Overview of parameters used for 2001 paper model and of modifications in 2002 model

J. Vermeulen, 8 June 2002

1 System Architecture

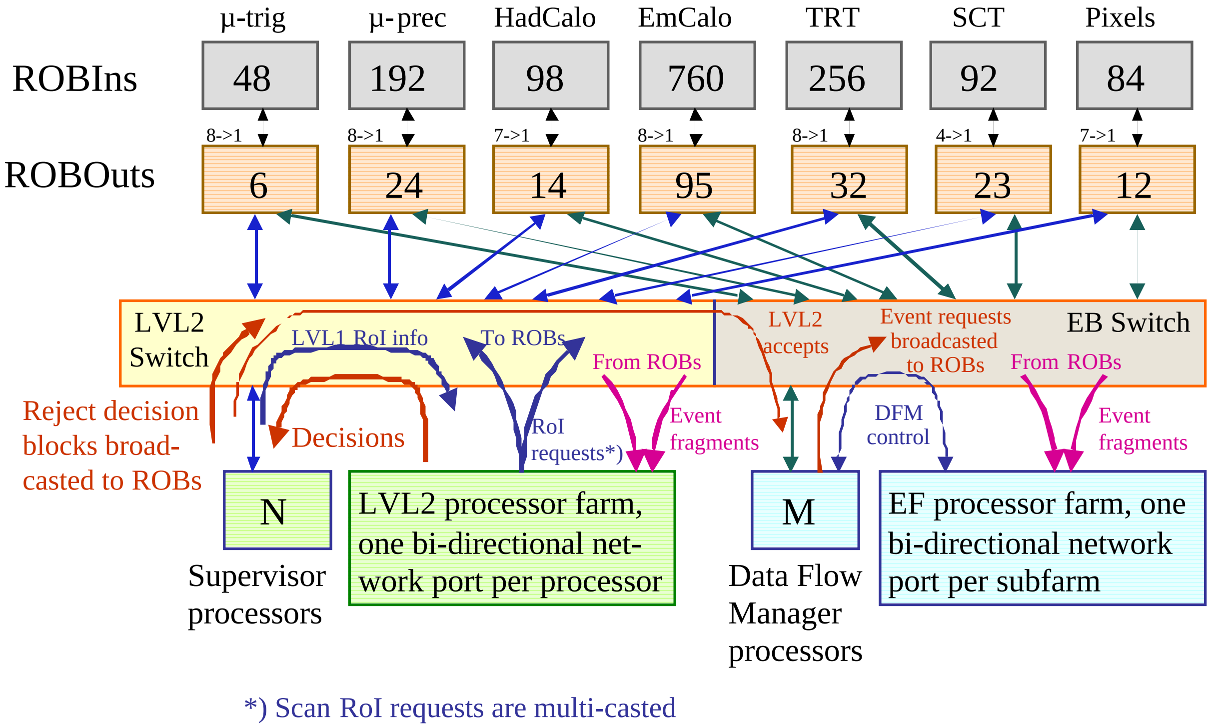

Figure 1. Overview of the system architecture and the data streams, for 2001 model. 2002; 120 ROBIns for pixels, concentration 6->1, blocks of 20 decisions (accepts and deletes) sent by DFM to ROBs, for each accepted event also a delete is sent (after completion of event building).

2 Most important components LVL1 Trigger menus:

Low lumi: 23 kHz MU6, 1 kHz 2 * MU6, 11.5 kHz EM20I, 1.6 kHz 2 * EM15I, 1.34 kHz TAU20 (+XE30); B-physics trigger runs at 5 kHz and produces on average 2.6 muon seeds for which data from the muon detectors are requested and 1.0 electron seed for which data from the calorimeters are requested.

Intermediate lumi = low lumi menu rates scaled up with a factor of 2 apart from B-physics rate
High lumi: 2001: 3.9 kHz MU20, 24.3 kHz EM30I, 4.9 kHz 2 * EM20I.

*2002 Modelling: Low lumi: mu pT threshold may be raised to 8 GeV, MU8 rate is 10 kHz in stead of 23 kHz. 2*MU6: 300 Hz. Intermediate lumi: apart from MU8 rate (19 kHz), 2 *MU6 (600Hz) and EM25I rate (12 kHz) everything assumed to be scaled by a factor of 2. High Luminosity; MU20I rate is 8 kHz in stead of 3.9 kHz., 2*MU6 rate = 3 kHz in stead of 4 kHz. B-physics trigger: scan no longer generates secondary RoIs.* (numbers for 2*MU6: S.Tapprogge, 27-2-2002)

3 Processing sequences and acceptance factors

RoI type	Data from	Action	Accepted	Comment
Muon	muon detectors	LVL1 confirmation	2001: 75 % 2002: 39 %, 8.6 %, high L	39 % for high L (MU20) was specified in mail 21 January S.Tapprogge, in mail of 11 January 8.6 % was specified
	inner detector	track matching	2001: 40 %, MU 6 and low L: 20 % 2002: 22 %, low+inter L	MU6 (MU8) determines B-trigger rate 2002: TRT not used, high L: inner detector not used
	calorimeters	isolation	2001: 4 % 2002: 1.8 %	only for high luminosity
e/gamma low and intermediate L	e.m. layer 2	LVL1 confirmation and refinement	50 %	used in fully sequential strategies
	e.m. layer 1		18 %	
	e.m. calorimeter	LVL1 confirmation and refinement	2001: 16 % 2002:19 %	*19 % comes from mail S.Tapprogge of 11 January*
	hadron calorimeter		2001: 13.5 % 2002: low L: 10 % inter L: 11%	
	inner detector	track find and match	2001: 1.8 % 2002: precision only: low L: 1.4 % inter L: 1.6 %	2002: TRT data requested for 1.4 % (low L) or 1.6 % (inter L) of e/gamma RoIs (final acceptance single em ROI (needed for fully sequential trigger):3.8 %)
e/gamma high L	e.m. layer 2	LVL1 confirmation and refinement	50 %	used in fully sequential strategies
	e.m. layer 1		20 %	
	e.m. calorimeter	LVL1 confirmation and refinement	2001: 18.3 % 2002: 16 %	*2002: for hadron calorimeter assume that it does not give further reduction, but that data is requested for 16 % of the e/gamma RoIs*
	hadron calorimeter		2001: 16.7 %	
	inner detector	track find and match	2001: 2.6 % 2002: precision only: 2.8 %	2002: TRT data requested for 2.8 % of e/gamma RoIs (final acceptance single em ROI (needed for fully sequential trigger):3.3 %)
jet	calorimeters	LVL1 confirmation	100 %	the worst case is considered
hadron	calorimeters	LVL1 confirmation	20 %	2002: TRT data requested for 4 % of hadron RoIs for high L, for 20 % of RoIs for low and inter. L
	inner detector	track find and match	8 % 2002: high L: precision only: 4 %	

Table 1. Processing steps and acceptance factors

4 RoI sizes, ROB mapping

See for a detailed specification the paper model document. In Table 2 the resulting average number of ROBIns per RoI can be found.

RoI type	mu prec	mu trig	total em	had cal	TRT	SCT	Pixels
muon	2.30	1.22	13.8	5.13	5.77	3.32	4.59
e/gamma	--	--	13.17	4.85	3.30	3.27	4.40
jet	--	--	25.00	8.90	--	--	--

Table 2. Average number of ROBIns per RoI, 2001

The mapping of the Pixels and of the SCT have changed, as well as the LVL1 muon grid, see Table 3 for changes in average number of ROBIns per RoI.

RoI type	mu prec	mu trig	total em	had cal	TRT	SCT	Pixels
muon	2.98	1.61	14.73	5.27	6.06	3.24	5.21
e/gamma	--	--	13.17	4.85	3.30	3.13	5.16
jet	--	--	25.00	8.90	--	--	--

Table 3. Average number of ROBIns per RoI, 2002

5 Event fragment sizes

The average sizes of the data fragments requested for an RoI from its associated ROBIns for each subdetector are listed in Table 4.

	Muon Precision	Muon Trigger	Calorimeters	TRT	SCT	Pixels
Low luminosity	800	380	1024* <u>752 / 344**</u>	300 <u>200</u>	250 <u>200</u>	300 <u>120</u>
Intermediate lumin.	800	380	1024* <u>752 / 344**</u>	560 <u>330</u>	370 <u>330</u>	350 <u>200</u>
High luminosity	800	380	1024* <u>752 / 344**</u>	1000 <u>1200</u>	1600 <u>1200</u>	800 <u>500</u>

*Table 4. Average raw data fragment size in Bytes. The * indicates the size after pre-processing for LVL2. Fragments passed to the event builder are 1800 Bytes (<u>2002: 752 Bytes or 344+8*26 = 552 Bytes (for Simion algorithm and design luminosity): source: mail W. Cleland, see below</u>). <u>2002: numbers are underlined</u>*

6 Processing times

ROBIn processing	Nominal time
Indexing per LVL1 trigger	5 µs
RoI request handling	10 µs
Decision block processing per block	50 µs
Decision handling per decision	1 µs

Table 5. Execution times for ROBIn processing

ROBOut processing	Nominal time
Fanning out an event request	2 µs
Decision distribution	2 µs
Input of an event fragment from ROBIn	2 µs
Receiving an event request	10 µs
Receiving a decision	10 µs
Output of an event fragment	10 µs
Speed of concatenating event data	160 MByte/s
Data size independent processing time for concatenating event data	10 µs

Table 6. Execution times for ROBOut processing

Detector	em/gamma RoI	hadron RoI	jet RoI	mu RoI
Pixels & SCT	2 1.20/1.94 ms	2 1.20/1.94 ms	--	2 1.20/1.94 ms
TRT	4 8.27/16.37 ms	4 8.27/16.37 ms	--	4 -- ms
Total calorimeters middle e.m. front e.m. back & presampler hadronic	20 45/70 µs 9 -- µs 4 -- µs 4 -- µs 4 12/18 µs	20 23/35 µs 9 -- µs 4 -- µs 4 -- µs 4 6/9 µs	500 240/390 µs	20 23/35 µs 9 -- µs 4 -- µs 4 -- µs 4 6/9 µs
Muon detector	--	--	--	250 200/250 µs

Table 7. Feature extraction processing times for design luminosity, 2002 numbers: mean or m_50 and m_95 values are specified, for muon isolation the FEX times for hadron RoIs have been used.

