

# Test-Beam Results from the ATLAS Level-1 Calorimeter Trigger Demonstrator

C. Bohm<sup>6</sup>, I. Brawn<sup>2</sup>, P. Bright-Thomas<sup>1</sup>, A. Connors<sup>1</sup>, J. Edwards<sup>5</sup>, E. Eisenhandler<sup>4</sup>, N. Ellis<sup>2</sup>, M. Engström<sup>6</sup>, P. Farthouat<sup>2</sup>, J. Garvey<sup>1</sup>, C.N.P. Gee<sup>5</sup>, A.R. Gillman<sup>5</sup>, P. Hanke<sup>3</sup>, R. Hatley<sup>5</sup>, S. Hellman<sup>6</sup>, S. Hillier<sup>1</sup>, E.E. Kluge<sup>3</sup>, M. Landon<sup>4</sup>, A.J. Maddox<sup>5</sup>, J.M. Pentney<sup>4</sup>, V.J. Perera<sup>5</sup>, U. Pfeiffer<sup>3</sup>, G. Schuler<sup>2</sup>, C. Schumacher<sup>3</sup>, T.P. Shah<sup>5</sup>, S. Silverstein<sup>6</sup>, R. Staley<sup>1</sup>, P. Watkins<sup>1</sup>, A. Watson<sup>1</sup>, M. Wunsch<sup>3</sup>

<sup>1</sup>School of Physics and Astronomy, University of Birmingham, UK

<sup>2</sup>CERN, Geneva, Switzerland

<sup>3</sup>Institut für Hochenergiephysik der Universität Heidelberg, Heidelberg, Germany

<sup>4</sup>Physics Department, Queen Mary and Westfield College, University of London, UK

<sup>5</sup>CLRC, Rutherford Appleton Laboratory, Chilton, Didcot, Oxon., UK

<sup>6</sup>Fysikum, University of Stockholm, Stockholm, Sweden

## Abstract

The ATLAS level-1 calorimeter trigger will utilise a number of advanced technologies, many of which have already been successfully demonstrated. To evaluate the different technologies associated with the important areas of high-speed data transport a large demonstrator system has been designed and operated during the last two years, using signals from prototype calorimeters in the ATLAS test-beam.

Using this system, inter-crate data transmission and reception have been demonstrated at over 1.4 Gbyte/s, with individual links running at up to 1.6 Gbaud. Operating with 160 Mbit/s signals across a transmission-line backplane, custom transceiver ASICs have achieved inter-module data fanout at peak rates above 800 Mbyte/s. With the addition of further modules, the system was extended to emulate a vertical slice through the ATLAS level-1 calorimeter trigger.

We present here the results from these tests, including measurements of bit-error rates across different data paths.

## I. INTRODUCTION

### A. The ATLAS trigger

ATLAS [1] is a general-purpose detector to be installed at the CERN Large Hadron Collider (LHC) which will produce collisions between proton bunches at 25 ns intervals at a centre-of-mass energy of 14 TeV. At a luminosity of  $\sim 10^{34} \text{ cm}^{-2}\text{s}^{-1}$  each bunch-crossing will produce an average of 18 proton-proton interactions, giving a total interaction rate of  $\sim 10^9$  Hz. The physics processes of interest (such as the production of Higgs bosons and other new physics) will only occur very rarely, so highly-efficient trigger processors will be essential to identify and select these events.

A three-level trigger system will be used, with level-1 employing fast hardware processing to reduce the 1 GHz interaction rate to a trigger rate of  $< 75$  kHz, as required by the detector electronics systems and the level-2 trigger. The trigger rate of selected events finally emerging from the level-3 event filter for storage and off-line analysis will be  $\sim 100$  Hz.

Detector data from  $\sim 10^7$  channels are stored in pipeline memories while a level-1 trigger decision is made, so to

minimise memory requirements the level-1 system must be designed to minimise latency. In ATLAS, the level-1 latency will be  $< 2.5 \mu\text{s}$ , which must include signal generation, transmission delays to and from the trigger system, as well as the trigger processing time.

