

The DAQ/HLT system of the ATLAS experiment

André dos Anjos*

University of Wisconsin, Department of Physics, Madison, USA

E-mail: andre.dos.anjos@cern.ch

on behalf of the ATLAS TDAQ Community (see Appendix A)

The DAQ/HLT system of the ATLAS experiment at CERN, Switzerland, is being commissioned for first collisions in 2009. Presently, the system is composed of an already very large farm of computers that accounts for about one-third of its final event processing capacity. Event selection is conducted in two steps after the hardware-based Level-1 Trigger: a Level-2 Trigger processes detector data based on regions of interest (RoI) and an Event Filter operates on the full event data assembled by the Event Building system. The detector read out is fully commissioned and can be operated at its full design capacity. This places the responsibility on the High-Level Triggers system to select only events of highest physics interest that will finally reach the offline reconstruction farms.

This paper brings an overview of the current ATLAS DAQ/HLT implementation and performance based on studies originated from its operation with simulated, cosmic particles and first-beam data. Its built-in event processing parallelism is presented and discussed.

XII Advanced Computing and Analysis Techniques in Physics Research

November 3-7, 2008

Erice, Italy

*Speaker.

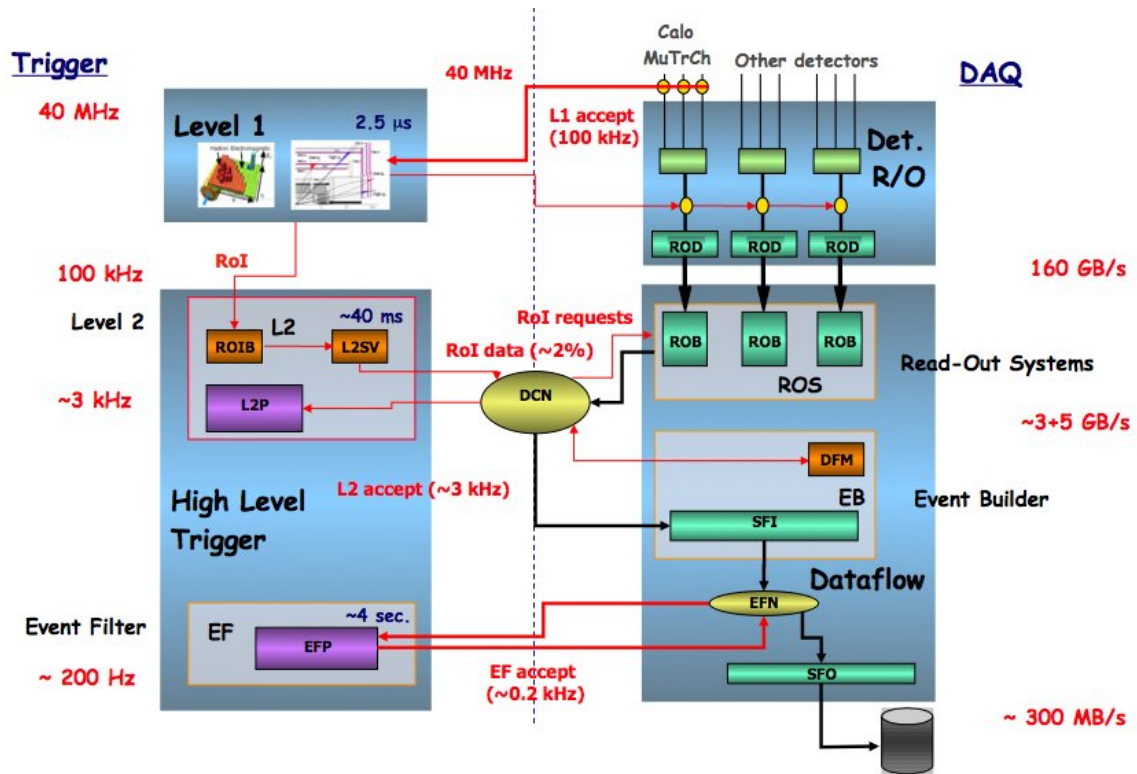


Figure 1: Overview of ATLAS Trigger and DAQ systems.

1. Introduction

The ATLAS experiment [1] at CERN, Switzerland counts on its Trigger and Data Acquisition (TDAQ) systems to select and store interesting physics reactions out of the Large Hadron Collider's (LHC) 14 TeV proton-proton bunch crossings, happening once every 25 ns (40 MHz). Each complete event read out in ATLAS is composed of roughly 1.6 Megabytes of data. Out of the large input rate only ~ 200 events per second can be recorded to permanent media, given the experiment's read out infrastructure (ROS) and offline computing capacity. An overview of the ATLAS TDAQ is depicted in Figure 1.

To achieve this design goal, the ATLAS Trigger is divided in three sub-levels. The First Level Trigger (L1) is built out of custom hardware, being capable of a rate reduction of three orders of magnitude. It uses signals from the ATLAS Calorimeters and Muon Detectors to reach up to 75 kHz (upgradeable to 100 kHz) output event rate. It does that by using simple scanning algorithms to detect high- p_T phenomena and to discard less interesting reactions. Such high- p_T objects have their region of interaction with the detector, or simply *Region of Interest* (RoI), annotated. In case of an accepted event, a summary of the L1 search known as the L1 Result (L1R) is passed to the High-Level Triggers system (HLT) and selection continues from this point. The L1 also triggers the read out system to buffer detector data while HLT refines its decision.

The HLT is composed of a Second Level Trigger (L2) and a third level or Event Filter (EF). By design, L2 delivers a maximum of 3 kHz output rate to EF of which only about 200 Hz survives

to tape. The L1 is connected to L2 through the means of a input gate known as the RoI Builder (RoIB). The RoIB has the task of fanning out the L1R to up to 20 L2 supervisors (L2SV) which, in turn, have access to up to 500 L2 Processors (L2Ps). Each L2P is equipped with the ATLAS Trigger Software and can perform selection, starting from the L1R and requesting RoI-centered data to the ROS for achieving a decision. Because only data around RoIs ($\sim 2\%$ of the total event data) is retrieved from the ROS, L2 is able to sustain such a high-input rate.

If the event is accepted by L2, a signal is sent to the Event Building system (EB) and data is collected from all ROS buffers and assembled before EF selection is performed. In case of a rejection, these aforementioned buffers are cleared and no event assembly happens. After event building, each EB processor known as Subfarm Input (SFI) waits for a transfer request from one of the 1800 EF processors (EFP).