Detector	em/gamma RoI	hadron RoI	jet RoI	mu RoI	scan
Pixels & SCT	2 0.38/0.68 ms	2 0.38/0.68 ms	--	2 0.38/0.68 ms	9 22.4/45 ms (102.4/298.72)* or 48.62/144.60* ms
TRT	4 3.35/5.55 ms	4 3.35/5.55 ms	--	4 -- ms	26 (53.7/103.84) 24.20/44.76 0 or 0 ms
Total calorimeters middle e.m. front e.m. back & presampler hadronic	20 34/50 µs 9 -- µs 4 -- µs 4 -- µs 4 9/13 µs	20 17/25 µs 9 -- µs 4 -- µs 4 -- µs 4 4/6 µs	500 240/390 µs	20 -- µs 9 -- µs 4 -- µs 4 -- µs 4 - µs	--
Muon detector	--	--	--	250 µs 200/250 µs	--

*Table 8. Feature extraction processing times for intermediate luminosity, 2002 numbers: mean or m_50 and m_95 values are specified, TRT scan time doubled with respect to low luminosity. *: for TRT scan, numbers between brackets: pT cut at 0.5 GeV*

Detector	em/had RoI	hadron RoI	jet RoI	mu RoI	scan
Pixels & SCT	2 0.28/0.52 ms	2 0.28/0.52 ms	--	2 0.28/0.52 ms	9 11.2/22.5 ms (51.2/149.36)* or 24.31/72.30* ms
TRT	2.74/4.20 ms	2.74/4.20 ms	--	4 -- ms	26 (26.85/51.92) or 12.10/22.38 or 0 ms
Total calorimeters middle e.m. front e.m. back & presampler hadronic	20 32/48 µs 9 -- µs 4 -- µs 4 -- µs 4 8/12 µs	20 16/24 µs 9 -- µs 4 -- µs 4 -- µs 4 4/6 µs	500 240/ 390 µs	20 -- µs 9 -- µs 4 -- µs 4 -- µs 4 -- µs	--
Muon detector	--	--	--	250 µs 200/250 µs	--

*Table 9. Feature extraction processing times for low luminosity, 2002 numbers: m_50 and m_95 values are specified. *: for TRT scan, numbers between brackets: pT cut at 0.5 GeV*

For calculating the time needed for data fragment merging, a merging speed of 160 MByte/s is assumed.

LVL2 processing	Time
RoI request Formulation : per RoI	10 µs
Decision processing	5 µs
Global step	50 µs

Table 10. Execution times for processing steps other than feature extraction

7 Modifications made in spreadsheet

- update of Pixel and SCT geometry and mapping: done, 2-2-2002
- update of Pixel and SCT fragment sizes: done, 2-2-2002
- update of reduction factors: done, 3-2-2002, updated: 24-2-2002
- update of feature extraction execution times: 3-2-2002, 12-2-2002, 21-2-2002, 23-2-2002
- change of calorimeter fragment size: done, 3-2-2002, revised 13-2-2002, using presentation W. Cleland: done, 19-2-2002
- B-physics trigger: new rates, processing times: 21-2-2002
- Update of menus, acceptance factors: 22-2-2002, 24-2-2002
- Grouping of accepts and rejects updated: 1-3-2002
- Information on RoIs generated by barrel LVL1 trigger updated: 14-4-2002
- Processing times and B-physics trigger strategy, if TRT scan is used, updated: 17-4-2002, bug in intermediate luminosity menu removed.

- Spreadsheet replaced by C++ program, final checks of results: 8-6-2002 (results for intermediate luminosity changed slightly)

8 Results

8.1 Number of LVL2 processors

Switch ROS: each ROBIn has its own network connection. Results are for calorimeter fragment size of 752 Bytes. Two numbers separated by a slash indicate the number of processors required with and without multi-casting of requests for the B-physics trigger.

				Number of LVL2 processors					
				No over-head	10 microsec		No over-head	10 microsec	
B-physics				bus / switch ROS	bus ROS	switch ROS	bus / switch ROS	bus ROS	switch ROS
MU6/ MU8	TRT scan	pT TRT scan 1.5 GeV/c	LVL1 rate (kHz)	mean or m_50			m_95		
None	n.a.	n.a.	16.4	11	16	20	16	21	25
MU6	No	n.a.	39.4	80	91/93	105/116	146	158/160	171/182
MU8	No	n.a.	26.4	42	50/51	58/63	74	82/83	90/95
MU6	Yes	No	39.4	429	442/446	468/492	1078	1091/ 1094	1116/ 1140
MU8	Yes	No	26.4	199	208/210	221/232	493	502/504	515/526
MU6	Yes	Yes	39.4	213	226/230	251/275	524	537/540	562/586
MU8	Yes	Yes	26.4	102	111/112	124/135	244	253/254	266/277

Table 11. Number of LVL2 processors, low luminosity

				Number of LVL2 processors					
				No over-head	10 microsec		No over-head	10 microsec	
B-physics				bus / switch ROS	bus ROS	switch ROS	bus / switch ROS	bus ROS	switch ROS
MU6/ MU8	TRT scan	pT TRT scan 1.5 GeV/c	LVL1 rate (kHz)	mean or m_50			m_95		
None	n.a.	n.a.	21.9	20	28	34	30	38	44
MU8	No	n.a.	40.9	130	143/145	157/167	243	256/258	271/280
MU8	Yes	No	40.9	728	743/746	767/788	1839	1854/ 1857	1878/ 1899
MU8	Yes	Yes	40.9	357	372/375	396/417	889	904/907	928/949

Table 12. Number of LVL2 processors, intermediate luminosity

LVL1 rate (kHz)	Number of LVL2 processors					
	No over-head	10 microsec		No over-head	10 microsec	
	bus / switch ROS	bus ROS	switch ROS	bus / switch ROS	bus ROS	switch ROS
	mean or m_50			m_95		
42.5	34	46	52	53	64	71

Table 13. Number of LVL2 processors, high luminosity

8.2 Fragment rate and data volume output per ROBIn

Event building rate is always 2.0 kHz, calorimeter event fragment size is 752 Bytes, in Table 14 and Table 19 the event fragment size if the algorithm described by Simion is applied is also shown.

Low luminosity, no B-physics		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total frag-ment rate per ROBIn (kHz)	Lower limit	2.01	2.01	2.07	2.28	2.02	2.04	2.04
	Average	2.01	2.02	2.42	2.59	2.03	2.08	2.10
	Upper limit	2.02	2.03	3.11	2.95	2.03	2.11	2.15
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	1.67	0.83	1.63 1.20*	1.79 1.27*	0.47	0.47	0.31
	Average	1.67	0.83	1.90 1.33*	2.03 1.39*	0.47	0.48	0.32
	Upper limit	1.68	0.84	2.44 1.59*	2.31 1.52*	0.47	0.49	0.33

*Table 14. Total fragment rate and data volume output per ROBIn, low luminosity. * algorithm described by Simion is used*

Low luminosity, B-physics, TRT scan, MU6		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total frag-ment rate per ROBIn (kHz)	Lower limit	2.18	2.51	2.07	2.28	7.22	7.40	7.43
	Average	2.37	2.80	2.42	2.59	7.22	7.59	7.69
	Upper limit	2.73	3.10	3.11	2.95	7.23	7.72	7.88
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	1.81	1.03	1.63	1.79	1.68	1.72	1.13
	Average	1.97	1.15	1.90	2.03	1.68	1.76	1.17
	Upper limit	2.27	1.28	2.44	2.31	1.68	1.79	1.20

Table 15. Total fragment rate and data volume output per ROBIn, low luminosity

Low luminosity, B-physics, TRT scan, MU8		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.08	2.23	2.07	2.28	4.36	4.44	4.46
	Average	2.17	2.36	2.42	2.59	4.36	4.55	4.61
	Upper limit	2.33	2.50	3.11	2.95	4.37	4.63	4.73
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	1.73	0.92	1.63	1.79	1.01	1.03	0.68
	Average	1.80	0.97	1.90	2.03	1.01	1.06	0.70
	Upper limit	1.94	1.03	2.44	2.31	1.01	1.07	0.72