The signatures used in the level-1 trigger to identify potentially interesting physics are:

- high- $E_T$  electrons/photons and hadrons/taus
- high- $E_T$  jets
- missing- $E_T$  and total scalar- $E_T$
- high- $p_T$  muons

All of these signatures (except for the muons) are derived from calorimeter data. 7300 trigger towers of typically  $0.1 \times 0.1$  in pseudo-rapidity  $\times$  azimuth space are formed on the detector by analogue summation of signals from both the electromagnetic and hadronic calorimeters, and the resultant signals are transported by 60 m shielded twisted-pair cables to the trigger processor electronics, where they are digitised at the LHC bunch-crossing rate of 40 MHz in the Preprocessor (figure 1).

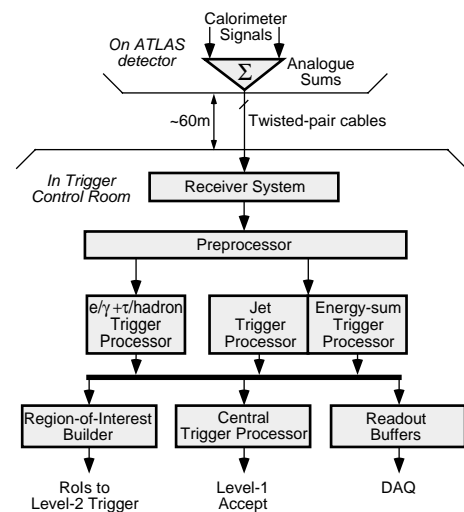


Figure 1: ATLAS level-1 calorimeter trigger architecture.

Considerations of Liquid Argon (LAr) calorimeter signal collection time and signal-to-noise performance require that

the raw calorimeter signals be shaped to produce pulses significantly wider than the 25 ns bunch-crossing interval, so an important function of the level-1 trigger system is the identification of the bunch-crossing (BCID) in which each signal originated. This is achieved by real-time digital analysis of the pulse shape. Before leaving the Preprocessor system, each trigger tower signal also undergoes pedestal subtraction, final  $E_T$  calibration and noise thresholding in look-up tables (LUTs).

The Preprocessor system transmits the data over high-speed serial links to two trigger processor systems. The  $e/\gamma$  and  $\tau$ /hadron Cluster Processor identifies and counts electromagnetic clusters consistent with isolated electrons and photons, and similar isolated hadronic clusters (produced possibly from  $\tau$  decays). The major challenge for this processor is to separate electromagnetic showers from the overwhelming jet background by the use of a two-dimensional cluster-finding algorithm and isolation vetoes. This algorithm is implemented in a cluster processor ASIC (CP ASIC). The Jet/Energy-sum Processor performs two separate tasks: identification of high- $E_T$  collimated jets of particles and calculation of their multiplicity, and calculation of the global missing- $E_T$  and total- $E_T$ .

Multiplicity information from these processors is fed to the level-1 Central Trigger Processor, where data from the muon trigger processor is also used to form an overall level-1 trigger decision for each bunch-crossing.

Region-of-interest (RoI) data are sent to the level-2 trigger processor system in order to highlight the regions of the detector that it must examine. For calibration and monitoring of trigger performance, input data and results are read out via the data acquisition system (DAQ) for off-line storage.

The implementation of the level-1 calorimeter trigger system outlined above is described in more detail in the ATLAS First-Level Trigger Technical Design Report [2].

## B. Technologies

The performance requirements of the first-level trigger are extremely demanding and require the use of several state-of-the-art technologies. To ensure their viability a large-scale demonstrator programme has been in operation for several years, some of the results from which have been reported earlier [3].

The technical challenges can be broadly grouped into the areas of data processing and data transport. The most demanding processing task is the operation of the  $e/\gamma$  and  $\tau$ /hadron cluster-finding algorithm, which must be implemented as a pipelined processor in order to minimise latency. A prototype ASIC [4], designed to evaluate this technique, was extensively tested in an earlier phase of the demonstrator programme. Digital techniques for bunch-crossing identification were also studied [5], [6].

The final phase of this work has concentrated on the challenging problem of data transport. Figure 2 shows the baseline solution for trigger-tower data transport from the ATLAS calorimeters to the  $e/\gamma$  and  $\tau$ /hadron cluster-processing logic, using several format transformations to match ASIC and module I/O requirements optimally at each stage. Data from each trigger tower are used in 16 concurrent

copies of the cluster-finding algorithm, implying extensive use of data serialisation to minimise pin-counts on ASICs and boards.