Each EFP, like in L2, is equipped with the HLT selection software. Based on the refined search by L2, it is able to reduce the rate of events to be recorded by a factor of 10, writing about 300 Megabytes per second to permanent storage. Each EFP is connected to a few Subfarm Output (SFO) processors, which take events that have survived the selection process and puts them into files separated by stream types [2]. ATLAS stream types are set by the HLT processors at L2 or EF, additively.

2. System performance

Except for the connection between detector front-end and the L1/RoIB to the ATLAS DAQ/HLT, detector data processing and messaging are implemented by means of commodity computing and network. Currently, this system is composed of Gigabit Ethernet switches and x86 multi-core computers running Linux.

The DAQ/HLT software was written to benefit from the multi-core architectures [3]. Where possible, DAQ applications make use of multiple threads of control for increased efficiency. L2 and EF HLT applications though, process events in multiple independent processes inside L2Ps and EFPs respectively. L2 Processing Unit (L2PU) is the name given to each instance of the L2 HLT application running on a L2P. EF Processing Task (EFPT) is the equivalent for EFPs. The number of L2PUs run by L2P is set to match the number of cores available in each machine; the same is valid for EFPTs with respect to the EFPs.

As far as concerns connectivity, detector data traffic routing is implemented using 3 separated logical networks: L2, EB and EF. Communication between nodes is configurable and can be run using either standard UDP or TCP over IP. The computing nodes are only connected to the relevant networks. For example L2Ps connect to the ROS through the L2 Network; SFIs build events through the EB Network, but are also connected to the EF Network in order to serve events to EFPs. Because of the unknown optimum balance factor between L2 and EF processing capacity that minimizes dead-time, a set of HLT processors is connected to both L2 and EF networks and can be setup to act either as L2Ps or EFPs. These are called Exchangeable Processing Units (XPUs). At the present state, there about 850 XPUs, 63 SFIs, 5 SFOs and all the 154 ROS computers required to read the ~ 1600 buffers which are connected directly to the detector read out. XPUs available are dual quad-core architectures and have 16 Gigabytes of RAM available to accommodate the HLT selection software. The remaining machines are either dual double or single-core.

In order to test this setup, it is possible to preload simulated or pre-acquired data simultaneously at the ROS computers and the L2SVs. In this case, the latter operates in self-triggered mode feeding the available L2Ps as fast as L1Rs can be analyzed. Resetting the system configuration to take triggers from L1 and data from the detector instead of preloading is equally simple. These setups are used in the following tests and explained in more details below.

2.1 Testing with simulated data

Testing with simulated data is necessary to stress the system limits, since no beam with the required structure (25 ns bunch-crossings) is available. For the presented results, the existing XPU's were split to make up a complete DAQ/HLT system composed of 360 L2Ps and 310 EFPs. This amounts to about, respectively, 70% and 17% of the L2 and EF final farm sizes. About 60 SFIs were used for EB and 5 SFOs were assuring the flow of events from the EF to disk. The L2 farm was managed by 4 L2SVs, each managing in turn ~ 90 L2Ps.

This apparently unbalanced configuration for the HLT is possible taking into account the rejection ratios programmed into the L2 system, to not overflow the available EB bandwidth. L2 and EF are both loaded with a trigger menu that is optimized for a $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity. The simulated data preloaded contains a mixture of events providing both signal and background events which is compatible with the expected collider production. The average event size in the mixture is 800 Kilobytes, mainly due to the lack of pile-up simulation.

The plots in Figure 2 summarize a few important results for tests done with simulated data. Figure 2a shows the combined L2 farm throughput. This shows a constant throughput 60 kHz during many hours of operation. Figure 2b the equivalent plot for the total event building throughput. Taking into consideration the average size in the preloaded sample, this rate is equivalent to about 3.5 Gigabytes per second bandwidth utilization. The spike in the rates, which is noticeable at around 1:30 a.m. at the plots, was due to disk-intensive cron jobs that were executed in parallel on all machines in the cluster. This problematic behavior has been fixed since these tests.

HLT thresholds were set to reject around 93% of the events in L2 and 62% of the events in the EF. The average execution time to reach a L2 reject is ~ 38 ms and ~ 125 ms for the EF as can be verified in Figures 2c and 2d. The average time it takes to accept events in L2 and EF are, respectively, ~ 136 ms and ~ 260 ms. The overall processing time for L2 and EF can be extrapolated from these numbers and are, respectively, ~ 45 ms and ~ 176 ms. These numbers are in agreement with design expectations of the ATLAS TDAQ. The low latency for EF event processing can be explained by the low luminosity condition in which the simulated events were generated. With increasing luminosity and, by consequence, detector activity, it is expected that the latency of event processing in EF grows to match typical reconstruction timings, which are in the order of a few seconds.

2.2 Testing with cosmics

Testing with the whole ATLAS experiment infrastructure is possible using muons originated from cosmic ray interactions. In this mode, the ATLAS detector is configured for normal operations an L1 is adjusted for triggering using either its muon and/or calorimetry infrastructure, opening up a reasonable number of configuration possibilities. Random L1 triggers may be mixed in to test

