

# OPTIONS FOR THE ATLAS LEVEL-2 TRIGGER

## ATLAS Level-2 Trigger Groups <sup>1</sup>

*Argonne National Laboratory, USA <sup>1</sup>; Bucharest IAP, Romania <sup>2</sup>  
CERN, Geneva, Switzerland <sup>3</sup>; Niels Bohr Institute, Copenhagen, Denmark <sup>4</sup>  
Henryk Niewodniczanski Institute of Nuclear Physics, Cracow, Poland <sup>5</sup>  
Dubna JINR, Russia <sup>6</sup>; University of Edinburgh, U.K. <sup>7</sup>  
University of Innsbruck, Austria <sup>8</sup>; University of California at Irvine, USA <sup>9</sup>  
Universität Jena, Germany <sup>a</sup>; Università di Lecce, Italy <sup>b</sup>  
University of Liverpool, U.K. <sup>c</sup>; University College London, U.K. <sup>d</sup>  
University of Manchester, U.K. <sup>e</sup>; Universität Mannheim, Germany <sup>f</sup>  
Marseille C.P.P.M, France <sup>g</sup>; Michigan State University, USA <sup>h</sup>  
Moscow State University, Russia <sup>i</sup>  
NIKHEF, Amsterdam, Netherlands <sup>j</sup>; University of Oklahoma, USA <sup>k</sup>  
Osaka University, Japan <sup>l</sup>; Oxford University, U.K. <sup>m</sup>  
Institute of Computer Sciences, Prague, Czech Republic <sup>n</sup>  
Universidade Federal do Rio de Janeiro, Brazil <sup>o</sup>  
Università di Roma 'La Sapienza', Italy <sup>p</sup>  
Royal Holloway and Bedford New College, University of London, U.K. <sup>q</sup>  
Rutherford Appleton Laboratory, U.K. <sup>r</sup>  
DAPNIA, CEA Saclay, France <sup>s</sup>; Universidad de Valencia, Spain <sup>t</sup>  
Weizmann Institute of Science, Rehovot, Israel <sup>u</sup>; University of Wisconsin, USA <sup>v</sup>*

---

<sup>1</sup> R. Blair, J. Dawson, J. Schlereth; <sup>2</sup> M. Caprini, F. Constantin; <sup>3</sup> R. Bock, A. Bogaerts, E. Denes, R.W. Dobinson, D. Francis, S. Haas, R. Hauser, R. Heeley, M. Liebhart, B. Martin, R. McLaren, P. Werner; <sup>4</sup> H. Bertelsen, M. Dam, J.R. Hansen, J.D. Hansen, B. Rensch; <sup>5</sup> Z. Hajduk, K. Korcyl; <sup>6</sup> I. Alexandrov, V. Kotov; <sup>7</sup> O. Boyle; <sup>8</sup> C. Hortnagl, D. Kuhn; <sup>9</sup> A. Lankford; <sup>a</sup> V. Dörsing, W. Erhard, P. Kammel, A. Reinsch; <sup>b</sup> P. Creti, O. Palamara, S. Petrera; <sup>c</sup> P. Maley; <sup>d</sup> R. Cranfield, G. Crone, N. Woolley; <sup>e</sup> R. Hughes-Jones, S. Kolya, R. Marshall, D. Mercer; <sup>f</sup> H. Högl, K. Kornmesser, A. Kugel, J. Ludvig, R. Männer, K.-H. Noffz, S. Rühl, M. Sessler, H. Simmler, H. Singpiel, R. Zoz; <sup>g</sup> F. Etienne, D. Laugier, Z. Qian, C. Rondot, F. Touchard; <sup>h</sup> M. Abolins, Y. Ermoline, D. Owen, B. Pope; <sup>i</sup> F. Rizatdinova, S. Sivoklov; <sup>j</sup> R.J. Dankers, J.C. Vermeulen; <sup>k</sup> P. Skubic; <sup>l</sup> M. Nomachi; <sup>m</sup> F. Harris, S. Hunt; <sup>n</sup> F. Hakl, M. Jirina, S. Rizek; <sup>o</sup> B. Kastrup, J. Seixas, R. Weber dos Santos; <sup>p</sup> A. Di Mattia, S. Falciano, L. Luminari, M. Mascagni, A. Nisati, L. Zanello; <sup>q</sup> A. Belias, M. Dobson, S. George, B. Green, W. Li, N. Madsen, J. A. Strong; <sup>r</sup> J. Baines, D. Botterill, R. Middleton, F. Wickens; <sup>s</sup> 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, B. Thooris; <sup>t</sup> A. Ferrer, V. Gonzalez, J.M. Lopez-Amengual, E. Sanchis; <sup>u</sup> L. Levinson; <sup>v</sup> D. Fasching.