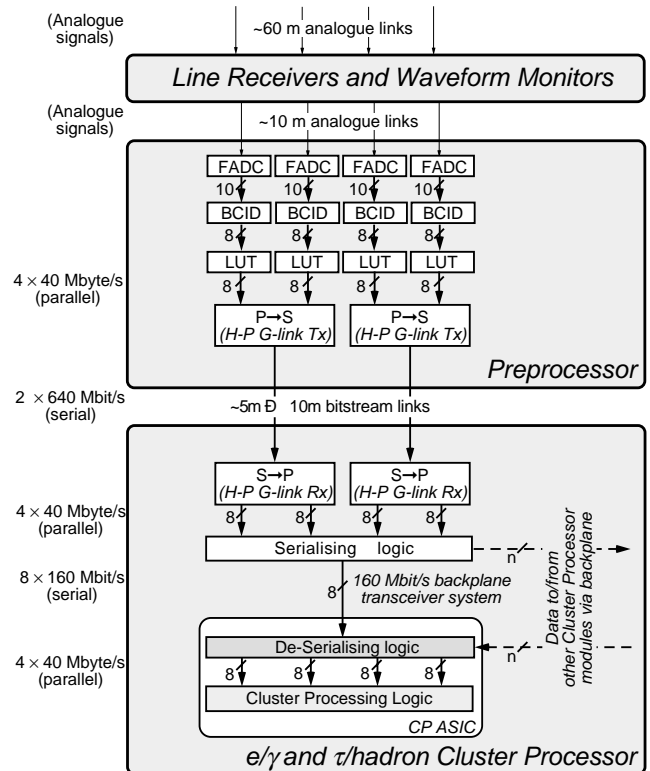


Figure 2: Data transmission chain —  $e/\gamma$  &  $\tau$ /hadron trigger processor

The key data transport technologies required for the trigger architecture, and therefore evaluated in this final part of the demonstrator programme, are as follows:

- Transmission of synchronous digitised data at  $>300$  Gbyte/s from the Preprocessor system to the trigger processors
- Data fan-out to multiple cluster-processing ASICs, using crate backplanes operating at 160 Mbit/s (single-ended)
- Serialised data input to the processing ASICs at 160 Mbit/s

With the addition of further modules the demonstrator system has also been successfully configured to emulate a complete vertical slice through the final first-level trigger.

As it is important to demonstrate these technologies in a realistic environment, emphasis has been placed on the operation of the demonstrator system in a test-beam with prototype electromagnetic and hadronic calorimeters.

## II. DEMONSTRATOR SYSTEM DESCRIPTION

### A. Hardware

The demonstrator system architecture was designed to be highly modular and hence scaleable. Alternative transmission technologies could be accommodated without major re-design work by the use of low-cost replaceable daughter-cards.

The basic data transmission chain (figure 3) transports analogue calorimeter trigger cell information ultimately to an electromagnetic cluster-finding ASIC. Following analogue transmission, the signals are digitised in a Flash ADC (FADC)

module and pre-processed in a Transmitter module (TXM). From there they are sent as serialised data to a Cluster Processing module (CPM), containing the cluster processing ASICs (CPASICs) and a 160 Mbit/s transceiver interface to the crate backplane for inter-module communication.

The full demonstrator system consists of a  $6 \times 6$  array of calorimeter trigger cells feeding nine sets of such data transmission chains, allowing the cluster-finding algorithm to fully process a  $3 \times 3$  array.

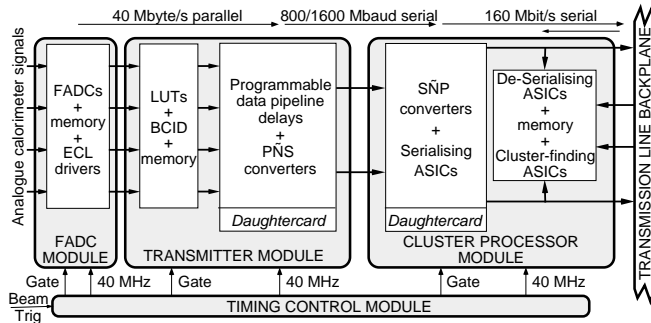


Figure 3: Data transmission chain — demonstrator system.

