designer rather than those of the algorithm developer or software engineer. Even the most cohesive development team has to be carefully managed to circumvent delays generated by red herrings and finger pointing.

To promote the full potential of this great technology, board vendors are now offering customers finished FPGA-based COTS products with the promise that customers can extend functions by reprogramming the FPGA to add their own IP. Not unlike opening Pandora's box, many trusting customers have unleashed a host of problems due to convoluted factory FPGA structures, a lack of in-house experience, and inadequate design information from the board vendor. This can lead to highly frustrated customers, a huge technical support load for the board vendor, and costly missed deadlines.

To head off these issues, board vendors developing COTS products suitable for customer reconfiguration must carefully architect the FPGA structures for this purpose, starting from the earliest phase of product definition. VHDL code must be organized into discrete modules dictated by specific FPGA functions that are thoroughly documented, fully tested, and supported with complete and consistent interface definitions. Full VHDL source code and accurate and informative comments are essential. Test benches must be developed to help with modeling and simulation across crucial modes of operation. This gives customers a better chance to successfully modify or replace only certain critical modules, while leaving the others intact.

The quality of the design team's FPGA skills is another essential requirement for projects involving FPGA code development by customers. No board sale should be consummated without a realistic assessment of staff capabilities, followed by a plan for any required training or outside consulting services. Commercially available IP cores and steady improvements in high-level design tools are helping. Lastly, the board vendor must provide a knowledgeable customer support staff to ensure a successful experience.

Because of the myriad benefits of FPGAs, companies that develop products for critical computing applications – and make investments in hiring the engineering talent, cultivating their design skills, and acquiring the development tools to exploit them – should be paid back handsomely. Provided these basic guidelines are followed, the answer to the question: "Are FPGAs really worth it?" is a resounding "yes." CS

__Rodger H. Hosking__ is vice president and cofounder of Pentek, which started in 1986. With more than 26 years of experience in the electronics industry, he is responsible for matching new technology to advanced signal processing applications and for the definition of new products. He designed the first commercial direct digital frequency synthesizer and holds patents in frequency synthesis and FFT spectrum analysis techniques. Rodger has a BS in Physics from Allegheny College and both BS and MS degrees in Electrical Engineering from Columbia University.

To learn more, contact Rodger at:

Pentek, Inc.
One Park Way • Upper Saddle River, NJ 07458
201-818-5900, Ext. 228
rodger@pentek.com • www.pentek.com