ATLAS Second Level Trigger Group

Status Report to PPESP (UK) - March 1995

R.E.Hughes-Jones, S.D.Kolya, R.Marshall, D.Mercer *Manchester University*
F.Harris, S.Hunt* *Oxford University*
D.R.Botterill, R.Hatley, J.Leake, R.P.Middleton, F.J.Wickens *RAL*
B.J.Green, Ll-M Mir, J.A.Strong *Royal Holloway College*
P.E.L.Clarke, R.Cranfield, G.J.Crone, P.Sherwood *University College London*
(* indicates student)

1. Introduction

The ATLAS-UK second level trigger group has, since October 1993, successfully developed ideas and prototypes from earlier R&D work into working components for a second level trigger demonstrator system. This status report describes the components, systems, test results, simulation and physics studies, and the evolution of the programme over the period and into the coming year. The milestones and use of funds are reviewed in this section, the following sections describe the system architecture, the work done (beam tests, system components, laboratory tests, modelling and physics studies) and the programme for the next year.

In the revised milestones submitted to the PPESP in October 1993, a beam test of a system with detectors was postponed, in view of the reduced funding available. Subsequently, however, the ATLAS technical proposal was delayed and it was judged that our technical arguments would be given greater weight if backed up by successful beam tests. Detailed consideration of the possibilities within the resource and technical constraints led to the decision that a small scale single detector system was possible and would produce worthwhile results. A large effort to achieve this goal was mounted by supplementing the PPESP funds with a great deal of equipment borrowed for the short duration of the tests or obtained via other funding sources. This also required considerable rescheduling of the laboratory tests, which have recommenced and are now well advanced. Progress in the global system, although substantial, was impacted by long delays in obtaining the Alpha SBCs and associated system software and to a lesser extent the SCI components. Table 1 lists the revised milestones and how well they have been attained. Table 2 compares the PPESP resources used, with the intended allocation of funds given to the PPESP with the revised milestones.

Milestone	Detail	Status
Complete prototype building block tests	T2B, Link Unit, FeX	Done
	SCI i/f's	C40, VME Done Completion of PCI Delayed
	Alpha SBC	Postponed by DEC. Replacement delayed, Standard WkStn used
	Supervisor	Replaced by RoI/T
Finish sub-system testing. Local + Global		Partially completed
Results from lab tests		Partially completed
Beam tests with detectors	Originally postponed	Done with single sub-detector

Table 1. Attainment of the revised milestones.

Following this work the entire ATLAS second level trigger community has been considering plans for further tests in 1995 and both larger scale beam and laboratory tests are proposed, depending heavily on the framework defined by the UK group. The UK contributions to this further work are described in this document and in the combined ATLAS-UK bid to the PPESP for 1995/6.

Item	Forecast	Actual	Comment
Alpha SBCs	16K	7K	Only 2 SBCs purchased to date
SCI i/f's	12K	12K	+4K earmarked to complete existing i/f's
C40 Commercial boards	18K	10K	
T2 Buffers	0K	6K	
Modsim (s/w lic)	0K	0K	
Crates, cables, etc.	0K	0K	
Test beam expenses	0K	1K	
Grand Total	46K	36K	Original total assumed 6K of non PPESP funds

Table 2. Use of PPESP funds.

2. Architecture

The second level trigger system (figure 1) is based on a local/global architecture. A processor in the local system is used to process data from a selected, small region of a sub-detector, whilst a global processor is used to process data derived from all the selected regions of all sub-detectors.

In the local part, guidance, in the form of region of interest (RoI) pointers, is taken from the level-1 system to enable fine grain raw data to be extracted from the second level buffers for processing in the local system's Feature Extractors (FeX). The resulting features (e.g. for a calorimeter: a cluster energy and position with some associated 'particle' classification) from all sub-detectors participating in the second level trigger system are then gathered together in a data concentration phase after passing through a network to the global sub-system. Here, features from different sub-detectors, corresponding to the passage of a particular particle or jet through the detector, are combined and a probable particle identification assigned. Certain physical quantities are then evaluated and used to classify the event according to topology and possible physics processes and the decision made whether to accept or reject the event.

The essential features to note are :-

- Localised processing of fine grain raw data
- Use of RoIs to reduce the required second level trigger system input bandwidth.
- Parallel processing in both local and global sub-systems to achieve a design decision frequency of 10^5 Hz, but with a processing time of about a millisecond for each event.
- System scalability to allow for the evolution of algorithms and physics goals.

