



CERN-LHCC-2013-018
ATLAS-TDR-023-2013

ATLAS

TDAQ System

Phase-I Upgrade
Technical Design Report



Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System

The ATLAS Collaboration

Issue: 1
Reference: CERN-LHCC-2013-018
ATLAS-TDR-023
Created: 21 September 2013
Last modified: 30 November 2013
Prepared by: ATLAS Collaboration

Executive Summary

The Phase-I upgrade of the ATLAS Trigger and Data Acquisition (TDAQ) system will allow the ATLAS experiment to efficiently trigger and record data at instantaneous luminosities that are up to three times that of the original LHC design while maintaining trigger thresholds close to those used in the initial run of the LHC. New Level-1 calorimeter feature extraction processors will be incorporated to allow finer granularity data from the Liquid Argon (LAr) Calorimeter to be used to improve electron, photon, and tau selection; more sophisticated and larger-area algorithms to be used to improve jet selection; and improved pile-up corrections to be used for missing momentum reconstruction. The finer granularity data will be optically transmitted from new dedicated LAr Calorimeter hardware which is described in a separate technical design report. The Phase-I TDAQ upgrade will also benefit from the construction of the New Small Wheel (NSW), described in a separate technical design report. The new signals from the NSW will be included in the Level-1 muon endcap trigger, significantly reducing the overall rate by rejecting a large fraction of fake triggers. In addition, signals from the outer layer of the extended barrel of the Tile Calorimeter will be made available to the Level-1 muon endcap trigger for reducing the fake trigger rate in the overlap region between barrel and endcap in the Muon Spectrometer. The upgraded Level-1 muon trigger electronics will provide input to the new Muon-to-Central-Trigger-Processor Interface (MUCTPI). The new calorimeter feature extraction processors and the new MUCTPI will produce object quantities (e.g. η , ϕ and momentum) that can be combined with each other as well as with signals from the existing calorimeter trigger in a new topological processor. The Central Trigger Processor will be expanded to allow up to 512 distinct triggers. The Data Acquisition and the High-Level Trigger (HLT) processing farm will be upgraded to allow full calorimetry information on a large fraction of the 100 kHz of accepted Level-1 events to be read out and processed. The processing of the HLT will be complemented by the new Fast Tracker (FTK), described in a separate technical design report, that provides full track reconstruction at the start of the HLT selection process at the full Level-1 accept rate. Improved dataflow software, and improved and refined HLT selection algorithms will allow the 100 kHz Level-1 accept rate to be reduced to less than 1 kHz for recording.

The ATLAS Collaboration

Argentina

Universidad de Buenos Aires, Buenos Aires

Universidad Nacional de La Plata, La Plata

Armenia

Yerevan Physics Institute, Yerevan

Australia

University of Adelaide, Adelaide

Research Centre for High Energy Physics, Melbourne University, Melbourne

University of Sydney, School of Physics, Sydney

Austria

Institut für Astro- und Teilchenphysik, University of Innsbruck, Innsbruck

Fachhochschule Wiener Neustadt (FHWN), Wiener Neustadt

Azerbaijan Republic

Institute of Physics, Azerbaijan Academy of Sciences, Baku

Republic of Belarus

Institute of Physics, National Academy of Sciences of Belarus, Minsk

National Centre for Particle & High Energy Physics, Minsk

Brazil

Universidade Federal de Juiz de Fora, Universidade Federal do Rio De Janeiro, COPPE/EE/IF, Rio de Janeiro, Universidade Federal de Sao Joao del Rei and Universidade de Sao Paulo

Canada

University of Alberta, Edmonton

University of Carleton, Carleton

University of Montreal, Group of Particle Physics, Montreal, Quebec

Department of Physics, McGill University, Montreal

Simon Fraser University, Burnaby, BC

Department of Physics, University of Toronto, Toronto

TRIUMF, Vancouver and York University, Toronto

Department of Physics, University of British Columbia, Vancouver

University of Victoria, Victoria

CERN

European Laboratory for Particle Physics (CERN), Geneva

Chile

Joint team from Pontificia Universidad Católica de Chile, Santiago and Universidad Técnica Federico Santa María, Valparaíso

China

Chinese cluster formed by IHEP Beijing, Nanjing, Shandong, Shanghai Jiao Tong and Hefei

Colombia

Universidad Antonio Narino (UAN), Bogotá

Czech Republic

Palacký University, Olomouc

Academy of Sciences of the Czech Republic, Institute of Physics and Institute of Computer Science, Prague

Charles University in Prague, Faculty of Mathematics and Physics, Prague

Czech Technical University in Prague, Faculty of Nuclear Sciences and Physical Engineering, Faculty of Mechanical Engineering, Prague

Denmark

Niels Bohr Institute, University of Copenhagen, Copenhagen

France

Laboratoire d'Annecy-le-Vieux de Physique de Particules (LAPP), CNRS-IN2P3, Annecy-le-Vieux

Laboratoire de Physique Corpusculaire, Université Blaise Pascal, CNRS-IN2P3, Clermont-Ferrand

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

Centre de Physique de Particules de Marseille (CPPM), CNRS-IN2P3, Marseille

Laboratoire de l'Accélérateur Linéaire (LAL), CNRS-IN2P3, Orsay

Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), Universités de Paris VI et VII, CNRS-IN2P3, Paris

Commissariat à l'Énergie Atomique (CEA), DSM/DAPNIA, Centre d'Études de Saclay, Gif-sur-Yvette

Georgia

Institute of Physics of the Georgian Academy of Sciences and Tbilisi State University, Tbilisi

Germany

Physikalisches Institut, University of Bonn, Bonn

Deutsches Elektronen-Synchrotron (DESY), Hamburg and Zeuthen

TU Dortmund, Experimentelle Physik IV, Dortmund

Technical University Dresden, Dresden

Fakultät für Mathematik und Physik, Albert-Ludwigs-Universität, Freiburg

Justus-Liebig-Universität, Giessen

Fakultät für Physik, II. Physikalisches Institut, Georg-August-Universität, Göttingen
Ruprecht-Karls-Universität Heidelberg, Kirchhoff-Institut für Physik and Zentrales Institut für Technische Informatik (ZITI), Heidelberg
Institut für Physik, Humboldt Universität, Berlin
Institut für Physik, Universität Mainz, Mainz
Sektion Physik, Ludwig-Maximilians-Universität München, München
Max-Planck-Institut für Physik, München
Fachbereich Physik, Universität Siegen, Siegen
Fachbereich Physik, Bergische Universität, Wuppertal
Julius-Maximilians-University, Würzburg

Greece

National Technical University of Athens, Athens
University of Athens, Athens
University of Thessaloniki, High Energy Physics Department and Department of Mechanical Engineering, Thessaloniki

