ATLAS Internal Note
DAQ-NO-055
27 June 1996

A Model for Sequential Processing in the ATLAS LVL2/LVL3 Trigger

Version 1.0

J. Bystricky, D. Calvet, J. Ernwein, O. Gachelin, T. Hansl-Kozanecka,
J.R. Hubbard, M. Huet, P. Le Dû, I. Mandjavidze, M. Mur,
M. Smizanska, and B. Thooris

CEA - DAPNIA, Saclay
91191 Gif-sur-Yvette, France

Abstract

A model of the ATLAS trigger is proposed, including preprocessing, data transfers, algorithm sequences and execution times. Latencies and processor and switch occupations are calculated, based on trigger menus for low luminosity (10^{33} cm^{-2} s^{-1}) described in a previous ATLAS note [1].

A sequential processing option is proposed in which all trigger selection algorithms are executed at LVL2. We show that this option requires fewer resources (network bandwidth and processors) than standard options in which trigger selection is split between LVL2 and LVL3 processing farms.

Table of Content

1. Introduction

The execution time and the latency of the LVL2 and LVL3 event selection depends on the initial (LVL1) and final (LVL3) trigger menus and on the trigger architecture and implementation. Many of the parameters - algorithm execution times, network transmission time, and overheads - can be measured on today's machines and extrapolated to future machines. In this note, we propose a model for sequential processing of the ATLAS trigger algorithms, and we estimate the parameters required by this model, based on the information we have today.

Two options for LVL2 and LVL3 processing are considered. Both options assume sequential processing of the trigger algorithms. In the first option, the trigger selection is separated into LVL2 algorithms and LVL3 algorithms as described in the ATLAS Technical Proposal [2]; LVL2 reduces the trigger rate to a few kHz, and LVL3 reduces the rate to about 100 Hz (or a bandwidth of about 100 MB/s). In the second option all selection algorithms are executed in the LVL2 processors, and LVL3 is reserved for physics analysis and data acquisition, with an event rate of about 100 Hz.

For each of these two options, we evaluate the event selection latencies and the required processing power and network bandwidth. A trigger menu for low luminosity (10^{33} cm^{-2} s^{-1}) described in an earlier ATLAS DAQ note [1] has been used for these studies.

The calculations contained in the present note have been labelled "paper modelling". This model is idealized with respect to the situation we will encounter in a real experiment. In real life, execution times will be distributed about some mean, with tails reaching out to execution times which are much longer than the average times estimated here. Furthermore, even though we have considered processor and network overheads in this note, we have not taken into account queuing delays. The results obtained here are necessarily optimistic with respect to the average parameters assumed, even though the parameters have been estimated conservatively.

2. Single Farm Architecture

The LVL2 system described in the Technical Proposal is based on a “local-global” scheme, as shown in Fig. 1. In this scheme, “feature extraction” algorithms are executed in parallel in “local” processor farms on “local” data from each subdetector and each RoI. The results of the local feature extraction are transferred to a “global” processor farm, which makes the overall LVL2 decision.

We propose a different LVL2 architecture with a single processor farm and a single network linking that farm to the Read-Out Buffers (ROBs). This “single-farm” scheme is motivated by the ATLAS physics requirements. Efficient use of resources (networks and processors) requires a sequential processing scheme, in which events can be rejected after any of the processing steps. In our scheme, this is accomplished by assigning a single processor to each event; this processor executes the trigger algorithms step-by-step, requesting data for the next processing step only if the analysis is still consistent with at least one set of trigger conditions.

The specific implementation described here is based on the functional model shown in Figs. 2 and 3. A group of ROBs is connected to a switching network via a ROB-to-Switch Interface (RSI). Similarly, a group of processors (typically 4-8) in the processor farm are connected to the same switching network via a Switch-to-Farm Interface (SFI). The LVL2 Supervisor assigns
an event to a given LVL2 processor, which requests data from the ROBs as needed for processing the sequential trigger algorithms. All of the control messages and all of the data transit through the same bi-directional switching network. The final LVL2 decision is sent to the LVL2 Supervisor by the LVL2 processor in charge of the event. The Supervisor broadcasts the decisions to all of the ROB's via the network.

If the event is accepted at LVL2, the ROBs send all their data to the LVL3 processor specified by the LVL2 (or LVL3) Supervisor; this process is called "Event Building". For certain types of events, partial data instead of full data may be sent to the LVL3 Event Builder, but this case will not be considered in this note.

3. Assumptions

In this paper we attempt to model the "final" ATLAS T/DAQ system, rather than today's demonstrators or prototypes. The set of working assumptions listed in this section are our "best-guess" values, based in part on measurements made with today's hardware and in part on estimates of future hardware and software performance. Data transfer times are taken to be those which can be obtained with today's ATM switching hardware and software performance. Algorithm execution times are taken to be five times faster than those obtained today.

Many people have contributed to our current understanding of the parameters, partly through discussions and partly by measuring certain parameters on today's hardware. Developping a better set of working assumptions requires more discussion, more measurements, and more complete modellization. We would appreciate constructive criticism of the assumptions used here.

3.1 Data Volumes

Our assumptions concerning the data volumes in the ROBs and the RSIs are based on the work by R. Bock and P. Le Dû on the interfaces with the front-end systems [3]. The data volumes (per event) are given in the following table for each subdetector, along with the number of ROBs, the number of RSIs, and $\Delta\eta \times \Delta\phi$ for the ROBs and the RSIs. There are two entries for the TRT data, one for low luminosity and one for high luminosity.

DATA VOLUMES PER EVENT IN ROBs AND RSIs

Subdetector	data	ROBs	$\Delta\eta$ x $\Delta\phi$	RSI's	$\Delta\eta$ x $\Delta\phi$
muon	185 kB	128 ROB	0.8 x 0.2	32 RSI	0.8 x 0.8
calorimeter	450 kB	256 ROB	0.4 x 0.4	64 RSI	1.6 x 0.4
TRT $L=10^{33}$	80 kB	256 ROB	0.8 x 0.1	64 RSI	0.8 x 0.4
TRT $L=10^{34}$	320 kB	256 ROB	0.8 x 0.1	64 RSI	0.8 x 0.4
SCT	320 kB	128 ROB	1.2 x 0.2	32 RSI	1.2 x 0.8
Total $L=10^{33}$	1035 kB	768 ROB		192 RSI	
Total $L=10^{34}$	1275 kB	768 ROB		192 RSI	

We assume that there has been zero suppression for the TRT data in the front-end Read-Out Drivers (ROD's). Without zero suppression, the TRT data would be 790 kB. The TRT data includes data from three beam crossings, but we assume that data from a single beam-crossing are used for feature extraction. The SCT data includes 55 kB of pixel data.

The number of ROBs in the table is only half of the standard number of ROBs. This is based on the assumption that we can input data to each ROB from two ROD's (i.e., two 1 Gbit/s links). This assumption needs verification. We will produce tables with the standard hypothesis (1 536 ROBs) in the next version of this note.

The data volume in each Region of Interest (RoI), as well as the size of the RoI for each trigger and/or for each subsystem, is given in the next table.

DATA VOLUMES IN RoIs AFTER PREPROCESSING

RoI type	system	RoI data	ROBs/RoI	RSIs/RoI	$\Delta\eta$ x $\Delta\phi$
muon trigger	MUON	1 kB	1 ROB	1 RSI	0.1 x 0.1
muon trigger	CALO	3 kB	3 ROB	2 RSI	0.4 x 0.4
e / γ / τ trigger	CALO	3 kB	3 ROB	2 RSI	0.4 x 0.4
jet trigger	CALO	3 kB	9 ROB	5 RSI	1.2 x 1.2
track $L=10^{33}$	TRT	0.1 kB	3 ROB	2 RSI	0. 8 x 0.2
track $L=10^{34}$	TRT	0.3 kB	3 ROB	2 RSI	0.8 x 0.2
track	SCT	0.2 kB	2 ROB	1 RSI	0.2 x 0.2
TRT scan $L=10^{33}$	TRT	16 kB	256 ROB	64 RSI	5.0 x 6.4
TRT scan $L=10^{34}$	TRT	64 kB	256 ROB	64 RSI	5.0 x 6.4
missing-E_T	CALO	16 kB	256 ROB	64 RSI	10.0 x 6.4
b-jet tag	SCT	3 kB	8 ROB	3 RSI	0.8 x 0.8

Note that we assume that only data inside the RoI is transferred to the LVL2 processors (after preprocessing in each ROB).The e / γ and τ algorithms are based on the same RoIs; they have been grouped together in the table. The single hadron trigger is a special case of the τ trigger.

3.2 Data Transfer

Data transfer rates and overheads have been measured at Saclay, both on ATM switching fabrics and on PCI busses. The nominal transfer rates are 19 MB/s (155 Mbits/s) for each link of the ATM switch, and 132 MB/s for the PCI bus; the maximum rates of useful data were measured to be about 15 MB/s for a ATM link, and 70 MB/s for the PCI bus. The data transfer latency was 50 μs through the ATM switch.

The ATM switches have dual links at each port, for data transfers in the two directions. The occupancy should be calculated for transfers in each direction. In our case, data transfers from the ROBs to the (LVL2 and LVL3) processors (designated ATM1) have much higher bandwidth than the control messages transferred from the processors to the ROBs (ATM2).

The processors suffer overheads for data transfers. The device-driver software overhead for sending a message has been measured to be about 30 μs; we assume that this overhead can be reduced to 10 μs in the future. The overhead for receiving a message has been measured to be about 50 μs; we assume that it will continue to be about 50 μs in the future. (The ATM measurements were made using a 50 MHz PowerPC with a LynxOS operating system and a 50 MHz Sparc20 workstation)

We assume that the data from the ROBs will be reordered in the RSIs before being transferred to the switching network, and that data arriving from different RSIs will again be reordered in the SFI's before transfer to the LVL2 processors. We assume that the data can be reordered at the rate of 50 MB/s.

