

The Digital Algorithm Processors for the ATLAS Level-1 Calorimeter Trigger

S. Silverstein for the ATLAS TDAQ Collaboration

Abstract— The ATLAS Level-1 Calorimeter Trigger identifies high- E_T jets, electrons/photons and hadrons, and measures total and missing transverse energy in proton-proton collisions at the Large Hadron Collider. Two subsystems – the Jet/Energy-sum Processor (JEP) and the Cluster Processor (CP) – process data from every crossing, and report feature multiplicities and energy sums to the ATLAS Central Trigger Processor, which produces a Level-1 Accept decision. Locations and types of identified features are read out to the Level-2 Trigger as regions-of-interest, and quality-monitoring information is read out to the ATLAS data acquisition system.

The JEP and CP subsystems share a great deal of common infrastructure, including a custom backplane, several common hardware modules, and readout hardware. Some of the common modules use FPGAs with selectable firmware configurations based on their location in the system. This approach saved substantial development effort and provided a uniform model for firmware and software development.

We present an in-depth description of the two subsystems as manufactured and installed. We compare and contrast the JEP and CP systems, and discuss experiences during production, installation and commissioning. We also briefly present results of recent tests that suggest an interesting upgrade path for higher luminosity running at the LHC.

Index Terms—triggering, level-1, first-level trigger, pipeline processing, real-time systems, programmable gate arrays, parallel architectures.

I. INTRODUCTION: THE ATLAS LEVEL-1 TRIGGER

THE ATLAS experiment at the CERN Large Hadron Collider (LHC) [1] is one of the most massive and complex detector systems ever built. To find evidence of rare, new physics processes, the LHC collides bunches of protons at a rate of 40 million bunch crossings per second in order to produce a handful of potentially interesting and rare events. To reduce the volume of readout data to manageable levels, the ATLAS sub-detectors store all sampled data in local pipeline buffers while reduced-granularity sets of data from the calorimeters and muon detectors are quickly sent to the Level-1 Trigger for fast analysis.

The Level-1 Trigger includes three systems: the Level-1 Calorimeter and Muon Triggers (L1Calo and L1Muon), and the Central Trigger Processor (CTP). L1Calo and L1Muon perform real-time pipelined analysis of detector data for each

25 ns bunch-crossing interval. Feature multiplicities and total and transverse energy results are forwarded to the CTP, which compares them against a trigger menu based on logical combinations of L1Calo and L1Muon result outputs, ultimately producing a Level-1 Accept (L1A) decision.

The total latency of the Level-1 Trigger – from when the collision occurs until the L1A reaches the detector – is fixed, and constrained by the on-detector pipelines to be less than 2.5 μ s; at the time of writing the measured latency is slightly greater than 2 μ s. A large fraction of this latency is due to signal transmission from and back to the detector front-end electronics.

The maximum L1A rate is nominally 75000 events per second, or less than 0.2% of the input rate. For these accepted events, data from all detector systems including the trigger are read out and transmitted to the data acquisition system (DAQ). The types and coordinates of features identified by L1Calo and L1Muon are separately read out to the Level-2 Trigger as so-called Regions-of-Interest (ROIs), which are used to seed selection of detector data for more detailed analysis.

A. Level-1 Calorimeter Trigger

The Level-1 Calorimeter Trigger (L1Calo) [2, 3] (Fig. 1) is a full-custom, fixed-latency, pipelined system located entirely in the electronics cavern adjacent to the main detector. It receives analog transverse-energy sums from approximately 7200 trigger towers, with typical granularity 0.1×0.1 in $\Delta\eta \times \Delta\phi$.

After reorganization and distribution through a set of patch panels, the sums are sent to an eight-crate PreProcessor (PPr) system that first conditions the analog signals and then digitizes them at 25 ns intervals with an ADC strobe that is adjustable to nanosecond precision. Digital ASICs (PPr-ASIC) then align the different channels with FIFOs, and a pipelined digital filtering algorithm including a peak-finder identifies the correct bunch crossing of the energy deposit. A lookup table is then applied for E_T calibration as well as thresholding for noise suppression, pedestal subtraction, and suppression of bad channels.

The PreProcessor finally sends its outputs over serial links to two digital algorithm processor subsystems operating in parallel. The four-crate Cluster Processor (CP) subsystem receives and processes the 0.1×0.1 granularity trigger-tower sums to identify isolated electromagnetic and τ /hadron cluster candidates, while the Jet/Energy-sum Processor (JEP) receives and analyses 0.2×0.2 transverse energy sums, or “jet elements”, to identify jet candidates and calculate the total and missing transverse energy for each event [4].

Manuscript received June 13, 2009. Corresponding Author: S. B. Silverstein, Stockholm University Department of Physics, AlbaNova University Centre, 106 91 Stockholm, Sweden (S. Silverstein may be contacted by phone: +46 (0)8 5537 8693; e-mail: silver@physto.se).

