

The e/γ and τ/hadron Processor System for the ATLAS First-Level Trigger

V. Perera, J. Edwards, C. N. P. Gee, A. Gillman, R. Hatley, A. J. Maddox, T. P. Shah

Rutherford Appleton Laboratory, Chilton, Oxon., UK

P. Bright-Thomas, A. Connors, J. Garvey, S. Hillier, R. Staley, P. Watkins, A. Watson

School of Physics and Astronomy, University of Birmingham, Birmingham, UK

E. Eisenhandler, M. Landon, J. M. Pentney

Queen Mary & Westfield College, University of London, London, UK

P. Hanke, E-E. Kluge, U. Pfeiffer, C. Schumacher, M. Wunsch

Institut für Hochenergiephysik der Universität Heidelberg, Heidelberg, Germany

K. Jakobs, G. Quast, U. Schafer

Institut für Physik, Universität Mainz, Mainz, Germany

C. Bohm, M. Engström, S. Hellman, S. Silverstein

Department of Physics, University of Stockholm, Stockholm, Sweden

Abstract

The e/γ and τ/hadron first-level trigger system for ATLAS will provide electron/photon and tau/hadron trigger multiplicity information to the Central Trigger Processor (CTP), and region of interest (RoI) information for the second-level trigger processor. The system will also provide intermediate results to the DAQ system. This paper will outline some of the more interesting and challenging technologies we have studied, ranging from ASIC design, through high-density packaging, to exploitation of commercial high-speed links for the final system.

1. INTRODUCTION

The e/γ and τ/hadron cluster processor (CP) will process 0.1×0.1 (eta \times phi) electromagnetic and hadronic trigger towers covering a pseudo-rapidity region of ± 2.5 . It will receive 6400 eight-bit trigger tower signals from the preprocessor system, serialised and BC-multiplexed [1] to transfer four trigger towers per link at 800 Mbit/s on 2080 coaxial cables. The CP system will provide electron/photon and tau/hadron trigger multiplicity information to the CTP, and region of interest information to the level-2 processor. Also the system will provide intermediate results as well as the final results to the DAQ system for monitoring and diagnostic purposes.

The trigger algorithms (see Algorithms below) use sliding and overlapping windows, and each trigger tower participates in sixteen different windows, implying a massive fan-out of signals. To keep the fan-out and the pin counts of modules and ASICs to a manageable minimum, careful system partitioning is necessary.

The CP system consists of four crates of cluster processing electronic modules, each processing a quadrant of the calorimeter in phi and a complete space in eta (± 2.5). This is carried out by using 13 cluster processor modules (CPMs) in

each crate, each CPM processing a $16 \times 4 \times 2$ area of the calorimeter.

The partial trigger multiplicity results from all the CPMs are then merged in a separate crate using cluster merger modules (CMMs) and the results transferred to the CTP. The RoI information and the DAQ data from all the CPMs are transferred to read-out driver modules (RODs) in a separate crate. Figure 1 shows the block diagram of the system.

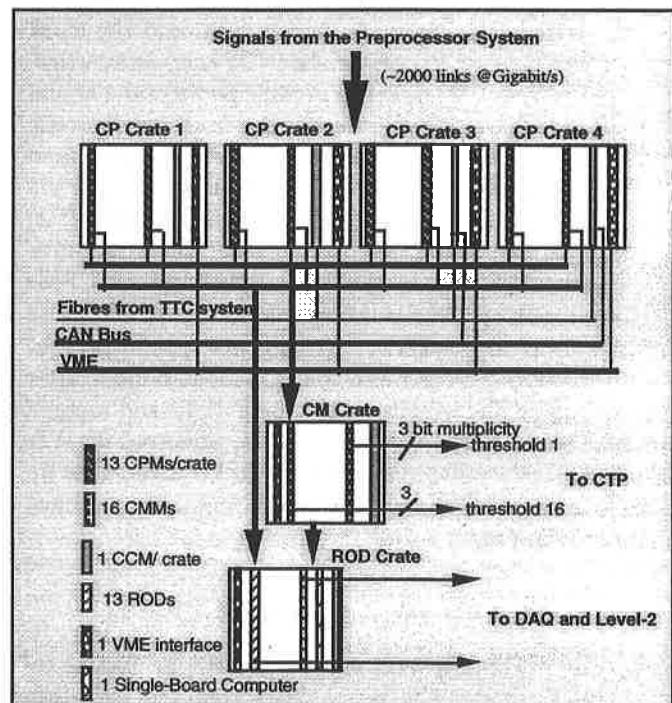


Figure 1: CP System

2. ALGORITHMS

Using a 4×4 sliding window and trigger towers of 0.1×0.1 in $\eta - \phi$ space (figure 2), the algorithm will search for isolated electromagnetic (e.m.) energy clusters and tau candidates and provide triggers and RoIs.