Israel

Department of Physics, Technion, Haifa
School of Physics, Tel Aviv University, Tel Aviv
Department of Particle Physics, The Weizmann Institute of Science, Rehovot

Italy

Dipartimento di Fisica dell' Università di Bologna e I.N.F.N., Bologna
Dipartimento di Fisica dell' Università della Calabria e I.N.F.N., Cosenza
Laboratori Nazionali di Frascati dell' I.N.F.N., Frascati
Dipartimento di Fisica dell' Università di Genova e I.N.F.N., Genova
Dipartimento di Fisica dell' Università di Lecce e I.N.F.N., Lecce
Dipartimento di Fisica dell' Università di Milano e I.N.F.N., Milano
Dipartimento di Scienze Fisiche, Università di Napoli 'Federico II' e I.N.F.N., Napoli
Dipartimento di Fisica Nucleare e Teorica della Università di Pavia e I.N.F.N., Pavia
Dipartimento di Fisica dell' Università di Pisa e I.N.F.N., Pisa
Dipartimento di Fisica dell' Università di Roma I 'La Sapienza' e I.N.F.N., Roma
Dipartimento di Fisica dell' Università di Roma II 'Tor Vergata' e I.N.F.N., Roma
Dipartimento di Fisica dell' Università di Roma III 'Roma Tre' e I.N.F.N., Roma
Dipartimento di Fisica dell' Università di Udine e I.N.F.N., Udine, I.C.T.P. e I.N.F.N., Trieste

Japan

Hiroshima Institute of Technology, Hiroshima

KEK, High Energy Accelerator Research Organization, Tsukuba

Kobe University, Kobe

Department of Physics, Kyoto University, Kyoto

Kyoto University of Education, Kyoto

Kyushu University, Kyushu

Nagasaki Institute of Applied Science, Nagasaki

Nagoya University, Nagoya

Faculty of Science, Okayama University, Okayama

Osaka University, Osaka

Faculty of Science, Shinshu University, Matsumoto

International Center for Elementary Particle Physics and Department of Physics, The University of Tokyo

Tokyo Institute of Technology, Tokyo

Physics Department, Tokyo Metropolitan University, Tokyo

Institute of Physics, University of Tsukuba, Tsukuba

Waseda University, Tokyo

Morocco

Faculté des Sciences Ain Chock, Université Hassan II, Casablanca, Université Mohamed Premier et LPTM, Oujda, Université Cadi Ayyat et LPHEA, Marrakech, CNESTEN et Université Mohamed V, Rabat

Netherlands

FOM - Institute SAF NIKHEF and University of Amsterdam/NIKHEF, Amsterdam

Radboud University Nijmegen and NIKHEF, Nijmegen

Norway

University of Bergen, Bergen

University of Oslo, Oslo

Poland

Institute of Nuclear Physics (IFJ PAN), Polish Academy of Sciences, Cracow

Faculty of Physics and Applied Computer Science, AGH University of Science and Technology and Marian Smoluchowski Institute of Physics, Jagiellonian University, Cracow

Portugal

Laboratório de Instrumentação e Física Experimental de Partículas (LIP), Faculdade de Ciências, Universidade de Lisboa, Centro de Física Nuclear da Universidade de Lisboa, Department of Physics, University of Coimbra, Departamento de Física, Universidade do Minho, Dep Física and CEFITEC of Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, University of Granada

Romania

Horia Hulubei National Institute of Physics and Nuclear Engineering (IFIN-HH), Institute of Atomic Physics, Bucharest, ITIM, Cluj Napoca, West University, Timisoara and University Politehnica Bucharest

Russia

Institute for Theoretical and Experimental Physics (ITEP), Moscow

P.N. Lebedev Institute of Physics, Moscow

Moscow Engineering & Physics Institute (MEPhI), Moscow

Moscow State University, Moscow

Budker Institute of Nuclear Physics (BINP), Novosibirsk

State Research Center of the Russian Federation - Institute for High Energy Physics (IHEP), Protvino

Petersburg Nuclear Physics Institute (PNPI), St. Petersburg

JINR

Joint Institute for Nuclear Research, Dubna

Serbia

Institute of Physics, University of Belgrade and Vinca Institute of Nuclear Sciences, Belgrade

Slovak Republic

Bratislava University, Bratislava, and Institute of Experimental Physics of the Slovak Academy of Sciences, Kosice

Slovenia

Jožef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana

South Africa

University of Cape Town, University of Johannesburg (UJ) and University of the Witwatersrand (WITS), Johannesburg

Spain

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

Physics Department, Universidad Autónoma de Madrid, Madrid

Instituto de Física Corpuscular (IFIC), Centro Mixto UVEG-CSIC, Valencia and Instituto de Microelectrónica de Barcelona, Bellaterra (Barcelona)

Sweden

Fysiska institutionen, Lunds universitet, Lund

Royal Institute of Technology (KTH), Stockholm

Stockholm University, Stockholm

University of Uppsala, Department of Physics and Astronomy, Uppsala

Switzerland

University of Bern, Albert Einstein Center for Fundamental Physics, Laboratory for High Energy Physics, Bern

Section de Physique, Université de Genève, Geneva

Taiwan

Academia Sinica, Taipei

Turkey

Department of Physics, Ankara University, Gazi University and TOBB ETU, Ankara

Department of Physics, Bogazici University, Dogus University and Gaziantep University, Istanbul

United Kingdom

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

University of Sussex, Brighton

Cavendish Laboratory, University of Cambridge, Cambridge

University of Warwick, Coventry

School of Physics & Astronomy, University of Edinburgh, Edinburgh

Department of Physics and Astronomy, University of Glasgow, Glasgow

Physics Department, Lancaster University, Lancaster

University of Liverpool, Liverpool

Department of Physics, Queen Mary and Westfield College, University of London, London

Department of Physics, Royal Holloway, University of London, Egham

Department of Physics and Astronomy, University College London, London

Department of Physics and Astronomy, University of Manchester, Manchester

Department of Physics, Oxford University, Oxford

Rutherford Appleton Laboratory, Science and Technology Facilities Council, Didcot

Department of Physics, University of Sheffield, Sheffield

United States of America

State University of New York at Albany, New York

Argonne National Laboratory, Argonne, Illinois

University of Arizona, Tucson, Arizona

Department of Physics, The University of Texas at Arlington, Arlington, Texas

Lawrence Berkeley National Laboratory and University of California, Berkeley, California

