JRA3 Data Acquisition System Overview

M. Warren[1]
(on behalf of UK DAQ groups: Cambridge, Imperial, Manchester, RHUL and UCL)

December 03, 2007

Abstract

The data-acquisition system for JRA3 aims to build a DAQ for multiple detector prototypes based on generic components, in a modular structure, using commercial hardware and existing protocols where possible. A description of the design and current state of the prototype are give. Information on hardware and interface choices is provided.

[1] University College London

1 Introduction

The data-acquisition (DAQ) system being developed for EUDET by the UK groups is designed to readout multiple calorimeter prototypes. The design is also sensitive to other detector requirements and attempts to satisfy these requirements where reasonably possible. As such it is modular – with most of the hardware shared by all detectors. To further take advantage of economies-of-scale the system aims to make use of commercially available components and connections with standard protocols where possible. Any detector specific customisations are via the firmware loaded into the modules.

A further goal of the DAQ project is to use PCs as the off-detector readout system. This is as an alternative to traditional crate based systems, although not at the exclusion of other environments. For example the use of serial protocols and bus standards will allow transfer of much of the work done here to a crate based architecture in future projects if desired.

2 Data Acquisition System

The DAQ is split between on-detector and counting room (off-detector) components. These are connected using gigabit optical data-links that can span a few hundred metres and provide electrical isolation. It is planned to transfer configuration, clock and control signals onto the detector over these optical links too, but this requires specialised protocols and development, so in addition copper connection is provided.

On-detector components are located as close to the detector ASICs as possible within the constraints of geometry and cooling.

The system comprises four main components (shown in figure 1) which ultimately connect to a *detector unit* which hosts ASICs and sensors. These components are: the detector interface (DIF) that connects the generic DAQ and services to the detector unit; the link/data aggregator (LDA) that connects multiple DIFs to the off-detector data-link and the off-detector receiver (ODR) - a PCI express card hosting the optical interface to LDA. The clock and control system is dealt with in a later section.

For ease of development USB interfaces are provided on the LDA and DIF modules. This allows stand-alone testing of individual components and subsets of the DAQ chain to be used without the full system.

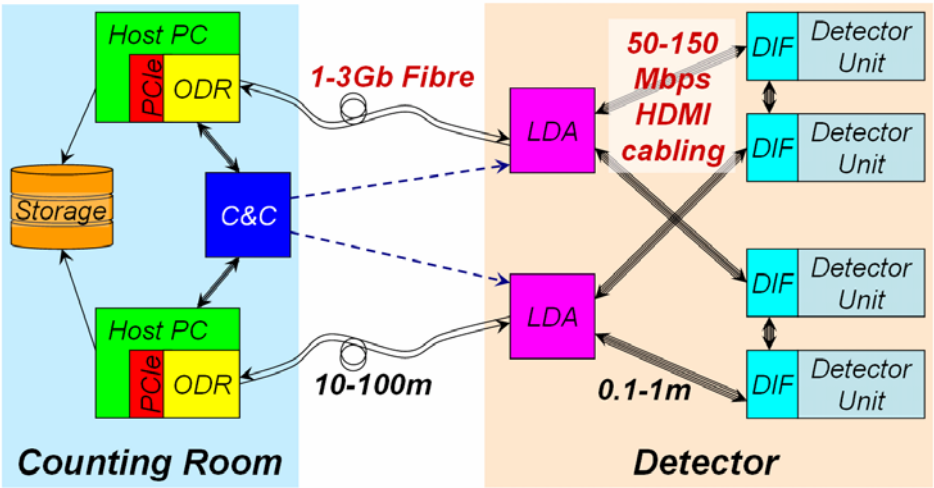

Figure 1. CALICE DAQ structure, showing the various modules of the system.

Hardware is also provided for low latency signals to be propagated asynchronously to the detector. Although this is certainly not required for the ILC, testbeam and other applications (e.g. cosmic tests) might need this.