3.3 Preprocessing

We assume that the raw data in the ROBs is preprocessed before sending it to the LVL2 processors. We assume that this preprocessing is performed separately for each ROB. The preprocessing algorithms have not yet been written (although the calorimeter preprocessing is rather trivial), so their execution times could not be measured. Our assumptions about the preprocessing algorithms and their execution times are given in the following table. Our model assumes intelligence equivalent to a 100 MIPS DSP in each of the ROBs.

PREPROCESSING EXECUTION TIMES PER ROB			
Trigger	system	preprocessing	algorithm
muon trigger	MUON	100 μs	Space points
e / γ / τ trigger	EMC	50 μs	Calibration only
e / γ / τ trigger	HAC	20 μs	Calibration only
jet trigger	EMC	50 μs	Calib.+Energy sums
jet trigger	HAC	20 μs	Calibration only
missing-E_T	EMC	70 μs	Calib. + Calculate E_x, E_y
missing-E_T	HAC	40 μs	Calib. + Calculate E_x, E_y
track L=10^{33}	TRT	200 μs	Data compression
track L=10^{34}	TRT	800 μs	Data compression
track	SCT	200 μs	Space points
b-jet tag	SCT	800 μs	Space points

3.4 Algorithm Execution Times

Stand-alone versions of the subdetector feature extraction algorithms have been written and benchmarked at several institutions. The pioneer work on algorithms and their execution times was performed at CERN [4]. More recently stand-alone muon feature extraction code has been written and benchmarked in Rome, using a 40 MIPS DSP with optimized Fortran code [5]. The calorimeter feature extraction code has been benchmarked at Saclay on several machines, with different levels of optimization. Detailed studies of the precision-tracker code have been carried out by S.Sivoklokov, R.Dankers, and J.Baines [6]. For the TRT, we use the benchmarks established by Hauser and Legrand [4], even though their algorithm has not yet been compared with the algorithm used in the ATRIG code. For the full TRT scan required for the B-physics algorithms, M.Smizanska has measured execution times on standard CERN processor farms [7]; execution times for optimized code have not yet been measured.

The assumptions about feature extraction execution times shown in the table below are based on extrapolations of the measured execution times to what we would expect using optimized code on a future 500 MHz processor. In general, we have divided the measured execution times by a factor five to account for the expected improvement in the speed of the processors, but we have multiplied the values obtained using stand-alone code by a factor two because we expect that additional data manipulation will be required in the final code. We have also reduced the execution time for the full TRT scan by a factor three because we believe that such an improvement can be obtained when we optimize the code for low-p_T tracks.

The algorithm for the b-jet tag has not been studied in the trigger group. Typically execution times for off-line code are "a few seconds", with long tails at larger execution times.

The estimate in the table below is 250 ms average execution time, using optimized code in the LVL2 processors. This assumes a factor five improvement in the processor speed, and another factor two due to optimization of the LVL2 code.

FEATURE EXTRACTION PROCESSING TIMES IN LVL2 PROCESSORS

Trigger	system	processing (μs)	measured execution times (μs)	
muon trigger	MUON	100	220	[5]
e / γ / τ trigger	CALO	100	125	
jet trigger	CALO	100		
track	TRT	600	700	[4]
track	SCT	800	1 840	[6]
TRT scan L=10^{33}	TRT	50 000	680 000	[7]
TRT scan L=10^{34}	TRT	200 000		
missing-E_T	CALO	100		
b-jet tag	SCT	250 000		

An important feature of the sequential processing model is that the LVL2 processors must be able to perform in a multi-task mode, with several events being processed at the same time. When a processor requests new data for the event being processed, these data will be available only after a delay of several 100 μs. In order to avoid losing this time, the processor must be able to switch to another event. The overhead due to this context switching has been measured at Saclay to be about 60 μs; we assume that the context switching will require 50 μs in the future.

In the more standard option for the ATLAS trigger, some of the event selection is performed at LVL3. Here we assume that any recalculation of the missing-E_T and any calculation of b-jet tags would be performed at LVL3 in this option. We expect that the LVL3 code would be similar to the off-line analysis code, and less optimized than the LVL2 code.

We assume that the processing time would be larger by a factor two at LVL3, compared to the processing time for the same algorithm at LVL2. Furthermore, we assume that the LVL3 analysis is not limited to Regions of Interest, and that the minimum time for the treatment of any event at LVL3 would be about 100 ms. We also assume that a full physics analysis can be performed in a few seconds, on average. These LVL3 execution times are summarized in the following table:

LVL3 ALGORITHM PROCESSING TIMES	
Algorithm	processing time (ms)
masses, multiplicities	100
B physics algorithms L=10^{33}	100
B physics algorithms L=10^{34}	400
b-jet tag	500
physics analysis	1 - 10 s

4. Processing Sequence

This chapter gives a quantitative model for sequential processing in a single-farm architecture. The timing is given for each step in the processing sequence. These estimated times are then used to calculate the event selection latencies and the resources required (network links and processors) for different LVL2 and LVL3 strategies. The average latency and the resources required are then calculated for a specific trigger menu and for a specific model for the sequential algorithms and for the step-by-step rejection factors.

Two important processing steps have been left out of the present study. For events accepted by the LVL2 trigger, LVL2 data produced by the LVL2 processors (features, masses,etc.) must be transfered through the switching network to the LVL2 trigger ROBs; this process has not been modelled here. Furthermore, the LVL2 and LVL3 processing must be monitored for efficiency and errors; unfortunately, we do not yet have a model for these control functions.

4.1 Processing Sequence Step-by-Step

The LVL2 and LVL3 processing sequence (for any given algorithm) is broken down into 13 processing steps, starting with the transfer of data from the front-end RODs to the ROBs, and ending with the transfer of the LVL3 decision to the supervisor. Any comparison of different event selection strategies must include all of the data transfers and all of the trigger selection algorithms (whether executed at LVL2 or at LVL3).

The execution times in this section are based on the assumptions listed in Chapter 3 of this note. When data transfers are required, an estimate of the data volume to be transferred is given in the last column, after the execution time. Execution times listed as "x/RSI" or "y/ROB" indicate that the execution times (x or y) should be multiplied by the number of RSIs or ROBs implicated in that particular data transfer. Execution times listed as "RoI data / z MB/s" indicate execution times of 1/z μs per byte of data in the RoI. The expressions "ROB", "RSI", and "RoI" refer to the number of ROB's, RSI's, or RoI's involved in that particular transfer of data. Preprocessing, feature extraction, and LVL3 algorithm execution times are denoted by "T_{pre}",

"T_{fex}", and "T_{alg}", respectively.

The sequential nature of the processing model described here is illustrated by the repetition indicated for processing steps 3 to 7. The sequential data requests and data transfers are described in these five processing steps. Each data request is preceded by a test of the compatibility of the event with the LVL2 trigger conditions.

The processing steps and their execution times are given below:

1)	SEND LVL1 DATA TO LVL2			
	ROD latency		100 µs	
	ROBIN	ROD -> ROB	10 µs / link	
	ROI Builder	RoIs -> SUPER	100 µs	
2)	ASSIGN LVL2 PROCESSOR			
	SUPER	code message	50 µs	
		send -> SFI	10 µs	
	ATM1	latency	50 µs	
		transfer	6 µs	0.1 kB
	SFI	receive	50 µs	
		send -> LVL2	10 µs	
	PCI	transfer	2 µs	0.1 kB
	LVL2	switch context	50 µs	
3)	REQUEST DATA FROM ROBs			
	LVL2	process RoI info	50 µs	
		send -> SFI	10 µs	
	PCI	transfer	2 µs	0.1 kB
	SFI	receive	50 µs	
		send -> RSI	10 µs	
	ATM2	latency	50 µs	
		transfer	6 µs / RSI	0.1 kB/RSI
	RSI	receive	50 µs / RSI	
		send -> ROB	10 µs / ROB	
	PCI	transfer	2 µs / ROB	0.1 kB/ROB
	ROB	receive	50 µs / ROB	
4)	PREPROCESS DATA IN EACH ROB			
	ROB	preprocess	T_{pre} / ROB	

5) SEND PREPROCESSED DATA TO LVL2

ROB	send -> RSI	10 μs / ROB	
PCI	transfer	RoI data / 70 MB/s	RoI data
RSI	receive	50 μs / ROB	
	reorder data	RoI data / 50 MB/s	
	send -> SFI	10 μs / RSI	
ATM1	latency	50 μs	
	transfer	RoI data / 15 MB/s	RoI data
SFI	receive	50 μs / RSI	
	reorder data	RoI data / 50 MB/s	
	send -> LVL2	10 μs	
PCI	transfer	RoI data / 70 MB/s	RoI data
LVL2	switch context	50 μs	

6) EXECUTE FEATURE EXTRACTION ALGORITHMS

LVL2	feature extraction	T_{fex} / RoI

7) TEST COMPATIBILITY WITH LVL2 TRIGGER

LVL2	reject / continue	50 μs

REPEAT STEPS 3 TO 7 AS REQUIRED FOR OTHER DETECTOR SYSTEMS AND/OR ROIs

---

8) TEST GLOBAL LVL2 TRIGGER CONDITIONS

LVL2	accept / reject	150 μs

9) SEND LVL2 DECISION TO SUPERVISOR

LVL2	send -> SFI	10 μs	
PCI	transfer	2 μs	0.1 kB
SFI	receive	50 μs	
	send -> SUPER	10 μs	
ATM2	latency	50 μs	
	transfer	6 μs	0.1 kB
SUPER	receive	50 μs	

10) BROADCAST LVL2 DECISION

SUPER	broadcast -> RSI	10 μs	
ATM2	latency	50 μs	
	transfer	6 μs / RSI	0.1 kB/RSI
RSI	receive	50 μs / RSI	
	send -> ROB	10 μs / ROB	
PCI	transfer	2 μs / ROB	0.1 kB/ROB
ROB	receive	50 μs / ROB	
	clear buffer	10 μs / ROB	

11) SEND DATA TO LVL3 PROCESSORS

ROB	send -> RSI	10 μs / ROB	
PCI	transfer	RoI data / 70 MB/s	RoI data
RSI	receive	50 μs / ROB	
	reorder data	RoI data / 50 MB/s	
	send -> SFI	10 μs / RSI	
ATM1	latency	50 μs	
	transfer	RoI data / 15 MB/s	RoI data
SFI	receive	50 μs / RSI	
	reorder data	RoI data / 50 MB/s	
	send -> LVL3	10 μs	
PCI	transfer	RoI data / 70 MB/s	RoI data
LVL3	switch context	50 μs	

12) EXECUTE LVL3 ALGORITHMS

LVL3	execute algorithms	T_{alg}

13) SEND LVL3 DECISION TO SUPERVISOR

LVL2	send -> SFI	10 μs	
PCI	transfer	2 μs	0.1 kB
SFI	receive	50 μs	
	send -> SUPER	10 μs	
ATM2	latency	50 μs	
	transfer	6 μs	0.1 kB
SUPER	receive	50 μs	

4.2 Latency and Occupation Step-by-Step

The execution times can be regrouped to indicate the overall latency at each step of the processing, as well as the total execution time for each of the hardware elements. The notation is the same as in the preceding section, except that the expression "max(ROB/RSI)' is used in the calculation of latencies to take account of the parallel operation of the groups of ROBs assigned to each RSI.

