7.7.3 Timescale

The detailed timescale given in the backup documents is based on an LHC start date of 2004. The trigger front end must be installed in 2002 with the LAr calorimeter, when the cluster processing section should be fully operational for detector commissioning and readout tests. Manufacture and testing of approximately 100 modules is expected to require two years, and should therefore start in 1999. The trigger concept must be frozen in 1997 to allow time for detailed design and final prototype testing before manufacture begins.

The resulting cost and engineering effort profiles are shown in tables 14 and 15.

Year	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004
Spend (UK k£)	115	80	35	186	431	840	1520	325	100	100
Accumulated (UK k£)	115	195	230	416	847	1687	3207	3532	3632	3732

Table 14: Level-1 requested spend profile.

Year	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004
Engineering staff (RAL S/Y)	2	2	3	5	6	6	4	4	2	1
Technical staff (RAL S/Y)	0.5	1	1.5	4	4	5	5	5	4	2

Table 15: Level-1 requested manpower profile.

7.7.4 Cost Sharing

Discussions on resource sharing are underway between the German, Swedish, and UK groups involved in the calorimeter trigger. Although detailed decisions cannot be reached until the trigger technology choices are made in 1997, it is already accepted that each country will contribute to infrastructure as well as having specific design, manufacturing and testing responsibilities.

Excluding CERN and commercial companies, from whom no capital contribution is expected, the UK groups account for significantly more than half of the signatories to the level-1 calorimeter trigger backup document to the ATLAS Technical Proposal. From the earliest stage of the ATLAS trigger work, UK groups have taken the leading role on physics simulation and architectural design studies, and test beam evaluations of trigger prototypes have been performed to date solely by the UK groups. To maintain this dominant role we feel that a UK share of 75% of the project is required.

7.8 Conclusions

The UK groups have established themselves over the last three years as the leading contributors to the ATLAS level-1 calorimeter trigger effort. We have carried out a very innovative and successful programme of trigger processor hardware development firmly based on simulation studies, and we continue to place particular emphasis on the importance of evaluating this hardware in the realistic environment of a test beam. Close working relationships have been developed with the calorimeter community, in particular in the area of front-end electronics where mutually-acceptable technical solutions must be found.

Our lengthy collaboration with the Stockholm group has been mutually beneficial in a number of ways, and we confidently expect to work closely with them after ATLAS has made technology choices in 1997. The collaboration with the Heidelberg group is more recent but is already proving very fruitful. With their strong background in hardware design and development they have made significant contributions to our current demonstrator programme, and together with their plans to build up expertise in ASIC design they would be well placed to contribute to the ATLAS front-end trigger electronics.

We therefore believe that the UK groups lead a very strong international team which is now accepted by ATLAS to be responsible for the design and construction of the level-1 calorimeter trigger system. With world-class electronics facilities, expertise and experience to draw upon, the UK groups are now ideally placed to capitalise on the work which we have done so far and to lead the project over the next decade. Financial and manpower resources commensurate with this leading role are thus sought from the PPESP.

8 The ATLAS Level-2 Trigger

8.1 Introduction

This section discusses the preferred ATLAS design for the level-2 trigger, describes the work performed to date in this area by the UK groups and presents the proposal for full participation in this area of ATLAS.

For the last several years, groups from Manchester, Oxford, RAL, RHBNC and UCL have been working to develop and demonstrate their ideas for the architecture of the level-2 trigger in ATLAS and they have played a leading role in defining the preferred architecture. From this and earlier activities, the UK groups have developed considerable experience and expertise in the areas proposed for future work.

This request for PPESP support would enable the UK groups to maintain their central position in the project and to design, construct and operate a significant part of the ATLAS level-2 system.

8.2 Physics Rates and Signatures

The level-2 trigger is required to reduce the level-1 trigger output rate to that which can be accommodated by the event builder. After allowing for uncertainties in those rates, the design criteria for the level-2 trigger at high luminosity (10^{34} cm^{-2}s^{-1}) are an input rate of 100 kHz and an output rate of 1 kHz. To achieve a rate reduction of this factor it is necessary to operate with new algorithms on full granularity data from the sub-detectors and to combine their information. It is not sufficient to repeat the level-1 algorithms with improved resolution.

Table 16 shows the input and output trigger rates expected from simulation work for inclusive trigger selections at level-2. Table 17 shows how these triggers can be used to select physics processes. Substantial improvements in selectivity are expected from applying topological constraints in the level-2 trigger and further studies will be necessary to examine this additional rejection power.

At low luminosity (10^{33} cm^{-2}s^{-1}), the 'high-luminosity' programme will be supplemented by measurements on the B system [30]. B-physics triggers will need to operate in parallel with the 'high-luminosity' triggers and the level-2 system will be designed to run them simultaneously.

8.3 ATLAS Level-2 Trigger System

8.3.1 Overview

At high luminosity, the total data rate from the detector after a level-1 trigger is around 150 GByte/s. To reduce the data transfer requirements into the level-2 trigger system, the overall philosophy adopted is to use guidance from level-1 to select the data to be processed. This guidance is in the form of pointers in $\eta - \phi$ space to regions of interest (RoIs) which correspond to significant transverse-energy deposits in the calorimeter or tracks in the muon trigger chambers. Only data from within a region of interest are transferred to the level-2 system for processing. This reduces the data volume to be transferred by up to two orders of magnitude.

As shown in figure 34, level-2 data processing is divided into two main steps.

- Feature extraction is used to reduce the data from one sub-detector from basic cell information to more physical quantities (e.g. for a calorimeter these would be cluster parameters).
- Features for all RoIs from all appropriate sub-detectors are combined and a global event decision is made.