2.1 Detector Interface (DIF)

The DIF straddles the boundary between a specific detector and a generic DAQ. It is physically connected to the detector unit and must match its geometry and connections, provide power and cooling and interface both physically and logically to the ASICs. The logic will be handled by an FPGA.

By defining a common interface, firmware can be written by each detector group to format its' data and controls such that they can integrate with the generic DAQ.

Provision is made on the DIF to connect to a neighbour. This will provide a link to the off-detector for a DIF even if its own connection has failed.

2.2 DIF-LDA Link

The DIF-LDA link connects the DIF to the rest of the DAQ. It must be small, of sufficient bandwidth and due to the number of links, low cost. The link provides clock, data up and down, fast-controls and a low-latency trigger using HDMI cabling.

HDMI [i] is a home entertainment system standard that defines cables and connectors with 4 shielded twisted pairs and 2 spare conductors in a small (13.9mm by 4.45mm) connector. These cables are commercially available, at low cost and can be sourced halogen-free. They are rated at more than 300Mb/s for data.

The data-link will run at a multiple of the machine clock (expected to be near 50Mbps) with the data encoded in a balanced protocol e.g. 8B/10B.

2.3 Link/Data Aggregator (LDA)

The LDA is a small FPGA based board tasked with connecting multiple DIFs to the off-detector. It has hardware to fan-out clocks, signals and data to the DIFs from the optical (and copper) up-link and aggregates data and monitoring from multiple DIFs, driving the off-detector optical link.

The LDA is positioned on the detector as close to the DIFs as possible for the shortest cable runs where convenient from a geometry point of view. It is expected that links will be between 30 and 100 centimeters in length.

The optimal number of DIFs per LDA must take into account the number of available pins on a low-cost FPGA, the bandwidth per DIF and the cost effective maximum bandwidth of the optical link. With the DIF link rated at 50Mbit, ~3Gbit optical links becoming standard and the number of usable pins on a small Xilinx Spartan 3 FPGA at around 200, 50 DIFs per LDA is the goal.

Physically the LDA has banks of HDMI connectors for connection to DIFs and a small form-factor pluggable (SFP) connector for the optical link off-detector.

The prototype version of the LDA is built on a commercial development board (*Enterpoint Broaddown2*[ii] with a *Xilinx Spartan3-2000* FPGA), with two purpose specific add-on boards. One provides the SFP and serialiser chipset for the optical link, and the second hosts 10 HDMI connectors along with clock fan-out hardware.

2.4 Off Detector Receiver (ODR)

The ODR is a PCI-Express card hosted in a PC. It collects data from an LDA where it is passed to (or fetched by) the PC and written to disk. The ODR is also responsible for sending data to LDAs and DIFs (and ultimately the ASICs) as well as clock and fast commands. Provision is also made for interfacing with fast-control (clock and synchronous command) sources.

We selected a commercially available development board for this function: the *PLDA XpressFX100*[iii] which is supplied with PCI-Express firmware. The board uses a *Xilinx Virtex4 FX100* FPGA and has 2 SFP gigabit interfaces already fitted, with the capability to add four more via plug-in modules.

The ODR FPGA firmware is responsible for receiving the data, and storing it in the card buffer memory. At present, the data stream is either data generated by the internal data generator (IDG) in firmware, or data received via on-board gigabit Ethernet interface.

The user interface to the ODR card contains two parts: a custom driver - mainly tasked with mapping card memory to the user space and providing direct memory access (DMA) support, and a "requester" program. The requester retrieves data from the ODR memory and stores it on the local disk. The requester is a fully programmable, multi-threaded application with separate threads for data transfer (DMA) from the ODR to the host memory, and IO threads for storing retrieved events to the local disk. The requester maximises data throughput. The current transfer rate as a function of the ODR data size is shown in figure 2 for the network, and internal data generator (ING) streams. A Maximum rate of 230 MB/s is achieved with no disk writes (i.e. ODR to memory transfer) and 123 MB/s when data is written to disk. In the case of ING, the upper curve shows transfer rate without data transfer to disk.

