Camera Data Acquisition for the Large Synoptic Survey Telescope

A. Perazzo, R. Herbst, M. Huffer, C. O'Grady, L. Sapozhnikov, E. Siskind, D. Tarkington, M. Weaver
Stanford Linear Accelerator Center
2575 Sand Hill Rd.
Menlo Park, CA 94025 USA

***Abstract*—The Large Synoptic Survey Telescope (LSST) is a proposed ground-based 8.4-meter telescope operating in the visible band. The LSST Camera Control System (CCS) will manage the activities of the various camera subsystems and coordinate those activities with the LSST Observatory Control System (OCS). The most demanding component of the CCS is the Science Array Data Acquisition Subsystem (SDS). Its principal responsibilities are to receive science data from the camera, perform cross-talk correction, reorganize, buffer and, finally, distribute these data to down stream clients. SDS will use System on Chip (SoC) technologies and high speed protocols like 10 Gb Ethernet. This paper will describe the SDS requirements, its conceptual design and the current development status.**

I. Introduction

The Large Synoptic Survey Telescope (LSST) is a proposed ground-based 8.4-meter telescope operating in the visible band [1]. It is scheduled to start taking data in early 2014 atop Cerro Pachon in northern Chile. The LSST camera will provide a field of view of ten square degrees which will be sampled by a three billion pixel array of sensors. The telescope will take ten seconds exposures and will be able to cover the available sky every three nights. Its ability to acquire movie-like windows on objects that change on rapid timescales will allow it to observe exploding supernovae and detect potentially hazardous near-earth asteroids. It will also be able to observe apparent distortions in the shapes of remote galaxies produced by lumps of dark matter.

II. The Camera

A. Sensors

The camera is composed by a mosaic of 201 CCD sensors. Each CCD is a square array of sixteen million pixels. One pixel worth of data is represented as a sixteen bits quantity, resulting in a CCD data size of 31.25 MB.

The CCD sensors are organized in groups called *rafts*. There are twenty one rafts populated with nine sensors and four rafts populated with three sensors.

B. Data Acquisition System

The Science Array Data Acquisition Subsystem (SDS) is the most demanding component of the Camera Control System (CCS). Data are carried to SDS through twenty five optical fibers, one per raft.

The front-end electronic of each raft readout system will send its data through this optical link to one of the Configurable Cluster Element (CCE) boards. The CCE is the core component of the DAQ system and will be described in the next paragraph.

Data are delivered to SDS in no more than one second. This implies that the data rate from a fully populated raft is 281.25 MB/s, resulting in a total aggregate data output rate of 6.28 GB/s.

After receiving data from the rafts, a CCE will perform cross-talk correction before reorganizing, buffering and distributing these data to a dedicated 10 Gb switch.

While SDS will be at base camp, the data center will live at about forty kilometers from the telescope site.

III. Configurable Cluster Element

The Configurable Cluster Element is based on System On Chip technology (SoC). Currently it is built around the FX60 device from the *Xilinx* Virtex-4 family [3].

In this family of devices, a wide array of hard core blocks complete the system. These blocks include one or more generic processors, 10/100/1000 Mb/s Ethernet Media Access Control blocks (Tri-Mode MAC), 622 Mb/s to 10 Gb/s serial transceivers blocks (MGT), dedicated DSP slices and high-speed clock management circuitry.

The generic processor block implements the PowerPC 405 architecture.

The FPGA fabric in the device interfaces to the memory subsystems, a JTAG debug port, a multi-function display and various I/O channels.

The estimated power consumption for a FX60 device is about 7 W plus 0.75 W for each MGT which is actually used.

A. Memory

There are four memory subsystems in a CCE:

1) Processor Memory: this subsystem uses *Micron* RLDRAM II devices [2] to provide 512 MB of RAM.

2) Platform Flash Memory: this subsystem is used to store the firmware code which is loaded into the FPGA fabric.