In general, feature extraction can take place in parallel for each RoI in each sub-detector. In some regions of the detector, however, it may be necessary for information from two sub-systems to be combined for satisfactory feature extraction, for example in the overlap area between barrel and forward sub-systems.

To achieve the required rate reduction, the level-2 trigger must use tracking information and event topology analysis. For example, the event vertex position is needed if mass cuts are to be effective and information from the precision tracker is required for this to be included in the global event decision.

8.3.2 Architecture

The basic level-2 architecture [1] is shown in figure 34. All data are moved off the detector after a level-1 trigger and transmitted to control rooms on the surface where they must reside in temporary buffers during the level-2 trigger decision time. Data from sub-detectors are grouped so that the rate into each buffer is approximately 100 MByte/s and the total number of buffers is around 1700.

Table 16: List of level-1 and level-2 trigger requirements and associated rates at high and low luminosity.

Trigger requirement	level-1 rate (kHz)	level-2 rate (kHz)
Luminosity $\mathcal{L} = 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$		
≥ 1 μ $p_{\mathrm{T}} > 20$ GeV ($1\mu_{20}$)	4	
≥ 1 isolated μ $p_{\mathrm{T}} > 20$ GeV ($1\mu_{20}$I)		0.2
≥ 1 μ $p_{\mathrm{T}} > 40$ GeV ($1\mu_{40}$)		0.1
≥ 1 isolated em cluster $E_{\mathrm{T}} > 30$ GeV ($1\mathrm{em}_{30}$)	20	
≥ 1 e $E_{\mathrm{T}} > 30$ GeV ($1\mathrm{e}_{30}$)		0.3
≥ 1 γ $E_{\mathrm{T}} > 60$ GeV ($1\gamma_{60}$)		0.1
≥ 2 μ $p_{\mathrm{T}} > 6$ GeV ($2\mu_{6}$)	1	
≥ 2 μ $p_{\mathrm{T}} > 10$ GeV ($2\mu_{10}$)		0.08
≥ 2 isolated μ $p_{\mathrm{T}} > 6$ GeV ($2\mu_{6}$I)		0.1
$\mathrm{B}^0_{\mathrm{d}} \to \mu^+\mu^-$, $\mathrm{B}^0_{\mathrm{s}} \to \mu^+\mu^-$		0.01
≥ 2 isolated em clusters $p_{\mathrm{T}} > 20$ GeV ($2\mathrm{em}_{20}$)	4	
≥ 2 e $p_{\mathrm{T}} > 20$ GeV ($1\mathrm{e}_{20}$) **or** ≥ 2 γ $p_{\mathrm{T}} > 20$ GeV ($2\gamma_{20}$)		0.2
≥ 1 jet $p_{\mathrm{T}} > 150$ GeV ($1\mathrm{j}_{150}$)	3	
≥ 1 jet $p_{\mathrm{T}} > 300$ GeV ($1\mathrm{j}_{300}$)		0.1
≥ 3 jets $p_{\mathrm{T}} > 150$ GeV ($3\mathrm{j}_{150}$)		0.04
Large E_{T}^{miss} (E_{T}^{miss})	1	0.1
Other prescaled triggers	5	0.1
Total	$\lesssim 38$	$\lesssim 1.4$
Luminosity $\mathcal{L} = 10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$		
≥ 1 μ $p_{\mathrm{T}} > 6$ GeV ($1\mu_{6}$)	8	
B-physics level-2 triggers		0.8
≥ 1 μ $p_{\mathrm{T}} > 20$ GeV ($1\mu_{20}$)		0.2
≥ 1 isolated em cluster $p_{\mathrm{T}} > 20$ GeV ($1\mathrm{em}_{20}$)	10	
≥ 1 e $p_{\mathrm{T}} > 20$ GeV ($1\mathrm{e}_{20}$)		0.2
≥ 1 γ $p_{\mathrm{T}} > 40$ GeV ($1\gamma_{40}$)		0.06
≥ 2 isolated em clusters $p_{\mathrm{T}} > 15$ GeV ($2\mathrm{em}_{15}$)	2.5	
≥ 2 e $p_{\mathrm{T}} > 15$ GeV ($2\mathrm{e}_{15}$)		0.01
≥ 2 γ $p_{\mathrm{T}} > 20$ GeV ($2\gamma_{20}$)		0.01
≥ 1 jet $p_{\mathrm{T}} > 100$ GeV ($1\mathrm{j}_{100}$)	0.9	
≥ 1 jet $p_{\mathrm{T}} > 200$ GeV ($1\mathrm{j}_{200}$)		0.06
≥ 3 jets $p_{\mathrm{T}} > 100$ GeV ($3\mathrm{j}_{100}$)		0.02
Large E_{T}^{miss} (E_{T}^{miss})	0.1	0.01
Other prescaled triggers	1	0.1
Total	$\lesssim 23$	$\lesssim 1.5$

Table 17: Representative list of physics processes and triggers through which they would be selected at high luminosity. The notation for the different triggers is defined in table 16.