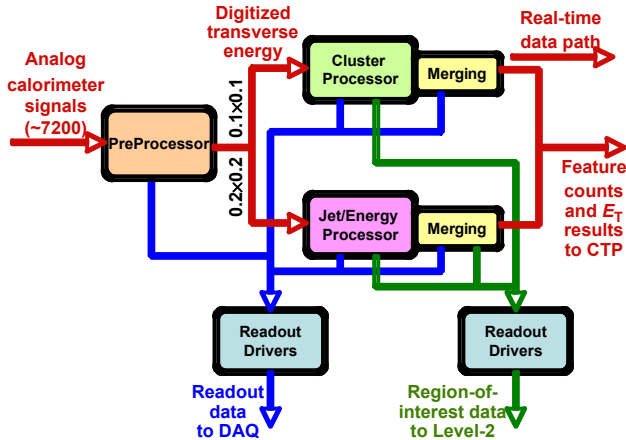


Fig. 1. Overview of the ATLAS Level-1 Calorimeter Trigger System. The real-time data path delivers results to the CTP every 25 ns, while events producing a L1A are read out to the data acquisition system and the Level-2 Trigger.

Real-time results from the two digital processor subsystems are sent on parallel LVDS cables to the CTP, and on a L1A detailed event data are read out to the DAQ and Level-2 trigger.

B. Outline of the paper

In this paper, we present a detailed description of the CP and JEP subsystems as produced and installed in the ATLAS electronics cavern. While different components of the two subsystems were developed by collaborators at different institutes, the architectures converged significantly over the specification and development phases of the project, and a great deal of common hardware was ultimately developed for use by both subsystems. Therefore, in the next two sections we first describe the important architectural features used in both systems, and then present each of the different boards in the system, their function and layout, and our experiences with them. Finally, we follow up with more general observations from production, installation and commissioning, as well as some recent test results that suggest an interesting upgrade path for future high-luminosity running.

II. SYSTEM ARCHITECTURE

The CP and JEP subsystems are both installed in modified VME crates outfitted with a common full-custom monolithic 9U backplane, water-cooled low-voltage DC supplies with heavy bus bars for power distribution, and custom hardware for supporting cables and interface boards at the back of the crate.

Besides the common crate and backplane, the CP and JEP use identical circuit boards, or “modules”, for the VME control/monitoring interface, clock alignment, timing distribution and control, and merging real-time data. In fact, the two subsystems differ only by the type of algorithm processing modules installed, and different FPGA firmware loaded for real-time data merging

Because the different modules share a number of common design features, we present in this section many of the key

features of the system, including signal handling, infrastructure and slow control. In the next section we describe the different hardware components in more detail.

A. Serial data input links

Serial data from the PreProcessor arrive at the CP and JEP over 11 m, shielded parallel pair cables using Low-Voltage Differential Signaling (LVDS). The cables are routed to vertical headers on the rear face of the backplane and passed via feed-through pins to the processor modules in front. There they are received and converted back to parallel by LVDS devices compatible with the 10-bit National Instruments Bus-LVDS serialisers used in the PreProcessor. Each link sends a 10-bit data word every 25 ns, plus two synchronization and framing bits, making a 480 MBaud serial stream.

The CP takes advantage of the fact that the peak-finder in the PreProcessor BCID algorithm never allows non-zero values in two consecutive bunch crossings. A so-called bunch-crossing multiplexing (BCMUX) scheme is used to multiplex two adjacent trigger towers in time and transmit them over a single link. In this way, the 2240 8-bit trigger tower sums to each CP crate use only 1120 link cables, with each link carrying 8 bits of transverse energy, one odd parity bit, and a tenth bit for BCMUX disambiguation. For the JEP, the jet elements are sums of four different trigger towers, so no BCMUX scheme is possible. Each link cable to the JEP carries a 9-bit electromagnetic or hadronic jet element plus one odd parity bit, for a total of 1408 links per JEP crate.

The LVDS deserialisers are somewhat sensitive to common-mode power supply noise, so they are powered by a “clean” 3.3 V DC supply separate from the main 5 V supply used for board logic.

B. Environment data sharing

Both the JEP and CP subsystems search for features with algorithms based on overlapping, sliding windows. To provide uniform algorithm coverage and avoid double counting of features it is important to share “environment” data across crate boundaries in ϕ and module boundaries in η . Sharing across crate boundaries is accomplished using duplicated links from the PreProcessor. Sharing between neighboring modules is done using short single-ended point-to-point links (so-called “Fan-In/Out” or FIO).

In the JEP subsystem, each module transmits and receives 33 10-bit transverse energy sums (electromagnetic plus hadronic) to and from its nearest neighbors. These sums are multiplexed to 80 Mbit/s, with the least significant bits sent first, requiring a total of 330 FIO signals between each pair of modules.

In the CP subsystem, neighboring modules need to both transmit and receive 120 8-bit trigger-tower energies, plus parity and BCMUX bits, with their neighbors. This is done at a data rate of 160 Mbit/s. With the modularity used, 320 FIO links are required between each pair of modules.

Low-voltage CMOS was shown to be suitable for data sharing at 80 and 160 Mbit/s during the prototyping phase of the project. For the production modules, the CP subsystem moved to the 2.5V SSTL2 signaling standard because of its lower drive currents and improved noise margin. The JEP

subsystem uses 1.5V CMOS for FIO signaling, with CMOS or HSTL sensing levels at the destination.

C. Real-time data merging

Real-time results from the individual processing modules are merged before the final results are sent to the CTP.