*Contact: J.R. Hubbard, SPP/DAPNIA, CEA Saclay, 91191 Gif-sur-Yvette, France*

This paper describes options under study for the ATLAS Level-2 trigger, based on commercial switches, general-purpose processor farms, and, in certain cases, fast FPGA processors. The demonstrator program designed to evaluate the various options will be described, and preliminary results will be presented.

*Key words:* Trigger architectures; Large systems; Switches

## 1 Introduction

The high rate of interactions in future LHC experiments places stringent demands on the trigger and data acquisition systems. The ATLAS experiment uses a three-level trigger [1]. Level 1 is based on special-purpose processors designed to reduce the trigger rate from the 40 MHz beam-crossing rate to below 100 kHz. After a Level-1 ACCEPT, all event data (1 MB per event, or up to 100 GB/s) are sent to readout buffers for temporary storage. Higher-level triggers are required to reduce the data flow to about 100 MB/s (100 Hz with full data) for permanent storage. These higher-level triggers will be implemented, if possible, using general-purpose processors and commercial switching networks. The volume of data transferred to the Level-2 processors is limited to regions of interest (RoIs) defined by the first-level trigger. The full event data (about 1 MB per event) is transferred to the Level-3 processors for events accepted at Level 2. This paper presents options for the ATLAS Level-2 trigger.

## 2 Demonstrator program

Level 2 presents many unresolved issues: push or pull architecture; separate or combined data and control networks; the use of parallelism within events; strategies to reduce the network load; control of sub-farms. These issues are under study in a demonstrator program centered around three Level-2 architectures (Fig. 1). Architectures A and B are optimized for parallel processing, with local processors extracting features in each of the detector systems and a global processor farm for the final Level-2 decision [1]; architecture A uses data-driven FPGA processors for fast local feature extraction (the *Enable++* implementation uses 24 highly-complex FPGAs with 12 MB of fast SRAM [2]); architecture B uses general-purpose processor farms. Architecture C is optimized for sequential processing, with all readout buffers connected through a switch to a single (global) Level-2 processor farm [3]. Hybrid architectures combining aspects of architectures A, B, and C are also under consideration. The demonstrator program will include tests with three types of switching

networks: ATM, DS-link, and SCI.

In all architectures, a supervisor receives trigger information from Level 1 and extracts global information on the trigger type as well as the characteristics of each RoI: RoI type ( $\mu, e/\gamma, hadron/\tau, jet$ ), thresholds passed, and position in  $\eta, \phi$ . The supervisor assigns processors to the event as needed - local and global processors for architecture B or a single processor for architectures A and C. The RoI information is transferred either to the readout buffers, for the ‘push’ data flow used in architectures A and B, or to the Level-2 processors, for the ‘pull’ data flow used by architecture C. Independent of the architecture, the final Level-2 decision (ACCEPT or REJECT) is sent back to the supervisor by the global Level-2 processors. Every 100 events or so, the supervisor sends a set of Level-2 decisions to the readout buffers. Rejected events are removed from the readout buffers, releasing the paged memory for new events. Accepted events are sent to the Level-3 processors.

The Level-2 system must satisfy present ATLAS physics requirements, while maintaining flexibility for unexpected future requirements. At the design luminosity ( $10^{34}/cm^2/s$ ) the trigger will be entirely guided by high-rate Level-1 RoIs, but B physics studies at the initial, lower luminosities require the entire tracking volume to be scanned for low- $P_t$  tracks, without Level-1 RoIs for guidance. The B-physics algorithm takes considerable time on general-purpose processors. This can be reduced by using fast FPGA processors for the initial track-finding in the Transition Radiation Tracker (TRT full scan). In one of the hybrid architectures to be tested (Fig. 1), a B physics candidate is first selected by a global Level-2 processor, then the TRT full scan is performed on an *Enable++* processor, and, finally, for selected events, precision tracking is performed by the global processor, leading to the final Level-2 decision.

Investigations also include optimization between Level 2 and Level 3, since some complex algorithms, such as the tagging of b-jets, may be more economically performed at Level 2, if the architecture allows direct access to data fragments. These algorithms require very limited additional data transfer if they are performed at Level 2; algorithms performed at Level 3 have greater precision, but they are slower, and they require transfer of the full event data.