Process	Level-1 trigger	Level-2 trigger
$(\mathrm{W} \rightarrow \ell\nu) + \mathrm{X}$	$1\mu_{20}$ $1\mathrm{em}_{30}$	$1\mu_{20}\mathrm{I}$ $1\mathrm{e}_{30}$
$(\mathrm{Z} \rightarrow \ell^+\ell^-) + \mathrm{X}$	$2\mu_6$ $2\mathrm{em}_{20}$ $1\mu_{20}$ $1\mathrm{em}_{30}$	$2\mu_6\mathrm{I}$ $2\mathrm{e}_{20}$ $1\mu_{20}\mathrm{I}$ $1\mathrm{e}_{30}$
$\mathrm{H} \rightarrow \gamma\gamma$	$2\mathrm{em}_{20}$	$2\gamma_{20}$
$\mathrm{H} \rightarrow \mathrm{ZZ}^* \rightarrow 4\ell$	$2\mathrm{em}_{20}$ $2\mu_6$ $1\mathrm{em}_{30}$ $1\mu_{20}$	$2e_{20}$ $2\mu_6\mathrm{I}$ $1\mathrm{e}_{30}$ $1\mu_{20}\mathrm{I}$
$\mathrm{A}^0 \rightarrow \tau^+\tau^-$	$1\mu_{20}$ $1\mathrm{em}_{30}$	$1\mu_{20}\mathrm{I}$ $1\mathrm{e}_{30}$
$\mathrm{t\bar{t}} \rightarrow \ell^+\ell^- + \mathrm{X}$	$2\mu_6$ $2\mathrm{em}_{20}$ $1\mathrm{em}_{30}$ $1\mu_{20}$	$2\mu_6\mathrm{I}$ $2\mathrm{e}_{20}$ $1\mathrm{e}_{30}$ $1\mu_{20}\mathrm{I}$
$\mathrm{t\bar{t}} \rightarrow \ell^\pm\nu$ + jets	$1\mu_{20}$ $1\mathrm{em}_{30}$	$1\mu_{20}\mathrm{I}$ $1\mathrm{e}_{30}$
SUSY $\rightarrow$ leptons	$1\mu_{20}$ $1\mathrm{em}_{30}$	$1\mu_{20}\mathrm{I}$ $1\mathrm{e}_{30}$
QCD jet studies	$1\mathrm{j}_{150}$	$1\mathrm{j}_{300}$
W', $\mathrm{Z}' \rightarrow$ jets	$1\mathrm{j}_{150}$	$1\mathrm{j}_{300}$
SUSY $\rightarrow$ jets	$1\mathrm{j}_{150}$ E_{T}^{miss}	$3\mathrm{j}_{150}$ E_{T}^{miss}
B-physics	$1\mu_{20}$ $2\mu_6$	$1\mu_{40}$ $2\mu_{10}$ $\mathrm{B}^0 \rightarrow \mu^+\mu^-$

Information from the different parts of the level-1 trigger system is gathered together in an RoI builder. This information is passed to the level-2 supervisor which is responsible for the overall control of the level-2 system. The supervisor sends the RoI information, together with any processor allocation information, to the level-2 buffers via an RoI distribution system, with a similar distribution system used for the final level-2 decisions received from the global processors.

Because data for one RoI in a sub-detector can be spread across a number of buffers, a network is required to gather the RoI data into one processor for feature extraction. After feature extraction, the data are passed through another network to one of a set of global processors. The feature extractors and the global processors are assumed to be programmable processors with a processing time long compared with 10 μs so that many of them are needed to cope with an input data rate of 100 kHz.

A critical parameter for the level-2 system is the latency. Increased latency implies that larger buffers are required to store the data during the level-2 decision, but, more significantly, that more processors are required in the system as the processing of more events overlap in time. To minimise the latency in the system, the overall philosophy adopted for high luminosity running is to divide processing into parallel tasks where possible. For example, RoIs in each sub-detector are processed in parallel in different feature extractors and each sub-detector is connected to an independent local processing system.

At low luminosity, the trigger rate and sub-detector cell occupancy are substantially lower, and a trigger which uses a sequential set of operations (and is compatible with the high luminosity architecture) is being studied. The process is to refine the muon trigger selection using the precision chambers and then to process the TRT data to find tracks above about 1 GeV/c in p_{T}. For regions of the detector around the TRT tracks, data from the silicon tracker and the calorimeter are then evaluated to provide information to the global processors

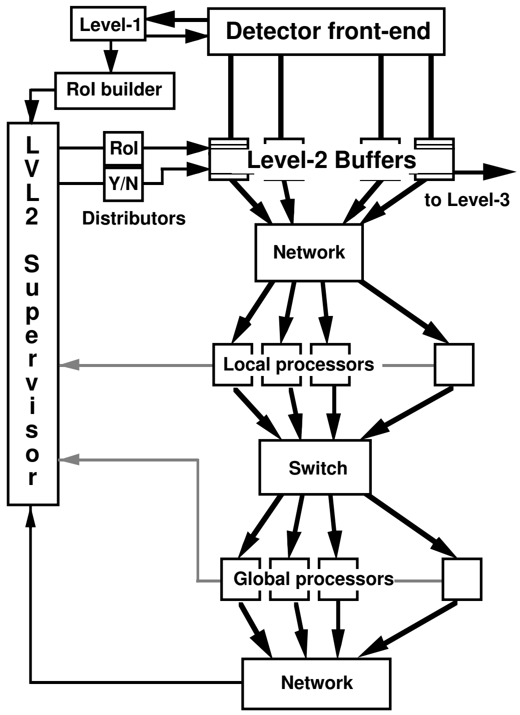

Figure 34: Block diagram of a level-2 trigger implementation.

where mass cuts can be applied to select the required physics signatures. As the luminosity is increased above 10^{33} cm^{-2}s^{-1} the muon threshold can be raised to maintain the trigger rate within the limits set by the level-2 system.

8.3.3 System components

Personality cards are required for each sub-detector, between the input data streams and the level-2 buffers, to format the data into a common frame appropriate for the level-2 trigger and event building. Liaison with sub-detector groups to define these cards is starting.