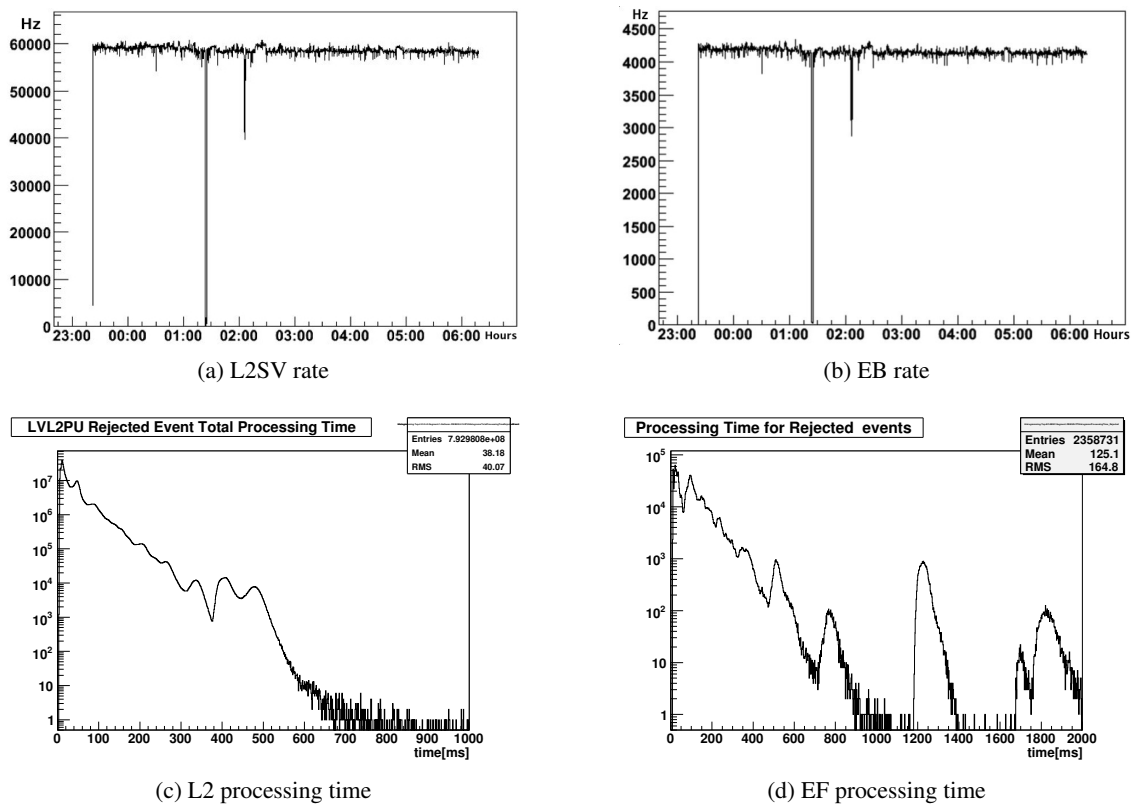


Figure 2: Results for timing tests using simulated data.

the trigger selectivity and stress the detector readout electronics. In this case, simple cosmic muon detection algorithms are placed into the HLT processors and used to discard random triggers.

The experiment has been operated in this condition for nearly 2 months without major problems. Figure 3 shows the number of events taken as a function of the time, in days. Each stream (trigger mode) is represented by a different color in such a plot. As can be verified, about 213 million events have been taken. The average size of these events is 2.1 Megabytes. The total size of data taken amounts to approximately 450 Terabytes.

3. The ATLAS HLT and multi-core technologies

3.1 Context

Event processing inside the ATLAS HLT is programmed sequentially for both L2 and EF. It is bootstrapped by the arrival of the precedent level's processing result and a set of feature extraction and hypothesis making algorithms is run one after the other to reach a decision as quickly as possible. The strategy is optimized for early rejection.

The development of the HLT infrastructure for ATLAS has seen major changes in the commodity computing industry with the development of multi-CPU (a.k.a. SMP) architectures and, since a few years, introduction of multi-core x86 CPUs. These phenomena have compensated for the so-called end of the *Frequency Scaling Era*, where programs created would just get faster by




Figure 3: Data streams during cosmics data taking.

waiting for the next generation of computer chips. In these new times, higher throughput should be obtained by simultaneous execution of many tasks. Because of the nature of triggering in High-Energy Physics, specially taking into consideration the sequential selection design by the ATLAS Trigger software, the most effective way to increase the event through by parallelization is by processing each event in a independent computer task.

3.2 Developments

Two main courses could be chosen taking into consideration the context described above, one which implements event processing in isolated processes (MP) and a second option that uses multi-threading (MT) to achieve parallelism. In the latter mode of operation, a single HLT process would be able to handle many events in parallel. Which to choose depends, as usual, on a number of trade-offs concerning application control limitations, configuration requirements and maintainability, just to mention a few.

Multi-threading allows the re-use of the current processing context (program variables) to implement synchronization and data transfer mechanisms, which makes programming easier and the resulting code typically faster. Unfortunately, the benefits of MT do not go much further. Presently there are very few tools to allow the detection of synchronization problems in MT programs and common knowledge in this area is very limited, which is of fundamental impact in large projects such as the ATLAS software.

Secondly, MT applications require that all code used in their context is *thread-safe*. A piece of code is thread-safe if it functions correctly during simultaneous execution by multiple threads, in particular considering multiple concurrent accesses to shared data. This is another drawback for large projects since all code produced and directly or indirectly used must be checked for conformity. In this context, an important aspect of the ATLAS Trigger design that comes into play

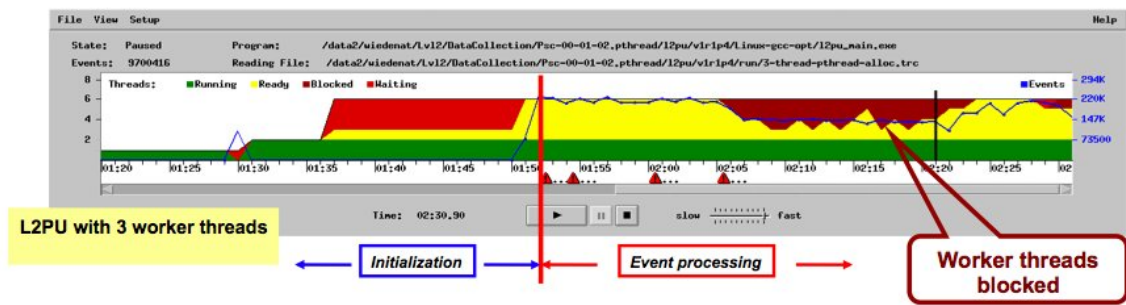


Figure 4: Thread inter-locking during event processing.

is code sharing. To simplify analysis and minimize maintenance [5], trigger and offline analysis share several tens of thousands lines of code, many of which were written during the Frequency-Scaling Era, when typical analysis applications were meant to be executed in a single-threaded way.

Still, thread safety does not satisfy a basic triggering condition which is time efficiency. Shared data may still be needed to pursue the processing of an event and its access may be blocked by other threads. This requirement has been called *Thread Efficiency*. Blocking concurrent threads for a resource represents a slow-down in the overall speed-up of the MT application leading to a sub-optimal result [6].