Physics Department of the University of Boston, Boston, Massachusetts

Brandeis University, Department of Physics, Waltham, Massachusetts

Brookhaven National Laboratory (BNL), Upton, New York

University of Chicago, Enrico Fermi Institute, Chicago, Illinois

Nevis Laboratory, Columbia University, Irvington, New York

University of Texas at Dallas, Dallas, Texas

Department of Physics, Duke University, Durham, North Carolina
Department of Physics, Hampton University, Virginia
Department of Physics, Harvard University, Cambridge, Massachusetts
Indiana University, Bloomington, Indiana
Iowa State University, Ames, Iowa
University of Iowa, Iowa City, Iowa
University of California, Irvine, California
Louisiana Tech University, Louisiana
University of Massachusetts, Amherst, Massachusetts
Massachusetts Institute of Technology, Department of Physics, Cambridge, Massachusetts
Michigan State University, Department of Physics and Astronomy, East Lansing, Michigan
University of Michigan, Department of Physics, Ann Arbor, Michigan
Department of Physics, New Mexico University, Albuquerque, New Mexico
Department of Physics, New York University, New York
Northern Illinois University, DeKalb, Illinois
Ohio State University, Columbus, Ohio
Department of Physics and Astronomy, University of Oklahoma
Oklahoma State University, Oklahoma
University of Oregon, Eugene, Oregon
Department of Physics, University of Pennsylvania, Philadelphia, Pennsylvania
University of Pittsburgh, Pittsburgh, Pennsylvania
Institute for Particle Physics, University of California, Santa Cruz, California
SLAC National Accelerator Laboratory, Stanford, California
Physics Department, Southern Methodist University, Dallas, Texas
State University of New York at Stony Brook, New York
Tufts University, Medford, Massachusetts
High Energy Physics, University of Illinois, Urbana, Illinois
Department of Physics, Department of Mechanical Engineering, University of Washington, Seattle, Washington
Department of Physics, University of Wisconsin, Madison, Wisconsin
Yale University, New Haven, Connecticut

Contents

Executive Summary	iii
1 Introduction	1
1.1 Changes for Detectors, Trigger and Data Acquisition	1
1.2 Upgrade Strategies	2
1.3 Present System and Planned Upgrades	3
1.4 Outline of this Report	4
2 Physics Motivation	7
2.1 Introduction	7
2.2 Rates and Performance	8
2.2.1 Electron and photon rates	8
2.2.2 Hadronic tau rate	9
2.2.3 Muon rate	9
2.2.4 Missing transverse momentum rate at high pile-up	11
2.2.5 Jet rate at high pile-up	11
2.2.6 Topological triggering	12
2.3 Level-1 Trigger Menus	13
2.4 HLT at Phase-I	16
2.4.1 HLT selections based on single electron and muon triggers	17
2.4.2 HLT selections based on hadronically decaying tau leptons	17
2.4.3 HLT selections based on hadronic triggers	18
2.5 Physics Studies	19
2.5.1 Higgs couplings and properties	19
2.5.2 Boosted objects	23
3 Level-1 Calorimeter Trigger	26
3.1 Introduction	26
3.2 Algorithms & Performance	27
3.2.1 Input data	27
3.2.2 Electron/photon trigger	27
3.2.3 Tau trigger	30
3.2.4 Jet trigger	31
3.3 System Evolution	34
3.3.1 Architecture	34
3.3.2 nMCM	37
3.3.3 CMX modules	39
3.4 Phase-I Architecture	41
3.4.1 Input interfaces	41
3.4.2 The Electron Feature Extractor (eFEX) module	43
3.4.3 The Jet Feature Extractor (jFEX) module	46
3.4.4 Output interfaces	49
3.4.5 Common modules	51
3.5 Latency Considerations	56
3.6 Firmware	56

3.7	Software	57
3.7.1	Online software	57
3.7.2	Offline software	57
3.8	Testing, Installation & Commissioning	58
3.8.1	Testing	58
3.8.2	Installation	59
3.8.3	Commissioning	59
3.9	System Performance	59
3.9.1	Validation	59
3.9.2	Calibration and monitoring	60
3.10	Optional Additions	61
3.10.1	The Global Feature Extractor (gFEX) module	61
3.11	Relation to the Phase-II Upgrade	64
3.12	Project Organisation	65
4	Level-1 Muon Trigger	67
4.1	Level-1 Muon Endcap Trigger	67
4.1.1	Introduction	67
4.1.2	Overview of the new system	67
4.1.3	Interface between the New Small Wheel and the New Sector Logic	69
4.1.4	The New Sector Logic board	71
4.1.5	Latency	72
4.1.6	Project planning	72
4.2	Level-1 Muon Barrel Trigger	73
4.2.1	Introduction	73
4.2.2	The new Sector Logic to MUCTPI interface board	73
4.2.3	Project planning	77
4.3	Level-1 Tile Muon Trigger	77
4.3.1	Tile hadronic calorimeter muon identification	78
4.3.2	Muon efficiency and fake rate reduction	78
4.3.3	System design	80
4.3.4	Project planning and milestones	81
5	Level-1 Central Trigger System	83
5.1	Central Trigger Processor (CTP)	83
5.1.1	Motivation	84
5.1.2	Trigger inputs	85
5.1.3	Trigger formation	86
5.1.4	Outputs to sub-detector partitions	87
5.1.5	Monitoring	87
5.1.6	Latency	88
5.2	Muon-to-CTP-Interface (MUCTPI)	88
5.2.1	Motivation	89
5.2.2	Interfacing of the existing MUCTPI to L1Topo in Run 2	89
5.2.3	Implementation	90
5.2.4	Interfaces	91
5.2.5	Latency	92

5.2.6	Software	92
5.2.7	Project planning	93
5.3	Level-1 Topological Processor	93
5.4	Level-1 Trigger Latency	96
6	Online and High-Level Trigger Software	99
6.1	Introduction	99
6.2	Description of the Run 1 and Run 2 Systems	99
6.3	Motivation	101
6.4	Baseline System, Assumptions and Options	102
6.5	Work Breakdown and Effort Estimates	103
6.5.1	High-Level Trigger selection software tasks	103
6.5.2	Data acquisition and control online software tasks	105
6.5.3	Effort estimates	107
6.5.4	Milestones	108
6.6	Summary	108
7	Data Acquisition and High-Level Trigger	110
7.1	Evolution during LS1	110
7.1.1	Read-out system	110
7.1.2	DAQ networks	113
7.1.3	Output to mass storage	115
7.2	Upgrades in LS2	117
7.2.1	Detector read-out	117
7.3	Evolution of HLT Compute Power	120
8	Resources, Organisation, and Workplan	124
8.1	Participating Institute Responsibilities	124
8.2	Management Organisation	127
8.3	Construction and Software Effort	128
8.4	Schedule and Milestones	128
8.5	Cost and Resources	128
A	Glossary	133
B	List of Authors	135

1 Introduction

When the LHC resumes operation in 2018 after the second long shutdown, it will be in many respects a new machine. The LHC performance will then by far surpass that achieved during the first pp run at $\sqrt{s} = 8$ TeV, with a peak luminosity of $0.77 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The upgraded LHC will produce collisions at a centre-of-mass energy at or near its ultimate design of 14 TeV; it will run with roughly double the number of proton bunches and, due to the upgrade of the injector complex, at similar or even higher bunch intensities, and therefore with significantly higher beam currents; and the beams will have smaller emittances with smaller β -functions at the interaction points, leading to a smaller luminous region. As a result, the LHC can be expected to reach instantaneous luminosities of $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This offers a major gain in physics potential through a substantial improvement in mass reach, with a dataset after three years of running totalling 300 fb^{-1} , more than ten times that which has been recorded to date.