For the electron/photon trigger, one of the two vertical or two horizontal sums in the electromagnetic calorimeter must be greater than a cluster threshold (eight thresholds), and separate electromagnetic and hadronic isolation threshold criteria are imposed (eight thresholds). Also, the central 0.2×0.2 region must be a local E_T maximum compared to its eight overlapping neighbours. This also performs the de-clustering and defines the RoIs.

The t/hadron trigger algorithm is very similar to the electron/photon trigger algorithm but uses both the electromagnetic and hadronic calorimeters for the core energy calculations. More details on the algorithms can be found in reference [1].

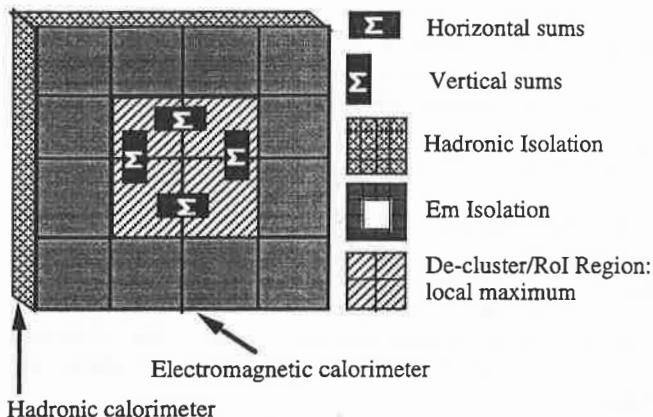


Figure 2: e/γ algorithm

3. TECHNOLOGY

The performance requirements of the ATLAS level-1 calorimeter trigger processor are extremely demanding in terms of advanced technologies. A great deal of work over the last few years has been involved in design studies of novel techniques and components, many of which are essential to the operation of the trigger and others of which could add significant improvements. In general, these studies have culminated in the design and fabrication of various items of hardware, most of which have been evaluated in a lengthy demonstrator programme. Whenever possible, the demonstrator system has been operated in the demanding environment of the ATLAS test-beam at CERN, and fed with signals from prototype calorimeters, although more detailed electronic studies and measurements have taken place in the laboratory.

Figure 3 shows the key technologies; such as Gigabit/s links, MCMs, 160 Mbit/s backplane, 40 MHz pipeline processing, etc., required within a 'data chain' from the preprocessor to the cluster processing electronics.

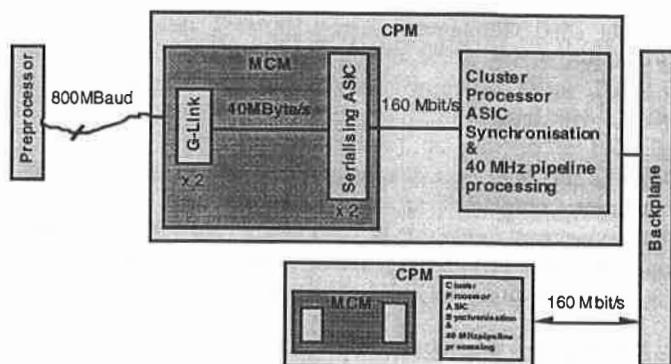


Figure 3: 'Data chains'

3.1 Serial Links.

The CP system has to process 6400 trigger towers. Given the algorithmic requirement to process overlapping windows, minimising fan-out implies maximising processing per module. If each CPM were to receive 8-bit parallel data from 160 trigger towers it would require 1280 connections. This would clearly be impractical, so it is proposed to transport the data serially. Serialisation at 320 Mbit/s would require one link per trigger tower, but by using Gigabit chip-sets (e.g. HP G-Link) and BC-multiplexing four trigger towers could share one link.

3.1.1 Hewlett Packard (HP) G-Links

HP G-Links (HDMP 1012/1014) [2] have been used in the trigger demonstrator system since 1995, driving optical fibres via Finisar [3] devices and driving coaxial cables. As the link lengths for the chosen architecture will be less than 10 m, the simple electrical solution is adequate and has been more extensively studied. Although designed to operate as ECL devices the G-Links may be operated in PECL mode, allowing direct interfacing to subsequent circuitry without additional conversion chips. Provision of clean power supplies, adequately filtered from the TTL supplies, is essential for reliable operation in this mode.