Table 16. Total fragment rate and data volume output per ROBIn, low luminosity

Intermediate luminosity, no B-physics		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.01	2.03	2.15	2.56	2.04	2.06	2.07
	Average	2.02	2.05	2.65	3.07	2.04	2.13	2.16
	Upper limit	2.04	2.06	3.71	3.71	2.05	2.18	2.24
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	1.67	0.84	1.68	2.01	0.74	0.74	0.48
	Average	1.68	0.84	2.08	2.41	0.74	0.77	0.50
	Upper limit	1.70	0.85	2.91	2.91	0.74	0.79	0.52

Table 17. Total fragment rate and data volume output per ROBIn, intermediate luminosity

Intermediate luminosity, B-physics with TRT scan		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.15	2.43	2.15	2.56	6.49	6.65	6.68
	Average	2.32	2.69	2.65	3.07	6.50	6.85	6.94
	Upper limit	2.63	2.95	3.71	3.71	6.51	6.98	7.13
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	1.79	1.00	1.68	2.01	2.35	2.41	1.55
	Average	1.93	1.11	2.08	2.41	2.35	2.48	1.61
	Upper limit	2.19	1.22	2.91	2.91	2.36	2.53	1.65

Table 18. Total fragment rate and data volume output per ROBIn, intermediate luminosity

High luminosity		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.11	2.31	2.06	2.23	2.01	2.19	2.22
	Average	2.23	2.49	2.72	2.63	2.01	2.40	2.49
	Upper limit	2.45	2.68	3.92	3.04	2.02	2.54	2.70
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	1.76	0.95	1.62 1.19*	1.75 1.25*	2.48	2.70	1.18
	Average	1.85	1.03	2.13 1.44*	2.06 1.41*	2.48	2.95	1.33
	Upper limit	2.04	1.10	3.07 1.89*	2.38 1.56*	2.49	3.13	1.44

*Table 19. Total fragment rate and data volume output per ROBIn, high luminosity. * algorithm described by Simion is used*

8.3 Fragment rate and data volume output per ROBOut

Event building rate is always 2.0 kHz, calorimeter event fragment size is 752 Bytes, in Table 20 and Table 25 the event fragment size if the algorithm described by Simion is applied is also shown.

Low luminosity, no B-physics		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.04	2.10	2.62	3.14	2.08	2.11	2.20
	Average	2.06	2.14	3.89	3.46	2.10	2.23	2.41
	Upper limit	2.11	2.22	6.69	4.01	2.14	2.35	2.59
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	13.35	6.64	13.01 9.57*	12.51 8.91*	3.76	1.89	1.86
	Average	13.38	6.67	15.19 10.62*	14.19 9.72*	3.76	1.93	1.92
	Upper limit	13.45	6.70	19.52 12.69*	16.16 10.66*	3.78	1.96	1.96

*Table 20. Total fragment rate and data volume output per ROBOut, low luminosity. * algorithm described by Simion is used*

Low luminosity, B-physics, TRT scan, MU6		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	3.24	5.43	2.62	3.14	7.28	7.85	8.32
	Average	4.15	6.83	3.89	3.46	7.29	8.32	9.17
	Upper limit	5.82	9.63	6.69	4.01	7.34	8.91	10.1
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	14.51	8.26	13.01	12.51	13.40	6.87	6.78
	Average	15.76	9.22	15.19	14.19	13.41	7.05	7.01
	Upper limit	18.18	10.23	19.52	16.16	13.42	7.17	7.18

Table 21. Total fragment rate and data volume output per ROBOut, low luminosity

Low luminosity, B-physics, TRT scan, MU8		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.56	3.55	2.62	3.14	4.42	4.68	4.94
	Average	2.97	4.18	3.89	3.46	4.43	4.95	5.43
	Upper limit	3.72	5.44	6.69	4.01	4.48	5.28	5.89
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	13.85	7.34	13.01	12.51	8.09	4.12	4.07
	Average	14.41	7.78	15.19	14.19	8.10	4.23	4.20
	Upper limit	15.51	8.23	19.52	16.16	8.12	4.30	4.30

Table 22. Total fragment rate and data volume output per ROBOut, low luminosity

Intermediate luminosity, no B-physics		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.07	2.20	2.98	4.03	2.13	2.18	2.33
	Average	2.13	2.28	4.87	4.61	2.15	2.36	2.66
	Upper limit	2.22	2.45	9.08	5.62	2.22	2.57	2.95
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	13.38	6.69	13.47	14.05	5.90	2.73	2.88
	Average	13.45	6.75	16.65	16.84	5.91	2.83	3.01
	Upper limit	13.60	6.81	23.29	20.38	5.95	2.89	3.11

Table 23. Total fragment rate and data volume output per ROBOut, intermediate luminosity

Intermediate luminosity, B-physics with TRT scan		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	3.07	4.95	2.98	4.03	6.58	7.08	7.54
	Average	3.85	6.16	4.87	4.61	6.61	7.56	8.41
	Upper limit	5.29	8.56	9.08	5.62	6.67	8.15	9.24
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	14.34	8.02	13.47	14.05	18.80	9.63	9.30
	Average	15.41	8.85	16.65	16.84	18.82	9.91	9.66
	Upper limit	17.50	9.72	23.29	20.38	18.85	10.11	9.93

Table 24. Total fragment rate and data volume output per ROBOut, intermediate luminosity

High luminosity		muon-MDT	muon-trig	e.m. cal	hadron cal	TRT	SCT	Pixels
Total fragment rate per ROBIn (kHz)	Lower limit	2.76	4.11	3.02	3.29	2.04	2.61	3.07
	Average	3.32	4.98	5.34	3.64	2.05	3.12	4.00
	Upper limit	4.36	6.70	10.38	4.25	2.07	3.73	4.87
LVL2+EB data volume per ROBIn (MByte/s)	Lower limit	14.05	7.62	12.95 9.54*	12.24 8.78*	19.82	10.79	7.09
	Average	14.82	8.21	17.06 11.51*	14.44 9.84*	19.84	11.80	7.96
	Upper limit	16.31	8.84	24.56 15.11*	16.67 10.90*	19.88	12.51	8.63

*Table 25. Total fragment rate and data volume output per ROBOut, high luminosity. * algorithm described by Simion is used*

9 Copies of relevant mails

Mail from Bill Cleland, 18 december 2001 (via Andy Lankford, 14 February 2002)

```
>>X-From_: cleland+@pitt.edu  Tue Dec 18 13:11:14 2001
>>Date: Tue, 18 Dec 2001 16:09:19 -0500
>>From: Bill Cleland <cleland+@pitt.edu>
>>Subject: Data compression
>>To: "Andrew J. Lankford" <lankford@lankford.ps.uci.edu>
>>Cc: Luc POGGIOLI <poggio@lapp.in2p3.fr>,
>>         Horst Oberlack <oberlack@mppmu.mpg.de>,
>>         Stefan Simion <simion@in2p3.fr>,
```

```
>>         Christian Zeitnitz <zeitnit@dhemze.physik.uni-mainz.de>,
>>         cleland <cleland@pitt.edu>, Mauro Citterio
>> <mauro.citterio@mi.infn.it>,
>>         Leonid Kurchaninov <kurchan@mppmu.mpg.de>,
>>         Lorenzo Moneta <moneta@mail.cern.ch>,
>>         Daniel LA MARRA <Daniel.LaMarra@physics.unige.ch>,
>>         Louis Fayard <lfayard@in2p3.fr>
>>X-Mailer: Mozilla 4.76 [en] (Windows NT 5.0; U)
>>X-Accept-Language: en
>>
>>
>>Dear Andy,
>>Last week, our Back End Electronics Subgroup (a subgroup of the liquid argon
>>stering group) discussed the question you raised concerning compression.
>>As you will recall, the uncompressed data block from a DSP servicing 128
>>channels has the following structure (32 bit words):
>>
>>                   header words    17
>>                   energy         128
>>                   time+chiquare   13
>>                                 ====
>>                          total   158
>>
>>Where time and chisquare are sent for only those channels above threshold,
>>which we assume to be 10% of the channels.
>>
>>Stefan's compression algorithm will be applied only to the energy block,
>>which will reduce its size to about 26 words, but the reduction factor will
>>depend on the luminosity, as it is sensitive to the noise. In order not to
>>require Level 2 to decode the energy, we would add the energy to the
>>time+chisqure block, so its size would double. Thus the compressed word
>>count would be
>>
>>                   header words    17
>>                   energy          26
>>                   time+chiquare   26
>>                                 ====
>>                          total    69
>>
>>Another suggestion to reduce the data volume has been made recently. Since
>>most of our cells have low energy, we could send the energy as a 16-bit word,
>>and append the cells with large energy to the end of the record. If we choose
>>the units to be MeV and add a sign bit, the threshold for appending an energy
>>value would be 32 GeV, so the rate would be exremely low. Since this scheme
>>would involve no compression, there would be no need to duplicate the energy
>>in the time+chisquare block. In this case the record would be:
>>
>>
>>                   header words    17
>>                   energy          64
>>                   time+chiquare   13
>>                                 ====
>>                          total    94
>>
>>So while this scheme does not accomplish the full benefit of compression
>>(59% vs 44% of the original record size), it accomplishes a good fraction
>>of it, it is luminosity independent, and it costs us very little in computing
>>power. The BESG therefore is comfortable in adopting this latter scheme as a
>>"baseline", whereas a decision on full compression can only be taken after we
>>understand better the computational demands on the DSP.
>>
>>                                 Best regards,
>>                                 Bill
```