For certain operations involving the LVL2 Supervisor and the broadcast of information by the LVL2 Supervisor, percentage occupations are listed in parentheses after the operation. These percentages are above 100%, so these operations will not be possible without special treatment. In step 1), the 200% WRITE occupation of the ROBs corresponds to the non-standard use of two 1 Gbit/s links for each ROB. The 600% Supervisor occupation in step 2) and the 500% Supervisor occupation in step 9) can be resolved by implementing the LVL2 Supervisor as a farm, with at least six bi-directional links connecting to the switching network. The 600% ROB and the 900% RSI occupations due to the broadcast by the LVL2 Supervisor to all of the ROBs

of the final LVL2 decisions can be resolved by grouping the decisions in the Supervisor, and broadcasting them together at fixed intervals (once every 1 ms); the occupation levels of the Supervisor, the RSI PCI busses, and the ATM2 links (Supervisor → ROBs) are also reduced to acceptable levels by this grouping of the final decision broadcasts.

The step-by-step latencies and the occupation of each of the hardware elements is given below:

---

1) SEND LVL1 DATA TO LVL2

Latency	100 μs	
SUPER occupation	10 μs	
ROB occupation	20 μs / ROB	(200% WRITE occupancy)
RSI PCI occupation	0	
RSI occupation	0	
ATM occupation	0	
SFI occupation	0	
SFI PCI occupation	0	
LVL2 processor occupation	0	

2) ASSIGN LVL2 PROCESSOR

Latency	228 μs	
SUPER occupation	60 μs	(600% occupation)
ROB occupation	0	
RSI PCI occupation	0	
RSI occupation	0	
ATM1 occupation	6 μs	
SFI occupation	60 μs	
SFI PCI occupation	2 μs	
LVL2 processor occupation	50 μs	

3) REQUEST DATA FROM ROBs

Latency	272 μs + 6 μs * RSI/RoI *RoI
	+ 12 μs * max(ROB/RSI)
SUPER occupation	0
ROB occupation	50 μs * ROB/RoI * RoI
RSI PCI occupation	2 μs * ROB/RoI * RoI
RSI occupation	10 μs * ROB/RoI * RoI
	+ 50 μs * RSI/RoI * RoI
ATM2 occupation	6 μs * RSI/RoI * RoI
SFI occupation	60 μs
SFI PCI occupation	2 μs
LVL2 processor occupation	60 μs

4) PREPROCESS DATA IN EACH ROB

Latency	$\max(T_{pre})$
SUPER occupation	0
ROB occupation	T_{pre} * ROB/RoI * RoI
RSI PCI occupation	0
RSI occupation	0
ATM occupation	0
SFI occupation	0
SFI PCI occupation	0
LVL2 processor occupation	0

5) SEND PREPROCESSED DATA TO LVL2

Latency	130 μs + RSI data / 70 MB/s + RSI data / 50 MB/s + 50 μs * max (ROB/RSI) + 50 μs * RSI/RoI * RoI + RoI * RoI data / 15 MB/s + RoI * RoI data / 50 MB/s + RoI * RoI data / 70 MB/s
SUPER occupation	0
ROB occupation	10 μs * ROB/RoI * RoI
RSI PCI occupation	RoI * RoI data / 70 MB/s
RSI occupation	50 μs * ROB/RoI * RoI + RoI * RoI data / 50 MB/s + 10 μs * RSI/RoI * RoI
ATM1 occupation	RoI * RoI data / 15 MB/s
SFI occupation	50 μs * RSI/RoI * RoI + RoI * RoI data / 50 MB/s + 10 μs
SFI PCI occupation	RoI * RoI data / 70 MB/s
LVL2 processor occupation	50 μs

6) EXECUTE FEATURE EXTRACTION ALGORITHMS

Latency	T_{fex} * RoI
SUPER occupation	0
ROB occupation	0
RSI PCI occupation	0
RSI occupation	0
ATM occupation	0
SFI occupation	0
SFI PCI occupation	0
LVL2 processor occupation	T_{fex} * RoI

7) TEST COMPATIBILITY WITH LVL2 TRIGGER

Latency	50 μs
SUPER occupation	0
ROB occupation	0
RSI PCI occupation	0
RSI occupation	0
ATM occupation	0
SFI occupation	0
SFI PCI occupation	0
LVL2 processor occupation	50 μs

REPEAT STEPS 3 TO 7 AS REQUIRED FOR OTHER DETECTOR SYSTEMS AND/OR ROI's

---

8) TEST GLOBAL LVL2 TRIGGER CONDITIONS

Latency	150 μs
SUPER occupation	0
ROB occupation	0
RSI PCI occupation	0
RSI occupation	0
ATM occupation	0
SFI occupation	0
SFI PCI occupation	0
LVL2 processor occupation	150 μs

9) SEND LVL2 DECISION TO SUPERVISOR

Latency	178 μs	
SUPER occupation	50 μs	(500% occupation)
ROB occupation	0	
RSI PCI occupation	0	
RSI occupation	0	
ATM2 occupation	6 μs	
SFI occupation	60 μs	
SFI PCI occupation	2 μs	
LVL2 processor occupation	10 μs	

10) BROADCAST LVL2 DECISION

Latency	218 μs	
SUPER occupation	10 μs	(100% occupation)
ROB occupation	60 μs / ROB	(600% occupation)
RSI PCI occupation	8 μs / RSI	(80% occupation)
RSI occupation	90 μs / RSI	(900% occupation)
ATM2 occupation	6 μs / RSI	(60% occupation)
SFI occupation	0	
SFI PCI occupation	0	
LVL2 processor occupation	0	

11) SEND DATA TO LVL3 PROCESSORS

Latency	330 μs + RSI data / 70 MB/s + RSI data / 50 MB/s + 50 μs * RSI + EVT data / 15 MB/s + EVT data / 50 MB/s + EVT data / 70 MB/s
SUPER occupation	0
ROB occupation	10 μs * ROB
RSI PCI occupation	EVT data / 70 MB/s
RSI occupation	50 μs * ROB + EVT data / 50 MB/s + 10 μs * RSI
ATM1 occupation	EVT data / 15 MB/s
SFI occupation	50 μs * RSI + EVT data / 50 MB/s + 10 μs
SFI PCI occupation	EVT data / 70 MB/s
LVL3 processor occupation	50 μs

12) EXECUTE LVL3 ALGORITHMS

LVL3	T_{alg}

13) SEND LVL3 DECISION TO SUPERVISOR

latency	178 μs
SUPER occupation	50 μs
ROB occupation	0
RSI PCI occupation	0
RSI occupation	0
ATM2 occupation	6 μs
SFI occupation	60 μs
SFI PCI occupation	2 μs
LVL2 processor occupation	10 μs

4.3 Detector Parameters for Processes at $L=10^{33}$

The latencies and occupations can be expressed more concisely if the variables are replaced by the detector parameters listed in Section 3. Here we treat the case of low luminosity, $L=10^{33}$ cm^{-2} s^{-1}. In the table below, we list the LVL2 parameters for the following processes (or algorithms):

MUON	=	stand-alone LVL2 muon trigger
e / γ / τ	=	electron / gamma / tau trigger in CALO system
JETs	=	jet trigger in CALO system
SCT	=	track trigger in SCT
TRT	=	track trigger in TRT
TRK	=	combined track trigger in SCT + TRT
MsE_T	=	missing-E_T recalculation at LVL2 (CALO system)
Scan	=	full TRT scan for B physics
b-tag	=	b-jet tag in SCT system

DETECTOR PARAMETERS FOR LVL2 PROCESSING AT L=10^{33} cm^{-2} s^{-1}

Detector parameters		MUON	e / γ / τ	JETs	SCT	TRT	TRK	MsE_T	Scan	b-tag
ROB/RoI		1	3	9	2	3	5	256	256	8
RSI/RoI		1	2	5	1	2	3	64	64	3
max(ROB/RSI)		1	3	4	2	3	3	4	4	4
ROB data	(kB)	1.0	1.0	0.3	0.1	0.03	0.06	0.1	0.1	0.4
RSI data	(kB)	1.0	1.5	0.6	0.2	0.05	0.10	0.3	0.3	1.0
RoI data	(kB)	1.0	3.0	3.0	0.2	0.10	0.30	16.0	16.0	3.0
PCI transfer at 70 MB/s										
ROB data/70	(μs)	14	14	4	1	1	1	1	1	6
RSI data/70	(μs)	14	21	8	2	1	1	4	4	14
RoI data/70	(μs)	14	42	42	2	1	4	229	229	42
Reorder data										
ROB data/50	(μs)	20	20	6	2	1	1	2	2	8
RSI data/50	(μs)	20	30	12	4	1	2	6	6	20
RoI data/50	(μs)	20	60	60	4	2	6	320	320	60
ATM transfer at 15 MB/s										
ROB data/15	(μs)	67	67	20	7	2	4	7	7	27
RSI data/15	(μs)	67	100	40	13	3	7	20	20	67
RoI data/15	(μs)	67	200	200	13	7	20	1 067	1 067	200
$max(T_{pre})$	(μs)	100	70	70	200	200	200	110	200	800
T_{fex}	(μs)	100	100	100	800	600	1 400	100	50000;250000	