It is a laborious task to detect all points in a program where a potentially inefficient piece of code is being used. It may be even harder to fix it as it may lie on libraries that are just being used by routines to which the programmer may have no access to. Figure 4 shows an example of lock between concurrent threads for a L2PU executing HLT selection code in 3 parallel threads - each thread processing an event. The remaining threads (a total of 6 is shown in the figure) represent the main program loop and other monitoring tasks. The horizontal axis describes the time since the start of the application and the vertical axis the number of threads available. The different colors mark the different state of each thread in the application. The green color on the bottom of this graphic means those threads are running while the yellow section marks how many threads are ready to run, but are not running, either because of system contingency or simply because the process may be paused (e.g. waiting the start of run command). Dark red indicates the referred threads are blocked waiting for a shared resource, while light red means the thread is waiting for a resource, which is not necessarily shared among the other threads (for example, waiting for disk access). After the process has been started, only one thread was executing at each time, with the other two waiting for a shared resource. Such a resource was later identified as a global memory manager used by the GNU C++ Standard Template Library (STL) [3]. These types of problems are not exclusive to any STL implementation and do happen, many times, in quite unexpected ways.

Multi-processing can cure problems observed with MT by relaxing on thread safety and efficiency requirements, at the expense of more computing power. Taking in consideration current dual quad-core platforms, going MP means increasing by 8 the number of applications to control and configure. More generally, MP imposes scalability requirements on TDAQ infrastructure which must follow the computer chip industry trends. If data sharing is not implemented between appli-

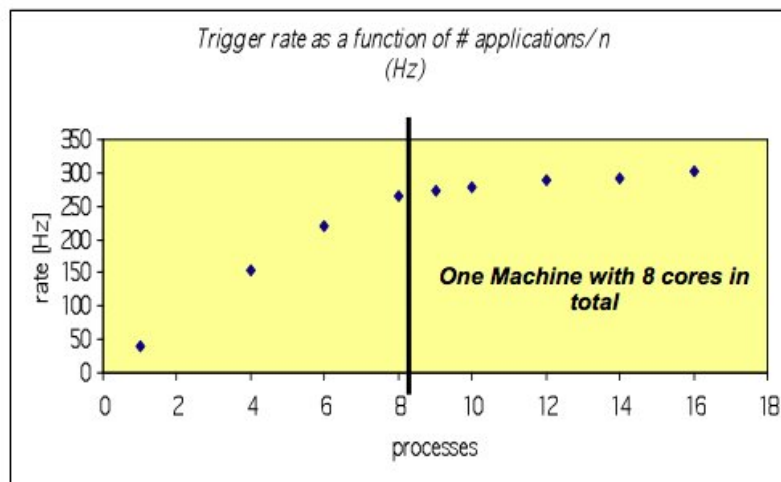


Figure 5: Speed-up of the ATLAS HLT as a function of the number of processes used in a single computing node.

cations executed in the same computing node, memory requirements are also expected to increase.

Because of the nature of the processing and the statistical independence of physics events flowing inside trigger systems, a perfect scaling is expected between the throughput achieved while executing code in a single task or multiple tasks excluded, of course, hardware limitations. Figure 5 shows the speed-up obtained by running the HLT code using an MP approach. The machine in question had 8 cores. As can be seen, the processing rate scales quite linearly as one increases the number of HLT processes running in the node. By running 9 or more processes the rate does not increase much more.

4. Conclusions

The ATLAS DAQ/HLT has reached the required maturity to take new physics, to be provided by LHC in mid-2009. It has been shown that this system can operate stably for many hours and fits perfectly under its design specifications. L2 can sustain up to 60 kHz which represents about 80% of its design capacity by counting only on 70% of the specified hardware. EB can assemble events stably at ~ 4 kHz for the same tests. Because of limiting interconnect conditions between the SFIs and the EF present in the system during these tests, only a fraction of the EB bandwidth could be utilized. Later tests with EF disconnected witnessed much higher EB bandwidth usage, above 6 Gigabytes/s (it was originally designed for 5 GB/s) with 60 SFIs, which represent only 60% of the expected hardware.

Tests with cosmics have been used to fine-tune interfaces between the DAQ/HLT, L1, the detector read out front-end and the offline analysis processing farms. Since the begin of operations in August 2008, nearly 1 Petabyte of data has been taken, 450 Terabytes from September 2008. These data were recorded by the SFO farm in nearly 400000 files. Streaming has been proven to work correctly and without major problems.

In particular for the HLT processing strategy, multi-processing is the current preferred solution and is deployed in both L2 and EF. For pure DAQ applications, like SFIs, L2SVs or SFOs, multi-threading is used to optimize data-acquisition to its maximum. The use of MP for HLT event processing is due to the lack of support tools for MT development, common expertise and support from software packages used in the project considering both thread safety and efficiency. MP has proven to scale well in current multi-core architectures while tests with the infrastructure have not shown major problems to control and configure the larger number of applications. Techniques for optimizing memory usage inside HLT processing nodes are currently under study.