The demonstrator system has 18 links operating with two channels per link at 800 Mbaud. Nine of the links may alternatively be run with four channels per link at 1600 Mbaud.

The links have been successfully operated at 800 Mbaud both in the lab and in the CERN test-beam environment. Using a purpose-built real-time hardware bit-error rate (BER) tester, BERs better than 5×10^{-13} have been measured with coaxial links up to 11 m in length. To avoid increasing the ATLAS trigger rate by $> 1\%$ a BER $< 10^{-9}$ per channel will be required. Link-lock was very robust, with no losses experienced.

In the 1600 Mbaud dual frame mode two 16-bit words are time-multiplexed within the 25 ns clock period, thereby achieving a density of four trigger towers per link and requiring only half the number of links and chip-sets. Corresponding module real-estate and power density requirements are also halved. Lab tests showed the links to be robust and error-free until VME accesses occurred in the crate housing the G-link transmitters, when lock was frequently

lost. At 1600 Mbaud these chips appeared very sensitive to +5V power supply noise, especially as the links were operated in PECL mode. HP now offers a version of the chip-set to operate with standard TTL logic. Also, the chip-sets are only specified to operate up to 1500 Mbaud (0°C to +85°C), although this may be extended to 1800 Mbaud over a reduced temperature range (0°C to +65°C).

The requirement to operate at 1600 Mbaud has now diminished since the BC-multiplexing scheme, which also packs four trigger towers into each link has been adopted.

3.1.2 GaAs Custom Tx ASIC

We have collaborated with the Middlesex University Microelectronics Centre to design a low-power high-speed chip-set. In the first instance, a transmitter ASIC has been designed to operate at 1600 Mbaud with a maximum power dissipation of 650 mW. This chip was designed to interface with the HP G-link 1014 Rx chip so that it can be tested within the existing test infrastructure. The Tx ASIC is under test at RAL at the time of writing this paper.

3.1.3 ECL "Simple Link"

This is a customised serial data link designed to run at 320 Mbit/s using differential ECL. At the transmitter end the eight bits (no error detection bits or clock encoding) are simply serialised at a 320 MHz rate. At the receiving end the bit-stream is latched on to two latches with anti-phase 160 MHz clocks to generate two 160 Mbit/s data streams. This scheme requires a clock and data alignment strategy as described in section 3.1.6. For testing, daughter cards have been designed to replace the G-link daughter cards in the demonstrator system. This scheme will not require any additional ASIC such as the serialising ASIC used with the G-Links, but additional circuitry would be needed for 'spying' on the data.

3.1.4 LVDS Chip-Sets

An alternative to the above scheme would be the use of commercial low-voltage differential signalling (LVDS) chip-sets, such as 'channel link' from National Semiconductors. As they are designed for the portable PC market to interface colour LCD displays, the modularity is three bytes (three colours) with four control bits, somewhat inconvenient for the trigger system. However at the time of writing this paper National Semiconductors have announced a new chip-set, which is a low power 10-bit serialiser/deserialiser which can operate at 40 MHz [4]. We will be investigating the possibilities of using this chip-set to reduce system power.

3.2 160 Mbit/s Serialiser ASIC

As shown in figure 3, in the demonstrator system the serial data are transmitted from the preprocessor system at 800 Mbaud using G-Link transmitter chips, and are received by the corresponding receiver chips and converted to parallel data. To reduce pin counts on backplanes and cluster processor ASICs, a dual function ASIC (RAL 163) was designed to demonstrate the interfacing between a G-link receiver chip and a cluster processor ASIC receiving data at 160 Mbit/s, as well as

module-to-module data transfers via the backplane at 160 Mbit/s.

The first function of this ASIC is to convert a 16-bit parallel word every 25 ns into four serial links operating at 160 Mbit/s. Hence two 160 Mbit/s links are required to transfer 8-bit data to a cluster processing ASIC within the 25 ns bunch crossing interval.

The 16-bit parallel word, together with the G-link error signal, can be captured in a 80-deep (2 µs) dual-port memory, to be read out following a level-1 YES decision. This memory can also be used in a playback scheme for test purposes, and for diagnostics.

