Contribution ID: **165** Type: **Oral Presentation**

Construction of high speed, massively parallel, ATCA based Data Acquisition Systems using modular components

Saturday 11 June 2011 14:40 (20 minutes)

Current generations of HEP Data Acquisition Systems either in production or development are differentiated from DAQ systems used in other disciplines by the significant amounts of data they must both ingest and process, typically at very high rates. In practice this has resulted in the construction of systems that are in fact massively parallel computing systems. They are distinguished from their commercial counterparts by the significantly greater amount of I/O capability required between computational elements as well as the unique and disparate I/O requirements required at their interfaces. Given their unique requirements, traditionally, such systems have been purpose built by individual experiments. However, it has long been recognized that all these systems share a large degree of architectural commonality. SLAC is currently embarked on a research project intended to capture this commonality in a set of generic components which can be used in the construction of arbitrarily sized DAQ systems, tailored to satisfy a variety of different experimental needs. The design and implementation of these components (the RCE and CI) will be described. The RCE is a generic computational building block based on SOC technology that provides arbitrary combinatoric logic, generic computational ability using both embedded CPU and DSP tiles, and many channels of generic, high speed I/O with an Ethernet interface capable of operation from 1-40 Gigabits/second. The embedded processor supports arbitrary operating systems, but comes bundled with the RTEMS Real-Time kernel. RTEMS is an Open Source kernel containing POSIX standard interfaces as well as a full TCP/IP network stack. A full suite of GNU cross-development tools as well as embedded software is provided, allowing the RCE to be easily configured to the specifications of varying applications. The CI is a low latency, Layer-3 compliant, 10G-Ethernet switch connecting together RCEs to form arbitrary computing Clusters. ATCA is an ideal platform to express these components and we describe those features that make its use particularly attractive for this project. As one example, a PICMG 3.8 compliant ATCA Front-Board containing a Cluster of 8 RCEs will be described. The cluster's Ethernet is connected to the Fabric at speeds up to 40 Gigabits/second. While compatible with any ATCA backplane, the board is optimized to take advantage of a full-mesh topology. A shelf with such a backplane, populated with 14 of these boards, would result in a system of 14 fully connected clusters. That system would be capable of absorbing up to 7 Terabits/second of arbitrary input data and providing more than 1 Terabit/second of external Ethernet. 600,000 DMIPS, 20 PetaMACs, and 500 Gigabytes of RAM would be available to process and buffer the data.

Author: THAYER, Gregg (SLAC National Accelerator Laboratory)

Co-authors: PERAZZO, Amedeo (SLAC National Accelerator Laboratory); O'GRADY, Chris (SLAC National Accelerator Laboratory); HALLER, Gunther (SLAC National Accelerator Laboratory); PANETTA, Jim (SLAC National Accelerator Laboratory); WEAVER, Matt (SLAC National Accelerator Laboratory); HUFFER, Michael (SLAC National Accelerator Laboratory); CLAUS, Ric (SLAC National Accelerator Laboratory); HERBST, Ryan (SLAC National Accelerator Laboratory); TETHER, Steve (SLAC National Accelerator Laboratory)

Presenter: THAYER, Gregg (SLAC National Accelerator Laboratory)

Session Classification: Trigger and DAQ Systems

Track Classification: Trigger and Data Acquisition Systems