A. The ATLAS TDAQ Community

Abolins M.,⁵⁶ Achenbach R.,³¹ Adorisio C.,¹⁵ Adragna P.,⁷⁶ Aharrouche M.,⁵⁷ Aielli G.,^{81,42} Al-Shabibi A.,³¹ Aleksandrov I.,²⁴ Alexandre G.,²⁶ Aloisio A.,^{68,38} Alviggi M.G.,^{68,38} Amorim A.,⁴⁹ Amram N.,⁸⁸ Anderson K.,¹⁶ Andrei V.,³¹ Anduaga X.,⁵⁰ Antonelli S.,^{10,36} Anulli F.,^{80,41} Aracena I.,⁸² Ask S.,⁵⁸ Åsman B.,⁸⁶ Avolio G.,⁹¹ Backlund S.,¹⁴ Badescu E.,¹¹ Baines J.,⁷⁷ Banerjee S.,⁹⁸ Barnett B.M.,⁷⁷ Bartoldus R.,⁸² Batraneanu S.,^{12,14} Battaglia A.,⁷ Bauss B.,⁵⁷ Beauchemin P.,⁷² Beck H.P.,⁷ Bee C.,⁵⁹ Behera P.,⁴⁵ Bell P.,⁵⁸ Bell W.H.,²⁸ Bellagamba L.,^{10,36} Bellomo M.,³⁹ Ben Ami S.,⁸⁷ Bendel M.,⁵⁷ Benhammou Y.,⁸⁸ Benslama K.,⁷⁸ Berge D.,¹⁴ Berger N.,⁴⁸ Bernius C.,⁹² Berry T.,³⁴ Bianco M.,^{52,37} Biau C.,⁹⁵ Biglietti M.,^{68,38} Bizjak I.,⁹² Blair R.E.,¹ Bogaerts A.,¹⁴ Bohm C.,⁸⁶ Boisvert V.,³⁴ Bold T.,^{91,20} Bondioli M.,⁹¹ Booth J.R.A.,⁸ Borer C.,⁷ Boscherini D.,^{10,36} Bosman M.,⁵ Boyd J.,¹⁴ Bracinik J.,⁸ Brawn I.P.,⁷⁷ Brelier B.,⁶² Bressler S.,⁸⁷ Brock R.,⁵⁶ Brooks W.K.,⁹⁷ Brunet S.,²¹ Bruni A.,^{10,36} Bruni G.,^{10,36} Bucci F.,²⁶ Buda S.,¹¹ Burckhart-Chromek D.,¹⁴ Buttar C.,²⁸ Camarri P.,^{81,42} Campanelli M.,⁹² Canale V.,^{68,38} Capasso L.,^{68,38} Caprini M.,¹¹ Caracinha D.,⁵³ Caramarcu C.,¹¹ Cardarelli R.,^{81,42} Carlino G.,³⁸ Casadei D.,⁶⁵ Casado P.,⁵ Cataldi G.,³⁷ Cerri A.,¹⁴ Charlton D.G.,⁸ Childers J.T.,³¹ Chiodini G.,^{52,37} Ciapetti G.,^{80,41} Cimino D.,⁴⁰ Ciobotaru M.,^{91,12} Coccaro A.,²⁷ Collins N.J.,⁸ Conde Muino P.,⁴⁹ Conventi F.,³⁸ Corradi M.,^{10,36} Corso-Radu A.,⁹¹ Costa M.J.,⁹⁶ Coura Torres R.,⁷⁹ Cranmer K.,⁶⁵ Crone G.,⁹² Crupi R.,^{52,37} Cuenca Almenar C.,⁹¹ Curtis C.J.,⁸ Czyczula Z.,⁶⁴ Dam M.,⁶⁴ Damazio D.,⁴ Darlea G.L.,¹⁴ Davis A.O.,⁷⁷ De Pedis D.,^{80,41} De Santo A.,³⁴ de Seixas J.M.,⁷⁹ DeAsmundis R.,³⁸ Degenhardt J.,⁷⁵ della Volpe D.,^{68,38} DellaPietra M.,^{68,38} Delsart P.-A.,⁶² Demers S.,⁸² Demirkoz B.,¹⁴ Di Mattia A.,⁵⁶ Di Nardo R.,^{81,42} Diaz M.,⁸⁴ DiCiaccio A.,^{81,42} Dionisi C.,^{80,41} Djilkibaev R.,⁶⁵ Dobson E.,⁷² Dobson M.,¹⁴ Dogaru M.,¹¹ dos Anjos A.,¹⁰⁰ Dotti A.,⁴⁰ Dova M.T.,⁵⁰ Drake G.,¹ Dufour M.-A.,⁶⁰ Eckweiler S.,⁵⁷ Ehrenfeld W.,^{30,21} Eifert T.,²⁶ Eisenhandler E.,⁷⁶ Ellis N.,¹⁴ Emelivanov D.,⁷⁷ Enoque Ferreira de Lima D.,⁷⁹ Ermoline Y.,⁵⁶ Etzion E.,⁸⁸ Falciano S.,⁴¹ Farthouat P.,¹⁴ Faulkner P.J.W.,⁸ Fedorko A.,¹⁴ Fedorko I.,¹⁴ Feng E.,¹⁶ Ferland J.,⁶² Ferrag S.,²⁸ Ferrari R.,⁷⁴ Ferrer M.L.,²⁵ Fischer G.,²¹ Flacher H.,¹⁴ Fleckner J.E.,⁵⁷ Flowerdew M.,⁵⁴ Föhlich F.,³¹ Fonseca-Martin T.,³⁴ Francis D.,¹⁴ Fratina S.,⁷⁵ Fukunaga C.,⁹⁰ Gadomski S.,²⁶ Gallacher M.P.,⁸ Garitaonandia Elejabarrieta H.,⁶⁹ Gaudio G.,⁷⁴ Gee C.N.P.,⁷⁷ George S.,³⁴ Ghete V.M.,⁴⁴ Giagu S.,^{80,41} Gillman A.R.,⁷⁷ Giusti P.,^{10,36} Goebel M.,²¹ Goncalo R.,³⁴ Gonçalo R.,³⁴ Gonzalez Pinto F.,¹³ Gonzalez Silva L.,¹³ Göringer C.,⁵⁷ Gorini B.,¹⁴ Gorini E.,^{52,37} Grabowska-Bold I.,^{91,20} Grancagnolo S.,³⁵ Green B.,³⁴ Groll M.,⁵⁷ Guida A.,^{52,37} Haas S.,¹⁴ Hadavand H.,⁸³ Hadley D.R.,⁸ Haller J.,^{30,21} Hamilton A.,²⁶ Hanke P.,³¹ Hansen J.R.,⁶⁴ Hasagawa S.,⁶⁷ Hasegawa Y.,⁸⁵ Hauschild M.,¹⁴ Hauser R.,⁵⁶ Hayakawa T.,⁴⁷ Head S.,⁵⁸ Heijboer A.,¹⁴ Heim S.,⁵⁶ Hellman S.,⁸⁶ Hershendorff A.,⁸⁷ Hid-