The second function of this ASIC is to convert the 160 Mbit/s data back into parallel data to be compatible with the demonstrator cluster finding ASIC (see section 3.3). This functionality demonstrates the operation needed for the front-end of the final cluster-finding ASIC by receiving the 160 Mbit/s serial data, de-serialising them, and synchronising the resultant parallel data to the 40 MHz system clock before they enter the algorithm processor, using 6.25 ns delay elements.

The dual-function ASIC was designed using the 0.7 µm CMOS technology from ATMEL-ES2, and has been tested successfully on the bench up to 176 MHz, as well as being used during the 1996 and 1997 beam tests in CERN.

3.3 Cluster Processor ASIC

The demonstrator cluster processor ASIC (RAL 114) was a cut-down version of the final CPASIC, processing only one trigger tower in the e.m. layer.

It was a 0.8 µm, 20,000 gates CMOS gate-array design to demonstrate the implementation of the cluster-finding algorithm using pipeline processing elements of adders and comparators.

When it was designed in 1993, the LHC bunch-crossing period was foreseen to be 15 ns, hence the ASIC was designed to work at 67 MHz.

This ASIC has been used since 1993 at various stages of the demonstrator program, and has been interfaced with the 160 Mbit/s data streams using the dual function ASIC (RAL 163) as described above, to convert the serial data to parallel data. In the final CPASIC, the data will be received on 160 Mbit/s serial links to minimise pin counts and maximise processing, with the 160 Mbit/s to 40 Mbyte/s conversion performed internally.

3.4 Timing and Synchronisation

In the final CP system, the G-link transmitters will use the 40 MHz LHC clock and the internal phase-locked loop circuits to multiply the clock to the required bit-rate clock. The 40 MHz clock will then be extracted at the receiving end by the G-link receiver chips, and phase-locked to the on-board 160 MHz clock to generate the 160 Mbit/s data streams. Unlike the high-speed serial links, the 160 Mbit/s data links do not have a clock recovery scheme. Therefore when the 160 Mbit/s data are received on the CPASICs they must be aligned to the 160 MHz clock. After conversion, the 40 Mbyte/s data then have to be synchronised to the 40 MHz LHC clock before entering the cluster finding logic.

In the 36 channel demonstrator system, clock alignment and synchronisation were performed manually using programmable delay lines, which was very time consuming, so to align 6400 trigger towers in the final system, an automated process will be required.

An ASIC (RAL 215) was designed to evaluate such a process. Using delay-locked loop (DLL) techniques the delay lines were built into the ASIC. The calibration logic on the ASIC scans through the 160 MHz clock phases (five 1.25 ns taps) and through a statistical process selects the appropriate clock phase to capture the calibration data. Then data are synchronised to the 40 MHz clock by using 6.25 ns delay elements. This process takes approximately 1.5 μ s. Figure 4 shows the block diagram.

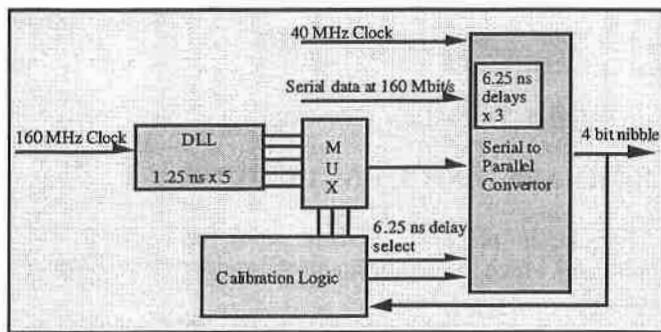


Figure 4. Clock Synchronisation Logic

3.5 Backplane

A 3U section of the 9U backplane on the demonstrator cluster processing crate is a high-speed transmission-line backplane operating at 160 Mbit/s single-ended. The high-speed backplane is required to fan-out trigger data within a crate to the neighbouring cluster processor modules. The signal transmission is point-to-point using ECL drivers, with path lengths ranging from two to eight slots. Figure 5 shows the multi-length fan-out from module to module. Nine double-width modules were used to fully process a 0.3 (η) x 0.3 (ϕ) calorimeter window.

Standard 2 mm, four-row Futurebus+ connectors are used, having 192 connections for signals and power.

The backplane was manufactured with 33Ω transmission-line tracks, using strip-line design in a 12-layer construction with four signal layers and eight power and ground layers. Ground guard tracks between signal tracks are used to minimise cross-talk.

Lab tests indicate that the backplane can be operated with a bit-error rate better than 10^{-13} .