For the LVL3 algorithms, the RoI is the full detector. The relevant detector parameters for low and high luminosity are given in the table below:

DETECTOR PARAMETERS FOR LVL3 PROCESSING			
		$L=10^{33}$	$L=10^{34}$
ROB		768	768
RSI		192	192
ROB/RSI		4	4
ROB data	(kB)	1.3	1.7
RSI data	(kB)	5.4	6.6
EVT data	(kB)	1 035	1 275
ROB data/70	(μs)	19	24
RSI data/70	(μs)	77	94
EVT data/70	(μs)	14 786	18 214
ROB data/50	(μs)	26	34
RSI data/50	(μs)	108	132
EVT data/50	(μs)	20 700	25 500
ROB data/15	(μs)	87	113
RSI data/15	(μs)	360	440
EVT data/15	(μs)	69 000	85 000
T_{pre}	(μs)	0	0
T_{alg}(min)	(μs)	10^5	10^5
T_{alg}(physics)	(μs)	10^7	10^7

4.4 Global Resources at $L=10^{33}$

Using the detector parameters listed above, we can calculate the resources required for any given trigger menu. The trigger menu must be decomposed into the trigger rate, the number of data requests for each subsystem and the number of RoIs of each type. Then the occupations can be derived from the equations shown below, where "RATE" is the relevant trigger rate, "REQ" is the number of data requests of a given type and "RoI" is the number of RoIs of that type. The total occupations are obtained by adding the total LVL2 occupations for the given trigger menu to the total LVL3 occupations. The split between LVL2 and LVL3 processing depends on the trigger strategy (option).

The total LVL2 occupations are obtained by adding the "minimal" LVL2 occupations to the contributions from each type of processing performed at LVL2 (RoIs and full TRT scan for both options, and missing-E_T and b-jet tags if all of the selection algorithms are executed at LVL2). The minimal LVL2 processing time includes the initialization steps 1) and 2) (sending the

LVL1 data to LVL2, and assigning a LVL2 processor), and the global decision steps 8) and 9). We assume that the final decisions are grouped together in the LVL2 Supervisor and broadcast to the ROBs once per millisec; this adds an average of 500 μs to the overall latency, but adds less than 10% to the ROB, RSI, and ATM occupancies. Only the 500 μs latency due to the broadcast of the final decisions has been added into the minimal LVL2 processing numbers listed below.

GLOBAL RESOURCES FOR LVL2 PROCESSING AT $L=10^{33}$

MINIMAL LVL2 PROCESSING (NO ROI DATA)					
Buffer occupation	RATE * [1	156 μs]			
ROB occupation		0			
RSI PCI occupation		0			
RSI occupation		0			
ATM1 occupation	RATE * [	6 μs]			
ATM2 occupation	RATE * [	6 μs]			
SFI occupation	RATE * [	120 μs]			
SFI PCI occupation	RATE * [	4 μs]			
LVL2 processor occupation	RATE * [	210 μs]			
MUON algorithms					
Buffer occupation	REQ * [	648 μs]	+	RoI * [	257 μs]
ROB occupation				RoI * [	160 μs]
RSI PCI occupation				RoI * [	16 μs]
RSI occupation				RoI * [	140 μs]
ATM1 occupation				RoI * [	67 μs]
ATM2 occupation				RoI * [	6 μs]
SFI occupation	REQ * [	70 μs]	+	RoI * [	70 μs]
SFI PCI occupation	REQ * [	2 μs]	+	RoI * [	14 μs]
LVL2 processor occupation	REQ * [	160 μs]	+	RoI * [	100 μs]
e / γ / τ algorithms					
Buffer occupation	REQ * [	759 μs]	+	RoI * [	514 μs]
ROB occupation				RoI * [	390 μs]
RSI PCI occupation				RoI * [	48 μs]
RSI occupation				RoI * [	360 μs]
ATM1 occupation				RoI * [	200 μs]
ATM2 occupation				RoI * [	12 μs]
SFI occupation	REQ * [	70 μs]	+	RoI * [	160 μs]
SFI PCI occupation	REQ * [	2 μs]	+	RoI * [	42 μs]
LVL2 processor occupation	REQ * [	160 μs]	+	RoI * [	100 μs]
JET algorithms					

Buffer occupation	REQ * [	790 μs]	+	RoI * [	682 μs]
ROB occupation				RoI * [	1 170 μs]
RSI PCI occupation				RoI * [	60 μs]
RSI occupation				RoI * [	900 μs]
ATM1 occupation				RoI * [	200 μs]
ATM2 occupation				RoI * [	30 μs]
SFI occupation	REQ * [	70 μs]	+	RoI * [	310 μs]
SFI PCI occupation	REQ * [	2 μs]	+	RoI * [	42 μs]
LVL2 processor occupation	REQ * [	160 μs]	+	RoI * [	100 μs]

SCT algorithms

Buffer occupation	REQ * [	782 μs]	+	RoI * [	875 μs]
ROB occupation				RoI * [	520 μs]
RSI PCI occupation				RoI * [	6 μs]
RSI occupation				RoI * [	184 μs]
ATM1 occupation				RoI * [	13 μs]
ATM2 occupation				RoI * [	6 μs]
SFI occupation	REQ * [	70 μs]	+	RoI * [	54 μs]
SFI PCI occupation	REQ * [	2 μs]	+	RoI * [	2 μs]
LVL2 processor occupation	REQ * [	160 μs]	+	RoI * [	800 μs]

TRT algorithms

Buffer occupation	REQ * [	840 μs]	+	RoI * [	722 μs]
ROB occupation				RoI * [	780 μs]
RSI PCI occupation				RoI * [	7 μs]
RSI occupation				RoI * [	302 μs]
ATM1 occupation				RoI * [	7 μs]
ATM2 occupation				RoI * [	12 μs]
SFI occupation	REQ * [	70 μs]	+	RoI * [	102 μs]
SFI PCI occupation	REQ * [	2 μs]	+	RoI * [	1 μs]
LVL2 processor occupation	REQ * [	160 μs]	+	RoI * [	600 μs]

TRK (TRT + SCT) algorithms

Buffer occupation	REQ * [	841 μs]	+	RoI * [	1 598 μs]
ROB occupation				RoI * [	1 300 μs]
RSI PCI occupation				RoI * [	14 μs]
RSI occupation				RoI * [	486 μs]
ATM1 occupation				RoI * [	20 μs]
ATM2 occupation				RoI * [	18 μs]
SFI occupation	REQ * [	70 μs]	+	RoI * [	156 μs]
SFI PCI occupation	REQ * [	2 μs]	+	RoI * [	4 μs]

LVL2 processor occupation	REQ * [	160 μs]	+ RoI * [	1 400 μs]

Missing-E_T recalculation at LVL2

Buffer occupation	REQ * [	6 120 μs]			
ROB occupation	REQ * [	43 520 μs]			
RSI PCI occupation	REQ * [	741 μs]			
RSI occupation	REQ * [	19 520 μs]			
ATM1 occupation	REQ * [	1 067 μs]			
ATM2 occupation	REQ * [	384 μs]			
SFI occupation	REQ * [	3 590 μs]			
SFI PCI occupation	REQ * [	231 μs]			
LVL2 processor occupation	REQ * [	260 μs]			

Full TRT scan

Buffer occupation	REQ * [	56 110 μs]			
ROB occupation	REQ * [	66 560 μs]			
RSI PCI occupation	REQ * [	741 μs]			
RSI occupation	REQ * [	19 520 μs]			
ATM1 occupation	REQ * [	1 067 μs]			
ATM2 occupation	REQ * [	384 μs]			
SFI occupation	REQ * [	3 590 μs]			
SFI PCI occupation	REQ * [	231 μs]			
LVL2 processor occupation	REQ * [	50 160 μs]			

b-jet tag calculation at LVL2

Buffer occupation	REQ * [	1 534 μs]	+ RoI * [	250 470 μs]
ROB occupation			RoI * [	6880 μs]
RSI PCI occupation			RoI * [	58 μs]
RSI occupation			RoI * [	720 μs]
ATM1 occupation			RoI * [	200 μs]
ATM2 occupation			RoI * [	18 μs]
SFI occupation	REQ * [	70 μs]	+ RoI * [	210 μs]
SFI PCI occupation	REQ * [	2 μs]	+ RoI * [	42 μs]
LVL2 processor occupation	REQ * [	160 μs]	+ RoI * [	250 000 μs]

The total LVL3 occupations are obtained in a similar way, by adding the “minimal” LVL3 occupations to the contributions from any special processing assigned to LVL3. In the standard option, the missing-E_T and the b-jet tag calculations are performed at LVL3. The missing E_T calculation can be absorbed in the “minimal” LVL3 occupations, since the missing-E_T algorithm is much faster than the minimal LVL3 processing assumed here.The b-jet tag, on the other hand, is slow, and must be taken into account explicitly.