Mail from Stefan Tapprogge, 11 January 2002

```
From Stefan.Tapprogge@cern.ch Fri Jan 11 08:53:50 2002
Received: (from root@localhost)
        by park.nikhef.nl (8.10.2/8.10.2) id g0B7rmJ24577
        for i73; Fri, 11 Jan 2002 08:53:48 +0100 (MET)
Received: from smtp3.cern.ch (smtp3.cern.ch [137.138.131.164])
        by park.nikhef.nl (8.10.2/8.10.2) with ESMTP id g0B7rfW24570
        for <i73@nikhef.nl>; Fri, 11 Jan 2002 08:53:42 +0100 (MET)
Received: from sunsci7 (sunsci7.cern.ch [137.138.90.187])
        by smtp3.cern.ch (8.12.1/8.12.1) with SMTP id g0B7qv5C007900;
        Fri, 11 Jan 2002 08:52:57 +0100 (MET)
X-Authentication-Warning: smtp3.cern.ch: Host sunsci7.cern.ch [137.138.90.187] c
laimed to be sunsci7
Sender: tapprogg@mail.cern.ch
Message-ID: <3C3E99D9.6D46@cern.ch>
Date: Fri, 11 Jan 2002 08:52:58 +0100
From: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Organization: CERN. European Lab. for Particle Physics
X-Mailer: Mozilla 3.01Gold (X11; I; SunOS 5.6 sun4m)
MIME-Version: 1.0
To: Jos Vermeulen <i73@nikhef.nl>, Chris Bee <Chris.Bee@cern.ch>
CC: Nick Ellis <Nick.Ellis@cern.ch>,
   Valerio Vercesi <Valerio.Vercesi@pv.infn.it>, abolins@pa.msu.edu
Subject: standard values / paper model parameters
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit
Content-Length: 4146
Status: RO

Dear Chris and Jos,

based on the short compilation from Jos on the paper model parameter
summary, which Chris has distributed, I have a first set of comments,
mostly concentrating on the accept fractions in table 1.

Before going into details, I would like to propose that we should aim
at a single document containing all the important relevant parameters,
both from the PESA side (and LVL1) -- i.e. rates, accept factor,
algorithm execution times etc. -- as well as more "technical" numbers
-- such as decision receiving, event request sending etc.. Maybe for
the former, Valerio, Nick and myself should make sure that we have
the correct information. For the latter, Chris, Jos and Maris are
obvious candidates.


Please note that in the following I am referring also to the results
as documented in the TP (thus I am not yet taking into account the
raising of two thresholds for the low luminosity scenario!).

---------------

muons:

  low lumi:     mu6    23 kHz LVL1
                                                         39%
                        9 kHz LVL2 muon confirmation
                                                         22%
                        5 kHz LVL2 combined muon
                                   (Inner Detector)

      (the mu20 which will be used also beyond B-physics is
       still in the noise of the rates)
```

```
  design lumi:  mu20i   8 kHz LVL1 (was 4 kHzin TDR)
                                                           8.6%
                        690 Hz LVL2 muon confirmation
                                                          >1.8%
                       >140 Hz LVL2 muon isolation
                                    (calorimeter)

    NOTE that for design L, no study has been made on the effect
         of the combined muon reconstruction for a further rate
         reduction

----------------------

electrons


   low lumi:     e20i   5.8 kHz LVL1 (11 kHz in TDR)
                                                          19%
                        1.1 kHz LVL2 e confirmation
                                     (calorimeter)
                                                          2.4%
                        140 Hz LVL2 track match
                                    (Inner Detector)


   design lumi:  e30i   22  kHz LVL1
                                                          16%
                        3.5 kHz LVL2 e confirmation
                                     (calorimeter)
                                                          2.1%
                        460 Hz LVL2 track match
                                    (Inner Detector)


----------------------

photons


   low lumi:     g40i   5.8 kHz LVL1 (11 kHz in TDR)
                                                          1.4%
                         83 Hz  LVL2 g confirmation
                                     (calorimeter)


   design lumi:  g60i   22  kHz LVL1
                                                          1.2%
                        250 Hz LVL2 g confirmation
                                     (calorimeter)

    NOTE that for the photon selection we are not performing
         any track search or veto at LVL2 (nor at present at
         the EF).

---------------------

At the moment I cannot comment on the sequential calorimeter
processing numbers, still have to dig out this information.

For the e/g part, the numbers which Jos has summarized and the
TP results are in good/reasonable agreement. In case of the
```

muons, my feeling is that in the paper model sometimes a too
large accept fraction is used. (Others: please verify/falsify
and comment).

For both signatures, there are two cases where the TP studies
lead to differences in the LVL1 rates wrt the ones documented in
the LVL1 TDR. For the muon, as far as I remember, we probably
have to take the higher rate (8 kHz instead of 4 kHz at design
luminosity) as this is the result of a more recent study. In
case of the em signature, there might be a normalisation problem
with the LVL1 rate in the TP study.

I will come back to the algorithm later in a second mail.

Cheers,

Stefan.

Mail from Stefan Tapprogge, 11 January 2002

Date: Fri, 11 Jan 2002 15:36:58 +0100
From: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Organization: CERN. European Lab. for Particle Physics
X-Mailer: Mozilla 3.01Gold (X11; I; SunOS 5.6 sun4m)
MIME-Version: 1.0
To: Jos Vermeulen <i73@nikhef.nl>, Chris Bee <Chris.Bee@cern.ch>
CC: Nick Ellis <Nick.Ellis@cern.ch>,
Valerio Vercesi <Valerio.Vercesi@pv.infn.it>, abolins@pa.msu.edu
Subject: standard values / paper model parameters (2)
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit
Content-Length: 2823
Status: RO

Dear Chris and Jos,

now for some comments on table 6, the LVL2 algorithm execution times.
We should keep in mind that there are still sizeable effects due to
e.g. data preparation for the Si space points to be taken into account.
All numbers are again taken from the TP, the values are the medians of
the timing distributions.

I did not do a complete scaling to a machine of a given number of
SI95, that's what we probably should do for the final documentation.
Either I just took the ratio of clock speeds or assumed 160 SI95 for
a 4 GHz machine (extrapolation goal).

muons:

muon algorithm (MUFAST)

2ms per RoI (10 SI95 machine) ===> 130 us

tracking (SCT+pixel, use electron number see below)

2.5 ms per RoI (500MHz PII) ===> 320 us

NOTE that the TRT has not been used so far for the
combined muon reconstruction at LVL2

```
       the matching between muon and ID track takes
       negligible time (20 us on 10 SI95, i.e. 1.3 us
       for the model), HOWEVER it is not clear whether
       there might be additional CPU time needed to evolve
       the track parameters to the position where the matching
       takes place.

   isolation (calorimeter)

      no numbers available, assuming the electron calo time
      seems to me a good approximation        ===>    20 us


electrons:


    calorimeter (1 step only)

       150 us per RoI (500MHz PII)            ===>    20 us

    SCT+pixel

       2.5 ms         (500 MHz)               ===>   320 us

    TRT

       31 ms          (500 MHz)               ===>     4 ms


photons:

    calorimeter (1 step only)

       200 us per RoI (500MHz PII)            ===>    25 us



jets:

    calorimeter

       1.9 ms per RoI (500MHz PII)            ===>   240 us


All the above numbers are for low luminosity. For the muon, we
do not have a design lumi number, for the jets it should be quite
similar (if a CONE algorithm is used, as it is the case).

For the electron/photon case, the corresponding numbers (extra-
polated to 4 GHz) are


      electron        calo       ====>     25 us

                     SCT/pixel   ====>    780 us

                     TRT         ====>     50 ms

      photon          calo       ====>     25 us


Numbers for the tau/hadron RoI have not been done for the TP, the
```

calorimeter part will not be too much different from the em part,
although here it is not obvious that a sequential processing will
make sense. Maybe the best would be to conservatively assume that
all calo cells are requested at once and a total CPU time of 25 us.
The tracking will be more complex than for the electron, as the
p_T threshold will be lower. This was never done for LVL2 using
a real algorithm with simulated data, so no information available.

Cheers,

Stefan.

From mail Stefan Tapprogge, 6 February 2002

Coming back to the issue of the TRT RoI processing time at design
luminosity. If I gather it correctly, you assume the execution of
both the precision and the TRT FEX with the same rate (i.e. the
LVL2 calo FEX output rate). Here one should include a further
sequential step, by first doing the precision FEX (which will get
you from 3.5 kHz to 620 Hz at design luminosity, and only then
executing the TRT and final matching to go down to 420 Hz). This
should reduce the impact of the TRT execution we were discussing
earlier and allow us to have a consistent picture (rates and CPU
times).