### 1) FADC module

A 4-channel 6U VME FADC module samples the calorimeter data at 40 MHz (the LHC bunch-crossing rate) with 8-bit resolution. (In ATLAS, the digitisation resolution will be increased to 10 bits.) Programmable pedestals allow observation of the full bipolar shape of the Liquid Argon (LAr) calorimeter pulses, and 256-byte scrolling memories provide for data capture, playback and injection of test data.

The real-time data are sent to the Transmitter modules as differential-ECL signals.

### 2) Transmitter module (TXM)

This is a 4-channel 6U VME module which pre-processes the incoming FADC data. Look-Up Tables perform pedestal-subtraction and calibration, and the BCID function is implemented in a Xilinx FPGA. Scrolling memories provide diagnostic and test facilities.

The data are delayed in programmable FIFOs by up to 255 system clock-ticks, then serialised by Hewlett Packard HDMP-1012 G-link transmitters [7] into 640 or 1280 Mbit/s bitstreams and transmitted to the Cluster Processor modules on 50  $\Omega$  cables. The G-link transmitters and associated logic are located on daughter-cards for easy evaluation of alternative link technologies.

### 3) Cluster Processor module (CPM)

This module converts the incoming TXM bitstreams to parallel data, and transports them via a 160 Mbit/s transceiver system to ASICs exercising the  $e/\gamma$  cluster-finding algorithm. It also fans these 160 Mbit/s data out to neighbouring CPMs via the high-speed point-to-point crate backplane, as required by the ASICs implementing the  $e/\gamma$  cluster-finding algorithm.

The Hewlett Packard HDMP-1014 G-link receivers [7], also located on daughter-cards, feed their data to RAL163 0.7  $\mu\text{m}$  CMOS ASICs [8] for serialisation at 160 Mbit/s, with two lines per trigger tower.

Although the final  $e/\gamma$  cluster-finding ASIC will receive serialised input data directly at 160 Mbit/s, the demonstrator system re-uses the prototype RAL114 cluster-finding ASICs [4] which receive only parallel data. Further RAL163 ASICs are therefore used to de-serialise the 160 Mbit/s data to parallel format for use by the RAL114 ASICs.

This 160 Mbit/s data transceiver system was designed for minimum clock and data skewing, to ensure that all channels of trigger tower data — both direct and from other CPMs via the backplane — arrive synchronously at the inputs of each cluster-finding ASIC. Diagnostic scrolling memories are provided for each channel.

### 4) Transmission-line backplane

In the demonstrator system, the fanout requirements of the  $e/\gamma$  cluster-finding algorithm demand that each CPM (occupying two crate slots) shares data with up to eight other CPMs in the crate, over path lengths ranging from two to eight slots. This is achieved with a high-speed point-to-point transmission-line backplane operating at 160 Mbit/s with single-ended ECL data.

Four 33  $\Omega$  layers accommodate signal striplines inside a 12-layer construction, with grounded guard tracks to minimise cross-talk. Standard 2 mm, four-row Futurebus+ connectors provide 192 connections for signals and power.

### 5) Timing and control system

A system-wide series of programmable-phase 40 MHz clocks advance the data along the processing pipelines under the control of the central Timing Control module (TCM), which also broadcasts control signals to freeze the scrolling memories for readout following a beam trigger [9].

## B. Software

Development of the necessary software to commission, control and monitor the real-time hardware, and in particular the use of object-oriented (OO) programming techniques, has proved valuable as a prototyping exercise.

### 1) Data acquisition

The DAQ software runs under LynxOS. Consisting of a number of UNIX processes centred around a buffer manager, it provides data readout and logging, readout-error tracking and simple histogramming. On-line analysis is limited to the monitoring and comparison of calorimeter pulse shapes along the processing pipelines to detect data transport errors.

### 2) Diagnostic tools

A user-friendly diagnostic software package was designed and implemented to perform the following major tasks: module debugging, DAQ software development, and module testing in the test-beam environment. The design requirements included multi-platform support, a graphical front-end, and module behaviour modelling.