The minimal LVL3 occupations are much higher than the minimal LVL2 occupations because we assume that all of the data are transfered to the LVL3 processors. We assume there is no preprocessing in the ROBs, and that the minimum processing time in the LVL3 ROBs is 100 ms. Finally, we assume that the LVL3 code is less optimal than the LVL2 code, so that algorithm execution times are longer by a factor two. The latencies and occupations (per event) for the minimal LVL3 processing and for the b-jet tag calculation are given below:

GLOBAL LATENCY AND GLOBAL RESOURCES FOR LVL3 PROCESSING AT $L=10^{33}$		
Minimal LVL3 processing per event		
Latency	214 779 μs	
ROB occupation	7 680 μs	
RSI PCI occupation	14 786 μs	
RSI occupation	61 020 μs	
ATM1 occupation	69 000 μs	
ATM2 occupation	6 μs	
SFI occupation	30 370 μs	
SFI PCI occupation	14 788 μs	
LVL3 processor occupation	100 060 μs	
b-jet tag calculation at LVL3		
Latency		RoI * [500 000 μs]
LVL3 processor occupation		RoI * [500 000 μs]

5. Sequential Processing

The model for sequential processing is a single LVL2 processing farm performing the full event selection, reducing the event rate from the initial 100 kHz LVL1 rate to the final 100 Hz LVL3 output rate. The supervisor assigns each event to a single LVL2 processor, which requests data from the ROBs as required for the event processing. Events can be rejected by the LVL2 processor at any time, if it is decided that the event does not satisfy the trigger selection criteria.

The processing sequence consists in the following steps:

1) confirm the LVL1 trigger using calorimeter and muon data from the trigger RoIs

2) verify muon, electron, and hadron triggers by matching the features in the calorimeter and muon systems to track parameters measured in the inner detector

3) perform the full TRT scan for events with a confirmed muon trigger

4) verify trigger muon isolation criteria in the calorimeter data

5) analyze non-trigger RoIs flagged by the LVL1.5 trigger menu [2]

6) recalculate the missing-E_T if required by the trigger menu

7) look for b-jet tags if required by the trigger menu

8) verify the global selection criteria by combining features from all RoIs satisfying the above selection criteria.

Events can be rejected at each step in the processing sequence.

5.1 Notation

The notation used for the trigger conditions is shown in the following table:

TRIGGER ALGORITHMS

Symbol	Subsystem	Description
mu	MUON	stand-alone muon algorithm
trk	TRT + SCT	track algorithm for RoI
calo	CALO	muon isolation algorithm
scan	TRT	full TRT scan
em	CALO	e / γ / τ algorithms
tau	CALO	τ / hadron algorithms
jet	CALO	jet algorithm
b-jet	CALO + SCT	jet + b-jet tag algorithm
me	CALO	missing-E_T calculation

5.2 Model for Sequential Processing

Trigger menus for low luminosity have been published in a preceding ATLAS note [1]. This trigger menu has been used to construct a model for the sequential processing described above. Rejection rates for each trigger algorithm have been estimated as best we can; a few details are given in ref.[1]. Each line in the table below corresponds to a different sequence of trigger algorithms (involving various threshold and isolation criteria) , some of which may be performed in the LVL3 processor farm.

The notation has been described in the previous section. Two columns of event rates are given: the first column is the rate for all events passing that set of algorithms, and the last column is the exclusive rate for that algorithm. The sum of the event rates in the last column should add up to the full LVL1 trigger rate.

SEQUENTIAL PROCESSING MODEL FOR $L=10^{33}$

Trigger algorithm sets	Inclusive rate	Exclusive rate
1 mu	8 000 Hz	601 Hz
1 mu + 1 trk	6 000 Hz	1 873 Hz
1 mu + 1 trk + 1 calo + scan	4 000 Hz	3 419 Hz
1 mu + 1 trk + scan + me	25 Hz	25 Hz
1 mu + 1 trk + 6 jet + scan	38 Hz	32 Hz
1 mu + 1 trk + 5 jet + scan	50 Hz	5 Hz
1 mu + 1 trk + 4 jet + scan	75 Hz	13 Hz
1 mu + 1 trk + 3 jet + scan	100 Hz	0 Hz
1 mu + 1 trk + 6 b-jet + scan	6 Hz	6 Hz
1 mu + 1 trk + 5 b-jet + scan	13 Hz	7 Hz
1 mu + 1 trk + 4 b-jet + scan	25 Hz	12 Hz
1 mu + 1 trk + 3 b-jet + scan	50 Hz	25 Hz
1 mu + 1 trk + 1 calo + 6 jet + scan	75 Hz	62 Hz
1 mu + 1 trk + 1 calo + 5 jet + scan	100 Hz	13 Hz
1 mu + 1 trk + 1 calo + 4 jet + scan	150 Hz	25 Hz
1 mu + 1 trk + 1 calo + 3 jet + scan	200 Hz	0 Hz
1 mu + 1 trk + 1 calo + 2 jet + scan	302 Hz	2 Hz
1 mu + 1 trk + 1 calo + 1 jet + scan	20 Hz	5 Hz
1 mu + 1 trk + 1 calo + 6 b-jet + scan	13 Hz	13 Hz
1 mu + 1 trk + 1 calo + 5 b-jet + scan	25 Hz	12 Hz
1 mu + 1 trk + 1 calo + 4 b-jet + scan	50 Hz	25 Hz
1 mu + 1 trk + 1 calo + 3 b-jet + scan	100 Hz	50 Hz
1 mu + 1 trk + 1 calo + 2 b-jet + scan	200 Hz	100 Hz
1 mu + 1 trk + 1 calo + 1 b-jet + scan	15 Hz	15 Hz
1 mu + 1 trk + 1 calo + scan + me	20 Hz	20 Hz
2 mu	400 Hz	155 Hz
2 mu + 2 trk	200 Hz	90 Hz
2 mu + 2 trk + 2 calo + scan	100 Hz	95 Hz
3 mu	20 Hz	13 Hz
3 mu + 3 trk	7 Hz	4.5 Hz
3 mu + 3 trk + 3 calo + scan	2.5 Hz	2.5 Hz
1 mu + 1 em	800 Hz	770 Hz
1 mu + 1 em + 2 trk	30 Hz	20 Hz
1 mu + 1 em + 2 trk + 1 calo + scan	10 Hz	10 Hz
2 mu + 1 em	40 Hz	32 Hz
2 mu + 1 em + 3 trk	8 Hz	6 Hz
2 mu + 1 em + 3 trk + 2 calo + scan	2 Hz	2 Hz
1 mu + 2 em	40 Hz	35 Hz
1 mu + 2 em + 3 trk	5 Hz	4 Hz
1 mu + 2 em + 3 trk + 1 calo + scan	1 Hz	1 Hz
1 mu + 1 tau	400 Hz	394 Hz
1 mu + 1 tau + 2 trk	6 Hz	2 Hz
1 mu + 1 tau + 2 trk + 1 calo + scan	4 Hz	4 Hz

1 em	10 200 Hz	6 133 Hz
1 em + 1 trk	1060 Hz	757 Hz
1 em + 1 trk + me	280 Hz	80 Hz
1 em + 1 trk + 6 jet + me	36 Hz	30 Hz
1 em + 1 trk + 5 jet + me	50 Hz	7 Hz
1 em + 1 trk + 4 jet + me	70 Hz	8 Hz
1 em + 1 trk + 3 jet + me	100 Hz	5 Hz
1 em + 1 trk + 2 jet + me	140 Hz	0 Hz
1 em + 1 trk + 1 jet + me	200 Hz	0 Hz
1 em + 1 trk + 6 b-jet + me	6 Hz	6 Hz
1 em + 1 trk + 5 b-jet + me	13 Hz	7 Hz
1 em + 1 trk + 4 b-jet + me	25 Hz	12 Hz
1 em + 1 trk + 3 b-jet + me	50 Hz	25 Hz
1 em + 1 trk + 2 b-jet + me	100 Hz	50 Hz
1 em + 1 trk + 1 b-jet + me	200 Hz	50 Hz
1 em + 1 tau	3 000 Hz	2 980 Hz
1 em + 1 tau + 2 trk	20 Hz	20 Hz
2 em	2 530 Hz	1 412 Hz
2 em + 2 trk	103 Hz	78 Hz
3 em	1 000 Hz	980 Hz
3 em + 3 trk	20 Hz	20 Hz
1 tau	5 000 Hz	3 960 Hz
1 tau + me	200 Hz	200 Hz
1 tau + 1 trk	40 Hz	40 Hz
2 tau	800 Hz	800 Hz
1 jet	8 000 Hz	5 100 Hz
2 jet	2 000 Hz	1 643 Hz
3 jet	200 Hz	157 Hz
4 jet	1 600 Hz	390 Hz
5 jet	800 Hz	145 Hz
6 jet	520 Hz	365 Hz
1 B-tag + 3 jet	500 Hz	202 Hz
1 B-tag + 4 jet	250 Hz	81 Hz
1 B-tag + 5 jet	150 Hz	120 Hz
2 B-tag + 2 jet	50 Hz	16 Hz
2 B-tag + 3 jet	25 Hz	7 Hz
2 B-tag + 4 jet	15 Hz	10 Hz
3 B-tag + 1 jet	10 Hz	4 Hz
3 B-tag + 2 jet	5 Hz	1 Hz
3 B-tag + 3 jet	3 Hz	2 Hz
4 B-tag	57 Hz	28 Hz
5 B-tag	29 Hz	11 Hz
6 B-tag	18 Hz	18 Hz
2 jet + me	80 Hz	72 Hz
3 jet + me	8 Hz	8 Hz
4 jet + me	200 Hz	160 Hz
5 B-tag + me	40 Hz	35 Hz
6 B-tag + me	5 Hz	5 Hz

me	30 Hz	30 Hz
TOTAL LVL1 trigger rate		34 270 Hz

5.3 Summary of Rates

The numbers needed to evaluate the latencies and occupations parameterized in section 4.4 of this note are the total rates for each of the trigger algorithms. These numbers are given in the table below.

The full TRT scan is required for the B-physics triggers. Only events with a confirmed muon trigger (confirmed in the stand-alone LVL2 muon system and in the inner tracking) are sent to the full TRT scan. Only about half of the events with a full TRT scan are still compatible with B-physics trigger conditions. New RoIs are defined for the tracks found in the TRT scan; on average, seven new RoIs are found for each event retained after the TRT scan. In this model we assume that only SCT algorithms and global algorithms are performed on these new RoIs, although it would be more correct to request calorimeter and muon data for some of them. In the tables below, the expression <SCT “scan"> refers to the execution of standard SCT algorithms on the RoIs defined in the TRT scan.

