MRSC 2008

A review of the FPGA high performance computing industry and the future role of FPGAs within data-centric processing architectures

Craig Petrie – Product Manager, Nallatech
c.petrie@nallatech.com, Tel. +44 1236 789 530

The computing industry's most significant period of change since the adoption of the PC is well and truly underway. The switch to parallel microprocessors is now a milestone recorded in history, but what impact has this irreversible step towards parallelism had upon the emerging FPGA high performance computing market and the role of FPGAs as accelerators in the future?

2007 was predicted by many to be the year when FPGA "technology push" would finally begin to translate into "market pull". There are significant indications that this is indeed the case.

Intel, recognizing the need for heterogeneality, honoured their promise to open up the Front Side Bus (FSB) interface to Altera and Xilinx as part of the "QuickAssist" accelerator initiative. This development has captured the imagination of the hardware world. FPGAs, despite being used extensively within markets such as Defense and Medical Imaging have often been dismissed as viable accelerators due to bandwidth and latency issues between the FPGA and the Host processor. Boasting bandwidths of up to 8.5GB/s and a latency of only 100ns, FPGA-enabled FSB accelerators are once again perceived as realistic alternatives for accelerating a wide range of compute-intensive applications.

In an attempt to bridge the gap between hardware and software environments, Intel also launched the "Accelerator Abstraction layer" (AAL) supporting FSB and PCI Express based accelerators. The AAL provides common protocols for communicating data and instructions to and from accelerators with common policies for managing memory and dealing with exceptions. This approach to defining a consistent interface for not only FPGAs, but accelerators in general is a significant step towards standardization – something that both evangelists and prospective customers within the embryonic FPGA high performance computing market have been desperate to see introduced.

Supporting this initiative, an eco-system of silicon vendors, C-to-FPGA compiler specialists and go to market experts, such as Nallatech, have been established to help enable and grow the FPGA high performance compute market. Encouragingly, this effort is not just a knee-jerk reaction to AMD's Torrenza program. It is part of a long term roadmap being developed by Intel, IBM and others called "Geneseo" that will embrace upcoming interface standards such as "Quickpath Interconnect" (QPI) and PCI Express 3.0.

In light of these promising developments and a renewed interest in FPGAs as accelerators, this presentation will look at the strengths and weaknesses of the FPGA in today's world of multi-core and many-core processors. In particular, it will discuss the role of the FPGA as a data-centric processing engine and explain, through illustration of example applications, how dramatic system-level improvements can be accomplished. It will also summarize the current state of the high level design tools and explain why optimized VHDL libraries will prove to be crucial to the adoption of FPGAs within high performance compute applications. Finally, this presentation will analyze business and technical aspects that make the FPGA a strong candidate for survival as an accelerator in the future.