Two groups of 25 single-ended output links from each processor module each carry 24 bits of real-time results, plus one odd parity bit, to two modules on either side of the crate for merging. The links vary in length, from about 2 cm for processor modules neighboring merger modules to over 40 cm. Both subsystems use 2.5 V CMOS to drive the merger links, but while the JEP drives these signals directly from the processing FPGAs, the CP merger outputs are buffered through discrete drivers followed by series termination resistors.

The crate-level merging of real-time data is performed in two large FPGAs (Xilinx XCV1000E), each of which accepts up to 400 data and parity signals from the processing modules. One additional pair of FPGAs per subsystem (also Xilinx XCV1000E) receives data from the crate-level mergers, and produces system-level results. These are transmitted to the CTP as parallel LVDS signals over shielded twisted-pair cable.

D. Readout

All modules in the real-time data path are read out when a L1A is generated. At the very minimum, all real-time inputs and outputs must be monitored for data quality monitoring and diagnostic purposes. In addition, the types and positions of identified features as well as other information potentially useful for higher level triggering are read out as Regions-of-Interest (ROIs) to the Level-2 Trigger.

L1Calo uses two independent but functionally similar systems for DAQ and ROI readout. Both are based on the Agilent HDMP 1022/1024 “G-Link” gigabit-rate transmit/receive chipset, which encodes up to 20 bits of user data into a 24-bit frame transmitted serially at 960 MBaud. The encoded G-link serial outputs are connected to fibre-optic converters for transmission to the readout crates. An important feature of the G-Link protocol is a Data Available (DAV) bit, which is part of the G-Link data frame.

Data to be read out are assembled into 20-bit wide formatted data streams and transmitted to the G-Link inputs at a rate of 40 Mbit/s to a Readout Driver (ROD) module. The DAV signal on the G-Link transmitter is asserted while valid data are being transmitted, and the ROD uses the transition of the received DAV signal to recognize the arrival of a new readout frame.

During quiescent periods the G-Link transmitter sends fill-frames to maintain the lock between the transmitter and receiver. Because of jitter issues with the distributed LHC clock, the readout G-Links are clocked asynchronously by a local 40 MHz crystal oscillator on each module. Extra registers between the readout logic and G-Link inputs are used to safely re-time the readout data between the two clock domains.

ROD modules in two separate VME64x crates receive and collect readout data from the different processor modules in

the CP and JEP subsystems. Large FPGAs on the RODs reformat, compress and buffer the data before sending them on to the DAQ or Level-2 using a CERN-developed optical link protocol (S-Link).

E. Timing

Timing and trigger signals are distributed throughout the experiment by the Timing, Trigger and Control (TTC) system developed at CERN. The TTC system encodes the 40 MHz LHC machine clock, L1A signals from the CTP, machine orbit synchronization, and an additional “B” channel for slow control and configuration commands, into a 160 Mbit/s bit-stream that is distributed optically to the different subsystems.

A common TTC daughter card solution was developed and used by every module in the two subsystems. An on-board TTCrx receiver ASIC receives the TTC bit-stream electrically via the backplane. Two “deskew” LHC clock signals are extracted that can be individually phase-adjusted within the TTCrx with 104 ps precision. Many other output signals of the TTCrx chip are also extracted and distributed, including the L1A strobe, the LHC bunch-number counter, and B-channel commands. PLL-based circuits on the daughter card can optionally be used for jitter reduction of the deskew clock outputs.

F. VME-based configuration and monitoring

A commercial 6U VME single-board computer is used for configuration and control of the modules in each crate. The full VME specification supports a variety of address and data bus widths, and different data transfer modes. However, due to pin count and space limits on the backplane it was necessary to reduce the size of bus. So a reduced VME specification (dubbed “VME--”) was adopted, where only a single master asserting A24D16 cycles is allowed. This reduces the bus width to only 43 pins. For signal compatibility with the module FPGAs, the VME-- bus uses 3.3 V LVTTL levels instead of 5 V.

The commercial VME computer (CPU) is mounted on a 9U VME Mount Module (VMM), which serves as an adaptor between standard VME and VME--. In addition to providing signal termination and 5 V to 3.3 V level translation, the VMM also includes circuitry for reset signal handling, and more importantly, protection against inadvertent use of address modes other than A24D16. If an inappropriate Address Modifier (AM) code is detected, the VMM terminates the cycle by asserting a bus error to the CPU.

The CP and JEP subsystems share an extended geographical addressing scheme, encoded in designated backplane pins, that allows all modules to determine their unique position in the system (i.e. the crate number and position number in the crate). The VME base address is automatically set by each module according to its address.

G. Slow control

The ATLAS Detector Control System (DCS) uses CANbus to collect and monitor power and temperature data from electronics racks and crates in a standard way. The L1Calo subsystems also monitor temperatures and voltages on each module as a protective measure to help avoid damage to the many large FPGAs. Each crate has an internal CANbus on

the backplane, and a CANbus node is implemented on each module in a standard microcontroller (Fujitsu MB90FG94) that includes a 10-bit ADC. CAN node identifiers for each module are derived from the module's geographic address.