Since the Technical Proposal [5] was written in 1994, the architecture has changed, and with the proposed ϕ -quadrant architecture as described in the Technical Design Report [1], each CPM will cover a quadrant in ϕ and 0.4 in η . The modules processing neighbouring regions in η will be in adjacent slots in the same crate, hence all backplane paths will be only one slot in length. The final backplane is therefore much less complex than the demonstrator version, which can already provide a performance considerably better than required.

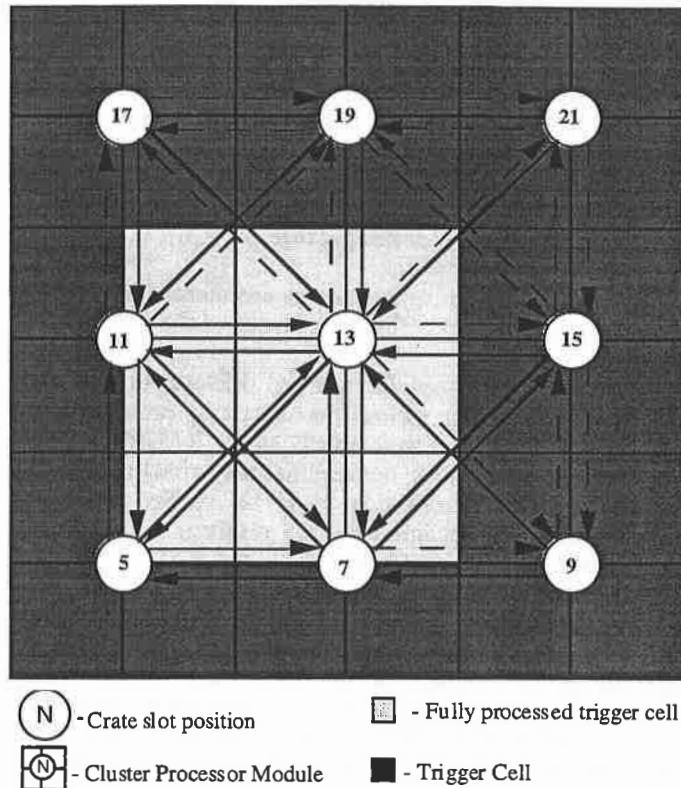


Figure 5. Backplane Multi-Length Fan-out

3.6 Multi-Chip Module (MCM)

The use of multi-chip module packaging technology brings many benefits such as:

- package efficiency (die size to package size), giving the capability of implementing many channels per module
- high-speed signals and interconnect routing confined to a small area (minimising the requirement of transmission lines on the PCBs)
- good EMC performance, etc.

The cluster-finding algorithms operating in the CPMs require a high degree of trigger tower fan-out to neighbouring modules, which should be minimised by maximising the processing in each module. Architectural studies indicate that data from 160 calorimeter trigger towers should be fed directly to each CPM from the preprocessor modules, thus demanding compact MCM packaging solutions.

Studies have been carried out at RAL into the use of MCM-L (laminate) technology with ball grid array packaging, and lead-less chip carrier packaging built-in to the substrate, as well as MCM-C (ceramic) technology. For reasons of thermal management (~5 W per MCM) and cost, MCM-C technology was chosen for the demonstrator design.

The demonstrator MCM incorporates two HP G-link dies (HDMP-1014D) to receive the trigger tower data at 800 Mbaud and convert them to 16-bit parallel words, and two serialiser ASIC dies (RAL163) which convert each 16-bit parallel word into four 160 Mbit/s bit-streams for the cluster-finding ASICs.

A summary of the MCM specification is given below

- Ceramic substrate (96% Alumina)
- Three metal (gold) layers

- Standard thick film technology
- Four dies (2 HP G-Links and 2 RAL163 ASICs) and printed resistors on substrate
- Capable of handling 200 ps rise-time signals on the substrate
- All components contained within a 30 mm \times 30 mm area
- All dies bonded to the MCM using 33 μ m gold wire
- 120-pin hermetically-sealed package
- 6 watts power dissipation

The complete specification, design and thermal modelling were carried out at RAL. The track layout of the MCM was carried out by industry.