SUMMARY FOR COMBINED LVL2/LVL3 PROCESSING

Trigger	System	Rate	Requests	RoIs
LVL1	-	34 270 Hz		
muon	muon	8 000 Hz	8 000 Hz	8 420 Hz
	track	6 000 Hz	6 000 Hz	6 256 Hz
	TRT scan	4 000 Hz	4 000 Hz	4 000 Hz
	SCT “scan”(*)	2 000 Hz	2 000 Hz	14 000 Hz
	EM calo	4 000 Hz	5 240 Hz	5 342 Hz
	jets	422 Hz	422 Hz	1 612 Hz
	b-jet tag	265 Hz	265 Hz	797 Hz
	missing-E_T	45 Hz	45 Hz	45 Hz
e / γ / τ	EM calo	17 700 Hz	17 700 Hz	25 030 Hz
	track	1 200 Hz	1 200 Hz	1 343 Hz
	jets	200 Hz	200 Hz	596 Hz
	b-jet tag	200 Hz	200 Hz	394 Hz
	missing-E_T	480 Hz	480 Hz	480 Hz

jet	jets	8 580 Hz	9 580 Hz	16 860 Hz
	b-jet tag	540 Hz	597 Hz	970 Hz
	missing-E_T	280 Hz	280 Hz	280 Hz
missing-E_T	missing-E_T	30 Hz	30 Hz	30 Hz
TOTAL	muon	8 000 Hz	8 000 Hz	8 420 Hz
	EM calo	21 700 Hz	22 940 Hz	30 372 Hz
	track	7 200 Hz	7 200 Hz	7 599 Hz
	TRT scan	4 000 Hz	4 000 Hz	4 000 Hz
	SCT “scan”(*)	2 000 Hz	2 000 Hz	14 000 Hz
	jets	9 202 Hz	10 202 Hz	19 068 Hz
	b-jet tag	1 005 Hz	1 062 Hz	2 161 Hz
	missing-E_T	835 Hz	835 Hz	835 Hz

(*) For the meaning of SCT “scan”, see text above.

5.4 Latency and Occupation at LVL2

The occupation of the trigger hardware can be calculated for each LVL2 or LVL3 process by plugging the numbers in Section 5.3 of this note into the equations in Section 4.4. The result for the LVL2 processes are shown in the tables below.

The (average) buffer occupations are obtained by multiplying the latencies by the event rates. Therefore, the average latencies are obtained merely by dividing the buffer occupation by the relevant trigger rate.

For the "hardware elements", the "number required" in the first column of numbers gives the required number of units for that hardware item (e.g., number of processors occupied at 100%). The second column gives the number of units in our model (see Section 3.1). The last column gives the occupation for those hardware items in percent.

We have assumed that we have a total of 192 RSIs, 192 SFI's, and 650 LVL2 processors. Thus we assume a 192x192 dual-link ATM switch. The physical ATM switch is assumed to have 512 full-duplex ports @ 155 Mbit/s. The 128 remaining ports are reserved for the LVL2 supervisor, the LVL1 and LVL2 trigger ROBs, and other units.

MINIMAL LVL2 PROCESSING (NO ROI DATA)

LVL1 trigger rate 34 270 Hz
LVL2 minimum latency 1.2 ms
Buffer length required 39.6 events

Hardware element	Number required	Number in model	Occupation
ROB	0	768	0
RSI PCI	0	192	0
RSI	0	192	0
ATM1	0.2	192	0.1 %
ATM2	0.2	192	0.1 %
SFI	4.1	192	2.1 %
SFI PCI	0.1	192	0.1 %
LVL2 processor	7.2	650	1.1 %

MUON stand alone algorithms

LVL1 muon trigger rate 8 000 Hz REQs 8 000 Hz RoIs 8 420 Hz
LVL2 latency 0.9 ms
Buffer length required 7.4 events

Hardware element	Number required	Number in model	Occupation
ROB	1.4	128	1.1 %
RSI PCI	0.1	32	0.4 %
RSI	1.2	32	3.7 %
ATM1	0.6	32	1.8 %
ATM2	0.1	32	0.2 %
SFI	1.2	192	0.6 %
SFI PCI	0.1	192	0.1 %
LVL2 processor	2.1	650	0.3 %

e / γ / τ calorimetry algorithms

LVL1 e / γ / τ rate 21 700 Hz REQs 22 940 Hz RoIs 30 372 Hz
LVL2 latency 1.5 ms
Buffer length required 33.0 events

Hardware element	Number required	Number in model	Occupation
ROB	11.8	256	4.6 %
RSI PCI	1.5	64	2.2 %
RSI	10.9	64	17.1 %
ATM1	6.1	64	9.5 %
ATM2	0.4	64	0.6 %
SFI	6.5	192	3.4 %
SFI PCI	1.3	192	0.7 %
LVL2 processor	6.7	650	1.0 %

JET calorimetry algorithms

LVL1 jet trigger rate 9 202 Hz REQs 10 202 Hz RoIs 19 068 Hz
LVL2 latency 2.3 ms
Buffer length required 21.1 events

Hardware element	Number required	Number in model	Occupation
ROB	22.3	256	8.7 %
RSI PCI	1.1	64	1.8 %
RSI	17.2	64	26.8 %
ATM1	3.8	64	6.0 %
ATM2	0.6	64	0.9 %
SFI	6.6	192	3.4 %
SFI PCI	0.8	192	0.4 %
LVL2 processor	3.5	650	0.5 %

TRT tracking algorithms

LVL2 tracking algorithms 7 200 Hz REQs 7 200 Hz RoIs 7 599 Hz
plus full TRT scan 4 000 Hz
LVL2 latency 1.6 ms for RoIs, 56.1 ms for full scan
Buffer length required 236.0 events

Hardware element	Number required	Number in model	Occupation
ROB	272.2	256	106.3 %
RSI PCI	3.0	64	4.7 %
RSI	80.4	64	125.6 %
ATM1	4.3	64	6.8 %
ATM2	1.6	64	2.5 %
SFI	15.6	192	8.1 %
SFI PCI	0.9	192	0.5 %
LVL2 processor	206.4	650	31.7 %

SCT tracking algorithms

LVL2 tracking algorithms 7 200 Hz REQs 7 200 Hz RoIs 7 599 Hz
plus full SCT scan 2 000 Hz REQs 2 000 Hz RoIs 14 000 Hz
LVL2 latency 1.7 ms for RoIs, 6.9 ms for full scan
Buffer length required 26.1 events

Hardware element	Number required	Number in model	Occupation
ROB	11.2	128	8.8 %
RSI PCI	0.1	32	0.4 %
RSI	4.0	32	12.4 %
ATM1	0.3	32	0.9 %
ATM2	0.1	32	0.4 %
SFI	1.8	192	0.9 %
SFI PCI	0.2	192	0.1 %
LVL2 processor	18.8	650	2.9 %

b-jet TAG USING SCT DATA at LVL2

LVL2 b-jet tag rate 1 005 Hz REQs 1 062 Hz RoIs 2 161 Hz
LVL2 latency 540.2 ms
Buffer length required 542.9 events

Hardware element	Number required	Number in model	Occupation
ROB	14.9	128	11.6 %
RSI PCI	0.1	32	0.4 %
RSI	1.6	32	4.9 %
ATM1	0.4	32	1.4 %
ATM2	0.04	32	0.1 %
SFI	0.5	192	0.3 %
SFI PCI	0.1	192	0.1 %
LVL2 processor	540.4	650	83.1 %

MISSING-E_T CALCULATION at LVL2

LVL2 missing-E_T rate 835 Hz
LVL2 latency 6.1 ms
Buffer length required 5.1 events

Hardware element	Number required	Number in model	Occupation
ROB	36.3	256	14.2 %
RSI PCI	0.6	64	1.0 %
RSI	16.3	64	25.5 %
ATM1	0.9	64	1.4 %
ATM2	0.3	64	0.5 %
SFI	3.0	192	1.6 %
SFI PCI	0.2	192	0.1 %
LVL2 processor	0.2	650	0.03%

The occupations due to the individual LVL2 algorithms shown above are far from saturation for most of the hardware elements. The exceptions are the TRT ROBs and RSIs and the LVL2 processors. The calorimeter RSIs approach saturation when the contributions from the individual algorithms are summed.

The TRT ROBs and RSIs are saturated in our model due to the 4 kHz rate of the full TRT scan needed for the B-physics algorithms. The saturation is slightly above 100% (125% for the RSIs). We take this as a warning, rather than an alarm, because the parameters of our model (and, in fact, the model itself) are not yet well-understood. On the other hand, our model indicates that it would be extremely difficult to execute the same B-physics algorithms at a luminosity three times higher than the luminosity considered here (i.e., at the intermediate luminosity, $3\ 10^{33}\ cm^{-2}\ s^{-1}$).

The LVL2 processors are near saturation (above 80%) for the b-jet algorithm, if this algorithm is performed in the limited LVL2 processor farm (650 processors). Since this algorithm is normally considered to be a LVL3 algorithm, the difficulty could be alleviated by

transferring some of the LVL3 processors to the LVL2 farm.

The calorimeter RSIs are near saturation (70%) if the occupation due to the missing-E_T calculations is added to those due to e / γ / τ events and the jets. (In the standard option, in which the missing-E_T calculation is performed at LVL3, the total RSI occupation is even higher.) Again, we take this as a warning, rather than an alarm. We need more work on the RSI functions, and we need to understand the probable evolution of the hardware and software performance better than we do today. For now, we consider that the missing-E_T calculations can indeed be performed at LVL2 at the rates obtained from our trigger menu. The warning is none-the-less serious, because we might easily be tempted to increase our jet trigger rates and our missing-E_T rates to increase our acceptance for SUSY events.