Object-oriented software is particularly applicable to this system, where there are general similarities, but specific differences, between the various modules and between registers and memories. C++ was therefore used to implement

a library of classes providing the functions of generic modules, memories and VME access, with subclasses describing the individual TXMs, CPMs and TCMs together with their registers. The modelling package was also implemented in C++ using ideas drawn from VHDL.

The Tcl/Tk scripting package was used to provide the user interface, and Tcl commands were written to control the class library. The user can open windows giving access to all the registers and memories of any module and allowing various levels of register and memory tests.

### III. DEMONSTRATOR SYSTEM PERFORMANCE

#### A. The H-P G-link system

For simple interfacing to the RAL serialising ASICs, the G-link receiver chips were operated in positive-ECL (PECL) mode, which demanded the use of clean power supplies, adequately filtered from the TTL supplies.

The full system of 18 links has been successfully operated at up to 800 Mbaud (640 Mbit/s of user data) per link, both in the laboratory and in the CERN test-beam environment. Link-lock was very robust, with no losses experienced.

Operation at 1600 Mbaud was also studied, as this offered the prospect of configuring four trigger towers per link in ATLAS, with the consequent savings in cost, power and module real-estate. Tests showed that data transport was reliable and error-free at this rate, but link-lock was frequently lost during VME accesses. The increased sensitivity to power supply noise observed in this speed regime was probably exacerbated by PECL mode operation.

#### B. The 160 Mbit/s transceiver system

The transceiver system, consisting of the 160 Mbit/s serialising/deserialising ASICs (RAL163) and the associated point-to-point backplane, was designed to demonstrate the viability of transporting 160 Mbit/s data into processor ASICs and between modules via a backplane.

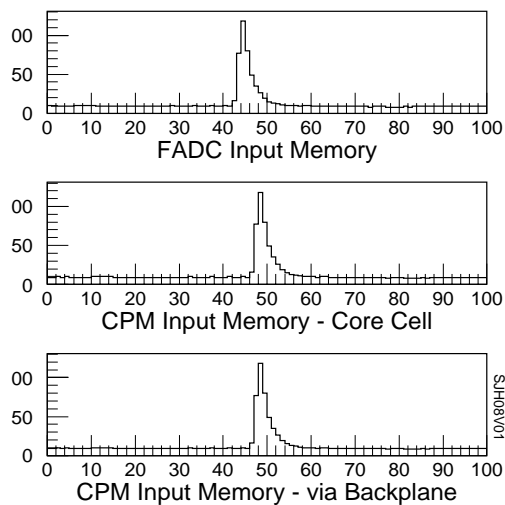


Figure 4: Tile calorimeter data — direct and via the backplane.

The RAL163 ASICs were successfully bench-tested up to 176 MHz, before incorporation into the CPMs. Performance measurements on the complete system then showed timing

margins of  $\sim 3$  ns, and signal crosstalk with four neighbouring signals switching to be below noise margins.

Figure 4 shows error-free propagation of a typical Tile calorimeter pulse from the FADC module to the inputs of the cluster-finding ASICs on two CPMs — one directly and the other fanned-out via the backplane.

#### C. Bit-error rate (BER) measurements

Data transmission errors within the trigger processor have the potential to cause unjustified additional trigger accept signals. The requirement that the rate of such false triggers should be low then imposes an upper limit to the transmission error rate. For the ATLAS calorimeter trigger, a maximum false trigger rate of  $\sim 1$  kHz corresponds to an average transmission error rate of no more than 1 part in  $10^9$ .

Measurement of errors at this level is very difficult by off-line analysis of recorded data, so a direct hardware solution was implemented. A purpose-built BER tester, which sampled the digital data from an FADC (at the start of the digital transmission chain) and at the cluster-processing ASIC inputs (at the end of the chain), was used in a configuration of five FADC-TXM-CPM chains, with the CPMs arranged to occupy five adjacent backplane positions. The central CPM therefore received data from modules either two or four backplane slots distant on either side. Any one of the 14 data channels ending at the central CPM could be studied, and for the chosen channel the tester compared all eight data bits on every 25 ns cycle. The chosen channel was populated with a pseudo-random data pattern covering the full dynamic range, and all other channels were driven with similar, but independent, pseudo-random data. Driving the rest of the system from a common source maximises the likelihood of channel-to-channel crosstalk.