3) Configuration Flash Memory: this subsystem uses a dedicated file system for storing software code and various configuration parameters. The user can select up to sixteen

Presented at 15th IEEE Real Time Conference 2007 (RT 07), 04/29/2007--5/4/2007, Batavia, Illinois

Work supported by US Department of Energy contract DE-AC02-76SF00515

separate images to bootstrap the device. The configuration memory uses *Samsung* K9F5608 chips to provide 128 MB of space.

4) Storage Flash Memory: this subsystem is interfaced through I/O channels using the Pretty Good Protocol (PGP). This is the same protocol adopted for the link between front-end electronics and the CCE and it will be described below. The storage flash memory subsystem features a low latency and a bandwidth of 1 GB/s. It uses *Samsung* K9NBG08 devices to provide up to 256 GB of storage space.

B. IP Cores

A high performance controller core for the processor memory has been developed to provide an overall bandwidth, shared between the I/O channels and the CPU, of 5.6 GB/s.

A generic DMA interface, called Packet Interface Core (PIC), has been designed as a set of VHDL blocks. The PIC eases the protocol implementation by providing a common interface, independent of the protocol itself, to efficiently transfer packets between networks and processor memory.

The PIC, in conjunction with the MGT blocks and the protocol cores, provides many channels of generic, high speed, serial I/O. In the current CCE design these channels will run PGP and 10 Gb Ethernet.

The PIC, in conjunction with the Tri-Mode MAC blocks, also provides commodity network interfaces (10/100 Mb/s and 1 Gb/s Ethernet).

C. Pretty Good Protocol

We developed the Pretty Good Protocol (PGP) for point-to-point full-duplex connectivity to/from the CCE. Its physical interface is serial with two LVDS pairs per lane. It features clock recovery ability, reliable frame transmission and reception, a deterministic latency and a small footprint. It is implemented as a VHDL protocol core interfaced to the PIC. It is extensible in both bit-rate and number of lanes.

D. Software

The CCE will use an open source real time kernel. The development language will be mainly C++ integrated with some C and assembly. The CCE will need a full TCP/IP network stack. In addition a specialized network stack for zero-copy UDP traffic will be developed.

IV. The Crate

The SDS crate will adopt an *Advanced Telecommunications Computing Architecture* backplane (ATCA) with sixteen 8U boards [4]. Two of the available slots will be used for network communication.

The first of these slots, called Fast Cluster Interconnect Module (FCIM), is a collection of managed 10 Gb switch devices. The second of these slots, called Slow Cluster Interconnect Module (SCIM), is a collection of commodity unmanaged switches.

The FCIM board will collect the data from the twenty five CCE boards through 10 Gb links running across the ATCA backplane. It will then distribute these data over 10 Gb links to the DAQ network clients.

The SCIM board will connect to the CCS network and it will exchange control traffic with the CCE boards over 1 Gb links running across the ATCA backplane.

V. Status and Conclusions

A. Firmware

At this time we have built the memory controller core, the flash memory interface, the PIC blocks, the PGP protocol core and the Tri-Mode MAC core. Next we will interface the PGP core and the Tri-Mode MAC core to the PIC. Finally, we will build the 10 Gb Ethernet core.

B. Software

We ported the real time kernel RTEMS [5] to our PMC evaluation board. We built a dedicated makefile system for software and firmware development. We coded the bootstrap loader, an exception model for debugging and a non-PIC network driver for the Tri-Mode MAC. Next we will code the PIC driver and all application code which is specific to the LSST.

C. Hardware

The firmware and software components mentioned above have been tested on a custom made PMC evaluation board. At the time of writing an ATCA based evaluation board with one CCE and no storage flash has just been built.

References

[1] http://www.lsst.org
[2] Data Sheet for the *Micron* 576 Mb CIO Reduced Latency (RLDRAM II), part number MT49H64M9, September 2005.
[3] http://www.xilinx.com
[4] http://www.picmg.org/v2internal/newinitiative.htm
[5] http://www.rtems.org