Figure 2 [1] shows schematically how such a scheme may be mapped into hardware.

3. Beam Test Demonstration

A major accomplishment of the group was the successful operation of a small system in the H8 beam line at CERN during September/October 1994. This provided an existence proof for an asynchronous local/global second level trigger system based on our architectural ideas and establishes a firm foundation from which future work can evolve.

The complete system (figure 3), contained two T2 buffers, one Link Unit, two Feature Extractors, a Global Gateway, a four node global network and a single global processor. The local system components were based on Texas Instruments TMS320C40 (C40) [2] digital signal processors (DSP), the global network was based on Scalable Coherent Interface (SCI) [3] and the global processor was a DEC Alpha [4] workstation. Parasitic operation in the beam line ensured minimal disruption to other parts of the ATLAS tests and enabled the work programme to remain independent of other activities. Invaluable experience was gained in integrating with a real detector providing real data.

Data was derived from a HIPPI spy unit which formed part of a router system [5] placed in the data readout path of the RD-6 Transition Radiation Detector (TRD) [6]. Two channels were equipped to feed data into two T2 buffers each under the supervision of their C40 based buffer manager. Collaborators from Cracow and Jena provided the interconnection to the RD-6 system. Data was fed through a single Link Unit to enable routing of data to the appropriate FeX. After feature extraction (a null algorithm in initial tests), the data were concentrated in the Global Gateway (another C40) and then sent via the SCI ring for final processing by a global processor.

The whole system was self-triggered by the passage of an event on the RD-6 HIPPI lines. In parallel to the data being sent to the T2 buffers, a message was sent to the RoI/T unit, which simulated the function of a T2 supervisor. This unit generated an event number and routing instructions for the event data to both the Feature Extractors and the global processor(s). Another function of the RoI/T was to generate special messages at the start and end of each burst, these were used to reset pointers in the Global Gateway and thus ensure synchronisation at the start of each burst. After all features from an event had been processed, an accept/reject signal was sent from the Global Gateway back to the T2 buffers to indicate that the data should be kept or overwritten.

The C40 sub-system was controlled from a PC through a daisy-chained JTAG interface. Nodes in the SCI sub-system were each self-configuring, with one of the nodes set to provide both the master clock for the ring and other ring management functions. The target Alpha processor with VxWorks [7] real-time kernel was hosted from another Alpha running OSF/1.

In the initial phase of the beam tests only the local part of the system was used. Since the prototype T2 buffers did not include a fast read-out port for the DAQ system, an option was included to pass the complete events through the Link Units and Feature Extractors to the Global Gateway. The data concentrated in the Global Gateway being recorded by the OS9 local DAQ system. In this phase some 50K events were successfully recorded to disk. An examples of an event recorded by the system is shown in figure 4.

In the second phase, the SCI ring and global processor were added. The Global Gateway node sent the data from each event to the Alpha node, with responses being returned. The other two SCI nodes were for diagnostic purposes only. In this configuration event data were recorded both from the Global Gateway by the OS9 system and by the Alpha VxWorks node to the OSF/1 Alpha. It is believed that this is the first time that 'live' detector data has been passed round an SCI ring at a beam line.

Figure 5 shows a SCI Move transaction (the mode used for data transfer in the test beam), consisting of 64 bytes of data plus a 16 byte header, between two nodes. The time taken for the transmitting node to send the move request onto the SCI ring and receive the response from the remote node was 2.1 µs, indicating that a node may transmit data at 30 MBytes/s. Note, however, that the SCI ring can contain other packets between other nodes during this time, up to a capacity of 125 MBytes/s at any point on the ring.

4. System Components

T2 Buffer

The basic T2 buffer design is shown in figure 6. Because of the early stage of the experimental design, no choice has been made on data transfer protocols from the sub-detectors and the current design accepts a 32-bit parallel word with a data valid signal and has a programmable input protocol which can be adapted to match sub-detectors requirements. An input FIFO allows the input data to be synchronised to its own clock. The 2 kByte input FIFO on the prototype board can operate at up to 50 MHz giving a peak input rate of 200 MBytes/s.

The 512 kByte memory is page managed with a nominal page size of 1 kByte corresponding to a transfer period of 10 µs at 100 MBytes/s. A new event starts a new page and an event can occupy more than one page. The page address is taken from the list held in the free page FIFO and the memory address is incremented until either the end of the page is reached when the next page is allocated or the end of the event is reached when the current address is written into the used page FIFO. The used page FIFO contains a list of used pages, a last page of event marker and the end of event address.