H. Firmware management

Several of the modules in L1Calo are multi-use, and require different FPGA configurations depending on their location in the system. To manage this, most modules use the Xilinx System Advanced Configuration Environment (System ACE) [5]. Firmware sets for each configuration option are stored on a CompactFlash card. On a system power-up or reset, the System ACE chipset can use three version-selection bits to select the appropriate firmware set, and configure the FPGAs via JTAG. The version select bits can be derived from the module's extended geographic address, allowing a replacement module installed anywhere in the system to correctly self-configure without user intervention.

III. HARDWARE

The CP subsystem is installed in four crates, each of which covers one ϕ quadrant of the calorimeter. Fourteen Cluster Processor Modules (CPMs) per crate each cover a "core" area of 4×16 trigger towers in η and ϕ , within a 7×19 "environment" of shared towers required for sliding-window searches for clusters whose positions are indexed by these towers. The CPMs occupy the most central positions in the crate, and are flanked by two Common Merger Modules (CMMs). The real-time results from each CPM are 16 3-bit multiplicities, eight of which represent electromagnetic cluster types, and the other eight of which can be either electromagnetic or τ /hadron clusters. These multiplicities are sent as two 24-bit parallel outputs to the CMMs, which sum crate- and system-wide multiplicities of these 16 cluster types.

The JEP is installed in two crates, each of which covers two opposing quadrants in ϕ . Eight Jet/Energy Modules (JEMs) cover each quadrant, for a total of 16 JEMs per crate. One JEM nominally covers a 4×8 core area of jet elements, within a 7×11 environment for the sliding-window based jet algorithm. The JEM real-time output includes either 8 3-bit jet multiplicities or 12 2-bit jet multiplicities (the latter for JEMs also handling forward calorimeter (FCAL) jets), as well as three 8-bit compressed-scale transverse-energy sums corresponding to ΣE_T , ΣE_x , and ΣE_y . The energy sums are sent to one CMM for total and transverse energy calculations, while the jet multiplicities are sent to the other CMM for global multiplicity summation and a jet multiplicity-based transverse-energy estimator known as "Jet E_T ".

In the following subsections we describe in more detail each of the components that comprise the CP and JEP systems.

A. Cluster Processing Module (CPM)

The CPM design (Fig. 2) evolved from an early concept based on G-Link receivers and two types of custom ASICs. The ASICs have since been replaced by FPGAs, but the production CPM broadly reflects the original concept, with many chips processing data in parallel, and a highly complex PCB providing connectivity between them.

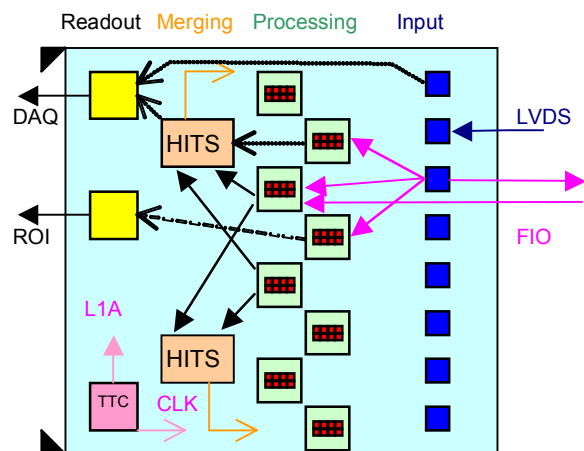


Fig. 2. Conceptual diagram of the Cluster Processor Module (CPM)




Fig. 3. Photograph of the production Cluster Processor Module (CPM)

The LVDS data links from the PreProcessor are received and deserialized individually by 80 LVDS receivers (National Semiconductor DS92LV1224). The 10-bit parallel outputs are then clocked into 20 "serialiser" FPGAs (Xilinx XCV100E), using the recovered 40 MHz clocks from the receivers.

The serialisers multiplex the input data to 160 Mbit/s and distribute them among the 8 CP chips (Xilinx XCV1000E) on the CPM, and to neighboring CPMs via the backplane FIO lines.

Each of the eight CP chips [6] covers a 4×2 core region surrounded by an environment of 7×5 trigger towers, so most of the trigger towers are shared between three CP chips on the same CPM, as well as up to three more CP chips on one of the neighboring CPMs. The 160 Mbit/s serialiser outputs are duplicated up to four times, with three of the outputs going to CP chips on the same CPM, and the fourth output to the backplane FIO. Same-CPM outputs are delayed by 3.125 ns to compensate for the later arrival time of the backplane signals. Sharing of FIO signals between multiple CP chips is necessary to reduce signal density, but reduces slightly the timing window for valid data.