végi A.,⁸⁶ Hillier S.J.,⁸ Hirayama S.,⁸⁹ Hod N.,⁸⁸ Hori T.,⁴⁷ Hryn'ova T.,⁴⁸ Hughes-Jones R.,⁵⁸ Huston J.,⁵⁶ Iacobucci G.,^{68,38} Idarraga J.,⁶² Iengo P.,^{68,38} Igonkina O.,⁶⁹ Ikeno M.,⁴⁶ Ishikawa A.,⁴⁷ Ishino M.,⁸⁹ Iwasaki H.,⁴⁶ Izzo V.,⁶⁸ Jain V.,⁴³ Jez P.,⁶⁴ Jimenez Otero S.,¹⁴ Johansen M.,⁸⁶ Johns K.,³ Jones G.,⁵⁸ Joos M.,¹⁴ Kadosaka T.,⁴⁷ Kajomovitz E.,⁸⁷ Kalinowski A.,⁷⁸ Kama S.,²¹ Kanaya N.,⁸⁹ Kanega F.,⁸⁹ Kaushik V.,³ Kawagoe K.,⁴⁷ Kawamoto T.,⁸⁹ Kazarov A.,⁷³ Kehoe R.,⁸³ Kessoku K.,⁸⁹ Khorauli G.,⁷⁸ Kieft G.,⁶⁹ Kirk J.,⁷⁷ Kiyamura H.,⁴⁷ Klofver P.,¹⁴ Klous S.,⁶⁹ Kluge E.-E.,³¹ Kobayashi T.,⁸⁹ Koeneke K.,²¹ Kolos S.,^{91,73} Kono T.,¹⁴ Konoplich R.,⁶⁵ Konstantinidis N.,⁹² Korcyl K.,¹⁹ Kordas K.,⁷ Kotov V.,²⁴ Kowalewski R.V.,⁹⁸ Krasznahorkay A.,^{14,22} Kreisel A.,⁸⁸ Kubota T.,⁸⁹ Kugel A.,³² Kurashige H.,⁴⁷ Kwee R.,³⁵ Landon M.,⁷⁶ Lane J.,⁵⁸ Lankford A.J.,⁹¹ Leahu L.,^{12,14} Leahu M.,¹² LeCompte T.,¹ Ledroit F.,²⁹ Lehmann Miotto G.,¹⁴ Lei X.,³ Lelouch D.,⁹⁹ Lendermann V.,³¹ Levinson L.,⁹⁹ Li S.,²¹ Liberti B.,^{81,42} Lilley J.N.,⁸ Lim H.,¹ Linne-
mann J.T.,⁵⁶ Lohse T.,³⁵ Losada M.,⁹ Lowe A.,⁴³ Luci C.,^{80,41} Luminari L.,⁴¹ Lundberg J.,¹⁴ Lupu N.,⁸⁷ Mackeprang R.,¹⁴ Maettig S.,^{30,21} Mahboubi K.,³¹ Mahout G.,⁸ Maltrana D.,⁸⁴ Mamuzic J.,²¹ Männer R.,³² Mapelli L.,¹⁴ Marchese F.,^{81,41} Marino C.,⁴³ Martin B.,¹⁴ Martin B.T.,⁵⁶ Martyniuk A.,⁵⁸ Marzano F.,⁴¹ Masik J.,⁵⁸ Matsushita T.,⁴⁷ McMahon T.,³⁴ Mcpherson R.,⁹⁸ Medinnis M.,²¹ Meessen C.,⁵⁹ Meier K.,³¹ Meirosu C.,¹² Merola L.,^{68,38} Messina A.,¹⁴ Middleton R.P.,⁷⁷ Miken-
berg G.,⁹⁹ Mincer A.,⁶⁵ Mineev M.,²⁴ Misiejuk A.,³⁴ Moa T.,⁸⁶ Moenig K.,²¹ Monticelli F.,⁵⁰ Mora Herrera C.,²⁶ Moreno D.,⁹ Morettini P.,²⁷ Morris J.D.,⁷⁶ Müller F.,³¹ Murillo Garcia R.,⁹¹ Nagano K.,⁴⁶ Nagasaka Y.,³³ Nakatsuka H.,⁴⁷ Navarro G.A.,⁹ Negri A.,^{74,39} Nemethy P.,⁶⁵ Neusiedl A.,⁵⁷ Nisati A.,⁴¹ Niwa T.,⁴⁷ Nomoto H.,⁸⁹ Nozaki M.,⁴⁶ Nozicka M.,²¹ Nurse E.,⁹² Ochi A.,⁴⁷ Oda S.,⁸⁹ Ohm C.,⁸⁶ Okumura Y.,⁶⁷ Oltmann B.,⁵⁷ Olvito D.,⁷⁵ Omachi C.,⁴⁷ Osculati B.,²⁷ Oshita H.,⁸⁵ Os-
una C.,⁵ Owen M.A.,⁵⁸ Özcan E.,⁹² Ozone K.,⁴⁶ Padilla C.,⁵ Panes B.,⁸⁴ Panikashvili N.,⁶¹ Parodi F.,²⁷ Pasqualucci E.,⁴¹ Pastore F.,¹⁴ Patricelli S.,^{68,38} Pauly T.,¹⁴ Perantoni M.,⁷⁹ Perera V.J.O.,⁷⁷ Perez E.,⁵ Perez Reale V.,¹⁸ Perrino R.,^{52,37} Pessoa Lima Junior H.,⁷⁹ Petcu M.,¹¹ Petersen B.,¹⁴ Petersen J.,¹⁴ Petrolo E.,^{80,41} Piegaia R.,¹³ Pilcher J.,¹⁶ Pinzon G.,⁹ Policicchio A.,¹⁵ Polini A.,^{10,36} Pope B.G.,⁵⁶ Potter C.,⁶⁰ Prieur D.P.F.,⁷⁷ Primavera M.,³⁷ Prokoshin F.,⁹⁷ Qian W.,⁷⁷ Quinonez F.,⁸⁴ Rajagopalan S.,⁴ Ramos Dos Santos Neves R.,⁴⁹ Reinherz E.,⁸⁸ Reinsch A.,⁷⁰ Renkel P.,⁸³ Rescigno M.,⁴¹ Rieke S.,⁵⁷ Risler C.,³⁵ Riu I.,⁵ Robertson S.,⁶⁰ Roda C.,⁴⁰ Rodriguez D.,⁹ Rog-
riquez Y.,⁹ Roich A.,⁹⁹ Romeo G.,¹³ Rosati S.,^{80,41} Rühr F.,³¹ Ruiz-Martinez A.,⁹⁶ Ryabov Y.,⁷³ Ryan P.,⁵⁶ Saavedra A.,⁹³ Sakamoto H.,⁸⁹ Salamon A.,^{81,42} Salnikov A.,⁸² Salvatore D.,¹⁵ Salva-
tore F.,³⁴ Sankey D.P.C.,⁷⁷ Santamarina C.,⁶⁰ Santonico R.,^{81,42} Sarkisyan-Grinbaum E.,⁹⁴ Sasaki O.,⁴⁶ Savu D.,¹⁴ Scannicchio D.,¹⁴ Schäfer U.,⁵⁷ Schiavi C.,²⁷ Schlereth J.,¹ Schmitt K.,³¹ Schmitt S.,²¹ Scholtes I.,¹⁴ Schooltz D.,⁵⁶ Schroer N.,³² Schultz-Coulon H.-C.,³¹ Schvienhorst R.,⁵⁶ Scott W.,⁷⁷ Segura E.,⁵ Sekhniadze G.,³⁸ Sherman D.,¹⁴ Shimojima M.,⁶⁶ Sidoti A.,³⁵ Silverstein D.,⁸² Silverstein S.,⁸⁶ Sinev N.,⁷⁰ Siragusa G.,⁵⁷ Sivoklokov S.,⁶³ Sjoen R.,⁶ Sjölin J.,⁸⁶ Sloper J.E.,¹⁴ Smizanska M.,⁵¹ Soffer A.,⁸⁸ Soloviev I.,⁷³ Spagnolo S.,^{52,37} Spiwoks R.,¹⁴ Staley R.J.,⁸ Stamen R.,³¹ Stancu S.,^{91,12} Steinberg P.,⁴ Stelzer J.,²¹ Stockton M.C.,⁸ Stradling A.,¹⁰⁰ Straessner A.,²³ Strom D.,⁷⁰ Su D.,⁸² Sugaya Y.,⁷¹ Sugimoto T.,⁶⁷ Sushkov S.,⁵ Sutton M.,⁹² Suzuki Y.,⁴⁶ Szy-
mocha T.,¹⁹ Taffard A.,⁹¹ Takahashi Y.,⁶⁷ Takeda H.,⁴⁷ Takeshita T.,⁸⁵ Tamsett M.,³⁴ Tan C.L.A.,⁸ Tanaka S.,⁴⁶ Tapprogge S.,⁵⁷ Tarem S.,⁸⁷ Tarem Z.,⁸⁷ Taylor C.,⁹² Taylor R.P.,⁹⁸ Teixeira-Dias P.,³⁴ Thomas J.P.,⁸ Thompson P.D.,⁸ Tokushuku K.,⁴⁶ Tollefson K.,⁵⁶ Tomoto M.,⁶⁷ Topfel C.,⁷ Torrence E.,⁷⁰ Touchard F.,⁵⁹ Tremblet L.,¹⁴ Tripiana M.,⁵⁰ Tsuno S.,⁴⁶ Ünel G.,⁹¹ Urquijo P.,²⁶ Urrejola P.,⁸⁴ Usai G.,¹⁶ Vachon B.,⁶⁰ Vandelli W.,¹⁴ Vari R.,^{80,41} Vaz Gil Lopes L.,¹⁷ Veneziano