After an initial set-up period, error-free transmission was observed throughout the system. The error rate limit is statistics-limited by the available running time. All 14 channels had no errors in 30-minute runs ( $BER < 2 \times 10^{-12}$ ) and all 11 tested in 16-hour overnight runs were also error free ( $BER < 6 \times 10^{-14}$ ).

In ATLAS, geometrical and latency considerations will limit link cable lengths to between approximately 5 m and 10 m, so studies were made of link operation under these conditions. Runs were initially taken with a short (2 m) cable carrying the 800 Mbaud TXM-CPM-G-link signal, but later studies used a range of longer cables formed by joining shorter lengths of 3 m, 4 m and 5 m. Five cells maintained error-free performance with 8 m cables, and the only two cells tested with 11 m cables were also error-free overnight. Errors were detected using 12 m cables, so further tests using longer cables without intermediate connectors are planned.

To ensure that the results were representative, two different CPMs were tested in the central CPM position, and three different TXM arrangements were tried. All showed error-free performance. The sensitivity of CPM timing was also checked, and it was found that no errors were produced until the CPM timing was advanced or retarded by more than 1.5 ns. This implies a stable region of  $> 3$  ns within the 6.25 ns (160 MHz) timing of the CPM backplane, which is more than adequate for robust performance of the system as a whole.

## D. Vertical slice operation

The addition of a Front-End module (FEM) [10], a Jet Processor module demonstrator (JPMD) and a Central Trigger Processor demonstrator (CTPD) [11] created a complete slice of the calorimeter trigger (figure 5). Using simulated calorimeter pulses, digitised data from 36 trigger towers were pre-processed in eight TXMs and one FEM and transported to nine CPMs, from where cluster hit flags were sent to the CTPD. The CPMs also passed real-time energy values to the JPMD, which in turn fed its jet hit flags to the CTPD. Using all these flags, the CTPD computed a trigger decision, indicated by a level-1 Accept (L1A) signal.

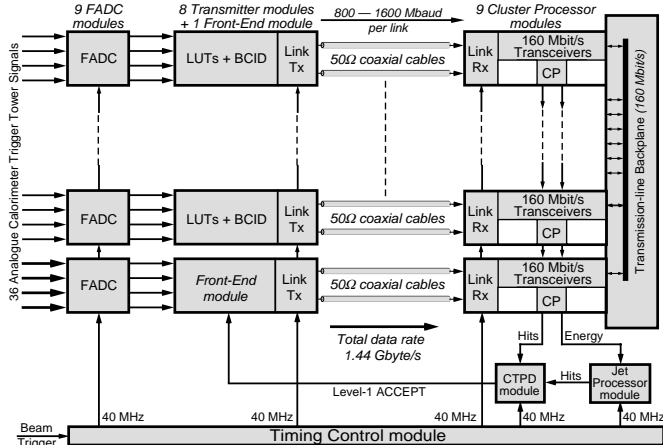


Figure 5: Complete slice demonstrator configuration.

Two modes of trigger operation were exercised, both emulating aspects of the ATLAS trigger system:

- Using the L1A signal as the event trigger to initiate readout. Reducing the simulated calorimeter pulse amplitudes caused the trigger rate to be cut to zero as cluster thresholds failed to be reached, demonstrating true trigger operation. The trigger saturated at the maximum readout rate allowed by the DAQ system.




Figure 6: BCID-processed data in FEM buffer memory.

- Using the L1A signal to initiate the fast pipeline readout in the FEM, where eight data timeslices surrounding L1A were captured in a double buffer without stopping the trigger pipelines. The buffer was read out when full, again without stopping the trigger pipelines. This “spy” mode allowed the L1A rate to be increased to several kHz, again limited by the DAQ system. Figure 6 shows a 2.5  $\mu$ s sample of the FEM buffer memory contents, where the BCID has zeroed seven in eight data timeslices. The pulse-height variation is due to asynchronous strobing of the analogue signal, producing samples at, or on either side of, the peak.