Mail from John Baines, 18 February 2002

Date: Mon, 18 Feb 2002 17:28:28 +0000 (GMT)
From: john.baines@rl.ac.uk
X-X-Sender: jtb@heplnx3.pp.rl.ac.uk
Reply-To: john.baines@rl.ac.uk
To: Jos Vermeulen <i73@nikhef.nl>
cc: abolins@pa.msu.edu
Subject: Paper model for B-trigger
Message-ID: <Pine.LNX.4.44.0202181614340.29080-100000@heplnx3.pp.rl.ac.uk>
MIME-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII
Content-Length: 2801
Status: RO

Hi Jos, Maris,

As I expect you know, we have been looking at updating the estimates of
execution times for the B-trigger. The following is the current status
(all execution times are for a 4GHz machine)

Muon Confirmation:

Muon FEX: ATLAS-TDAQ-2000-036 is still the only timing measurement
available. It could be argued that the histogram in that note is more
significant with 0.2ms than 0.25ms, but thats a detail.

TRT FEX: 4ms is for the TBTREC FEX, the time for the TRT-LUT FEX is 0.23ms.
However, we would like to consider the scenario where the TRT is not used
at all for muon confirmation (in fact the rate estimates are based on
Si+Muon only).
SCT+PIX: 2ms is the time for the for the (old) SCT Hough FEX with a
generous 1ms for data preparation. A better estimate would be 0.2ms
based on the SctKalman FEX and a more realistic estimate for data
preparation.

Scan:

Pixel-scan : The pixel-scan time needs to be increased to 8ms to allow for the data-preparation step.
SCT: 4ms is correct for a pT>1.5 GeV threshold including data preparataion time
TRT: 26ms is correct for a pT>0.5 GeV threshold, but we would also like to consider a 1.5 GeV threshold which takes 13ms, and also the case of omitting the TRT scan altogether.

So a minimum estimate for the LVL2 cpu requirements for B-physics would be, as a function of muon pT threshold and luminosity:

	mu6@10^33			mu8@10^33			mu8@2x10^33		
	Rate	Time	cpu	Rate	Time	cpu	Rate	Time	cpu
LVL2 MU	23	0.2	4.6	10	0.2	2.0	19	0.2	3.8
MU+Si	9	0.2	1.8	4	0.2	0.8	8	0.4	3.2
pixelscan	5	8	40	2.5	8	20.	5	16	80.0
SctKalman	5	4	20	2.5	4	10	5	8	40.0
Total LVL2			66			33			127

This is without any TRT information (which means no J/psi(ee) trigger).

Given the financial constraints, one option we would like to pursue is that given in the middle group of columns, i.e. a 8 GeV muon threshold and an assumption that we would only run the B-triggers when the luminosity has fallen to 10^33. Would it be possible to re-run the paper model for this scenario and with the above cpu times to determine the cpu requirements?

Cheers
John

Mail from John Baines, 18 February 2002

From: john.baines@rl.ac.uk
Date: Wed Feb 20, 2002 11:42:44 AM Europe/Amsterdam
To: Jos Vermeulen <i73@nikhef.nl>
Cc: abolins@pa.msu.edu
Subject: Re: Secondary RoIs
Reply-To: john.baines@rl.ac.uk

Hi Jos,
Sorry for the delay in replying.

I realise now that I am being pessimistic in the scaling of the execution time for the Si FEX in muon RoI. The search is seeded by the muon, so it is not the same as doing a full-scan in the RoI. We should use the same time for the high pT electron FEX.

I think we should drop the RoIs created by the ID scan and extrapolated outwards, and instead assume that these will me covered by LVL1 RoI i.e. 2 muon triggers (mu6mu6), or a muon-ecal (mu6e5) trigger. I'm not sure if you have the information necessary to include these in the model? We have asked Leandro for the di-muon rates. I'm don't think we have an estimate for an e5 rate, so we proably have to omit that for the time being.

Cheers
John

Mail from Stefan Tapprogge, 21 February 2002

From: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Date: Thu Feb 21, 2002 05:11:37 PM Europe/Amsterdam
To: Jos Vermeulen <i73@nikhef.nl>
Cc: Valerio Vercesi <Valerio.Vercesi@pv.infn.it>
Subject: paper model parameters

Hi Jos,

please find below a compilation of the parameters for the paper model (so far concentrating on the high pT) part, which I have checked and also new ones, esp. for the 2*10^33 scenario.

I started to produce numbers for the total amount of CPU's needed for various scenarios. Not being the expert on the model, I would appreciate very much if you were able to take the input given below and produce the number of CPU's needed for the three luminosity scenarios (and both for the m_50 and m_95 values of the execution times), as well as for a 40 kHz input rate and the scaling to 75 kHz - so that there is a solid cross-check of my modifications and results. Please leave out the full scan for the time being, here we should wait for some new input from John et al.

If you have any questions on the numbers below please let me know.

One question: is there any assumption used at the moment in the paper model on the execution time needed for data preparation?

See you next week at CERN, where we would like to understand better what constitutes the OS part of the CPU time (beyond context switching).

Cheers,

Stefan.

10^33 m_50 (m_95 values in brackets) [in ms]

algorithm / RoI type	muon	calo	SCT/pixel	TRT
MU	0.13 (0.83)		0.28 (0.64)	
EM		0.032 / 0.008 (0.048 / 0.012)	0.28 (0.64)	3.44 (23.33)
TAU		0.016 / 0.004 (0.024 / 0.006)	0.28 (0.64)	3.44 (23.33)
J		0.24 (0.39)		

the two numbers for the calo part indicate the sequential processing first in em. compartment, next in hadronic comp.

note that the numbers for the EM are doubled wrt to the version 5alpha to take into account the fact that we will have possibly (slightly) different algorithms for electrons and photons.

10^33 acceptance factors

EM into SCT/pixel: 0.10

EM into TRT: 0.014 (done only after SCT/pixel)

TAU into SCT + TRT: 0.2

MU into SCT/pixel: 0.39

2*10^33 m_50 (m_95 values in brackets) [in ms]

algorithm	muon	calo	SCT/pixel	TRT
RoI type				
MU	0.13 (0.83)		0.34 (0.75)	
EM		0.034 / 0.009 (0.05 / 0.013)	0.34 (0.75)	8.6 (37.4)
TAU		0.017 / 0.004 (0.025 / 0.006)	0.34 (0.75)	8.6 (37.4)
J		0.24 (0.39)		

2*10^33 acceptance factors

EM into SCT/pixel: 0.11

EM into TRT: 0.016 (done only after SCT/pixel)

TAU into SCT + TRT: 0.2

MU into SCT/pixel: 0.39

rates (only the ones which have been changed wrt to the 10^33 menu are given, the others have been left untouched -- simplified approach)

```
  MU6    19000   (this actually correspond to a  8 GeV selection now)
 2MU6     2000
 EM20I   12000   (this actually correspond to a 25 GeV selection now)
2EM15I    3200
```

giving a total rate of 39248 Hz

```
10^34   m_50  (m_95 values in brackets)    [in ms]


   algorithm        muon         calo         SCT/pixel       TRT

RoI type


  MU           0.13 (0.83)                    0.78 (1.59)

  EM                        0.045 / 0.012     0.78 (1.59)  50 (150)
                           (0.07  / 0.018)

  TAU                       0.023 / 0.006     0.78 (1.59)  50 (150)
                           (0.035 / 0.009)

  J                           0.24 (0.39)


10^34 acceptance factors

  EM into SCT/pixel:    0.16

  EM into TRT:          0.028   (done only after SCT/pixel)

  TAU into SCT:         0.2

  TAU into TRT:         0.04

  MU into SCT/pixel:    0.39

     -------------------------------------------------------------
```

Mail from John Baines, 3 April 2002

```
Date: Wed, 3 Apr 2002 11:56:51 +0000 (GMT)
From: john.baines@rl.ac.uk
X-X-Sender: jtb@heplnx3.pp.rl.ac.uk
Reply-To: john.baines@rl.ac.uk
To: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: Re: m_95 values for the B-physics algorithm execution times

Hi Jos,

  The definitions come from Saul Gonzalez, they are simply the points
with respectively 50% and 95% of execution times less than that value
(sorry I think I confussed "mode" and "median" in my previous email).

The times I have sent you in the past are the mean values. As Saul points
out m_50 is less prone to tail effects than the mean, but the values
should be similar. What I will do is measure mean, median (m_50) and
m_95.

              Cheers
                    John
```

Mail from John Baines, 4 April 2002, morning

Date: Thu, 4 Apr 2002 10:20:14 +0000 (GMT)
From: john.baines@rl.ac.uk
X-X-Sender: jtb@heplnx3.pp.rl.ac.uk
Reply-To: john.baines@rl.ac.uk
To: Jos Vermeulen <i73@nikhef.nl>
cc: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Subject: Re: m_50 and m_95 times for scan

Hi Jos,
Sorry, I didn't give you the full table since the other numbers are just calculated assuming that scan execution time is proportional to occupancy and hence luminosity. I have given the full tables below.

I fear that a full set of mean execution times are not available for the RoI based studies as Saul chose to use m_50 for the electron trigger studies since this is a more robust measurement. The means can be very susceptible to one or two events in the tail. I suspect that for the calorimeter algorithms, the difference between m_50 and mean will be small. This is probably not true for the RoI-based Si and TRT algorithms, but I think we probably have to make the assumption that mean~=m_50 for the time being.