S.,^{80,41} Ventura A.,^{52,37} Vercesi V.,⁷⁴ Vermeulen J.C.,^{69,2} Von Der Schmitt J.,⁵⁵ Vorwerk V.,⁵ Wang M.,⁷⁸ Watkins P.M.,⁸ Watson A.T.,⁸ Weber P.,³¹ Weidberg T.,⁷² Wengler T.,⁵⁸ Werner P.,¹⁴ Werth M.,⁹¹ Wessels M.,³¹ Wheeler-Ellis S.,⁹¹ Whiteson D.,⁹¹ Wickens F.J.,⁷⁷ Wiedenmann W.,¹⁰⁰ Wielers M.,⁷⁷ Wildt M.,⁵⁷ Wilkens H.,¹⁴ Winklmeier F.,¹⁴ Wu S.-L.,¹⁰⁰ Wu X.,²⁶ Xella S.,⁶⁴ Yamazaki Y.,⁴⁷ Yang U.,⁵⁸ Yasu Y.,⁴⁶ Zanello L.,^{80,41} Zema F.,¹⁴ Zhang J.,¹ Zhao L.,⁶⁵ Zobernig H.,¹⁰⁰ zur Nedden M.³⁵

¹ Argonne National Laboratory, Argonne, Illinois

² University of Amsterdam, Amsterdam

³ University of Arizona, Tucson, Arizona

⁴ Brookhaven National Laboratory (BNL), Upton, New York

⁵ Institut de Física d'Altes Energies (IFAE), Universitat Autònoma de Barcelona, Bellaterra (Barcelona)

⁶ University of Bergen, Bergen

⁷ Laboratory for High Energy Physics, University of Bern, Bern

⁸ School of Physics and Astronomy, The University of Birmingham, Birmingham

⁹ Universidad Antonio Narino, Bogotá, Colómbia

¹⁰ Università di Bologna, Dipartimento di Fisica, Bologna

¹¹ National Institute of Physics and Nuclear Engineering, Bucharest

¹² UPB - 'Politehnica' University of Bucharest, Bucharest

¹³ Universidad de Buenos Aires, Buenos Aires

¹⁴ European Laboratory for Particle Physics (CERN), Geneva

¹⁵ Dipartimento di Fisica dell' Università della Calabria e I.N.F.N., Cosenza

¹⁶ University of Chicago, Enrico Fermi Institute, Chicago, Illinois

¹⁷ Laboratório de Instrumentação e Física Experimental de Partículas, Coimbra

¹⁸ Nevis Laboratory, Columbia University, Irvington, New York

¹⁹ Institute of Nuclear Physics, Polish Academy of Sciences, Cracow

²⁰ Faculty of Physics and Applied Computer Science, AGH University of Science and Technology, Cracow

²¹ Deutsches Elektronen-Synchrotron (DESY), Hamburg and Zeuthen

²² University of Debrecen

²³ Technical University Dresden, Dresden

²⁴ Joint Institute for Nuclear Research, Dubna

²⁵ Laboratori Nazionali di Frascati dell' I.N.F.N., Frascati

²⁶ Section de Physique, Université de Genève, Geneva

²⁷ Dipartimento di Fisica dell' Università di Genova e I.N.F.N., Genova

²⁸ Department of Physics and Astronomy, University of Glasgow, Glasgow

²⁹ Laboratoire de Physique Subatomique et de Cosmologie de Grenoble (LPSC), IN2P3-CNRS-Université Joseph Fourier, Grenoble