The level-2 buffer accepts data from the sub-detector via the personality card and stores them during the level-2 decision time. On receipt of a message giving the event identifier and relevant information from the level-2 supervisor via the RoI distributor, the buffer pushes the trigger data into the level-2 system. A level-2 accept signal causes the buffer to push the full event data to the level-3 network and then to release the buffer space for further events. On a level-2 reject signal, the buffer space is released. By including the capability to pre-process the trigger data, in or close to the level-2 buffer, substantial savings can be made in the network bandwidth required. For calorimeter jet data, for example, there is little point in transferring basic cell data to the feature extractor.

The local network must route data for one RoI from a number of buffers for one sub-detector to a single feature extractor. Among critical issues for the networks in the system are latency and data integrity. The network for each sub-detector has to handle events containing around 5 RoIs at 100 kHz event rate. However the performance of any part of the network depends on the connectivity demanded, e.g. the average rate at which one data buffer in the calorimeter system is asked to supply data is less than 5 kHz. The volume of data to be passed to the feature extractor for one RoI depends on the sub-detector and the degree of pre-processing in the data buffer. For an electromagnetic cluster trigger, the calorimeter RoI data volume is about 1.5 kBytes but this is spread over several buffer sources (the number depends on the position of the RoI relative to the buffer coverage) and the network must be capable of handling a significantly higher data volume per RoI. This leads to a throughput requirement for a sub-detector local network of between one and five GByte/s.

Feature extraction is performed in the local processors which need to be fast with good I/O response.

Measurements have been made on a number of systems for current TRT, SCT and calorimeter algorithms [31] but the overall performance of the processors in a realistic situation must be tested. For a 300 μs execution time of the complete algorithm including handling the I/O and any data formatting required, with 5 RoIs and 50% utilisation, 300 local processors are required for one sub-detector. For the calorimeter algorithm, current processors with about 100 Mips equivalent performance can satisfy this requirement. For the TRT and SCT algorithms, processors with hardware accelerators are needed. The estimated total computing power in the local processing system is between one and five times 10^5 Mips equivalent distributed over about one thousand processors. Improvements in processor computing power and their I/O handling, in particular, would lead to a dramatic reduction in the number of processors required. Commercial boards with RISC processors and digital signal processors (DSP) are being evaluated.

The global network transfers information for all the RoIs in the event from the feature extractors for the different sub-detectors active for that event to a global processor. The network must connect all feature extractors to all global processors. A high performance network is required to handle the feature data of typically 1–2 kBytes per event at 100 kHz. The global network must also be fault tolerant, and must be flexible enough to allow expansion of the processor farm without latency penalties or excessive reorganisation. At present there is no proven technology which will address the high bandwidth, low-latency, many-to-one problem in the transmission of many small feature records to a single global processor. Technically the most promising solution is Scalable Coherent Interface (SCI) [32] which uses high speed point-to-point links, and where the total bandwidth scales roughly with the number of nodes. This technology is being used for supercomputer applications and is being developed commercially for various computer connections.

Global processors are organised as a farm operating with a 100 kHz decision rate, running algorithms developed in offline trigger studies and thus requiring high performance general purpose processors. Assuming a total time of 400 μs is required for the tasks of collecting and organising the feature records, and making the global decision, the global system will need a farm of approximately 80 processors for 50% occupancy. Wherever possible data collection, formatting and reorganisation should be performed in hardware, and we are developing solutions based on FPGAs. Algorithm tests with current RISC processors do not achieve the full speed required, but with performance increases of a factor of 5–10 likely by the time the final system is required this number of processors should satisfy the requirements. The level-2 decisions from the global processors are passed to the level-2 supervisor via part of the global network.

The RoI trigger builder is required to provide the guidance information to level-2. Information from the level-1 calorimeter, muon and central trigger processors are combined to provide the information on event identification and regions of interest for use by the level-2 trigger. The information is passed to the level-2 supervisor. A design study is in progress to establish the full functional requirements [33]. A circuit based on FPGAs could accomplish this task, but a test using a DSP [34] indicates that a programmable processor may provide a more flexible option if some increase in performance can be achieved.

The level-2 supervisor controls the operation of the level-2 trigger. It receives the event information from the level-1 system, adds, if required, information on routing (including global processor allocation) and passes the modified information to the RoI distributor. Results from the global processors are used to initiate data transfers from the level-2 buffers to the event builder or free the buffer space. Information on transfer and processing errors are passed to the supervisor which also monitors time-outs in the processing chain. A design study is in progress to establish the full functional requirements [35].

The RoI and level-2 decision distributors sit between the level-2 supervisor and the buffers. The RoI distributor guides messages to the appropriate buffers. In this way each buffer does not have to respond to all RoI messages for every event but only to those which require it to take action. All buffers have to respond to the level-2 decision message at 100 kHz, but in most cases (99%) this is a reject signal which results in the relatively simple operation to free buffer space. An RoI message requires the buffer to organise the movement of data to the level-2 system and a level-2 accept signal requires the buffer to move data to the event builder.

8.4 Current Status of UK Level-2 Activities

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 [36]. Invaluable experience was gained by integrating with a real detector providing data. This demonstration provided an existence proof for an asynchronous local/global level-2 trigger system based on our architectural ideas and establishes a firm foundation from which future work can evolve. Much study has also been carried out under test conditions in the laboratory and has started to cross-fertilise architectural modelling work.

In much of the work described below, the UK group has led the field and increasingly other groups have started to contribute.