For the detectors, the substantial increase in physics potential comes with considerable challenges. The increased luminosity will be accompanied by much increased levels of pile-up. The mean number of interactions per bunch crossing, $\langle \mu \rangle$, which during the 2012 run averaged 21 and peaked at around 40, is expected to go up to 80. The bunch spacing reduction from 50 ns to 25 ns, while helping to contain the increase in in-time pile-up, will nevertheless increase out-of-time pile-up, and also increase beam-induced fake trigger rates, particularly in the muon system. Most importantly, both the production cross sections for electroweak-scale particles (W , Z , H), and the QCD jet backgrounds under which they are buried, will rise by factors of 2–4 as the LHC energy increases from 8 TeV to 14 TeV.

1.1 Changes for Detectors, Trigger and Data Acquisition

The ATLAS detector has to be upgraded to exploit the new physics reach, preserve the physics acceptance, and to meet the challenges of the high-luminosity environment. The upgrades that are being prepared for ATLAS operations after 2018 are referred to as Phase-I upgrades.¹ Among the main upgrades to the detector systems are those to the Liquid Argon (LAr) calorimeter electronics [1.1], and the replacement of the inner stations of the endcap muon system with a new muon detector, the New Small Wheel (NSW) [1.2]. A new hardware-based Fast Tracker (FTK) [1.3] is being added to process data from the silicon tracking detectors, aiming to reconstruct charged particle tracks within a latency that is short enough for the tracks to be used at the start of the High-Level Trigger (HLT) processing.

The Trigger and Data Acquisition (TDAQ) system not only needs to adapt to the changes made to the detectors, but itself faces event rates and pile-up levels much higher than the original design values. The present TDAQ system was conceived in the late 1990s and has operated successfully through the end of the current data-taking (Run 1), while undergoing only a limited evolution. This has involved rolling replacements of commodity hardware, CPUs, networking, etc., as well as continuous improvements to the Trigger and DAQ software, which together have enabled the TDAQ system to operate reliably over a wide range of running conditions.

In contrast, the custom electronics in the Level-1 calorimeter and muon triggers, the Central Trigger Processor, as well as the read-out system, have remained unchanged, with only

¹Phase-I and Phase-II refer to the LHC machine and detector upgrades anticipated to be completed in 2018 and 2022 respectively. The terms LS1, LS2, LS3 refer to the long shutdowns of the LHC in 2013-14 and anticipated in 2018 and 2022, during which the detector upgrades occur. Run 1, Run 2, Run 3 refer to the periods of data-taking operation of ATLAS before and between the long shutdowns.

modifications in firmware since their original construction. In the face of the high-luminosity conditions, these parts will now have to be upgraded as well, and will be able to take advantage of major advances in technology in the past 15 years. This document reports on the design of the Phase-I upgrade of the TDAQ system.

Some LHC machine parameters, such as the 25 ns bunch spacing and the near-design centre-of-mass energy, are expected to be reached after long shutdown one (LS1), i.e. for operation during Run 2, albeit at lower luminosities. The upgrades for Run 2 present an important stepping stone towards the Phase-I TDAQ system and are included in the present document.

1.2 Upgrade Strategies

High levels of pile-up degrade the calorimeter resolution, soften the trigger turn-on and compromise isolation of single particles, thereby decreasing trigger efficiency. A major handle to counteract this is the use of increased granularity in calorimeter trigger data. The 0.1×0.1 trigger towers² used in the present Level-1 Calorimeter Trigger are too coarse to handle the increased levels of pile-up, and the primary aim of the LAr electronics upgrade is to make finer-granularity data available to the trigger electronics. The improved granularity will give rise to better isolation, which in turn can preserve lower energy thresholds. The upgraded Level-1 Calorimeter Trigger will make full use of the finer granularity, using new electron/photon and jet feature extraction modules, referred to as eFEX and jFEX, respectively.

Beam-induced backgrounds, primarily from slow particles (low-momentum protons) emanating from the endcap toroid or shielding, are an important source of fake triggers in the muon endcap. The key to suppressing these backgrounds in the muon endcap trigger is to require coincidence with the inner stations in the muon detector. The main objective of the muon New Small Wheel is to provide such coincidences in position and direction. This requires new Level-1 muon trigger electronics to interface to the Central Trigger Processor, as well as new sector logic for the muon endcap and changes to the interface of the muon barrel trigger. With these changes, the rates can be brought to manageable levels without restricting the acceptance and consequential loss of physics efficiency. In addition to the above, the shorter bunch spacing will increase fake triggers in the $1 < |\eta| < 1.3$ transition region. One way to suppress these is to use the outermost layer of the Tile calorimeter in the muon trigger to form a coincidence.

As luminosities increase, the physics coverage can be improved by complementing inclusive selections with more exclusive ones. The present central trigger relies on combinatorial logic, using information on object energies and momenta but with no data on their relative location within the ATLAS detector. The new topological trigger processor, being prepared for Run 2, will allow selections based on angles, effective masses etc. A new core processor within the central trigger is also being prepared. This will provide double the number of inputs, and give substantially more flexibility in choice and number of trigger menu items. The new muon interface to the central trigger will also provide full granularity Region-of-Interest (ROI) information to the topological trigger processor.

In the High-Level Trigger, the increased pile-up will demand more sophisticated corrections, as is already common in offline reconstruction. Many of these corrections involve tracking, which is the most CPU-intensive task performed by the HLT software. The main strategy

²Regions in the detector are given in units of pseudo-rapidity η and azimuth ϕ , and are specified as $\Delta\eta \times \Delta\phi$. Thus, an area quoted as 0.2×0.1 spans 0.2 in η and 0.1 rad in ϕ .

is to make more use of tracks in the HLT, using the fast hardware tracking provided by the FTK. Another strategy is an increased reliance on full calorimeter read-out. In general, as more sophisticated techniques are moved upstream from the HLT into the Level-1 trigger, the HLT will in turn have to more extensively adopt offline processing algorithms to maintain the required rejection factors. The upgrades to the HLT hardware and software are designed to provide the necessary resources to make this possible.