The MCMs were manufactured and delivered to RAL in September 1997. After various manufacturing problems had been resolved the devices were tested. As the RAL163 ASICs were designed with built-in test facilities they could be rapidly checked, but G-link operation could be verified only by observing the lock condition and the resultant data via the RAL163 ASICs. Communication with the RAL163 ASICs worked correctly, but the G-links did not lock reliably. After a thorough investigation the cause of the problem was revealed to be excessive noise on the 5 V supply lines, which the G-links use as the reference. In this design they operate with positive ECL (PECL) logic levels to enable direct interfacing to the RAL163 CMOS ASICs. At design time only the ECL version was available from HP but a TTL version now exists. With extra de-coupling capacitors added to the substrate the noise levels have been reduced and the two channels lock successfully. Further tests need to be carried out to eliminate minor problems such as loss of lock when DAQ system intervenes. These problems are still related to power supply noise, due to poor power supply track layout on the MCM substrate.

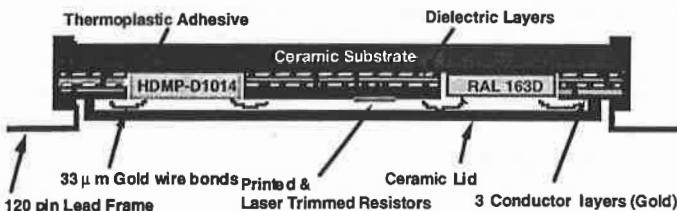


Figure 4: MCM

4. Demonstrator Programme

The ATLAS test beam has been used every year since 1993 to test various components and technologies required for the final ATLAS trigger system, and in 1997 a complete system 'slice test' was carried out [1]. The strategy for the demonstrator programme was to build a small scale version of the trigger system to evaluate critical components required for the final system. Except for the MCM all other technologies required for the final trigger system have been tested in the CERN test beam environment. Figure 5 shows various milestones achieved during 1993 to 1997.

Further to the test beam work extensive lab tests have been carried out, including bit-error rate measurements on the complete 'data chain'. The measured BER is better than 5×10^{-13} , which should be compared to the $<10^{-9}$ BER limit imposed by restricting the trigger rate increase to $<1\%$.

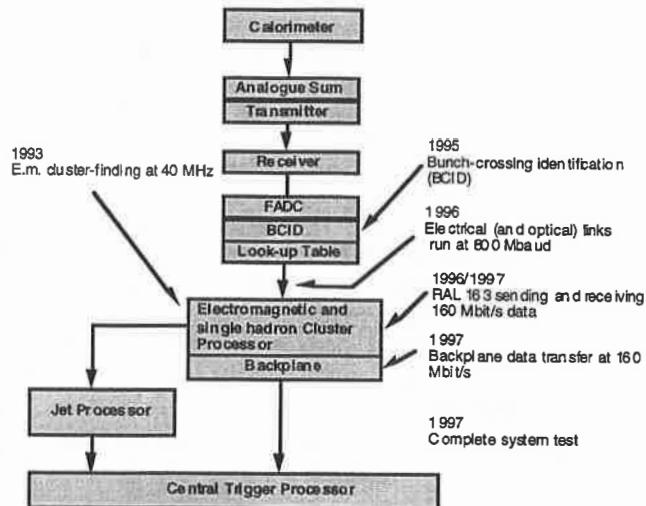


Figure 5. Technology Milestones

5. SUMMARY AND CONCLUSIONS

The cluster processor system relies on several key technologies which the demonstrator programme was designed to evaluate.

The first phase of the project studied the implementation of the cluster-finding algorithms on semi-custom ASICs at full LHC speeds. In the latter part of this phase the implementation of digital bunch-crossing identification logic was studied.

The second phase of the demonstrator programme concentrated on the 'data chain' from the detector through the FADCs and into the processing crates, and then the fan-out on a transmission-line backplane.

Most of the key technologies have now been successfully tested, and proven solutions exist in all areas to support the final system design as described in the Technical Design Report [1].

The final CP system architecture includes two ASIC designs, a MCM design and four module designs.

6. REFERENCES

- [1] ATLAS Level-1 Trigger Technical Design Report, ATLAS Level-1 Trigger Group, ATLAS TDR-12, 24 June 1998.
- [2] Low Cost Gigabit Rate Transmit/Receive Chip Set, Technical Data, Hewlett Packard.
- [3] FTM -8510 low Cost Gigabit Optical Transmitter/Receiver, Finisar Corporation, CA 94025.
- [4] 16-40 MHz 10 Bit Bus LVDS Serializer and Deserializer, National Semiconductor, September 1998
- [5] ATLAS Technical Proposal, CERN/LHCC/94-43, 1994