A C40 DSP is used as the buffer manager as its six communications ports can be used for all I/O transactions apart from the main data input. In the current design, memory reads are under the control of the manager which can access the memory at a maximum of one in four clock cycles. For the 40 MHz devices currently used this corresponds to a read rate of up to 40 MBytes/s and a memory write rate of at least 120 MBytes/s. The communications ports are capable of operating at up to 16MBytes/s within a total rate of 40 MBytes/s. The current design has all the hardware associated with the buffer on a board which can plug into a commercial board holding the buffer manager DSP. The commercial board, a DBV44 from Loughborough Sound Images (LSI) [8], can hold up to four C40s sited on TIM-40™ modules [9]. One of the TIM modules is replaced by the buffer board which then connects to the global bus of the C40 on another TIM site.

The Buffer Manager software on the C40 has to manage the memory and handle both data requests and second level trigger decisions. Event indexing, page address and global decision handling must be performed at the average first level trigger rate of ~100 kHz. RoI handling can proceed at a much lower rate (~10kHz) and read-out to the DAQ proceeds at the second level trigger accept rate (~1kHz). All communications on the second level trigger side takes place over C40 links.

Since it is unlikely that anything would be gained by letting buffer manager functions interrupt one other, we have avoided the complexity and overhead of interrupts altogether, using instead a simple loop in which input sources and output queues are polled in turn and handled immediately if they require servicing. Code was originally written in ANSI C for portability e.g. to other DSP families, but has been rewritten in assembler to maximise speed. All the current code and variables fit into the 2 kiloword on-chip RAM with only the index table held in external memory.

Link Unit, Feature Extractors, Global Gateway and RoI/T

The Link Unit, Feature Extractors, Global Gateway and RoI/T were implemented using DBV4x commercial modules. In all of these the software was written in ANSI C and, as in the T2 buffer, the code used polling loops rather than interrupts. The Link Unit used header information in the data received from the T2 buffers to route the data on to the requested Feature Extractor. For the beam tests the Feature Extractors merely passed the data on to the Global Gateway. The Global Gateway as well as acting as the only communication point between the local and global parts of the system, sending data from the Feature Extractors to the global processor and the event accept/reject signals back to the T2 buffers, also stored the data it received in buffers to be read by the local DAQ. The RoI/T acting as a system supervisor, on receipt of a trigger signal, generated 'RoI' tags, event number and processor allocations, and at start and end of burst generated signals for system synchronisation.

SCI Modules

All SCI interfaces were based on the CMOS SCI NodeChip™ (designed by Dolphin [10] and manufactured by LSI Logic [11]) which implements the SCI protocols and drives SCI links at 125 MByte/s. The NodeChip interfaces to external logic with a 64-bit bi-directional synchronous bus called the Cbus.

Two approaches were used. A Daughter Card (figure 7) was designed to interface SCI to a simple synchronous bus (DBus) and a series of interfaces were made to connect the DBus to VME and the C40. The Daughter Card contains a NodeChip, high performance FIFOs and control logic. In parallel, a design based on the work of the CERN RD-24 collaboration was fabricated. This interfaced a SCI Mezzanine card from Dolphin [12] to an R3000 processor in a RIO VME module from CES [13].

The operation of the Daughter Card centres on 4 synchronous FIFOs in the data path between the Cbus and the DBus; one each for the outward requests, inward requests, outward responses and inward responses, thus ensuring deadlock free operation. In operation, the NodeChip translates incoming SCI packets to Cbus packets and, assuming no other Cbus traffic, these are written to the relevant inward FIFO. The FIFO status Flags can be polled or generate vectored interrupts over the DBus. The external host can then subsequently empty the FIFO over the same DBus. Outward request or response packets are written to the relevant FIFO and then a word-count posted to the relevant down-counter in the CSR space. Writing these counters initiates the emptying of the FIFO over the Cbus into the NodeChip.

The VME interface to the Daughter Card was through the DBus and was used both to commission the Daughter Card and to interface DEC Alpha and Motorola 68K processors. The board is wire-wrap and based on the Newbridge SCV64 (Single Chip VME) [14] which provides high performance pipelined slave access operating in both D32 and D64. An on board DMA engine can be used to 'push' data from the Daughter Card directly to VME memory.