1.3 Present System and Planned Upgrades

Figure 1 shows a schematic overview of the TDAQ system in 2012. A detailed description of the original system can be found in Ref. [1.4]. Event data move from the ATLAS on-detector electronics into the front-end buffers at the bunch crossing rate of 40 MHz. Starting from this huge flow of data, the task of the Trigger and DAQ system is to select a few hundred events per second for recording to permanent storage.

The DAQ system is responsible for the transport and assembly of the event data all the way from the front-end buffers to the logging to disk. Along this path, event fragments are processed in parallel by the sub-detector Read-Out Drivers (RODs), which perform varying levels of feature extraction. The data pass to the Read-Out System (ROS), where they are buffered until requested by the High-Level Trigger farm, where they are assembled into events (event building) and ultimately recorded to disk. The two-tier trigger system orchestrates these stages, with each level accepting only a subset of data to be passed on to the next. The Level-1 Trigger is clock-driven, the High-Level Trigger event-driven.

The Level-1 trigger, built in custom electronics, processes data from the calorimeters and the muon detectors, searching for signatures such as large electromagnetic energy deposits or high- p_T muon tracks. Once identified, information on these features is collected and sent to the central trigger, where it is combined and the result is compared to 256 programmable trigger items. If the trigger conditions are met, a Level-1 Accept is issued, initiating read-out and processing in the RODs, and subsequent transfer to the ROS. At the same time, Regions-Of-Interest (ROIs) are built. These are small data packets describing the feature locations and energies, which are sent to the Level-2 Trigger and used to steer the processing. The rate of the Level-1 Accepts reached about 70 kHz during Run 1.

The High-Level Trigger is implemented in software, and operates on a large farm of commercial computer processors. It executes chains of reconstruction and signature algorithms that analyse the properties of the events. The original architecture employed two distinct levels, running on physically separate sub-farms. The Level-2 trigger generally requested fragments defined by the Level-1 ROIs, but at full granularity for the detector regions concerned, amounting to approximately 10% of the data for each event. The Level-2 nodes delivered a decision within a few tens of milliseconds and were typically configured to accept about one in ten events for further processing. The accepted events were then fully built, prior to processing in the Event Filter, in which a complete reconstruction of the event was performed, running essentially offline algorithms to make a final decision within approximately one second. At this stage, another rejection factor of ten was realised, resulting in a recording rate of a few hundred events per second.

Many of these original components will be retained in the Phase-I system. Figure 2 shows the system after the Phase-I upgrade. The main parts to be upgraded are the Level-1 Calorimeter Trigger electronics, including the new eFEX and jFEX processors, the RODs and optical plant, the Level-1 Muon Trigger sector logic for the endcap, the Central Trigger Processor

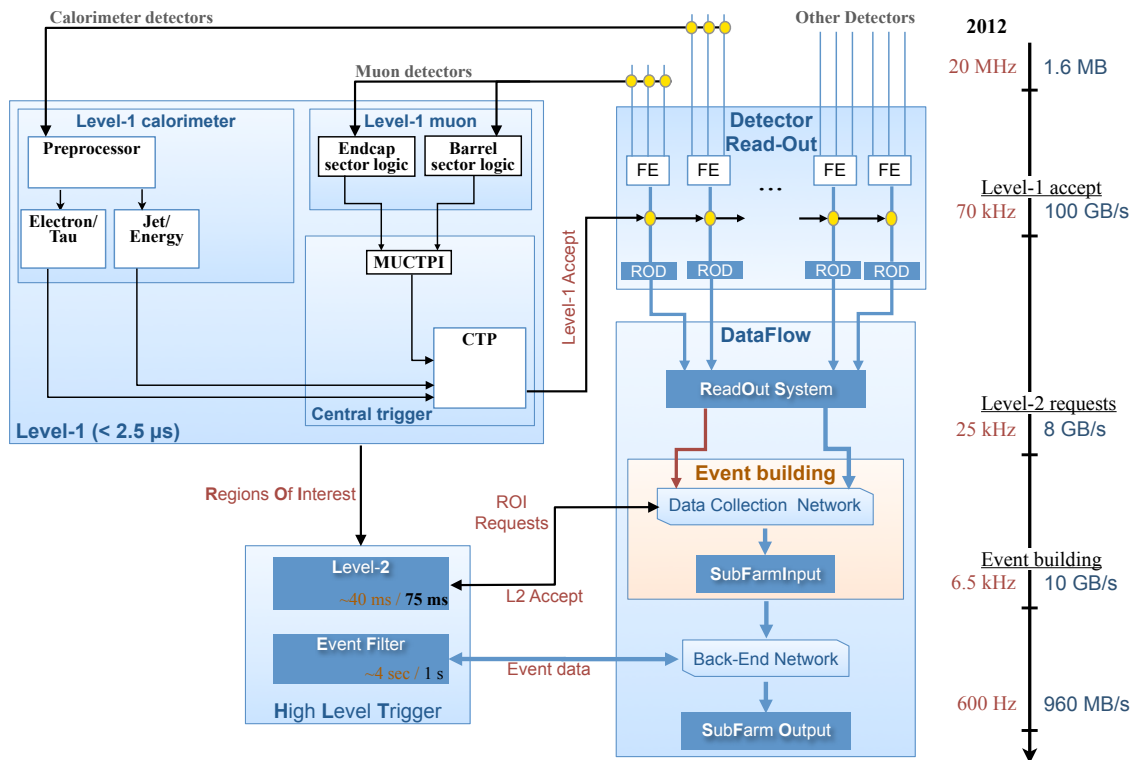


Figure 1: Schematic overview of the Trigger and DAQ system in 2012 (Run 1)

(CTP), and the muon-to-CTP interface (MUCTPI). In addition, there will be a new Level-1 Topological Processor (L1Topo), a Tile Calorimeter muon trigger, a New Small Wheel trigger processor, and the new Fast Tracker.

Trigger latency is a critical parameter for the Level-1 Trigger upgrades, as the majority of the detector front-end electronics will not be replaced before LS3. All the new trigger functions have to be performed within the existing latency limit of $2.2 \mu\text{s}$.

In the DAQ/HLT system, the ROS hardware will be replaced. Wide-ranging software upgrades will be needed to core software, dataflow, configuration and control software, trigger menus and algorithms, as well as possible use of new technologies. The Level-2, event building and Event Filter will be merged during LS1 into a single processing unit running in a homogeneous HLT farm. Finally, the HLT CPU farm and online infrastructure will be upgraded and extended, to meet the high-luminosity performance requirements. The evolution of the HLT compute power is based on the arguments presented in the physics motivation section and the improvements expected from upgraded trigger algorithms.