For the muon algorithm, I think the execution time is in any case a (by eye?) mean? I must admit I'm very supprised by the m_95 value for the muon execution time. As I understand it, the source is DAQ-2000-036 and the distribution is shown in Figure 16. It shows a mean of ~3ms @ 10 SI95, but nothing above 4.6ms @ 10 SI95. A m_95 value of 0.83ms @ 160 SI95 corresponds to 13ms @ 10 SI95, which doesn't seem consistant with this Figure.

	mu6@10^33			mu8@10^33			mu8@2x10^33		
MEANS	Rate	Time	cpu	Rate	Time	cpu	Rate	Time	cpu
LVL2 MU	23	0.18?	4.1	9.5	0.18?	1.7	19	0.18?	3.4
MU+Si	9	(0.28)	(2.5)	3.7	(0.28)	(1.0)	7.4	(0.34)	(2.5)
pixelscan	5	7.1	35.5	2.1	7.1	14.9	4.2	14.2	59.6
SctKalman	5	3.9	19.5	2.1	3.9	8.2	4.2	7.8	32.8
			61.6			25.8			98.3

	mu6@10^33			mu8@10^33			mu8@2x10^33		
m_50	Rate	Time	cpu	Rate	Time	cpu	Rate	Time	cpu
LVL2 MU	23	0.13?	3.0	9.5	0.13?	1.2	19	0.13?	2.5
MU+Si	9	0.28	2.5	3.7	0.28	1.0	7.4	0.34	2.5
pixelscan	5	5.3	26.5	2.1	5.3	11.1	4.2	10.6	44.5
SctKalman	5	3.2	16	2.1	3.2	6.7	4.2	6.4	26.9
			47.3			19.8			76.4

	mu6@10^33			mu8@10^33			mu8@2x10^33		
m_95	Rate	Time	cpu	Rate	Time	cpu	Rate	Time	cpu
LVL2 MU	23	0.83!	19.1!	9.5	0.83!	7.9!	19	0.83!	15.8!
MU+Si	9	0.64	5.8	3.7	0.64	2.4	7.4	0.75	5.6
pixelscan	5	17.0	85	2.1	17.0	35.7	4.2	34.0	142.8
SctKalman	5	7.7	38.5	2.1	7.7	6.7	4.2	15.4	64.7
			148.4			52.7			228.9

Cheers
John

Mail from John Baines, 4 April 2002, evening

Date: Thu, 4 Apr 2002 16:18:57 +0000 (GMT)
From: john.baines@rl.ac.uk
X-X-Sender: jtb@heplnx3.pp.rl.ac.uk
Reply-To: john.baines@rl.ac.uk
To: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: m_50 and m_95 times for scan

Hi Jos,

Everything I have given you so far is for pixel-scan + SCT with a 1.5 GeV threshold and based on the assumption of a B-physics trigger that does not use the TRT-scan (i.e. no J/psi(ee) trigger).

For the TRT-LUT scan the times are :

	10^33			2x10^33		
	mean	m_50	m_95	mean	m_50	m_95
TRT-scan (pT>1.5 GeV)	12.1	11.0	22.4	24.2	22.0	44.8
TRT-scan (pT>0.5 GeV)	26.9	24.1	51.9	53.8	48.2	103.8

I don't know if you also wanted the times for extrapolating these tracks inwards into the SCT and pixels (i.e. TRT-seeded reconstruction):

	10^33			2x10^33		
	mean	m_50	m_95	mean	m_50	m_95
SiKalman (pT>1.5 GeV)	24.3	15.1	72.3	48.6	30.2	144
SiKalman (pT>0.5 GeV)	51.1	32.4	149.3	102.2	64.8	298.6

These times are substantially longer than those for the pixel-seeded reconstruction because the TRT-LUT algorithm gives many more tracks to be followed into the Si than the pixel-scan. This is why we have changed the baseline to be pixel-seeded reconstruction.

Cheers
John

Mail sent to John Baines, 4 April 2002

Hi John,

> For the TRT-LUT scan the times are:

Thanks very much !

> I don't know if you also wanted the times for extrapolating these tracks
> inwards into the SCT and pixels (i.e. TRT-seeded reconstruction):

I have not been using these times until now. I have built the following into the paper model for low luminosity:

Choice between:
- no B-physics trigger
- MU6 + Pixel scan + SCT for 1.5 GeV threshold
- MU8 + Pixel scan + SCT for 1.5 GeV threshold
- MU6 + Pixel scan + SCT for 1.5 GeV threshold + TRT scan for 0.5 GeV threshold
- MU8 + Pixel scan + SCT for 1.5 GeV threshold + TRT scan for 1.5 GeV threshold

For intermediate luminosity for the B-physics trigger the second and fourth item of the above list (the items with a MU6) are not considered.

I understand now that the last two items in the list probably better are
replaced by :

- MU6 + TRT scan for 0.5 GeV threshold + SiKalman (pT > 0.5 GeV)
- MU6 + TRT scan for 1.5 GeV threshold + SiKalman (pT > 1.5 GeV)
- MU8 + TRT scan for 0.5 GeV threshold + SiKalman (pT > 0.5 GeV)
- MU8 + TRT scan for 1.5 GeV threshold + SiKalman (pT > 1.5 GeV)

Is this correct?

Cheers, Jos

Mail from John Baines, 5 April 2002

Date: Fri, 5 Apr 2002 09:38:06 +0000 (GMT)
From: john.baines@rl.ac.uk
X-X-Sender: jtb@heplnx3.pp.rl.ac.uk
Reply-To: john.baines@rl.ac.uk
To: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: SiKalman
In-Reply-To: <200204042019.g34KJwh26347@bilbo.nikhef.nl>
Message-ID: <Pine.LNX.4.44.0204050930220.19890-100000@heplnx3.pp.rl.ac.uk>
MIME-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII
Content-Length: 1188
Status: RO

Hi Jos,

> - no B-physics trigger
> - MU6 + Pixel scan + SCT for 1.5 GeV threshold
> - MU8 + Pixel scan + SCT for 1.5 GeV threshold
>
> - MU6 + TRT scan for 0.5 GeV threshold + SiKalman (pT > 0.5 GeV)
> - MU6 + TRT scan for 1.5 GeV threshold + SiKalman (pT > 1.5 GeV)
> - MU8 + TRT scan for 0.5 GeV threshold + SiKalman (pT > 0.5 GeV)
> - MU8 + TRT scan for 1.5 GeV threshold + SiKalman (pT > 1.5 GeV)

Just to be clear, SiKalman is the algorithm used for the SCT or
SCT & pixels in all cases.

> Is this correct?

Yes. But I guess we could question whether we need to retain all these
possibilities. I think its fairly clear that the TRT-seeded option with a
0.5 GeV threshold is too resource hungrey, so we could drop this, unless
you want to keep it for comparison with old results (where this was
the baseline).

Cheers
John

Mail from John Baines, 10 April 2002

From: john.baines@rl.ac.uk
Date: Wed Apr 10, 2002 12:56:43 PM Europe/Amsterdam
To: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Cc: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: m_50 and m_95 times for scan
Reply-To: john.baines@rl.ac.uk

```
Hi Stefan,
  I'm not quite clear what numbers you mean. For the muon FEX you have
m_50 = 0.13ms, m_95=0.83ms. The 0.13ms is consistant with
2ms * 10SI95/160SI95 (2ms is given on p17 of ATL-DAQ-2000-036) but the
0.83ms seems suprisingly large. Are you saying that in fact the 0.83ms
includes the sct? An if so, is the inference that the muon m_95 time
should be 0.83-0.64=0.19ms? That would be 3ms @ 10SI95, and a m_95-m_50
of 1ms @ 10SI95 seems consistant with  the width (hwhm~0.5ms @ 10SI95) of
the distribution shown in figure 16 of ATL-DAQ-2000-036.

So we would have :
           m_50  m_90
 10SI95    2ms    3ms
160SI95    0.13ms 0.19ms

Although this does leave the mystery of why Figure 16 looks more
a mean of 3ms and a m_95 of 4ms

                     Cheers
                        John
```

Mail from John Baines, 10 April 2002

```
From: john.baines@rl.ac.uk
Date: Wed Apr 10, 2002  03:33:38 PM Europe/Amsterdam
To: Jos Vermeulen <i73@nikhef.nl>
Cc: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Subject: Re: Re: m_50 and m_95 times for scan
Reply-To: john.baines@rl.ac.uk

Hi Jos,


Shouldn't we take then (for 160SI95) : m_50 = 0.2 ms, m_95 = 0.25 ms ?

On the face of it, that would seem a safer bet. But I don't have any
justification for believing the histogram is correct and the text is wrong
in Leandro's note, and it was the 2ms (=>0.13ms) that went in the TP. I
sent Leandro an email asking about the discrepancy some weeks ago, but
received no reply. I was going to raise it when I next saw him.

               Cheers
                    John
```