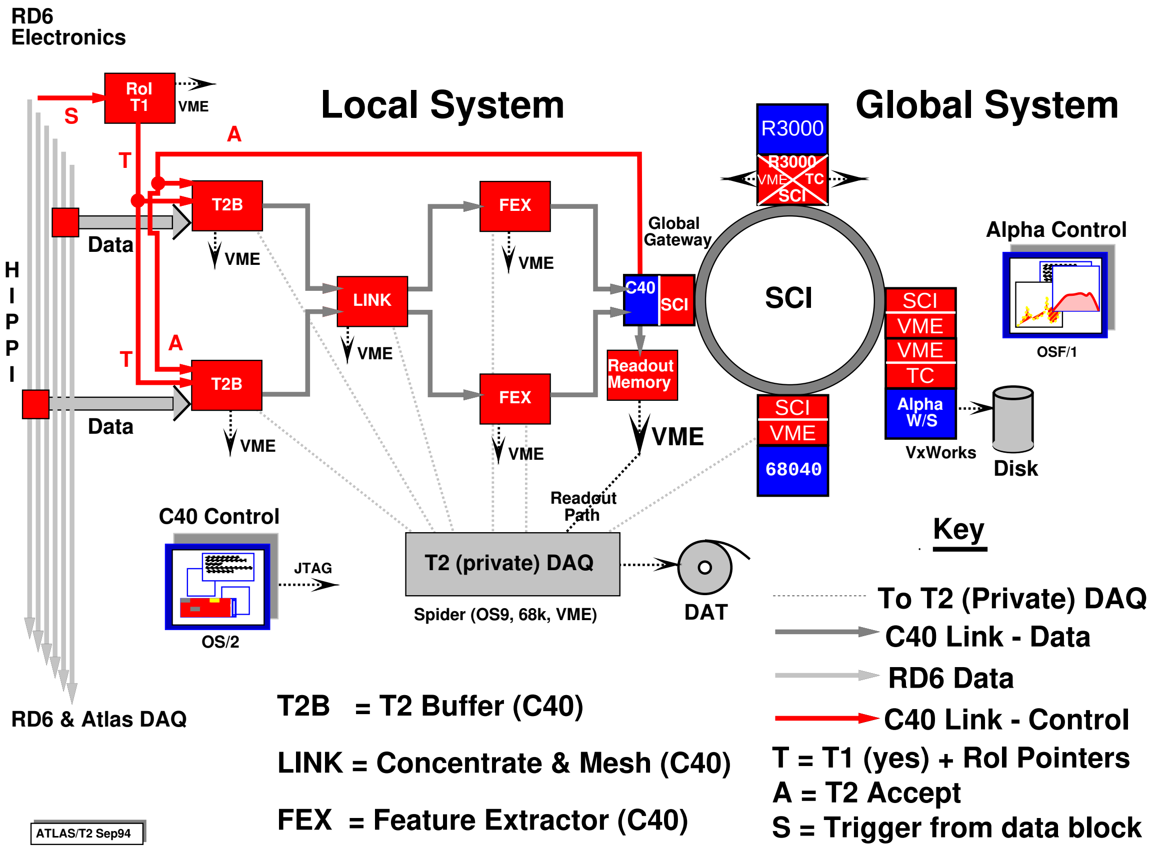

Figure 35: Test beam set-up, autumn 1994.

8.4.1 Beam test demonstration

The beam test system (figure 35), contained two level-2 buffers [37], one 'link unit', two feature extractors, a 'global gateway', a four node global network and a single global processor. The 'link unit' provided the local network and the 'global gateway' acted as the interface from the local to the global system. The local system components and the 'global gateway' were based on Texas Instruments TMS320C40 (C40) DSP, the global network was based on SCI and the global processor was a DEC Alpha workstation. The buffers and the SCI node interfaces were designed and built 'in-house'. Data were derived from a HIPPI spy unit which formed part of a router system [38] placed in the data readout path of the RD6 Transition Radiation Detector (TRD) [39]. The whole system was self-triggered by the passage of an event on the RD6 HIPPI lines. 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.

In the initial phase of the beam tests only the local part of the system was used. In this phase some 50K events were successfully recorded to disk. An example of an event recorded by the system is shown in figure 36. In the second phase, the SCI ring and global processor were added. It is believed that this is the first time that 'live' detector data has been passed round an SCI ring at a beam line.

8.4.2 Laboratory tests

The level-2 buffer is controlled by buffer manager software. Timing tests of this program have been made with a software emulation of the buffer hardware. Additional DSPs acted as data sources and sinks, so that the I/O environment seen by the buffer manager was essentially the same as in a real system. The performance of the

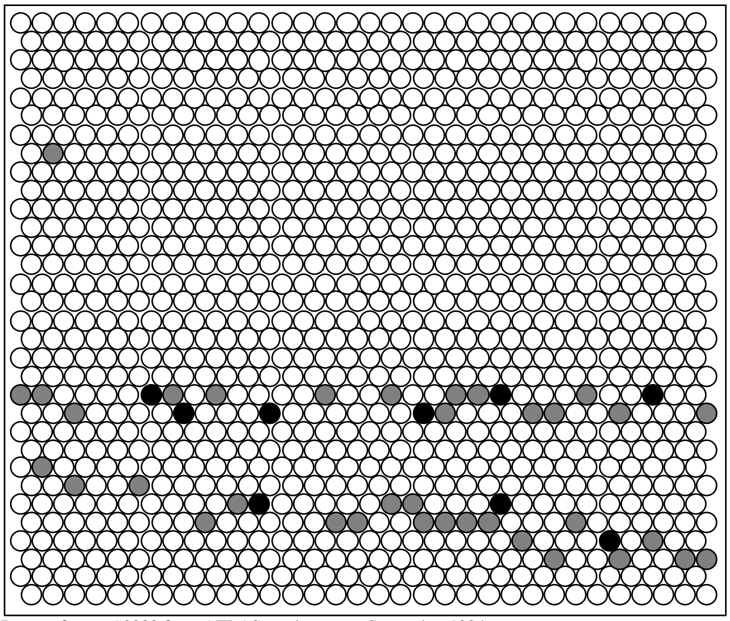

Dump of event #0089 from ATLAS test beam run September 1994

Figure 36: Example event recorded by the test beam system.

assembler-coded buffer manager was comfortably above the design goal of 100 kHz event rate. The results of these software-based timing tests are consistent with those obtained using a hardware level-2 buffer driven by a SLATE [40] event generator board.