1.4 Outline of this Report

Following this introduction, the document is organised in eight principal sections. Section 2 presents the physics case for the TDAQ Phase-I upgrade. The basic constituents of the current trigger menu, single-electron and single-muon triggers, taus, jets and missing transverse energy triggers, are examined in the context of their current thresholds and rates. The evolution of the rates with the upgraded LHC is then discussed, particularly how low-threshold single-lepton triggers – and HLT rejection factors – can be preserved, despite the much higher

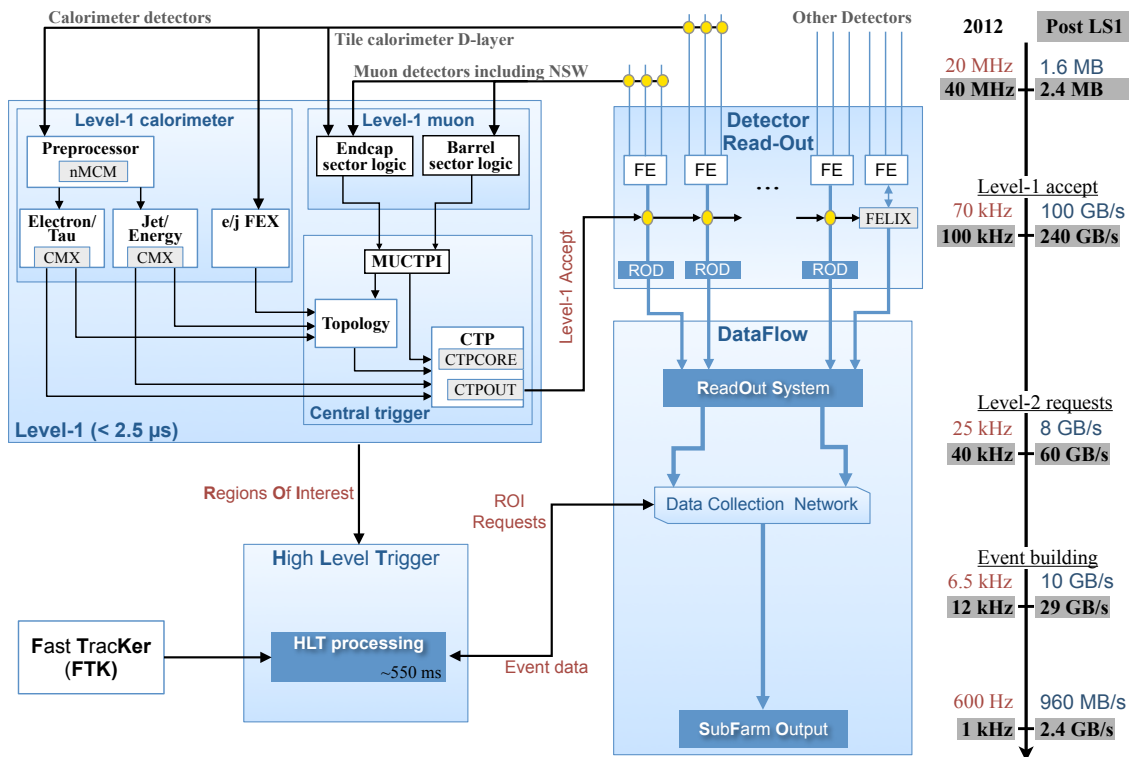


Figure 2: Schematic overview of the Trigger and DAQ system after the Phase-I upgrade

particle rates by taking advantage of the upgraded detector and trigger. Physics studies are presented that review the most important physics channels.

The following sections then cover details of the upgrade. Section 3, the longest of this document, describes the upgrades to the Level-1 Calorimeter Trigger, including examples of new algorithms and their performance, the eFEX and jFEX processors with their interfaces and infrastructure, firmware and software upgrades, and the plan for testing, commissioning and validation. The upgrades of the Level-1 Muon Trigger are presented in Section 4, which is divided into three parts: the muon endcap trigger and new sector logic for the NSW; the new muon barrel trigger interface modules; and the design of the Tile calorimeter D-layer coincidence electronics. Section 5 details the upgrades to the Level-1 Central Trigger, specifically the CTP, the new MUCTPI, and the L1Topo trigger processor. A summary of the Level-1 Trigger latency in the new system is presented here as well.

In Section 6, the evolution of the online and High-Level Trigger software is discussed. Section 7 focuses on the hardware of the DAQ and HLT systems, covering the upgraded Read-Out System as well as the DAQ networks after LS1, followed by a description of a proposed upgrade for the Phase-I read-out. The final part of the document, Section 8, summarises the participating institutes and their areas of interest, the management organisation, work breakdown structure and schedule, as well as the cost and resources.

The glossary in appendix A defines many of the common acronyms and terms used in this document.

The strategy presented in this technical design report provides ATLAS with a robust configuration for higher-luminosity running, enabling the experiment to take full advantage of the accelerator upgrades.

References

- [1.1] ATLAS Liquid Argon Calorimeter Group, *ATLAS Liquid Argon Calorimeter Technical Design Report: Trigger Electronics for the LHC Phase-I Upgrade*, Tech. Rep. CERN-LHCC-2013-???. ATLAS-TDR-??, CERN, Geneva, (*in preparation*).
<https://twiki.cern.ch/twiki/bin/view/LAr/LArPhaseITDR>.
- [1.2] ATLAS Collaboration, *New Small Wheel Technical Design Report*, Tech. Rep. CERN-LHCC-2013-006. ATLAS-TDR-020, CERN, Geneva, Jun, 2013.
<https://cds.cern.ch/record/1552862>.
- [1.3] ATLAS Collaboration, *Fast TracKer (FTK) Technical Design Report*, Tech. Rep. CERN-LHCC-2013-007. ATLAS-TDR-021, CERN, Geneva, Jun, 2013.
<https://cds.cern.ch/record/1552953>.
- [1.4] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, **JINST 3 (2008) S08003**.

2 Physics Motivation

2.1 Introduction

The ATLAS integrated luminosity near $\sqrt{s} = 14$ TeV is expected to total more than 300 fb^{-1} by the end of Run 3, with the bulk of the data taken after the installation of all the components of the Phase-I upgrade. This data set will allow ATLAS to make precise measurements of Higgs production rates and properties and, equally importantly, search in a very large phase space for physics Beyond the Standard Model (BSM). These data will also allow ATLAS to improve the measurement of many Standard Model (SM) physics processes and to improve the foundation on which the Higgs and BSM studies are built. To satisfy these demands, the Phase-I upgrade of the Trigger and Data Acquisition (TDAQ) system of ATLAS must provide comprehensive and efficient coverage of all triggers needed for Higgs and SM physics, while at the same time serve the very diverse needs of the searches for BSM physics.