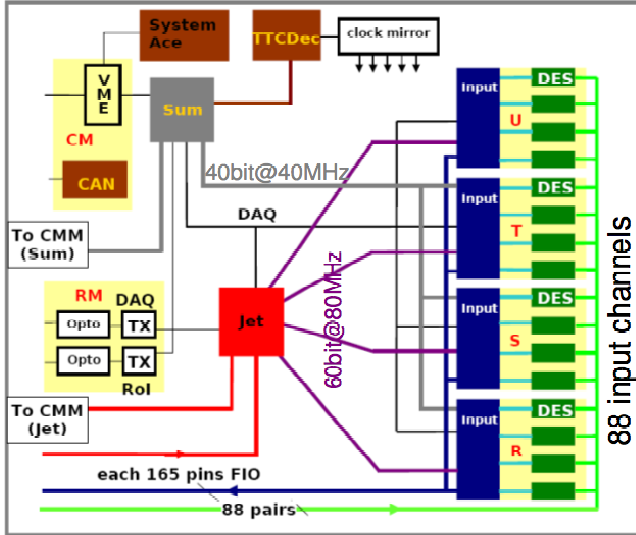


Fig. 4 Conceptual diagram of the Jet/Energy Module (JEM)

The serialisers also pipeline and read out the serial link input data to DAQ if and when an L1A is issued. It should be noted that while the data streams to the CP chips are still BCMUX and parity encoded for signal density reasons, the serialisers internally decode the trigger tower data before they enter the readout pipelines.

The CP chip receives the 160 Mbit/s streams from the serialisers, de-multiplexes them to the 25 ns LHC clock speed, decodes the BCMUX data and checks parity. The data then enter the pipelined CP algorithm to identify cluster ROIs.

Because the CP algorithm requires a feature to have a local maximum energy in overlapping 2×2 trigger-tower regions, each CP chip can identify a maximum of two ROIs in a single event. Therefore the CP chip real-time output contains 16 2-bit multiplicities divided into two 16-bit “hit words”. Each CP chip also pipelines and reads out the fine positions and passed thresholds of the two potential clusters as ROIs for use by Level-2.

The 2-bit cluster multiplicities from each CP chip need to be merged before they are sent to the CMM. Two “hit count” FPGAs (Xilinx XCV100E) each produce 3-bit multiplicity sums of eight cluster types, and then transmit these results as a 24-bit data word plus one odd parity bit through discrete buffers to the backplane merger lines. The real-time module outputs generated in the hit count FPGAs are pipelined for readout to DAQ.

Other important components on the CPM include two readout controller (ROC) FPGAs, which coordinate and retime readout streams from the different FPGAs, as well as extract and append the event’s 12-bit bunch-crossing identifier to the data. A PLD mediates communication with the VME-bus, and a microcontroller reports voltages, currents and FPGA temperatures to the slow-control system.

The final result is a highly complex 18-layer PCB (Fig. 3) with signal line widths down to 0.0762mm, over 20,000 solder joints and around 200 integrated circuits.



Fig. 5 Photograph of a production Jet/Energy Module (JEM)

B. Jet/Energy Module (JEM)

In contrast with the CPM, the JEM design (Fig. 4) began at a later point in the project, and underwent significant changes over several prototype and pre-production iterations. The final production JEM (Fig. 5) uses a smaller number of larger, more modern FPGAs, and many of the complex components are implemented on daughter cards to reduce the complexity of the main PCB [7].

The 88 LVDS data links from the PreProcessor are received and deserialised on four input daughter cards, each with four six-channel deserialiser devices (National Semiconductor SCAN921260) and an “input processor” FPGA (Xilinx XC2V1500). The input processor FPGAs produce 10-bit electromagnetic plus hadronic E_T sums for each jet element, multiplex the data to 80 Mbit/s and distribute them to a “jet-processor” FPGA (Xilinx XC2V3000) on the same JEM, as well as to neighbouring JEMs via the backplane. Sums of E_T , E_x and E_y for the “core” jet elements in each input processor are produced and sent to the same-board “sum-processor” FPGA (Xilinx XC2V2000). Both main processor FPGAs are mounted directly on the main board.

The jet-processor FPGA covers the JEM’s 4×8 core region surrounded by an environment of 7×11 jet elements, so any jet element is duplicated no more than twice. The 80 Mbit/s data distribution gives a much wider valid data window than for the CPM, and the same TTC clock can be used to time in both same-JEM and backplane data to the jet-processor FPGA. The 25-bit output of parity-protected jet multiplicities is driven directly from the FPGA to the jet CMM as LVCMOS levels.

The sum-processor FPGA produces final transverse and vector E_T sums of the JEM’s 4×8 core region. The 25-bit parity-protected energy-sum output is also driven directly from the FPGA to the energy-sum CMM using LVCMOS levels.

The jet- and energy-processor FPGAs have functionally identical ROC logic blocks to pipeline and readout results and

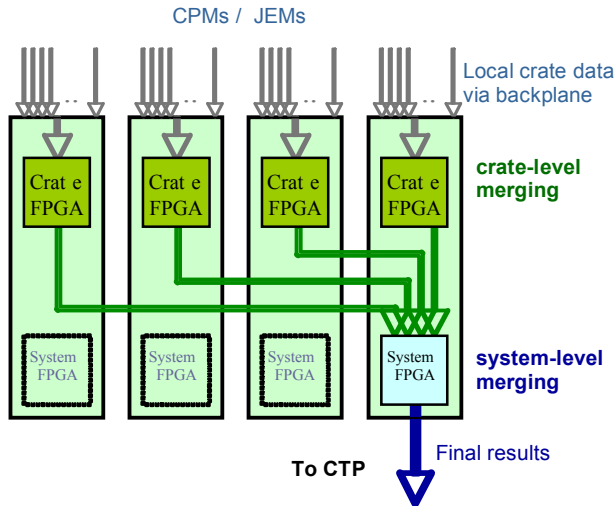


Fig. 6 Crate and system level merging by the CMMs

ROI data to DAQ and Level-2, respectively. Local readout sequencer blocks on the input processors receive control signals from the sum processor ROC and add their data to the DAQ readout streams. The G-Link inputs are driven directly from the input and processor FPGA outputs. Both G-Links and their respective fibre-optic transmitters are implemented on a daughter card near the front panel of the JEM.