Mail sent to Daniel Pomarede et al, 10 April 2002

```
From: Jos Vermeulen <i73@nikhef.nl>
Date: Wed Apr 10, 2002  12:24:05 PM Europe/Amsterdam
To: Daniel Pomarede <Daniel.Pomarede@cern.ch>
Cc: falciano@roma1.infn.it, aleandro.nisati@roma1.infn.it,
stefano.veneziano@roma1.infn.it, abolins@pa.msu.edu, Jos Vermeulen
<i73@nikhef.nl>
Subject: Re: Re: Barrel LVL1 muon trigger: RoI positions

Dear Daniel.

> I am not at all the expert on RoI definition.

From the documentation on the LVL1 muon barrel trigger I learned that the
possible RoI positions(not the extent, this is something that can be defined
by the LVL1 supervisor) are determined by the coincidence matrices in
combination with the PAD logic. My understanding is that each PAD outputs 2
```

bits, indicating which CM (of the four it is receiving input from) fired. So what I need to know are the dimensions in eta-phi space of the areas "seen" by the CM matrices. Your web pages in any case give the impression that you have access to this information. With the information I found on your web pages and with the engineering drawings I have have been able to calculate the following:

Large sectors, pivot plane is at r = 7051:

CM	z extent	high z	high eta	eta extent
0	850	1000	0.13	0.13
1	850	1850	0.24	0.11
2	850	2700	0.35	0.11
3	850	3550	0.46	0.10
4	850	4400	0.56	0.10
5	850	5250	0.65	0.10
6	732	5982	0.73	0.08
7	488	6470	0.78	0.05
8	732	7202	0.85	0.07
9	488	7690	0.90	0.05
10	745	8435	0.97	0.07
11	745	9180	1.03	0.06
12	480	9660	1.07	0.04

Small sectors, pivot plane is at r = 8396:

CM	z extent	high z	high eta	eta extent
0	850	1000	0.12	0.12
1	850	1850	0.22	0.10
2	730	2580	0.30	0.08
3	1100	3680	0.43	0.12
4	730	4410	0.50	0.08
5	730	5140	0.58	0.08
6	730	5870	0.65	0.07
7	1055	6925	0.75	0.10
8	980	7905	0.84	0.09
9	730	8635	0.90	0.06
10	730	9365	0.96	0.06

The z extents are a little bit different from the values I found in the documentation, I used for the high z values numbers that I found in the engineering drawings and then computed what the z extent should be, assuming that there are no overlaps (whichanyway are "removed" by the LVL1 trigger). For the large sectors, for CM 6 and 7, and CM 8 and 9, I also used information from the "The Level-1 Trigger System of the ATLAS Barrel Muon Spectrometer"" , p.9 on the number of input channels used (16 vs. 24, so I concluded that the z extent should be for CM6 and 8 1.5 that for CM 7 and 9).

For the phi ranges I get:

phi extent	low phi	high phi	type
0.122	0.000	0.122	LARGE
0.100	0.122	0.223	LARGE
0.071	0.223	0.293	SMALL
0.099	0.293	0.393	SMALL
0.099	0.393	0.492	SMALL
0.071	0.492	0.563	SMALL
0.100	0.563	0.663	LARGE
0.122	0.663	0.785	LARGE
0.122	0.785	0.907	LARGE
0.100	0.907	1.008	LARGE
0.071	1.008	1.079	SMALL
0.099	1.079	1.178	SMALL

0.099	1.178	1.278	SMALL
0.071	1.278	1.348	SMALL
0.100	1.348	1.449	LARGE
0.122	1.449	1.571	LARGE
0.122	1.571	1.693	LARGE
0.100	1.693	1.793	LARGE
0.071	1.793	1.864	SMALL
0.099	1.864	1.963	SMALL
0.099	1.963	2.063	SMALL
0.071	2.063	2.134	SMALL
0.100	2.134	2.234	LARGE
0.122	2.234	2.356	LARGE
0.122	2.356	2.478	LARGE
0.100	2.478	2.579	LARGE
0.071	2.579	2.649	SMALL
0.099	2.649	2.749	SMALL
0.099	2.749	2.848	SMALL
0.071	2.848	2.919	SMALL
0.100	2.919	3.020	LARGE
0.122	3.020	3.142	LARGE
0.122	3.142	3.264	LARGE
0.100	3.264	3.364	LARGE
0.071	3.364	3.435	SMALL
0.099	3.435	3.534	SMALL
0.099	3.534	3.634	SMALL
0.071	3.634	3.704	SMALL
0.100	3.704	3.805	LARGE
0.122	3.805	3.927	LARGE
0.122	3.927	4.049	LARGE
0.100	4.049	4.149	LARGE
0.071	4.149	4.220	SMALL
0.099	4.220	4.320	SMALL
0.099	4.320	4.419	SMALL
0.071	4.419	4.490	SMALL
0.100	4.490	4.590	LARGE
0.122	4.590	4.712	LARGE
0.122	4.712	4.834	LARGE
0.100	4.834	4.935	LARGE
0.071	4.935	5.006	SMALL
0.099	5.006	5.105	SMALL
0.099	5.105	5.205	SMALL
0.071	5.205	5.275	SMALL
0.100	5.275	5.376	LARGE
0.122	5.376	5.498	LARGE
0.122	5.498	5.620	LARGE
0.100	5.620	5.720	LARGE
0.071	5.720	5.791	SMALL
0.099	5.791	5.890	SMALL
0.099	5.890	5.990	SMALL
0.071	5.990	6.061	SMALL
0.100	6.061	6.161	LARGE
0.122	6.161	6.283	LARGE

I have assumed that the overlap regions are divided into two halves with equal phi extent, with one half assigned to a LARGE sector and the other half assigned to a SMALL sector.
I like to hear whether these numbers are reasonable. Note that the numbers are not used for computing trigger efficiency, they are used for estimating RoI request rates and all holes are neglected for these estimates. The rates for some of the ROBs therefore will be somewhat too high, but I think this is acceptable given the uncertainties in the LVL1 trigger rates,

Regards, Jos

Mail from Stefan Tapprogge, 11 April 2002

From: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Date: Thu Apr 11, 2002 07:07:52 PM Europe/Amsterdam
To: john.baines@rl.ac.uk
Cc: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: m_50 and m_95 times for scan

Hi John,
hi Jos,

unfortunately I have not yet been able to sort things out, and I will not be able to do so during the next days (i.e. before the T/DAQ week starts).

For time being, let's stick to the values Jos proposes for the muons and take the SCT/pixel numbers which we have for the electron selection.

Cheers,

Stefan.

Mail from John Baines, 12 April 2002

From: john.baines@rl.ac.uk
Date: Fri Apr 12, 2002 07:25:25 PM Europe/Amsterdam
To: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Cc: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: m_50 and m_95 times for scan
Reply-To: john.baines@rl.ac.uk

Hi Jos,
I have done some benchmarking for EM RoI in jet events (this is what we are also using as the time for MU RoI). I have made timings for jets with and without pile-up at design lumi and interpolated to get the time for 2x10^33, which I think is the same proceedure Stefan used.

Execution times (ms) @ 160 SI95

	jets (low lumi) mean (m_50 / m_95)	jetmb (10^34) mean (m_50 / m_95)	estimate 2x10^33 mean (m_50 / m_95)
pix + sctKalman	0.61 (0.48 / 1.03)	2.38 (1.98 / 5.02)	0.80 (0.65 / 1.48)
sctHough	0.28 (0.25 / 0.52)	1.20 (1.07 / 1.94)	0.38 (0.34 / 0.68)
TRT-LUT	0.77 (0.83 / 1.11)	1.05 (1.12 / 1.35)	0.80 (0.87 / 1.13)
SctKalman	0.63 (0.43 / 1.83)	3.03 (2.57 / 6.66)	0.90 (0.66 / 2.37)
xKalman-TRT	2.74 (2.55 / 4.20)	8.27 (6.85 /16.37)	3.35 (3.03 / 5.55)
SctKalman	0.47 (0.33 / 1.36)	1.37 (1.08 / 2.96)	0.56 (0.41 / 1.54)

sctHough in ATRIG is what was used for Saul's original timings (the numbers Stefan gave you). My timings are from sctHough in CTrig and include the data-preparation step. I don't have a time for TBTREC (which was used for Saul's timings) but I have instead times for TRT-LUT and the TRT part of xKalman. Again timings from CTrig. I've also given the time for pix-scan(within RoI)+sctkalman for comparison.

So my conclusion is that the mean and m_50 values are indead quite

similar. My timings for sctHough are more-or-less in agreement with the
times Stefan gave you, and the times for xKalman-TRT arn't that
different from the TRT times you have, so probably no need to
change the numbers you are using.

However there is obviously potential for significant saving if we use
TRT-LUT, but the performance is as yet un-proven.

Cheers
John

Mail sent to John Baines, 15 April 2002

Date: Mon, 15 Apr 2002 12:09:45 +0200
Subject: FEX processing times for Inner Detector
Content-Type: text/plain; charset=US-ASCII; format=flowed
Mime-Version: 1.0 (Apple Message framework v481)
Cc: Stefan.Tapprogge@cern.ch, Jos Vermeulen <i73@nikhef.nl>
To: john.baines@rl.ac.uk

Hi John,

Thanks for the new numbers. For the TRT they differ considerably from
the numbers Stefan provided:

low: 2.74/4.20 (in stead of 3.44/23.33 ms
2*10**33 : 3.35/5.55 in stead of 8.6/37.4 ms
design : 8.27/16.37 in stead of 50/150 ms

Numbers above are average/m_95 for your numbers, m_50/m_95 for Stefans
numbers. Although you concluded that the times for xKalman-TRT are about
the same as the TRT times specified by Sefan, I think that the numbers
above show that there is a considerable difference

For the PIxels and the SCT, we get for pix + sctKalman (sctHough):

low: 0.61/1.03 (0.28/ 0.52) in stead of 0.28 / 0.64
2*10**33 : 0.80/1.48 (0.34/ 0.75) in stead of 0.38/0.68
design : 2.38/5.02 (0.78/1.59) in stead of 1.2/1.94

We also need to take into account processing of the Pixels data, which I
guess is not included in sctHough.