Much of the trigger phase space is covered by the decays of electroweak particles such as the W , Z and Higgs bosons, which have masses between 80 and 125 GeV. W bosons are ubiquitous, not only in SM decays, such as those of the top quark, but also in BSM decays, for example, in the decay of supersymmetric particles. The decays of these electroweak particles are mainly to leptons and jets resulting in detector objects which predominately exceed an offline transverse momentum, p_T , of about 30 GeV. The most useful Level-1 trigger signatures for these *electroweak-scale* particles are from isolated electrons and muons which produce samples highly enriched in W and Z bosons³.

The main goal of the Phase-I upgrade is to preserve the low- p_T trigger thresholds at Level-1 for single electrons and muons that were used successfully in the Run 1 physics program. The ability to maintain unprescaled Level-1 triggers on these objects not only will maximise data samples, for example for Higgs studies, but will also produce data samples that can be thoroughly understood, which will lead to high quality physics results with small systematic uncertainties. An added benefit of the Phase-I upgrade is that the threshold for multi-lepton triggers can be reduced, which is helpful for SUSY scenarios with small mass splittings.

Another important goal of the Phase-I upgrade is to maintain the ATLAS experiment's ability to do physics with electroweak-scale particles that produce hadronically-decaying tau leptons, jets and missing transverse momentum, E_T^{miss} (from neutrinos and other non-interacting particles). The signatures of these objects are less distinct than those of electrons and muons. To efficiently trigger on these decay products of electroweak-scale particles, it is desirable to combine several trigger objects at Level-1 (including very low- p_T electrons and muons, if appropriate) in order to reach rates that can be accommodated by the ATLAS TDAQ system. At Level-1, these combined objects greatly benefit from the proposed topological trigger, which is particularly helpful in rejecting background from di-jet events. The ability of the trigger system to accommodate these combined trigger signatures relies on having low Level-1 single-electron and -muon rates leaving adequate bandwidth for the combined objects.

The last goal of the Phase-I upgrade is to significantly improve the High-Level Trigger (HLT) to keep the output rate reasonable. The improvements are needed because after the Phase-I upgrade the 100 kHz of accepted Level-1 triggers will be enriched in interesting physics objects, compared to Run 1, by a factor of 10 or more and will contain many more

³Some events with W and Z bosons may need to be rejected in the HLT, as described in Section 2.4.1. The full detector information available at the HLT allows these events to be rejected with minimal loss to the ATLAS physics program.

pile-up interactions. This requires a combination of software improvements, expansion of the HLT computing farm and the addition of the Fast Tracker (FTK).

In Section 2.2 the Level-1 trigger rates expected before and after the Phase-I upgrade for electrons or photons, taus, muons, jets and E_T^{miss} are discussed. In Section 2.3, a possible evolution of the Level-1 trigger menu is presented. This is followed by a discussion of the expected physics performance of the High-Level Trigger in Section 2.4, including the impact of the FTK processor. In Section 2.5, the physics performance of the Phase-I upgrade is evaluated for specific processes.

2.2 Rates and Performance

In this section, the expected Level-1 trigger rates for $\sqrt{s} = 14 \text{ TeV}$ and $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity, with (Run 3) and without (Run 1 and Run 2) the Phase-I upgrade are discussed. In Run 1, constraints external to the trigger limited the total Level-1 rate to approximately 65 kHz. During Run 2 and Run 3, the ATLAS Level-1 rate is expected to reach 100 kHz. Bandwidth allocations are based on detailed arguments and agreed in the menu coordination group. Typically roughly equal shares are given to electrons and muons, with a large share of bandwidth reserved for jets, E_T^{miss} , taus and multi-object triggers, implying that rates for single electrons and muons would each be limited to approximately 25 kHz or less.

2.2.1 Electron and photon rates

Rates of electron and photon (EM) triggers, which are indistinguishable at Level-1, can be controlled through both isolation cuts and by raising p_T thresholds. If no isolation cuts were applied, the rate at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for an offline p_T threshold of 50 GeV ($\sim 40 \text{ GeV}$ trigger threshold) would be approximately 33 kHz and the trigger would only be sensitive to the highest- p_T electrons from boosted W bosons. During Run 2, the isolation cuts will be made using the $0.1 \times 0.1 \eta - \phi$ trigger towers in the current system. Electrons are found as maxima in a 0.1×0.2 or 0.2×0.1 region in the electromagnetic calorimeter. A hadronic energy isolation cut of $\leq 1 \text{ GeV}$ is applied on the sum of hadronic energy in the 0.2×0.2 region immediately behind the core of the EM region. The EM isolation cut of $\leq 2 \text{ GeV}$ is made on the twelve 0.1×0.1 EM towers surrounding the 0.2×0.2 EM core region. These isolation requirements, in conjunction with η -dependent thresholds that account for energy loss before the LAr calorimeter, allow the EM trigger threshold to be 30 GeV (38 GeV offline cut) with a rate of 14 kHz, but with a loss of efficiency of $\sim 3\%$ relative to efficiencies before the changes.

After the Phase-I upgrade it will be possible to reduce the thresholds by applying a cut on R_η instead of the EM isolation cut. R_η is the ratio of energy deposited in a $0.075 \times 0.2 \eta - \phi$ region to a 0.175×0.2 region. Figure 3b shows the expected rates and additional rejection achieved if an R_η cut with a similar efficiency to the Run 2 EM isolation cut is used instead of the Run 2 EM isolation cut.⁴ Figure 3a shows the efficiency versus offline p_T for a fixed rate of $\sim 14 \text{ kHz}$. The rates in Figure 3b have been normalised to the observed rates in Run 1 at $\sqrt{s} = 8 \text{ TeV}$ and extrapolated to $\sqrt{s} = 14 \text{ TeV}$ using Pythia 8.165 [2.1] with the MSTW2008LO [2.2] leading-order parton distribution function (PDF) set with the AU2 underlying-event tune [2.3].

More sophisticated isolation and shower shape requirements are under study and have been discussed in Ref. [2.4]. In addition, improvements in energy resolution from the use of optimal weights when combining LAr Calorimeter layers will result in better energy resolution

⁴Offline thresholds are typically set higher than the online thresholds at the point at which the Level-1 efficiency reaches $\sim 90\%$ of its plateau value.

and can sharpen the turn-on curves shown in Figure 3a, allowing the offline p_T threshold to be significantly reduced [2.4].