³⁰ University of Hamburg, Germany

³¹ Kirchhoff Institut für Physik, Universität Heidelberg, Heidelberg

³² ZITI Ruprecht-Karls-University Heidelberg

³³ Hiroshima Institute of Technology, Hiroshima

- ³⁴ Department of Physics, Royal Holloway, University of London, Egham
³⁵ Institut für Physik, Humboldt-Universität zu Berlin
³⁶ I.N.F.N. Bologna
³⁷ I.N.F.N. Lecce
³⁸ I.N.F.N. Napoli
³⁹ I.N.F.N. Pavia
⁴⁰ Dipartimento di Fisica dell' Università di Pisa e I.N.F.N., Pisa
⁴¹ I.N.F.N. Roma
⁴² I.N.F.N. Tor Vergata
⁴³ Indiana University, Bloomington, Indiana
⁴⁴ Institute for Astro- and Particle Physics, Leopold-Franzens-Universität Innsbruck
⁴⁵ Iowa State University, Ames, Iowa
⁴⁶ KEK, High Energy Accelerator Research Organisation, Tsukuba
⁴⁷ Kobe University, Kobe
⁴⁸ Laboratoire d'Annecy-le-Vieux de Physique des Particules (LAPP), IN2P3-CNRS, Annecy-le-Vieux
⁴⁹ Laboratório de Instrumentação e Física Experimental, Lisboa
⁵⁰ Universidad Nacional de La Plata, La Plata
⁵¹ Department of Physics, Lancaster University, Lancaster
⁵² Università degli Studi del Salento, Dipartimento di Fisica, Lecce
⁵³ University Católica-Figueira da Foz and University Nova de Lisboa, Lisbon
⁵⁴ University of Liverpool, Liverpool
⁵⁵ Max-Planck-Institut für Physik, München
⁵⁶ Michigan State University, Department of Physics and Astronomy, East Lansing, Michigan
⁵⁷ Institut für Physik, Universität Mainz, Mainz
⁵⁸ School of Physics and Astronomy, University of Manchester, Manchester
⁵⁹ Centre de Physique des Particules de Marseille, IN2P3-CNRS, Marseille
⁶⁰ Department of Physics, McGill University, Montreal
⁶¹ University of Michigan, Department of Physics, Ann Arbor, Michigan
⁶² University of Montreal, Montreal
⁶³ Moscow State University, Moscow
⁶⁴ Niels Bohr Institute, University of Copenhagen, Copenhagen
⁶⁵ Department of Physics, New York University, New York
⁶⁶ Nagasaki Institute of Applied Science, Nagasaki
⁶⁷ Nagoya University, Nagoya
⁶⁸ Università di Napoli 'Federico II', Dipartimento di Scienze Fisiche, Napoli
⁶⁹ Nikhef National Institute for Subatomic Physics, Amsterdam
⁷⁰ University of Oregon, Eugene, Oregon
⁷¹ Osaka University, Osaka
⁷² Department of Physics, Oxford University, Oxford
⁷³ Petersburg Nuclear Physics Institute (PNPI), St. Petersburg
⁷⁴ Università di Pavia, Dipartimento di Fisica Nucleare e Teorica, Pavia
⁷⁵ Department of Physics, University of Pennsylvania, Philadelphia, Pennsylvania

- ⁷⁶ Physics Department, Queen Mary, University of London, London
- ⁷⁷ STFC Rutherford Appleton Laboratory, Harwell Science and Innovation Campus, Didcot, Oxon
- ⁷⁸ University of Regina, Regina
- ⁷⁹ Universidade Federal do Rio de Janeiro, COPPE/EE/IF, Rio de Janeiro
- ⁸⁰ Università di Roma I 'La Sapienza', Dipartimento di Fisica, Roma
- ⁸¹ Dipartimento di Fisica dell' Università di Roma II 'Tor Vergata'
- ⁸² Stanford Linear Accelerator Center (SLAC), Stanford
- ⁸³ Department of Physics, Southern Methodist University, Dallas, Texas
- ⁸⁴ Pontificia Universidad Católica, Santiago
- ⁸⁵ Faculty of Science, Shinshu University, Matsumoto
- ⁸⁶ Fysikum, Stockholm University, Stockholm
- ⁸⁷ Department of Physics, Technion, Haifa
- ⁸⁸ School of Physics and Astronomy, Tel Aviv University, Tel Aviv
- ⁸⁹ International Center for Elementary Particle Physics, University of Tokyo, Tokyo
- ⁹⁰ Physics Department, Tokyo Metropolitan University, Tokyo
- ⁹¹ Department of Physics and Astronomy, University of California, Irvine, California
- ⁹² Department of Physics and Astronomy, University College London, London
- ⁹³ University of Sydney, Sydney
- ⁹⁴ Department of Physics, The University of Texas at Arlington, Arlington, Texas
- ⁹⁵ Université de Technologie de Belfort-Montbéliard
- ⁹⁶ Instituto de Física Corpuscular (IFIC) Universidad de Valencia
- ⁹⁷ Universidad Técnica Federico Santa María, Valparaíso
- ⁹⁸ University of Victoria, Victoria
- ⁹⁹ Department of Particle Physics, The Weizmann Institute of Science, Rehovot
- ¹⁰⁰ Department of Physics, University of Wisconsin, Madison, Wisconsin
- ¹⁰¹ Università di Napoli 'Parthenope'

References

- [1] The ATLAS Collaboration, G. Aad et al., *The ATLAS Experiment at the CERN Large Hadron Collider*, **2008** JINST 3 S08003
- [2] A. Battaglia et al, *The Data-Logging System of the Trigger and Data Acquisition for the ATLAS Experiment at CERN*, **2008**, Nuclear Science Symposium and Medical Imaging Conference
- [3] S. Gadomski et al, *Experience with multi-threaded C++ applications in the ATLAS dataflow software*, **2003**, Conference for Computing in High-Energy and Nuclear Physics
- [4] A. dos Anjos et al., *The Second Level Trigger of the ATLAS Experiment at CERN's LHC*, **2004**, IEEE Transaction on Nuclear Science, volume 51, number 3, pages 909-914
- [5] W. Wiedenmann et al., *Studies for a common selection software environment in ATLAS : from the Level-2 Trigger to the offline reconstruction*, **2004**, IEEE Transactions on Nuclear Science, volume 51, number 3, pages 915-920
- [6] A.S. Tanenbaum, *Modern Operating Systems*, **1992**, Prentice Hall