For me it is not entirely clear now what I should use for the paper
model. It is maybe best to use the worst case numbers, i.e the results
for the "pix + sctKalman" and xKalman-TRT" algorithms:

Pixels and SCT:
low: mean = 0.61, m_95 = 1.03
2*10**33 : mean = 0.80, m_95 = 1.48
design : mean = 2.38, m_95 = 5.02

TRT FEX

low: mean = 2.74, m_95 = 4.20
2*10**33 : mean = 3.35, m_95 = 5.55
design : mean = 8.27 m_95 = 16.37

Would you agree with this approach ?

Cheers, Jos

Mail from John Baines, 15 April 2002

From: john.baines@rl.ac.uk
Date: Mon Apr 15, 2002 02:33:46 PM Europe/Amsterdam
To: Jos Vermeulen <i73@nikhef.nl>
Cc: john.baines@rl.ac.uk, Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Subject: Re: FEX processing times for Inner Detector
Reply-To: john.baines@rl.ac.uk
Hi Jos,

> Although you concluded that the times for xKalman-TRT are about
> the same as the TRT times specified by Sefan, I think that the numbers
> above show that there is a considerable difference

Sorry, I was concentrating on low lumi, you are right that for high lumi the TBTREC times are very much higher than the numbers I gave.
A summary of all the TRT measurements is as follows:

	10^33	10^34	
TBTREC/Atrig	3.88/26.25	50/150	(m_50/m_95 Saul)
CTrig/TRT-LUT (no o/head)	0.23/0.61	0.43/0.78	(m_50/m_95 Saul)
CTrig/TRT-LUT (no o/head)	0.29/0.48	0.47/0.72	(mean/m_95 John)
CTrig/TRT-LUT (incl. o/head)	0.77/1.11	1.05/1.35	(mean/m_95 John)
CTrig/XKalmanTRT	2.74/4.20	8.27/16.37	(mean/m_95 John)

(the overheads are the times taken to convert event input and output data from one format to another. I have been conservative and included this time.)
Now the problem is that only TBTREC has been used for physics performance studies, which is why Saul has, very reasonably, used this algorithm for the baseline TRT execution time. However I suspect that to use the 50ms time for Atrig/TBTREC at high lumi is to be overly conservative. To take the CTrig/XKalmanTRT times seems reasonable to me, but its a difficult decision in the absence of the corresponding physics performance results.

> For the PIxels and the SCT, we get for pix + sctKalman (sctHough):
>
> low: 0.61/1.03 (0.28/ 0.52) in stead of 0.28 / 0.64
> 2*10**33 : 0.80/1.48 (0.34/ 0.75) in stead of 0.38/0.68
> design : 2.38/5.02 (0.78/1.59) in stead of 1.2/1.94

SctHough INCLUDES the pixels and is the algorithm used for all the e/gamma trigger studies, so I think this is what we should use as the SCT+Pixel baseline for the RoI triggers. In general my times for CTrig/SctHough are fairly consistant with Saul's, but its true that the CTrig times are significantly faster at design lumi. It seems resonable to me to take the CTrig times.

So my prejudice would be for :

	MU, EM and TAU RoI		
	10^33	2x 10^34	10^34
	mean/m_95	mean/m_95	mean/m_95
CTrig/TRT-LUT (incl. o/head)	0.77/1.11	0.80/1.13	1.05/1.35
CTrig/Sct-Hough(SCT & Pixels)	0.28/0.52	0.38/0.68	1.20/1.94

As we discussed, the muon FEX would be 0.2/0.35 at all luminosities.

For the SCT & Pixel full-scan:

	10^33	2x 10^34
	mean/m_95	mean/m_95
CTrig/idscan (SCT+pixels)	11.2/22.5	22.4/45.0

I have chosen to change to the idscan algorithm as the baseline since the pixel-scan algorithm time contains a speed-up factor assumed for part of the code, whereas this is the actual time measure for the idscan. In any case, the times are very similar to what I sent you before for the pixel-scan.

For the TRT-guided full-scan (pT>0.5) (same numbers as in earlier email):

	10^33 mean/m_95	2x 10^34 mean/m_95
CTrig/TRT-LUT(pT>0.5)	26.85/51.92	53.70/103.84
CTrig/sctkalman	51.20/149.36	102.40/298.72

For the TRT-guided full-scan (pT>1.5) (same numbers as in earlier email):

	10^33 mean/m_95	2x 10^34 mean/m_95
CTrig/TRT-LUT(pT>1.5)	12.10/22.38	24.20/44.76
CTrig/sctkalman	24.31/72.30	48.62/144.60

Mail from John Baines, 16 April 2002

Date: Mon, 15 Apr 2002 16:55:44 +0100 (BST)
From: john.baines@rl.ac.uk
X-X-Sender: jtb@heplnx3.pp.rl.ac.uk
Reply-To: john.baines@rl.ac.uk
To: Jos Vermeulen <i73@nikhef.nl>
cc: Stefan Tapprogge <Stefan.Tapprogge@cern.ch>
Subject: Re: FEX processing times for Inner Detector

Hi Jos,

> > As we discussed, the muon FEX would be 0.2/0.35 at all luminosities.
>
> I thought that we decided on 0.2/0.25 (3/2 * 0.13 / 4/3 * 0.19) ?

Sorry, that was a mistype, your numbers are correct

Cheers
John

Mail from John Baines, 16 April 2002

From: john.baines@rl.ac.uk
Date: Tue Apr 16, 2002 02:15:30 PM Europe/Amsterdam
To: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: FEX processing times for Inner Detector
Reply-To: john.baines@rl.ac.uk

Hi Jos,

> 1) for the TRT you suggested that it would be reasonable to use the
> numbers for CTrig/xKalmanTRT, but you specified the CTrig/TRT-LUT values
> as your prejudice. I am inclined to use the CTrig/xKalmanTRT times. Is
> this OK ?

Sorry, I cut n' pasted the wrong line. I agree CTrig/xKalmanTRT is better since this is similar to the TBTREC algorithm that was used for the physics performance measurements, whilst we have no experience of TRT-LUT for electron ID.

> 2) the CTrig/sctkalman times (associated with the TRT scan) are longer
> than the times specified for SCT & pixel scan. Is this reasonable ?

Yes. You get 3-4 times more seed tracks from xkalman-TRT than
from the pix-scan for the same pT threshold. This is partly due to
material in the SCT and pixels causing secondaries and partly to the
fact that the TRT is in 4 separate parts. It is time-consuming to
link together track-segments from different parts, so this has not been
done. It may be we can get some improvement in the future.

I have put tables of the full scan timing results on:

http://hepunx.rl.ac.uk/atlasuk/simulation/level2/Bphys/bmark.html

Cheers
John

Mail from John Baines, 17 April 2002

From: john.baines@rl.ac.uk
Date: Wed Apr 17, 2002 11:55:35 AM Europe/Amsterdam
To: ATLAS B physics trigger mailing list <atlas-trig-bphysics@atlas-lb.cern.ch>
Cc: Jos Vermeulen <i73@nikhef.nl>, Monika Wielers <Monika.Wielers@cern.ch>
Subject: Update to Benchmarking web page
Reply-To: john.baines@rl.ac.uk

Hi,

I've updated the B-trigger benchmarking web page. I have redone all
the timings to measure mean, median (m_50) and 95% (m_95) values. The
results are on:

http://hepunx.rl.ac.uk/atlasuk/simulation/level2/Bphys/bmark.html

In order to get estimates of the execution time for muon confirmation
in muon RoI, I also ran benchmarking on 0.2x0.2 EM RoI in jet events. The
results are on:

http://hepunx.rl.ac.uk/atlasuk/simulation/level2/eid/bmark.html

These should be treated with care since only the SCT-Hough algorithm
has been used in RoI guided physics performance studies.

There are a lot of numbers there, and it is possible I have mistyped.
Please let me know if you notice anything that looks suspicious.

Cheers
John

Mail from Aleandro Nisati, 18 April 2002

```
From: Aleandro Nisati <Aleandro.Nisati@cern.ch>
Date: Thu Apr 18, 2002  01:58:57 PM Europe/Amsterdam
To: Jos Vermeulen <i73@nikhef.nl>
Subject: Re: Copy of returned mail concerning barrel LVL1 mu trigger


Dear Jos,

first of all please accept my apologies to react so late to your message.
We are very busy these days in Rome becasue we are preparing all the
hardware/software needed to test our Asics use for the level1 muon
trigger in the barrel. This explain why none of my colleagues in rome
has reacted.
So far, we haven't re-evaluated the RoIs dimensions: we will do it
as soon as possible; in the meantime please use the values you extracted
from the web pages.
As soon as we will have the new RoI grid we will let you know.

Cheers Leandro
```