The link to the local C40 network uses a DBV42 VME board (from LSI) with two C40 based TIM modules and user expansion via a proprietary bus, the dBeX32. Translation from dBeX32 bus to DBus is performed with an interface implemented in wire-wrap on a double Eurocard board. Since the dBeX32 and the DBus match rather well the translation can be accomplished by a small number of programmable logic devices (22V10 PALs/GALs).

Global Processors

At the time of the PPESP proposal it was planned to use a particular Alpha based single board computer (SBC) being developed by DEC. Subsequently this particular SBC was postponed and plans were modified to use an alternative SBC (EB66 - also from DEC [15]) - although as this replacement was not available in time for the beam tests, use was made of a standard Alpha workstation with a VME interface to connect to SCI. The EB66 evaluation board has the advantage of using a second generation Alpha chip (21066) with on-chip PCI controller [16] and has a PC motherboard format. Whilst the performance of the 21066 is slightly less than the original 21064 chip, it is significantly more cost effective, being targeted at the commodity PC market. The SCI interface to the EB66 is via the PCI bus and design work is well advanced on this.

The Chorus operating system has been evaluated and is functionally well suited to our needs. However, commercial support is primarily for OEMs and to date no commercial port exists for the Alpha processor. At present we are using VxWorks and have implemented SCI packet handling code within it.

The global system supervisor function has been implemented, for simplicity, in the RoI/T. It is foreseen to implement a dedicated supervisor for future tests.

5. Laboratory tests

Buffer and Link Unit tests

Timing tests of the Buffer Manager have been made with a software emulation of the buffer hardware, using additional DSPs as data sources and sinks, so that the I/O environment seen by the Buffer Manager is essentially the same as in a real system. The performance of the assembler coded Buffer Manager running on a 40 MHz DSP with 1 kByte Region of Interest (RoI) records and 1 kByte DAQ events is given in Table 3. It is seen that the average event rate was comfortably above 100 kHz and is roughly twice that obtained from earlier tests using modular, portable code written in C.

Fraction of events with RoIs	20%	5%	5%
Fraction of second level trigger "yes" decisions	1%	1%	4%
Maximum mean event rate	127 kHz	134 kHz	131 kHz

Table 3. Buffer Manager maximum mean event rate as a function of RoI and DAQ transfers.

The results of these software-based timing tests are consistent with those obtained using a hardware T2 buffer driven by a SLATE [17] event generator board. In these hardware-based tests, a 33 MHz DSP (with no DMA transfers) sustained a continuous event rate of 125 kHz.

For the Link Unit tests, one C40 was used to send RoI records randomly to each of a set of Link Unit processors. The Link Units were interconnected to represent the relevant part of a mesh (see figure 8). The record destination was determined as either another Link Unit or a Feature Extractor by the RoI generator and entered into the record header. With only a limited number of units available to us, we have wrapped the interconnections to simulate, to a limited extent, a fully connected mesh system.

Tests have been run using two and three interconnected Link Units. Fifteen million events with event sizes between 1 and 2 kBytes were transferred through the system with zero errors. Results (see Table 4) show that for these events the throughput of a Link Unit is approximately 20 MBytes/s.

Event size	0.25	0.5	1.0	2.0	4.0	kBytes
RoI records	36.10^3	31.10^3	18.10^3	10.10^3	5.10^3	per second

Table 4. Link Unit RoI record transfer rate as a function of record size.

Tests of global system components

Laboratory testing of the SCI components has confirmed the correct operation of the SCI Daughter Card when driven by either the VME or C40 interfaces with data rates of several MBytes/s - much higher than that used in the beam-line tests. For example a 2 node ringlet with each node sending data to itself has achieved data rates on the ring of up to 20 MBytes/s demonstrating correct data transfer in spite of packet collisions. It should be noted that the maximum data rate that the current NodeChips can inject into the ring is about 30 MBytes/s, the rates given above are limited by the prototype driving interfaces and are very similar to those achieved elsewhere [18]. In addition to checking the performance, the tests have allowed measurements of the operational parameters of the LSI NodeChip beyond information given by the manufacturers or in the very limited amount of published data available. An understanding of these details is important input to the simulation studies, for the determination of the ultimate performance that can be achieved with currently available NodeChips and for optimal upgrades to the current designs.