5.5 Latency and Occupation at LVL3

The latency and occupation due to trigger algorithms executed at LVL3 (except for B-physics) are shown in the following table. We assume that LVL3 shares the same 192 RSIs and the same 192 SFIs as LVL2, but we assume LVL3 has its own farm of 2 000 LVL3 processors.

MINIMAL LVL3 PROCESSING PER EVENT (TRANSFER ALL DATA TO LVL3)

LVL2 trigger rate	1 627 Hz	(except B-physics)
LVL3 minimum latency	214.8 ms	
Buffer length required	349.5 events	

Hardware element	Number required	Number in model	Occupation
ROB	12.5	768	1.6 %
RSI PCI	24.1	192	12.5 %
RSI	99.3	192	51.7 %
ATM1	112.3	192	58.4 %
ATM2	0.01	192	0.0 %
SFI	49.4	192	25.7 %
SFI PCI	24.1	192	12.5 %
LVL3 processor	162.8	2 000	8.1 %

b-jet TAG USING SCT DATA at LVL3

LVL3 b-jet tag rate	1 005 Hz	RoIs	2 161 Hz
LVL3 latency	1 075.1 ms		
Buffer length required	1 080.5 events		

Hardware element	Number required	Number in model	Occupation
LVL3 processor	1 080.5	2 000	54.0 %

We note that the average occupation of the RSIs and ATM1 are rather high (above 50%). The occupation of the LVL3 processors is also above 50% because of the calculation of the b-jet tags. We will discuss these items further in the next section.

5.6 Sequential Processing Options

The first of our two processing options splits the algorithms into LVL2 algorithms and LVL3 algorithms. In the model presented here, the b-jet tags and the missing-E_T are calculated at LVL3. The resources required at LVL2 and at LVL3 and the sum of the resources required in LVL2 and LVL3 are shown in the tables below.

TOTAL OCCUPATION FOR LVL2 PROCESSING (without b-jet tag and missing-E_T)

LVL1 trigger rate	34 270 Hz
LVL2 average latency	10.6 ms
Buffer length required	363.2 events

Hardware element	Number required	Number in model	Occupation
ROB	318.9	768	41.5 %
RSI PCI	5.8	192	3.0 %
RSI	113.7	192	59.2 %
ATM1	15.3	192	8.0 %
ATM2	3.0	192	1.6 %
SFI	35.8	192	18.6 %
SFI PCI	3.4	192	1.8 %
LVL2 processor	244.7	650	37.6 %

TOTAL OCCUPATION FOR LVL3 PROCESSING with (b-jet tag and missing-*E*)

LVL2 trigger rate	1 627 Hz	(except B-physics)
LVL3 latency	878.9 ms	
Buffer length required	1 430.0 events	

Hardware element	Number required	Number in model	Occupation
ROB	12.5	768	1.6 %
RSI PCI	24.1	192	12.5 %
RSI	99.3	192	51.7 %
ATM1	112.3	192	58.4 %
ATM2	0.01	192	0.0 %
SFI	49.4	192	25.7 %
SFI PCI	24.1	192	12.5 %
LVL3 processor	1 243.3	2 000	62.2 %

TOTAL OCCUPATION FOR LVL2 PLUS LVL3 PROCESSING

LVL1 trigger rate	34 270 Hz	
LVL2 + LVL3 latency	52.3 ms	(average for all LVL1 events)
Buffer length required	1 793.2 events	(LVL2 + LVL3 buffers)

Hardware element	Number required	Number in model	Occupation
ROB	331.4	768	43.2 %
RSI PCI	29.9	192	15.6 %
RSI	213.0	192	110.9 %
ATM1	127.6	192	66.5 %
ATM2	3.0	192	1.6 %
SFI	121.0	192	63.0 %
SFI PCI	27.5	192	14.3 %
LVL2+LVL3 processor	1 488.0	2 650	56.2 %

The total number of processors required for the trigger selection algorithms with this processing option is 1 488, more than half of the total number of processors (LVL2 plus LVL3) in our model. The switching network is heavily loaded (66% of the maximum bandwidth for data transfers from the ROBs to the processors), and the RSIs are saturated (average occupation 110%).

In our second processing option, all of the algorithms, including the b-jet tags and the missing-E_T calculations, are executed at LVL2. The resources required in this case are given in the following tables:

TOTAL OCCUPATION FOR LVL2 PROCESSING (including b-jet tag and missing-E_T)

LVL1 trigger rate	34 270 Hz
LVL2 average latency	26.6 ms
Buffer length required	911.2 events

Hardware element	Number required	Number in model	Occupation
ROB	370.1	768	48.2 %
RSI PCI	6.5	192	3.4 %
RSI	131.6	192	68.5 %
ATM1	16.6	192	8.6 %
ATM2	3.3	192	1.7 %
SFI	39.3	192	20.5 %
SFI PCI	3.7	192	1.9 %
LVL2 processor	785.3	650	120.8 %

MINIMAL LVL3 PROCESSING (all selection at LVL2)

LVL2 trigger rate	107 Hz	(except B-physics)
LVL3 latency	214.8 ms	
Buffer length required	23.0 events	

Hardware element	Number required	Number in model	Occupation
ROB	0.8	768	0.1 %
RSI PCI	1.6	192	0.8 %
RSI	6.5	192	3.4 %
ATM1	7.4	192	3.8 %
ATM2	0.00	192	0.0 %
SFI	3.2	192	1.7 %
SFI PCI	1.6	192	0.8 %
LVL3	10.7	2 000	0.5 %

TOTAL OCCUPATION FOR LVL2 PLUS LVL3 PROCESSING (all selection at LVL2)

LVL1 trigger rate	34 270 Hz	
LVL2 + LVL3 latency	27.3 ms	(average for all LVL1 events)
Buffer length required	934.2 events	(LVL2 + LVL3 buffers)

Hardware element	Number required	Number in model	Occupation
ROB	370.9	768	48.3 %
RSI PCI	8.1	192	4.2 %
RSI	138.1	192	71.9 %
ATM1	24.0	192	12.5 %
ATM2	3.3	192	1.7 %
SFI	42.5	192	22.1 %
SFI PCI	5.3	192	2.8 %
LVL2 + LVL3 processors	796.0	2 650	30.0 %

The total number of processors required for this option is only 796, nearly 700 less than those required in the standard processing option, where the processing is split between LVL2 and LVL3. Furthermore, the load on the switching network is reduced by a factor five. The occupation of the RSIs is still high (above 70%), but significantly lower than the saturation level observed in the standard option (110%).

5.7 Sequential vs. Parallel Processing

The advantage of sequential processing of the trigger algorithms can be seen in the two tables below, which summarize the number of events, the number of data requests, and the number of RoIs that must be treated in each of the detector subsystems.

SUMMARY FOR SEQUENTIAL PROCESSING

Algorithm	Subsystem	Rate	REQs	RoIs
muon	MUON	8 000 Hz	8 000 Hz	8 420 Hz
EM calo	CALO	21 700 Hz	22 940 Hz	30 372 Hz
track	TRT + SCT	7 200 Hz	7 200 Hz	7 599 Hz
jets	CALO	9 202 Hz	10 202 Hz	19 068 Hz
b-jet tag	SCT	1 005 Hz	1 062 Hz	2 161 Hz
missing-E_T	CALO	835 Hz	835 Hz	835 Hz

SUMMARY FOR PARALLEL PROCESSING

Algorithm	Subsystem	Rate	REQs	RoIs
muon	MUON	8 000 Hz	8 000 Hz	8 420 Hz
EM calo	CALO	25 700 Hz	25 700 Hz	34 690 Hz
track	TRT + SCT	25 700 Hz	25 700 Hz	34 690 Hz
jets	CALO	24 980 Hz	24 980 Hz	82 420 Hz
b-jet tag	SCT	20 500 Hz	20 500 Hz	73 820 Hz
missing-E_T	CALO	14 030 Hz	14 030 Hz	14 030 Hz

The full TRT scan has not been included in the summaries above because the B-physics algorithms used in our rate calculations require new RoIs to be formed after the full TRT scan. This is only possible using sequential processing. At least three alternatives are possible in a parallel processing architecture:

- Parallel processing could be maintained by performing the full TRT scan on the full 8 kHz of LVL1 muon triggers, and executing the SCT algorithms at LVL3.

- The B-physics events could be considered as an exception to the parallel processing scheme, and the sequential algorithms could be executed at LVL2.

- A hybrid solution could be used, in which special data-driven processors perform the TRT full scan and determine the coordinates of the new RoIs. Note that special data-driven processors would be needed in our model if we want to extend the B-physics algorithms to intermediate luminosities ($3\ 10^{33}$ cm^{-2} s^{-1}) without tightening the selection requirements.

The number of muon and EM calorimeter RoIs that have to be processed is reduced only slightly by using sequential rather than parallel processing. On the other hand, the number of tracking RoIs and the number of jet RoIs are reduced by a factor 4 when sequential processing is used. Note that the advantage of sequential processing is obtained only if the LVL2 processors operate in a multi-task mode, so that they can switch to another event when they post a request for additional data for the event in hand. This multi-task operation has been modelled successfully at Saclay, and the overhead has been measured and included in the model described here.

Missing-E_T and b-jet tags can be calculated at LVL2 using sequential processing, but this is not possible using parallel processing. This processing is only possible with sequential processing because the number of candidate events can be reduced by large factors: the number of missing-E_T events to be processed at LVL2 is reduced by a factor 16 in our model using sequential processing, and the number of b-jet tags is reduced by a factor 34.

6. Conclusions

One of the purposes of this study was to establish a method for comparing different LVL2 and LVL3 architectures without the machinery required for full modellization. These "paper models" can be used for initial evaluations and preliminary optimization of proposed trigger architectures. Meaningful comparisons will require the use of similar models and comparable sets of detector parameters for the different architectures. Final evaluation and final optimization will require the full modelling studies with fully-simulated data.