An additional daughter card carries a CPLD for VME control, as well as a microcontroller for monitoring and reporting supply voltages and board temperatures to the slow control system.

The main board of the final JEM was held to only 14 layers by transferring much of the dense routing around the input processors to the four input daughter cards (12 layers). The decision to place G-Link and CAN/Control circuitry on additional daughter boards was because these designs had not been finalized at the time that the main board was completed. The end result is a less monolithic board with lower signal density than the CPM, but with more subcomponents as well as board-to-board connectors in the real-time data path.

C. Common Merger Module (CMM)

The CMM is a single hardware design for collecting and merging real-time cluster, jet and energy-sum data in the two subsystems. Two large FPGAs are responsible for nearly all of the functionality, including real-time crate- and system-level results merging as well as DAQ and ROI readout (see Fig. 6). A separate CPLD (Xilinx XCR3384XL) is responsible for the VME interface and firmware management, and a third FPGA (Xilinx XCV100E) provides an interface with the TTC daughter card.

Each subsystem crate contains two CMMs (Fig. 7), each of which receives up to 400 bits of data and parity from the processing modules. The “crate” FPGA (Xilinx XCV1000E) on each CMM receives all of these signals and produces up to 50 bits of crate-level results.

The two CMMs in one of the subsystem crates are designated as “system mergers”. These collect crate-level results from their own crate FPGAs, as well as those from




Fig. 7 Photo of a production Common Merger Module (CMM)

remote crates via 40 Mbit/s parallel LVDS cables. The “system” FPGA (also Xilinx XCV1000E) produces and reports up to 75 bits of system-level results via 40 Mbit/s parallel LVDS cables to the CTP. In CMMs that are not system mergers, the system FPGA is loaded with “dummy” firmware versions without merging functions, but which maintain uniform VME and readout behaviour.

A passive rear-transition module with three high-density cable connectors behind each CMM provides connectivity for the system-level merging. Two of the cable ports are bidirectional, so that they can be used either for transmitting or receiving data. Two output cable connectors on the front panel send system results to the CTP.

The Xilinx System ACE chipset is used to maximum advantage on the CMM. Depending on the data being merged (em, hadron/ τ , jet, or energy-sum) and whether or not the crate is a system merger, one of 8 different firmware sets will be loaded. During power-up, the CPLD automatically finds the exact geographic location of the CMM in the system from the backplane, and asserts the appropriate version selection bits to the System ACE chipset. Since all possible CMM firmware loads can be stored on the Compact Flash card, all CMMs are completely interchangeable without user reconfiguration.

D. Timing and Control Module (TCM)

The Timing and Control module (TCM) provides primary clock distribution in the two subsystems, as well as an interface to the ATLAS detector control system (DCS).

In its primary role, the TCM receives and converts a single optical signal from the TTC system to differential PECL format, and then duplicates and distributes it in a star configuration to the other modules in the crate.

Its second major role is to provide a CANbus processing node, with interfaces to DCS as well as the internal crate CANbus. A microcontroller on the TCM gathers temperature and voltage readings over the internal crate CANbus from the different modules and creates a single table of all values. DCS then gathers data and receives alarms directly from the TCM in each crate. The TCM microcontroller can also store data in




Fig. 8 Rear view of custom backplane with hardware.

a dual-port memory shared by the VME interface, allowing the local single-board CPU to read temperatures and voltages without connection to DCS.

A third, ancillary function of the TCM is to provide a diagnostic VME bus display, which is especially important for monitoring the VME bus in the custom processor backplane.

E. Processor Backplane (PB)

The processor backplane serves as the backbone of both subsystems. It is built on a full custom, monolithic PCB with 9U height and 21 module positions, and is populated mainly with male, 5+2 row, 2mm-pitch Hard Metric connectors.

The CP backplane requirements are a subset of those of the JEP, so the backplane includes slots for up to 16 processor modules, each accommodating up to 88 LVDS input cables, and 330 single-ended FIO links between nearest neighbors. Eight hundred diagonal point-to-point links carry merger data to the two CMMs, and the 43-line VME bus, TTC fanout from the TCM and a differential CANbus are distributed to all modules. The production backplane is a 4.9 mm thick, 18-layer board with eight signal layers. Traces are separated by typically 2 mm, and each signal layer is sandwiched between two ground planes, providing uniform impedance and minimal cross-talk. The CPM/JEM and CMM positions each have 1148 signal and ground pins, with a local signal-to-ground ratio as low as 4:3 when both ground-shield rows are used.