For the 'link unit' tests, one C40 was used to send RoI data randomly to each of a set of link unit processors. With only a limited number of units available, interconnections have been wrapped around to simulate 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 yielding a link unit throughput of a approximately 20 MByte/s.

In the laboratory, tests of the SCI components have been made at data rates approaching the limits of our first generation SCI hardware [41]. In two node tests we have achieved ring occupancy of over 20%, and have investigated the latency effects associated with packet collision on the SCI ring. The correct operation of the custom SCI interfaces have been confirmed in this regime, and it has been shown how performance will degrade as the system scales. An understanding of these details is important input to the modelling studies and for the determination of the best ways to incorporate and use this technology in future interfaces.

Investigations have also started into the way the PCI local bus is driven by the DEC Alpha processor (21066), and the best way to access the processor memory system for optimal flow of feature data from the global network. Including PCI overheads, data transfer at a rate of 80 MByte/s has been observed — implying an overhead of about 25 μs to transfer all of the feature data for an event from the SCI interface into the memory.

8.4.3 Modelling

Members of the UK group have played a key part in setting up and directing a broad collaboration of people to model the performance of multi-level DAQ/Triggering systems based on networks and multi-processing. In addition, they have taken on the responsibility of developing and maintaining the core of the software.

A modelling system, SIMDAQ [42], has been developed based on the object-oriented simulation language MODSIM II. Input files of data generated from physics simulations and of detector readout segmentation allow architectures to be studied under realistic conditions.

The first implementation of this system has been applied to simulations of level-2 trigger architectures of the type studied by the UK groups. The performance of the level-2 trigger has been evaluated for three ATLAS sub-detectors, the electromagnetic and hadron calorimeters, and a tracking detector, the TRT. SIMDAQ has proved flexible and studies using it have provided valuable information on the performance requirements of system components. Figures 37 and 38 show example results from models of the beam test system and a larger ATLAS system.

Because test systems will be constructed on a scale small compared with ATLAS, modelling will provide the only way of assessing full scale architecture performance. Modelling, therefore, will remain a central component of the UK activity. Measurements on test systems and components will be used to provide input to models and results from the modelling of full scale systems will be used to assist in the specification of component parts.

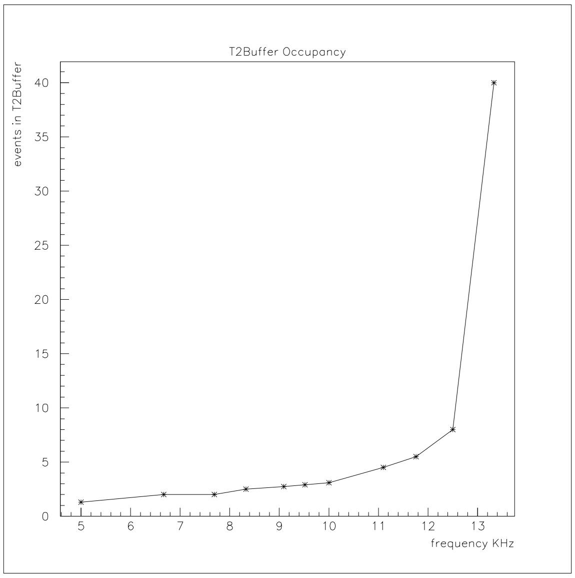

Figure 37: Example result from modelling the test beam system. The occupancy of a level-2 buffer is shown as a function of the buffer hit-rate. Note that in the full ATLAS system this hit-rate is expected to be between 1 and 5 kHz.

8.5 UK Responsibilities

Discussions on the division of responsibilities in the level-2 trigger are only just starting and the programme described below takes into account our best evaluation of other group's preferences. The UK groups support and complement each other in their range of interests, experience and expertise. Based on these strengths, we propose to work on the development of the local level-2 trigger system for the inner detector, the global level-2 trigger system, physics studies in relevant areas (including B-physics triggers) and system modelling for the trigger and data acquisition.

In this proposal we bid for the construction of 75% of the level-2 local system for the inner detector and for the construction of the level-2 global system. The RHBNC and UCL groups will take responsibility for the buffers, local processing, physics studies and tracking algorithm development. The Manchester and RAL groups will take responsibility for the networks, global processing and global algorithm developments. The Oxford group will lead the modelling activity and participate in the algorithm development. Table 18 shows further details.

The significant UK participation in the inner detector makes it sensible and efficient for part of the level-2 activity to be concentrated in this area. The work, in collaboration with the inner detector group, will be on the development of off-detector hardware for the inner detector and on the algorithms for track recognition for both triggering and offline analysis. UK groups have been working on high performance networks with tightly integrated RISC processors, and are now starting to investigate how this promising technology can be optimised for global processing. These activities will lead naturally to the UK responsibilities outlined above, where the scale of the proposed activity is commensurate with the size and strength of the UK groups involved.

8.6 Work Programme

The work programme is divided into three phases. The current phase is to prove architectural ideas and examine generic problems and will lead to an architecture choice at the end of 1997. This choice will be followed by a period of technology development and assessment resulting in technology selection for the level-2 system components around 2000. The third stage is the final design, construction and operation of the system.