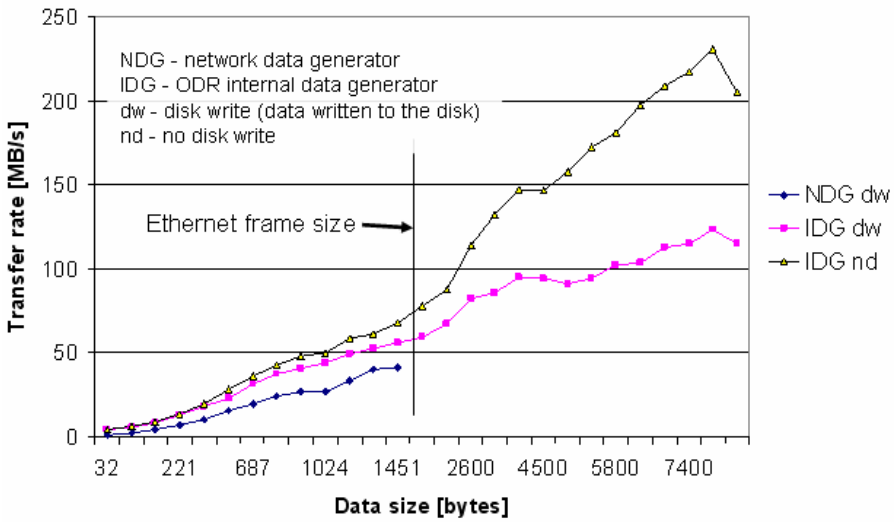

Figure 2. ODR through-put performance. All measurements with a single requester thread, single IO thread (disk write). Data is written to a disc-array, with Events stored in groups of 10000 events written to the same file.

2.5 Operation and Software

The DAQ software will be based on an existing system where possible. As the system is trigger-less our requirements are closer to those of a slow-control system: after configuration only train-start signals need to be sent to the detector. The ASICs do event selection and the DIF automatically co-ordinates sending of the data.

We are still in the selection phase, having examined EPICS[iv], ACE[v] and DOOCS[vi]. The latter is the most likely candidate as it is open source, in active development, provides slow and fast controls tools and is already used by ILC community.

3 Clock and Control

The goal of the project is to deliver clock and control signals using the data distribution system with commercial hardware and protocols. Commercial networking hardware is not suited to this task as it most efficient when it can buffer data and provides no guarantees on delivery times (by design). Similarly networking hardware built into modern FPGAs suffers from varying latency, as do serialisers in some devices.

In this case a custom protocol has advantages and our development will focus on at least fixing the byte-stream latency. This allows a single fibre uplink between the ODR and LDA. Similarly a single copper connection from LDA to DIF would be the best solution. In both cases an auxiliary means of connecting these signals is provided.

The clock and control (C&C) module must interface with the machine and provide stand-alone signal and clock generation. It then distributes these signals to the detector, either via the ODR or LDA. It will also receive a busy signal. The HDMI cables configuration used for DIF-LDA link is re-used here.

The C&C module will be a commercial FPGA board, with USB for communication with run-control, connected to a custom board providing hardware for a good clock, fail-over clock switching, low skew/latency fan-out and signals I/O connectors. The FPGA firmware will provide command encoders, busy filtering and signals enables.

Acknowledgement

This work is supported by the Commission of the European Communities under the 6th Framework Programme “Structuring the European Research Area”, contract number RII3-026126.

References

[i] HDMI – High Definition Multimedia Interface Consortium, http://www.hdmi.org
[ii] Enterpoint Broaddown2, http://www.enterpoint.co.uk/moelbryn/broaddown2.html
[iii] PLDA XPressFX, http://www.plda.com/prodetail.php?pid=48
[iv] Experimental Physics and Industrial Control System, http://www.aps.anl.gov/epics/
[v] Adaptive Communication Environment, http://www.cs.wustl.edu/~schmidt/ACE.html
[vi] Distributed Object Oriented Control System, http://tesla.desy.de/doocs/doocs.html