## IV. SUMMARY AND CONCLUSIONS

The overall success of this programme has demonstrated the viability of the baseline level-1 calorimeter trigger architecture for ATLAS. In addition, it has highlighted several areas where the original proposed design could be significantly improved, and some of these changes have already been incorporated into the ATLAS First-Level Trigger Technical Design Report [2].

Wide experience has been gained in the successful use of the H-P G-link chipset at 800 Mbaud. However, 1600 Mbaud operation, running in the PECL configuration described above, appeared insufficiently robust for the ATLAS trigger system using several thousand links.

Operating the G-links at this rate in an ECL configuration (or using the current TTL chipset) was not studied, as an alternative technique for doubling link channel density was proposed. Bunch-crossing multiplexing utilises the 25 ns timeslice immediately following an identified bunch-crossing, which the BCID algorithm guarantees to be empty.

The thermal management of the serial links ( $\sim 15$  kW system-wide dissipation) could be considerably simplified if a low-power alternative to G-link were available. The recently-announced low-voltage differential signalling (LVDS) chipset from National Semiconductor [12] appears promising and is being evaluated.

Experience of setting up the demonstrator transceiver system timing has resulted in significant architectural simplifications, whereby each CPM will share data with only its two nearest backplane neighbours (1-dimensional c.f. 2-dimensional fanout), leading to a considerably simpler and more uniform backplane design. The excellent BER results already achieved — considerably better than required in ATLAS — using the relatively complex demonstrator system architecture provide convincing proof that a robust error-free system can be designed for ATLAS.

## V. ACKNOWLEDGEMENTS

We would like to thank the ATLAS Liquid Argon and Tile calorimeter communities for their invaluable assistance in many parts of the programme.

This work was supported in the UK by the Particle Physics and Astronomy Research Council, in Germany by the Bundesministerium für Bildung, Wissenschaft, Forschung und Technologie and in Sweden by the Swedish Natural Science Research Council and the Swedish Council for Planning and Coordination of Research.

## VI. REFERENCES

- [1] ATLAS Collaboration, Technical Proposal for a General-Purpose pp Experiment at the Large Hadron Collider at CERN, CERN/LHCC/94-43 (1994).
- [2] ATLAS First-Level Trigger Technical Design Report, TDR 12, CERN/LHCC/98-14 (1998).
- [3] I.P. Brawn et al., Beam tests of a prototype level-1 calorimeter trigger for LHC experiments, Nucl. Instrum. Methods A349 (1994) 356.
- [4] V. Perera, Data on the LHC electromagnetic cluster-finding ASIC RAL 114, RAL (1992).
- [5] I. Brawn et al., Bunch-Crossing Identification for the ATLAS First-Level Calorimeter Trigger, CERN/LHCC/95-36, First Workshop on Electronics for LHC Experiments, Lisbon, 1995.
- [6] A. Mass et al., Design and Test of an ASIC for Bunch Crossing Identification in the ATLAS Level-1 Trigger, CERN/LHCC/95-36, First Workshop on Electronics for LHC Experiments, Lisbon 1995.
- [7] Low Cost Gigabit Rate Transmit/Receive Chip Set, Technical Data, Hewlett-Packard Company, 3000 Hanover Street, Palo Alto, CA 94304-1185.
- [8] J. Edwards and V. Perera, User guide for the first level calorimeter trigger processor demonstrator dual function ASIC RAL163, RAL (1994).
- [9] B. Hatley and J. Leake, Programming model for the trigger timing control module used in the phase-2 ATLAS bit-parallel calorimeter trigger demonstrator, RAL (1996).
- [10] U. Pfeiffer et al., Performance of the Front-End Demonstrator System of the ATLAS Level-1 Calorimeter Trigger, CERN/LHCC/97-60, Third Workshop on Electronics for LHC Experiments, London, 1997.
- [11] I. Brawn et al., A Demonstrator for the ATLAS Level-1 Central Trigger Processor, CERN/LHCC/97-60, Third Workshop on Electronics for LHC Experiments, London, 1997.
- [12] 16-40 MHz 10 Bit Bus LVDS Serializer and Deserializer, National Semiconductor, September 1998.