Work is also progressing to investigate and understand the way the PCI local bus is driven by the DEC Alpha chip (21066) on the SBC and the best way to access the Alpha memory system for optimal flow of trigger data from the C40 local system. Figure 9 gives a logic analyser trace showing a DMA access from a PCI peripheral to the Alpha memory and secondary cache. The PCI transaction starts when the peripheral asserts the FRAME signal, when the Alpha chip is willing to accept data it asserts the TRDY signal and data is then transferred on each clock cycle. Including PCI overheads the figure shows a measured data transfer rate of 60 MBytes/s. By transferring all of the data for an SCI packet in one PCI transaction this would increase to a data transfer rate between the Alpha memory and the PCI-SCI interface of 89 MBytes/s.

6. Modelling

A DAQ/Trigger modelling system, SIMDAQ, has been developed for evaluating the performance of multi-level DAQ/Triggering systems based on switching systems and multi-processing. The first implementation of this system has been made with a view to evaluating designs for the LHC/ATLAS experiment and, in particular, second level trigger architectures of the type studied by the UK groups. Members of the group have played a major part in setting up and directing a broad collaboration of people to study the problem. In addition, they have taken on the responsibility of developing and maintaining the base program.

SIMDAQ is based on the object oriented simulation language MODSIM II [19]. It allows the model under study to accept both input generated from physics simulations and the definition of the detector readout segmentation. Hence SIMDAQ allocates data to detector readout buffers based upon real physics event distributions. Models of the switching elements required in the trigger architectures have been developed for ATM, SCI and a generic switch. Models of other network technologies, e.g. FiberChannel and Transputer switches, will be integrated into the SIMDAQ framework as they are developed.

The simulation has been used to evaluate second level trigger system performance for 3 key ATLAS sub-detectors, the electromagnetic (EMC) and hadron calorimeters (HAC), together with a tracking detector (TRT). First simulation runs have been made corresponding to 10000 events through the system, with evaluations of deadtime and latencies, the latter providing guidance as to required depths of buffers in the system. Figure 10 shows results for a system using SCI as the switching elements. The system has proved flexible and very useful in the first evaluations of the requirements in performance of the system components.

Using the framework described above, models have been developed of the basic components used by the UK T2 demonstrators. Results from a simulation of the beam test have been delayed by the necessity to produce results for the technical proposal.

7. Physics studies

The trigger activity is driven by the need to accept interesting physics and reject background with high efficiency and reliability. The current architecture was developed to deal with high luminosity data and we still have much work to do, both to prove our ideas and to develop them into a working system.

The study of B physics with ATLAS at low luminosity requires a new approach as the first level trigger is a single muon and no other regions of interest are generated by the first level trigger system. In order to study the second level trigger requirements, two of our group have been working as part of the B-physics group [20] and we will use the expertise they have gained to guide our studies.

8. Future Programme

The future work of the group has two motivations: (i) to maintain the central role we have attained as a driving force of the ATLAS level 2 trigger development and (ii) to perform, on a timescale commensurate with the start date of ATLAS, a set of more specific studies and tests necessary to prove the proposed architecture.

Overview

The UK group, in promoting the beam tests in 1994, gained benefits in experience of leading edge technologies, but also attracted many others to participate in further work. The beam tests proposed for 1995 are more comprehensive and will involve most of the other groups active in second level trigger within ATLAS. The proposed configuration is shown in figure 11. The aim is to demonstrate the principle of a common (generic) trigger system taking data from up to 4 different sub-detectors. Equally importantly, however, is that this will bring together experts from different sub-detector groups, establishing working links and a better mutual understanding which will be essential in order to specify and build the final trigger.

The more specific, but equally important goals of the 1995 programme are:

- Beam tests:
  - Connecting to different sub-detectors and handling their data in the feature extractors.
  - Event building and association of data from different sub-detectors in the global system.
  - Demonstrating RoI guidance principles

- Laboratory tests (using high rate simulated data sources):
  - Measurement of SCI characteristics under high load in a 'many to one' configuration. (Needed for SCI development and selection between competing technologies).
  - Tests of performance improvements with hardware assisted SCI interfaces.
  - Construct 2nd generation prototype ATLAS level 2 buffer.
  - Test a C104 'DS link' switching fabric with the MACRAME collaboration in CERN.

- Simulation programme
  - The simulation programme will continue to be given high priority, since only by comparison of prototypes with their simulations can the final system be simulated with confidence. The prototypes will also continue to supply vital parameters for the simulations.