This study has already helped us to locate some of the problem areas in the single-farm architecture, and to begin the optimization procedure. More work is needed to refine the model and to determine the parameters of the final ATLAS system.

The main purpose of this study was to investigate different processing options for the single-farm LVL2 trigger architecture. The striking conclusion is that the trigger selection should not be split between LVL2 and LVL3. The number of processors and the network bandwidth required are reduced if all the selection algorithms are performed at LVL2. The number of processors is reduced because LVL2 uses preprocessed data and fully optimized code. The bandwidth is reduced because only a small fraction of the data is required to complete the selection algorithms at LVL2, whereas the full event must be sent to LVL3.

This study has also allowed us to compare sequential and parallel processing schemes, even though we do not have a complete model for the parallel processing scheme. It is well known that a parallel architecture does not provide a scheme for processing our B-physics algorithms. Furthermore, the number of tracking RoIs and the number of jet RoIs that must be processed is much higher (by a factor 4) in the parallel scheme than in the sequential scheme, because secondary RoIs must be processed, in addition to trigger RoIs, for all events. Finally, b-jet tags and missing-E_T must be calculated at LVL3 in the parallel scheme, whereas resources can be saved by performing these operations at LVL2 in the sequential scheme. There seem to be clear advantages to using sequential processing rather than parallel processing for the LVL2 trigger.

These studies will be extended to intermediate and high luminosities as soon as the corresponding trigger menus are available.

References

[1] J. Bystricky, J. Ernwein, T. Hansl-Kozanecka, J.R. Hubbard, P. Le Dû, and M. Smizanska, Trigger Menus for the ATLAS Trigger at Luminosity 10^{33} cm^{-2} s^{-1}, ATLAS Internal Note DAQ-N0-054.

[2] ATLAS Technical Proposal, CERN/LHCC/94-43, 15 December 1994.

[3] Trigger & DAQ Interfaces with Front-End Systems: Requirement Document, ATLAS Internal Note DAQ-NO-049.

[4] R. Hauser and I. Legrand, Algorithms in second-level triggers for ATLAS and benchmark results, ATLAS Internal Note DAQ-NO-027.

[5] S. Falciano, A. Nisati, O Palamara, S. Petrera, and L. Zanello, The ATLAS Muon Trigger Algoritm for Level-2 Feature Extraction, ATLAS Note DAQ in preparation.

[6] S. Sivoklokov, R. Dankers, and J. Baines, Second Level Triggering in the Forward Region of the ATLAS Inner Tracker, ATLAS Internal Note INDET-NO-111.

[7] M. Smizanska, Second Level TRT Trigger for B-Physics, ATLAS Internal Note PHYS-NO-089.

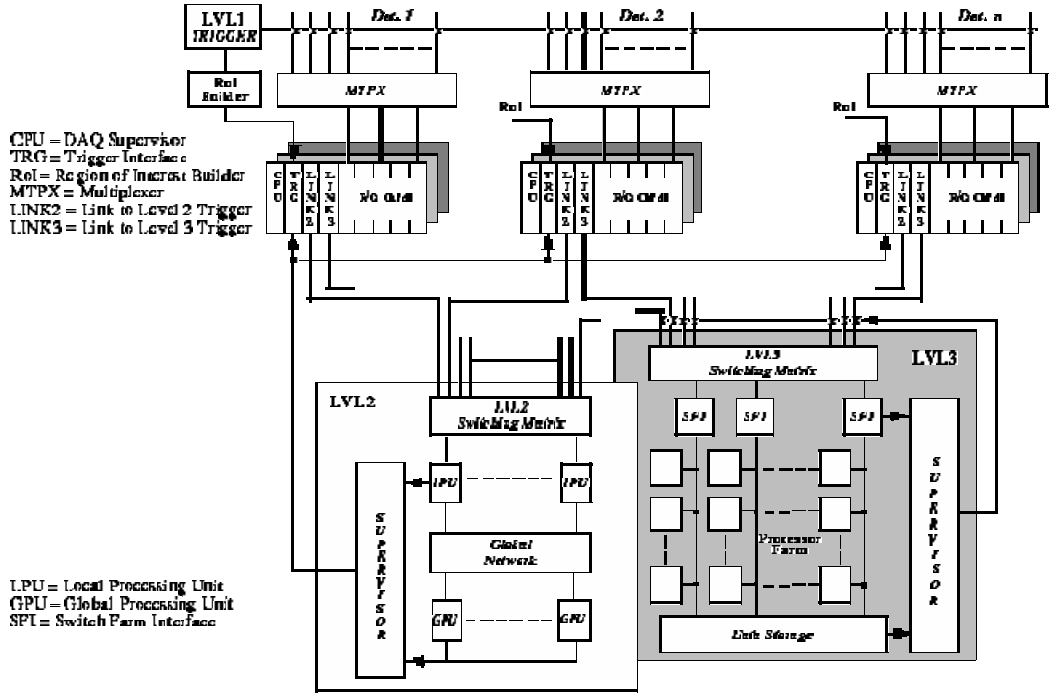

Figure 1. Overall trigger/DAQ architecture (from the ATLAS Technical Proposal, p. 150).

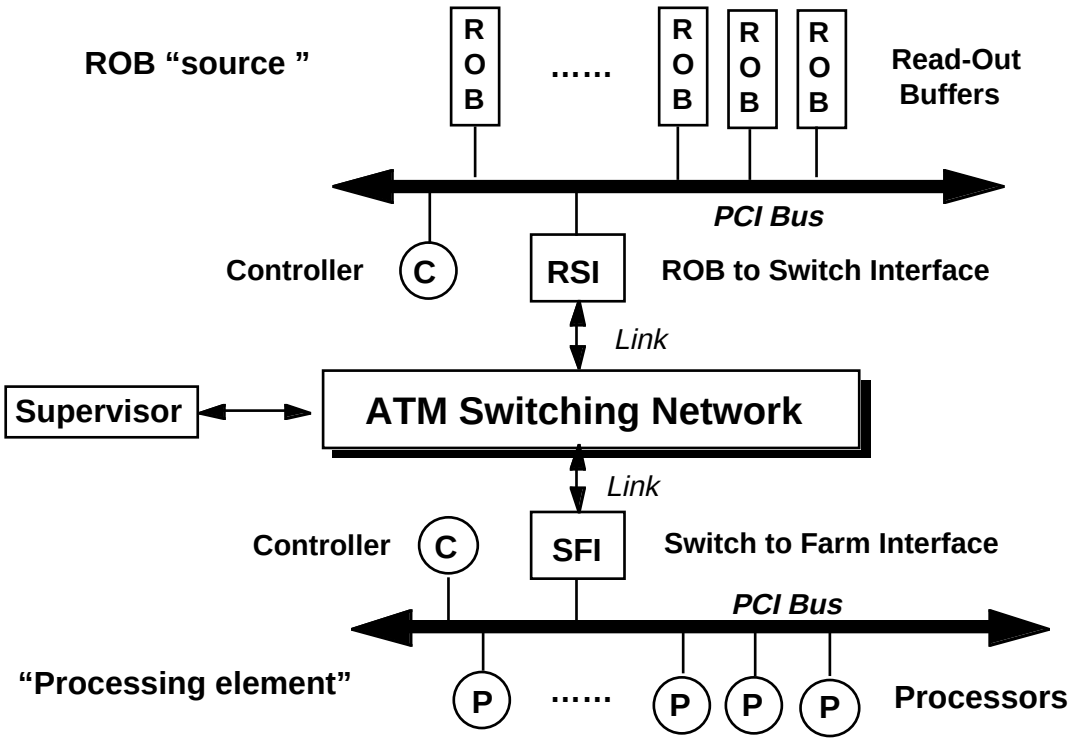

Figure 2. Functional Model for a Single Farm Architecture.

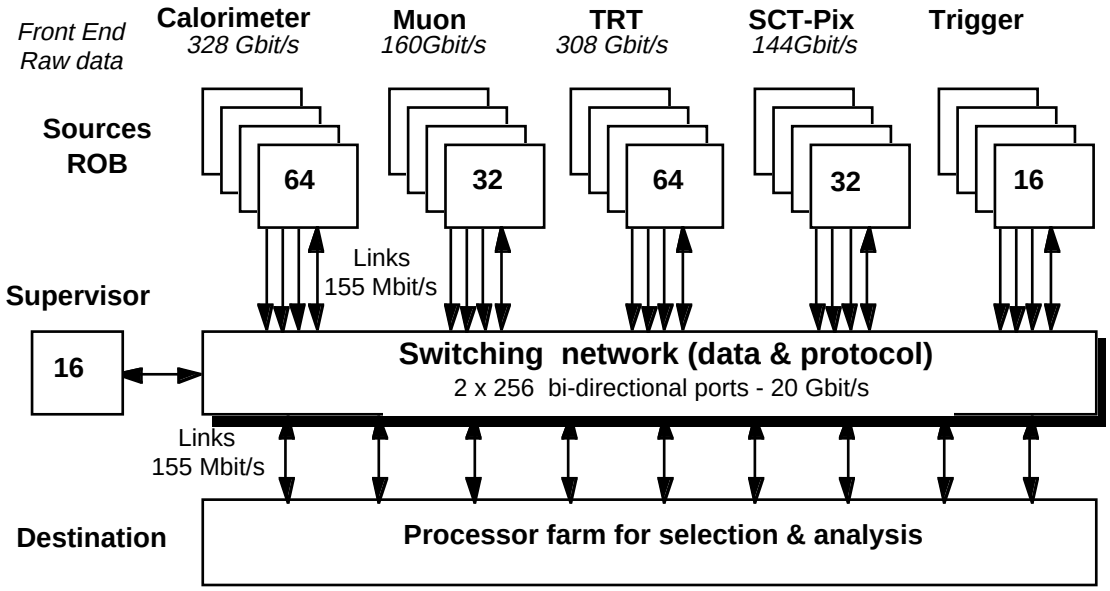

Figure 3. Single Farm Architecture