The very high module pin-counts lead to large insertion and extraction forces, typically around 450 N for a properly aligned module. Connector misalignment and bowing of the PCB greatly increase the forces, so a guide pin at the top of each module is used to ensure correct alignment, and six vertical reinforcement ribs made from solid brass reduce the maximum displacement during insertion or extraction to less than 1 mm. Mixed pin heights provide a small additional reduction in the insertion force, and durable, solid aluminum inject/eject handles compliant with the IEEE 1101.10 mechanics standard provide the needed leverage to insert and remove modules with thumb pressure.

The vertical reinforcement ribs also serve as mounting points for custom retention hardware (Fig. 8) that secure the

LVDS input cables to their input shrouds, as well as heavy copper bus bars for low voltage power distribution. Three high-current DIN connector pins at the bottom of each module position provide up to 20 A of 3.3 V and 5 V DC power, with a common signal-ground return. The backplane, retention hardware and bus bars are assembled in a single unit, expediting replacement or repairs.

Extended geographic addresses are encoded in every module slot by designated pins that are either grounded or left floating. The modules provide pull-up resistors and power to read the address from these pins. Three bussed crate-address lines are set with a rotary switch on the backplane to ensure that every module knows its unique position in the system for VME, CAN and TTC addressing as well as firmware selection.

F. Clock Alignment Module (CAM)

As described previously, processor modules in each crate share FIO data with their neighbours across high-speed point-to-point links. These links do not contain timing information, but rather rely on the relative phase adjustment of the TTC clocks on each module. As the TTCrx receivers on each module show some non-uniform behaviour and temperature dependence, the CAM was developed to help maximize the safe data timing windows within a crate by directly measuring phase differences between the deskew clocks in each module, as well as monitoring the phase of TTC clocks in other crates.

Each CPM or JEM provides a single-ended source-terminated CMOS clock signal on a front-panel coaxial connector. The CAM accepts up to 16 AC-coupled inputs from the processor modules, and two 16-way multiplexers can select any pair of clocks for comparison. A high-stability local clock is also recovered from the backplane TTC feed using very fast PECL logic to extract an 80 MHz clock signal from the raw TTC signal, and converting it to 40 MHz with a fast CPLD. This clock is available locally, and is also brought out to a fibre optic transmitter on the front panel. The CAM also has three optical receivers, allowing clock signals from other crates to be received.

At the time of writing, calibration and optimization procedures and software using the CAM have been demonstrated to work well in the timing-critical CP crates, yielding a notable improvement in the safe FIO timing windows. While the 80 MHz FIO links in the JEP crates are less critical, it is also planned to introduce CAM-based timing in that subsystem as well.

G. Readout Driver (ROD)

L1Calo has two crates of readout electronics, one for DAQ and the other for ROI readout to Level-2. Both crates are VME64x and populated with 9U Readout Driver (ROD) modules as well as a commercial VME CPU and a readout-crate version of the TCM that distributes the TTC signal over a custom J0 connector.

The ROD modules each receive up to 18 readout links on nine dual optical receivers, which are then converted to 40 MHz parallel data by G-Link deserialisers. Five large input FPGAs (Xilinx XC2VP20) collect and process the data, including parity checking, appending the correct bunch-crossing number, and data compression or zero suppression.

The processed data are stored in FIFOs until all required information has been received. A “switch” FPGA (Xilinx XC2VP30) then assembles the formatted data into a complete ATLAS event fragment and transmits the output to the readout-system and Level-2 over four S-Links at up to 160 Mbytes/s per link.

Like the CMM, the RODs are a single FPGA-based hardware design that uses different firmware to process readout data streams from different types of modules. The different firmware versions are stored on Compact Flash memory cards, and as for the CMM the Xilinx System ACE chipset selects appropriate firmware based on the module’s geographic address. Because the VME64x backplane only supports local and not system-level geographic addressing, the J0 connector used by the TCM also includes custom pins for designating the crate number.

IV. DEVELOPMENT AND COMMISSIONING EXPERIENCE

A. Common architecture and hardware

Over the development period of L1Calo, the architectures of the two digital processor subsystems have evolved substantially to the point that the same architectural features, and in many cases the same hardware, are used in both. An initial motivation for this was to avoid duplicated engineering effort between the member institutes, but other benefits became clear over time.

From a hardware perspective, common modules mean fewer types of components in the system, greatly simplifying spares policies. Multi-use, FPGA-based modules have enough built-in flexibility that new functionality can be added to the system if needed, such as threshold rate and parity-error histograms recently added to the CMM for online monitoring.

A large amount of software development was also necessary to configure and control the hardware and integrate the two subsystems into ATLAS TDAQ. Common hardware solutions also yielded similar or identical register maps and programming models for components of the two subsystems, substantially reducing the software development effort required.

B. Hardware development of CPM vs. JEM

The CPM and JEM followed different development paths. As we discussed earlier, the CPM followed a smooth development path from an early ASIC-based design, and the final production module maintained a similar overall architecture, with a complex, monolithic PCB populated with a larger number of smaller FPGAs. The JEM design was begun later and experienced several substantial revisions, eventually using two large FPGAs per board for energy and jet processing, and migrating much of the board complexity, including the input data handling, onto daughter cards.