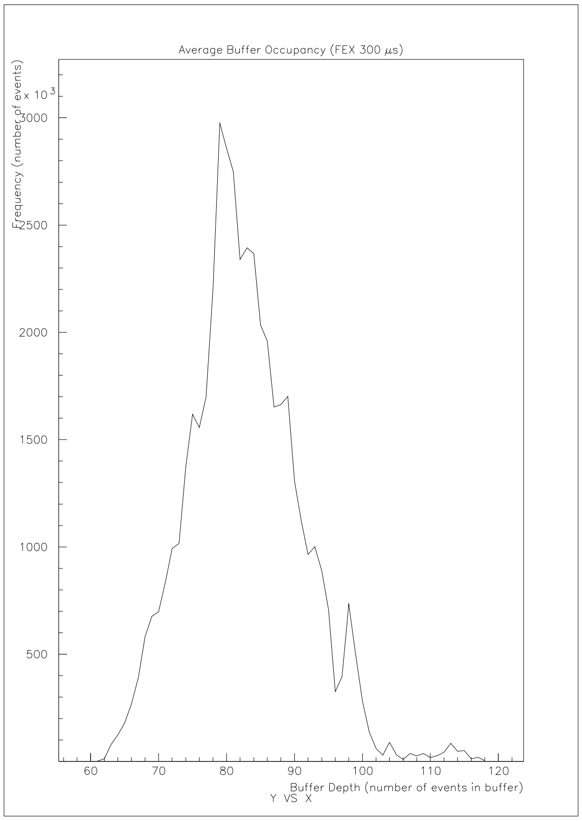

Figure 38: Example result from modelling a level-2 system with an SCI network. The distribution of level-2 buffer occupancy is shown for a trigger rate of 100 kHz.

The UK groups will continue with the development and exploitation of their current systems during the first phase and will move to improved devices and better technologies if and when these become available.

8.6.1 Short term (1995–97) — evaluation and design

For a system of the complexity of level-2, which depends on the interaction of several thousand processors, the monitoring and control of the system performance is vital. Each part of the system must be testable independently, and as a whole, without modification of any of the connections in the system. Fault tolerance must be built-in. With data spread across many memories and processors, the system must be designed to tolerate the loss of single components without serious compromise to the overall performance. Rapid reconfiguration of the system is required to circumvent failed components. For areas where reconfiguration does not or cannot help, high reliability of components will be essential. These issues are central to the ability of the system to operate reliably.

With many processors active on many different events simultaneously, the performance of the system as a whole must be considered. It is insufficient to make measurements on each component of the system in isolation. It is imperative that complete systems are modelled and tested in the laboratory before serious architecture and technology comparisons can be made. For this reason, the twin activities of modelling and laboratory (and test beam) demonstrations are vital elements of the level-2 trigger programme in both the short and medium terms.

Over the next 3 years the work will continue with a balanced programme of modelling, simulation, and tests both in the laboratory and at a test beam. Many issues need to be studied including RoI guidance, integration of multiple detectors, fault tolerance, error-recovery, optimisation of the use of network bandwidth and compute power, initialisation procedures, monitoring and control. In parallel, activities concentrated on the definition of data formats, the development of prototype local and global algorithms and the specification of level-2 supervisor functionality will be conducted. At all stages the results of laboratory tests will be used as

Item	Responsible	Other Contributors
ID interface	UCL	RHBNC, ID groups
Buffers	RHBNC/UCL	
Local processing	RHBNC/UCL	
Node interfaces	Manchester	RAL
Networks	Manchester/RAL	
Global processing	RAL	Manchester
Modelling	Oxford	RAL, RHBNC, UCL
Trigger simulation	RHBNC/UCL	Oxford

Table 18: Level-2 areas of responsibility.

input to the modelling activity. The modelling will seek to simulate the performance of the test system and to estimate the performance of full scale systems.

Preparations for this test work will include the development of a front end to the current prototype level-2 buffer to accept data from the silicon detectors. This will be done in collaboration with the ATLAS UK groups working on the inner detector. Hardware solutions to the specific problems of the global networking interfaces will be developed, especially hardware assistance to accelerate the management of the feature data.

This work will be conducted in close collaboration with many non-UK partners working on complementary parts of the level-2 system. At the end of this phase, decisions on details of architecture and specific technologies will be made by ATLAS and will be adopted by the UK group for subsequent phases.

8.6.2 Medium term (1998–2000) — large scale prototype

In this phase a large scale prototype will be constructed with the chosen architecture and candidate technologies will be subjected to rigorous tests and optimisation. This will include evaluation of the physics performance of the full system design and the definition of diagnostic procedures.

This work will take the form of construction of a complete slice through the level-2 trigger system (both local and global). The large scale prototype will be tested both in the laboratory and in the test beam in conjunction with the level-1 demonstrator system, the inner detector prototype and the DAQ system. The connection to the inner detector prototype will provide the level-2 group with real data and therefore realistic tests.

The prototype tests will be accompanied by further modelling work to ensure optimum and scalable system performance. During the last year of this phase final technology choices will be made and the full system design completed.

8.6.3 Long term (2001–04) — production and integration

The final phase is sub-divided into procurement or construction, module tests, integration and initial operation. Table 19 shows the number of modules that will be required and indicates the full scale of the project.

The UK expertise in high performance networks and tightly integrated processors will be brought to bear in the specification of modules to be purchased commercially or to be built 'in-house'. The latter being necessary to achieve the fully optimised performance required.

All modules must undergo test procedures to ensure quality and reliability of the final system. To this end, these procedures will be defined as appropriate.

Integration of components will not only involve bringing together parts of the local and global sub-systems, but also the connections to sub-detectors, to the level-3 event building and DAQ system and to the level-2 supervisor. The final local and global trigger algorithms will be implemented and put to test.