Although the final technology choices do not have to be made for several years, there is much to be done in the intervening time. It is important to obtain a better understanding of the potential architectures and technologies. Key areas, such as low latency networks and processors with well integrated I/O, are being developed commercially, but tailoring to our requirements will be needed and in some cases developments are susceptible to our influence. Strategies for fault tolerance and for monitoring and control of the large number of processors need to be developed. Work needs to be done in collaboration with the sub-detector groups to define and demonstrate the interconnection between their front-end electronics and the trigger/DAQ system. This work, needed to ensure a suitable interface to the common second level buffer and a mutually convenient partitioning of the data, must be completed before the sub-detector designs are frozen. We firmly believe that the best way to address all of these issues is a well balanced programme of the form outlined.

Specific Developments

C40 Embedded board

To gain more detailed knowledge on the operation of embedded processors we have built a prototype board containing a single C40 DSP. This board was designed to examine a range of problems but the final design is a good match to our requirements for a buffer manager and we intend to build more for use as buffers for future tests. The commercial boards are better matched for use as Link Units and Feature Extractors and will be employed for these functions in future.

Global System

In addition to the existing SCI nodes a new PCI-SCI interface will be developed from the existing Daughter Card design. The current C40-SCI design will also be adapted to use with the embedded C40 board so that several can be provided more economically and some optimisations made to reduce the overhead per SCI packet. Improvements to the interfaces are planned with the development of hardware assistance for the specific applications of the interfaces. These will include hardware to handle SCI packets and hardware to build the features of each event directly into the memory of each global processor. SCI data sources and other diagnostic tools will be constructed to allow further investigations of SCI performance at high rates.

Modelling

We intend to continue the development of our modelling activities within the general ATLAS framework. Specific areas requiring work in the immediate future are: an object protocol agreement to facilitate code integration; a durable scheme for connecting to the physics simulation activities; a detector readout architecture database; development of component and algorithm models and functional extension to a wider range of system components and detector configurations.

9. Conclusions

The UK group has been instrumental in setting the preferred ATLAS second level trigger architecture and has developed modules and test systems in line with this architecture. We have developed the critical components needed to enable us to demonstrate our ideas using current technology. We have successfully connected together a small asynchronous local/global trigger system and demonstrated it coupled to the prototype RD-6 detector in the H8 test beam at CERN. Some critical performance characteristics of our modules have been measured in the laboratory and at the test beam. We have made major contributions to the general ATLAS modelling activity and this is providing us with a framework for the simulation of our current and future test activities.

Primarily as a result of our beam test demonstration, many other groups working in the second level trigger area are eager to collaborate with us. We are developing a detailed programme of the tests and developments needed to establish the criteria for architecture and technology selection for the ATLAS second level trigger on the expected timescale for the LHC.

References

[1] ATLAS Collaboration Technical Proposal. CERN/LHCC/94-43, LHCC/P2, 15 December 1994
[2] TMS320C40 Users Guide, Texas Instruments Inc.
[3] "SCI, Scalable Coherent Interface", IEEE standard 1596-1992.
[4] Alpha Architecture Reference Manual, ISBN 1-55558-098-X,
Digital Equipment Corporation.
[5] CERN/EAST note 92-09.
[6] RD-6 Collaboration, CERN/DRDC/P8.
[7] VxWorks, Wind River Systems Inc.
[8] DBV42 & DBV44 Technical Reference Manuals,
Loughborough Sound Images Ltd, Loughborough, UK.
[9] TIM-40, Texas Instruments Inc.
[10] Dolphin Interconnect Solutions AS, PO Box 70 Bogerud, Oslo, Norway.
[11] L64601 SCI NodeChip Technical Manual, LSI Logic Corporation.
[12] CMOS Mezzanine Card Users Guide, Dolphin Interconnect Solutions
[13] RIO 8260 User Guide, Creative Electronic Systems (CES), Geneva, Switzerland.
[14] SCV User Manual, Newbridge Microsystems.
[15] EB66: DECchip 21066 and DECchip 21068 Alpha AXP Evaluation Board
[16] Peripheral Component Interconnect (PCI) Specification,
PCI SIG, c/o Intel Corporation, Hillsboro, Oregon, USA.
[17] SLATE (Second Level Architecture Test Environment), CERN/EAST 91-02, 91-13, 92-20.
[18] RD-24 Collaboration, Private communication.
[19] CACI Products Co, La Jolla, California, USA.
[20] B Physics in ATLAS, Ll-M Mir, P Sherwood in P Eerola et al., NIM A351,(1994)84

List of Figures