Both approaches had their own risks and drawbacks. The high layer count and small feature size of the CPM made PCB manufacture a bigger challenge, and substantial hardware upgrades are basically ruled out. The JEM’s daughter card-based design led to less complex PCBs, and leaves open the possibility to upgrade many of the board components. But the use of board-to-board connectors can also lead to reliability

issues and potential impedance mismatches in the real-time data path.

Ultimately, both approaches were successful [8]. Although the CPM and JEM both encountered problems during production, the problems were due to manufacturing errors and not the module designs themselves. And both modules have been successfully installed and commissioned at ATLAS, and thus far have proved to work reliably. If there is any distinction to be drawn between the two modules, the CPM design – having been frozen earlier – uses older and slower FPGAs, limiting the performance potential for future upgrades.

C. Cabling

Using National Instruments Bus-LVDS serial links to transmit data from the PreProcessor has many advantages over earlier G-Link based designs, including lower cost and significantly lower power consumption. But one consequence was a doubling of the number of link cables to more than 7000. A massive cable plant (Fig. 9) of approximately 1900 4-pair cable assemblies needed to be carefully routed between the eight PreProcessor crates and the six digital processor subsystem crates. Custom supports and strain relief hardware were engineered to accommodate and support the cables, and the installation effort took months and significant manpower. But at least partly due to careful engineering and installation effort the cable plant has proved robust, with very few problems noted.

D. Backplane

The processor backplane provides the infrastructure for a relatively compact processor design, with features such as same-crate data merging and back-of-crate installation of the large numbers of input cables. Male Hard Metric connectors specified by international standard IEC 1076-4-101 provide the needed signal and ground pin densities and excellent signal characteristics. Disadvantages include high insertion forces and alignment issues as we have already noted. Another significant issue that arose during development and installation was occasional damage to individual male pins from accidental module misalignment, or unnoticed debris between the module and backplane.




Fig. 9 Cabling for the two JEP subsystem crates

To resolve this issue, we have investigated repair and replacement procedures. For single pin damage, Tyco/AMP produces a toolkit for field replacement of individual pins, which we have successfully used on multiple brands of male connectors. For more extensive damage the backplane should be removed from the system and entire connectors replaced by an external company.

V. OUTLOOK

Although at the time of writing L1Calo and the entire ATLAS experiment are still preparing for first collision data, planning for system upgrades is already underway. A two-phase luminosity upgrade for LHC is expected over the program lifetime. Phase 1, currently planned for 2014, aims to double or triple the $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ design luminosity through maximum exploitation of existing infrastructure. Phase 2, currently planned for 2018, will be a major upgrade, with a luminosity target of $10^{35} \text{ cm}^{-2} \text{ s}^{-1}$. The trigger and data acquisition systems will need to be upgraded to cope with the higher background rates without losing valuable physics.

While Phase 2 will require an entirely new L1Calo architecture, the timescale for Phase 1 precludes any major changes to the current infrastructure. Simply raising trigger thresholds in the existing hardware to reduce trigger rates would likely degrade interesting physics signals, so we are currently investigating how to augment the existing trigger with topological algorithms at Level-1 using ROIs in the real-time data path.

Recent transmission tests across the long data-merging lines of the processor backplane show that they can be run much faster than 40 Mbit/s. Using CMOS drivers and parallel termination at the destination (which is feasible for the JEM) very clean eye diagrams are seen at 320 Mbit/s. Series

termination at the source (as for the CPM) considerably degrades the signal, but a data rate of at least 160 Mbit/s appears feasible.

By increasing the CPM/JEM to CMM data volume at least four-fold, we see an opportunity to use the real-time data path for topological, in addition to just multiplicity data. The current CMM could be replaced with a new module that would gather all crate-level ROIs and transmit them over high-speed optical links to a new global merger subsystem that can execute both the existing multiplicity-based algorithms and a variety of additional, topological triggers.

REFERENCES

- [1] G. Aad et al., *The ATLAS Experiment at the CERN Large Hadron Collider*, 2008_JINST_3_S08003.
- [2] R. Achenbach et al., *The ATLAS Level-1 Calorimeter Trigger*, 2008_JINST_3_P03001.
- [3] J. Garvey et al., *The ATLAS Level-1 Calorimeter Trigger Architecture*, IEEE Trans Nucl. Sci., Vol. 51, p356-360, 2004
- [4] E. Eisenhandler, *ATLAS Level-1 Calorimeter Trigger Algorithms*, ATLAS note ATL-DAQ-2004-011, 2004
- [5] Documentation on Xilinx System ACE may be found at http://www.xilinx.com/support/documentation/system_ace_solutions.htm
- [6] J. Garvey et al., *Use of an FPGA to Identify Electromagnetic Clusters and Isolated Hadrons in the ATLAS Level-1 Callorimeter Trigger*, Nucl. Inst. Meth., A512, p506-516, 2003
- [7] J. Garvey et al., *ATLAS Level-1 Calorimeter Trigger: Subsystem tests of a Jet/Energy-sum Processor Module*, IEEE Trans Nucl. Sci., Vol. 51, p2356-2361, 2004
- [8] R. Achenbach et al., *Commissioning experience with the ATLAS Level-1 Calorimeter Trigger System*, IEEE Trans Nucl. Sci., Vol. 55, p99-105, 2008