The small-scale demonstrators will measure the technological limits of each of the critical system elements, such as processor allocation and multi-task operation, data collection, and switch performance. Scalability will be studied on a 1024-node DS-link emulator, on which the different architectures can be emulated [4]. Final system performance and scalability will be evaluated using modelling studies [5]. All of these studies will be used to produce improved cost estimates. Preliminary results for the various architectures studied will be presented.

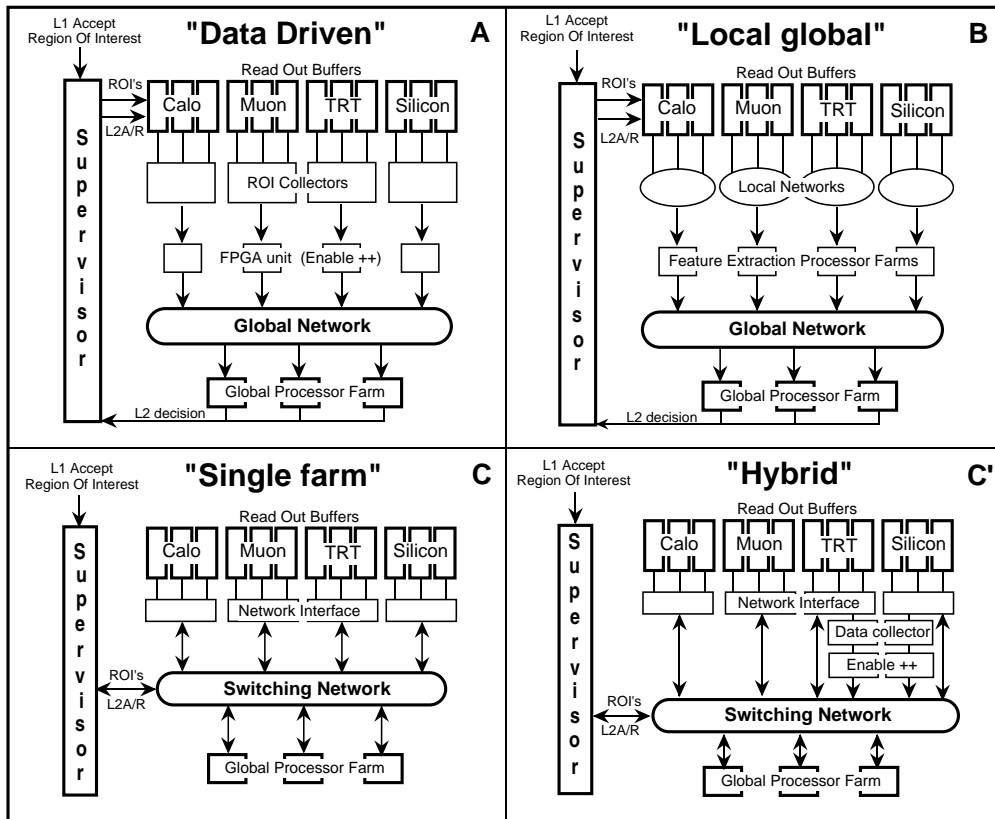


Fig. 1. Block diagrams for architecture A, with local data-driven FPGA processors and a global processor farm, architecture B, with local and global processor farms, architecture C, with a single switch and a single farm, and hybrid architecture C', with FPGA processors for the initial track finding.

## References

- [1] ATLAS Technical Proposal, CERN/LHCC/94-43, 15 December 1994.
- [2] K.-H. Noffz, et al., "Enable++: A General-Purpose L2 Trigger Processor"; Proc. IEEE Nucl. Sci. Symp., San Francisco, CA (1995) 667-671.
- [3] J. Bystricky, et al., "A Sequential Processing Strategy for the ATLAS Event Selection," to be published in IEEE Trans. on Nucl. Sci. Vol. 44, 3 June 1997.
- [4] R.W. Dobinson et al., "The MACRAME 1024 Node Switching Network: Status and Future Prospects," Proceedings of the 2nd International Data Acquisition Workshop on Networked Data Acquisition Systems (DAQ96), 13-15 November 1996, RCNP, Osaka, Japan
- [5] S. Hunt, et al., "SIMDAQ - a System for Modelling DAQ/trigger Systems," IEEE Trans. on Nucl. Sci. 43 (1996) 61.