Final debugging and test of the full system will be during the year 2003 and will include the implementation of full scale initialisation, monitoring, control and diagnostic procedures. This should ensure a reliable level-2 system for production operation in 2004, able to efficiently trigger on candidate physics processes.

8.7 Resource Tables

Tables 20 and 21 show the costs of the proposed level-2 items and their spend profile. Table 22 show the effort available for this work and table 23 shows the additional electronic engineering effort requested at RAL.

Item	**number**
Input to Inner Detector Buffers (GByte/s)	20
Number of Inner Detector Buffers	240
Number of Feature Extractor Processors	350
Local Processor power (MIPs/processor)	500
FeX processing time (μs)	350
Local network throughput (MByte/s)	2000
Global Processors	80
Global Processor power (MIPs/processor)	1000
Global processing time (μs)	400
Global network throughput (MByte/s)	200

Table 19: Estimates of system size and requirements for the inner detector local system and the global system.

Item	**total**	**fraction**	**total**
Inner detector buffers	1440	75%	1080
demonstrator	140	75%	110
local system	3600	75%	2700
demonstrator	360	75%	270
Total (kCHF)	5540	75%	4160
global system	1530	100%	1530
demonstrator	230	100%	230
Total (kCHF)	1760	100%	1760
Grand Total (kCHF)	7300		5920

Table 20: CORE costings for the inner detector buffers and local system, and the global system.

Year	pre-1996	1996	1997	1998	1999	2000	2001	2002	2003	2004
Spend (UK k£)	82	56	61	83	61	278	1056	1333	167	111
Accum. (UK k£)	82	138	199	282	343	621	1677	3010	3177	3288

Table 21: Level-2 spend profile.

Year	1996	1997	1998	1999	2000	2001	2002	2003	2004
Physicists	10	11	11	12	12	13	13	13	13
Engineers	6	6	7	7	7	7	7	7	7
Technicians	4	5	5	6	6	7	7	5	5

Table 22: University and RAL PPD staff for level-2 (excluding students).

Year	pre-1996	1996	1997	1998	1999	2000	2001	2002	2003	2004
Electronics SYs	-	2	3	3	4	8	8	8	8	6
Accumulated SYs	-	2	5	8	12	20	28	36	44	50

Table 23: Level-2 RAL Engineering Staff request.

8.8 Conclusions

The UK groups have made substantial contributions to the work on the level-2 trigger within ATLAS and play a leading role in this activity. The programme outlined allows the groups to maintain a significant profile during the design and development stages and to provide a substantial section of the level-2 trigger. It addresses both the use of tracking information and the application of event topology constraints, which are vital to achieve the rate reduction required of level-2. The work programme is well matched to the experience and expertise in the UK groups and provides an interesting and varied programme of work covering many aspects of the trigger activity. The proposed spread of trigger activities, spanning from level-1 through local level-2 to the global level-2 system, gives the UK activities breadth and cohesion. By concentrating the local system work on the connection to the inner detector, the coherence of the UK efforts is further strengthened.

9 Summary of Requests

The following tables summarise the requests to the PPESP for the three major sub-projects proposed by the ATLAS UK community, and the level of university manpower available for ATLAS construction. Additional information is given in Appendix B and Appendix C on ATLAS software activities in the UK and forward toroid design, respectively. Although no bid for resources is made for these activities in this proposal, they represent a significant part of the UK contribution to ATLAS.

9.1 Human Resources

Institute	Physicists	Eng./Appl.Phys./Support Phys.	Prog./Phys.Prog.	Technical Support
Birmingham	8.25	1.5	1.5	3.0
Cambridge	2.5	2.5	1.5	3.0
Edinburgh	1.0	–	1.0	–
Glasgow	3.75	1.5	1.0	3.0
Lancaster	3.0	1.0	0.5	2.0
Liverpool	4.75	3.0	1.0	4.0
Manchester	3.25	3.5	–	7.5
Oxford	6.5	4.5	2.0	6.0
QMW	3.0	3.0	1.5	3.5
RAL PPD	7.0	12.5	6.0	2.0
RHBNC	4.0	1.0	1.0	2.5
Sheffield	2.5	–	–	2.0
UCL	3.0	2.25	2.5	2.5

Table 24: Overall university and RAL(PPD) manpower effort (f.t.e.) during the construction phase.

9.2 Profile of Spend and Additional Manpower Requested

Item	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	Total
Inner Detector											
Spend (UK k£)	170	650	1080	3400	3570	2800	1850	700	300	100	14620
Manpower (S/Y)	19	33	48	52	49	50	39	30	23	16	359
Level-1 Trigger											
Spend (UK k£)	115	80	35	186	431	840	1520	325	100	100	3732
Manpower (RAL Tech S/Y)	2.5	3	4.5	9	10	11	9	9	6	3	67
Level-2 Trigger											
Spend (UK k£)	82	56	61	83	61	278	1056	1333	167	111	3288
Manpower (RAL Tech S/Y)		2	3	3	4	8	8	8	8	6	50
Common Fund Contribution											
Estimated total and profile (UK k£)		431	526	589	3543	3690	704	578	231	221	10513
Total profile											
Spend (UK k£)	367	1217	1702	4258	7605	7608	5130	2936	798	532	32153
Manpower (S/Y)	21.5	38	55.5	64	63	69	56	47	37	25	476

Table 25: Profile of requested spend, assuming 1.8 CHF/ £, and additional manpower.

Note that throughout this document an exchange rate of 1.8 CHF/ £ has been used for consistency.

Estimates of our Travel and Subsistence requirements, based on our experience at LEP and HERA, indicate a peak level of approximately £ 1M per annum. More detailed analysis of the requirements are being prepared.