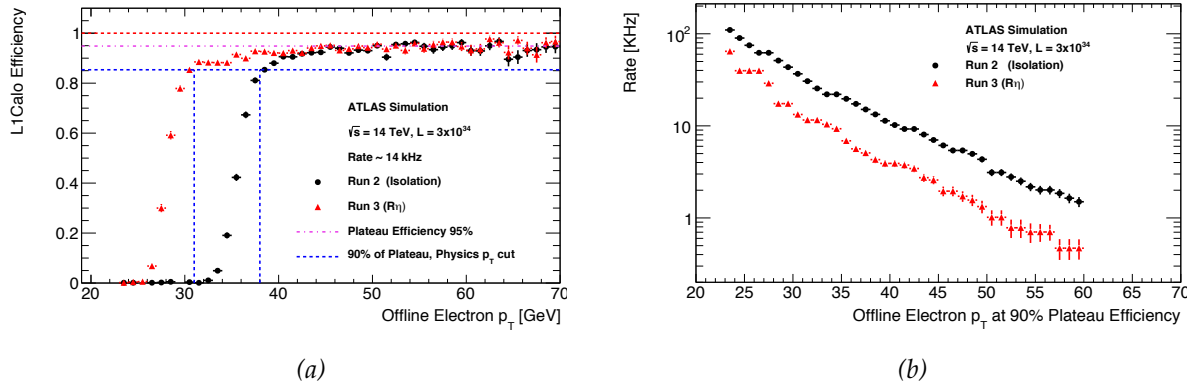


Figure 3: (a) Efficiency vs reconstructed electron p_T using the existing system in the Run 2 configuration with an isolation cut (black circles), and in Run 3 (after the Phase-I upgrade) with the R_η cut (red triangles). For both systems, the E_T threshold value varies with η to correct for variations in material. (b) EM trigger rates for Run 2 (black circles) and Run 3 (red triangles) systems. The E_T scale is the offline electron p_T for which the E_T thresholds used would give 90% of the plateau efficiency.

2.2.2 Hadronic tau rate

Individual hadronically-decaying tau leptons will be identified at Level-1 with a method very similar to that used for electrons, except that the hadronic energy behind the EM core is included in the energy sum.

The rate of these tau triggers will be reduced using electromagnetic isolation. For example, a variable similar to R_η has been shown to give a rejection of 1.3 to 1.4 in the eFEX. In the jFEX, a wider area isolation region can be used.

Because of the energy carried away by the neutrinos in tau decays, the thresholds for hadrons from taus must be lower than those for muons and electrons produced directly from W boson decays. However, the rate at Run 2 for a hadronic tau trigger threshold of 15 GeV is nearly 1 MHz, dominated by multi-jet background, making it impossible to trigger on single tau leptons produced by electroweak-scale objects. Instead, di-object triggers are used together with topological cuts.

2.2.3 Muon rate

Without changes to the trigger system, the rate of single-muon triggers ($p_T > 20$ GeV) would be more than 50 kHz at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ as can be seen in Table 1. The bulk of this rate originates from beam backgrounds that produce relatively low energy protons in the forward muon spectrometer which in turn produce both random and correlated hits. In contrast to the electron and hadronic tau triggers, where the relative p_T resolution of the calorimeter improves as p_T increases, the relative momentum resolution of the muon spectrometer degrades as p_T increases. Because of the limited spatial resolution of the present Level-1 muon system, significantly raising the nominal p_T threshold above 20 GeV produces a slow turn-on curve that is passed by a large number of muons below the threshold, but at the same time still has $\sim 10\%$ inefficiency well above the threshold [2.5]. The strategy for Run 2 is to control the

Online $p_T > 20 \text{ GeV}$ $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$	No Change	TGC EIL4 + (TGC FI or NSW)	TGC EIL4 + (TGC FI or NSW) + Tile Cal.	TGC EIL4 + (TGC FI or NSW) + Tile Cal. + low field mask
	Rate [kHz]	Rate [kHz]	Rate [kHz]	Rate [kHz]
Run 2 (pre NSW)	51	34	31	28
Run 3 (post NSW)		17	15	13

Table 1: Expected single-muon Level-1 rate (based on 2012 data and 8 TeV) with bunch spacing 25 ns and instantaneous luminosity $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for online $p_T > 20 \text{ GeV}$. The first column shows the rate with no added coincidences, the second with a coincidence between EI and FI chambers and outer chambers in the region $1.0 < |\eta| < 2.4$, the third adds a coincidence between the outer layer of tile calorimeter and outer chambers in the region $1.0 < |\eta| < 1.3$, the fourth column adds a mask that removes some small regions in ϕ with $|\eta| \sim 1.6$, that would result in an efficiency loss of 1%.

single-muon rate by adding additional coincidence requirements to the trigger, and failing that, reducing the acceptance in η to exclude the regions with the highest fake rates. Multi-object triggers are largely insensitive to the beam-induced background and can be used for all η in all scenarios.

During Run 1, only TGC chambers furthest from the interaction point were used to generate triggers. In Run 2, the trigger rate will be reduced by introducing a coincidence with a doublet of TGC chambers located upstream of the endcap toroids. The EIL4-TGC doublets partially ($\sim 50\%$) cover the rapidity region $1.0 < |\eta| < 1.3$, while the EI-TGC fully cover the region $1.3 < |\eta| < 1.9$. The coverage of the EIL4-TGC doublets can be increased by adding a coincidence with the outer-most layer of the ATLAS Tile Calorimeter (see Section 4.3).⁵ With these improvements, the trigger rate for $p_T > 20 \text{ GeV}$ is not expected to exceed 28 kHz, when excluding the low field regions ($< 5\%$ acceptance loss) in Run 2. Additional reduction could be obtained by excluding the region $1.9 < |\eta| < 2.4$.

In Run 3, the New Small Wheel (NSW) [2.5] reduces the trigger rate through the use of trigger chambers that will provide two displaced space points for muons, allowing a constraint on both the position and direction of the candidate muon at the location of the NSW. The two trigger space points are each determined from four sTGC trigger chambers allowing for both a local coincidence and for redundancy in case of chamber inefficiency. The trigger chambers are complemented by eight MicroMegas (MM) chambers which will avoid the degradation that occurs in the present offline muon reconstruction due to the high rate of beam-induced hits in the muon drift tubes (MDTs) and in addition provide an independent Level-1 trigger with both a position and direction constraint. The NSW will provide hermetic coverage in the range $1.3 < |\eta| < 2.4$ and reduce the total single-muon rate to a tolerable 17 kHz at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Including the planned reduction in the region $1.0 < |\eta| < 1.3$ from the Tile calorimeter the rate will be further reduced to 15 kHz at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. In Run 3 it is likely the low field region will not need to be masked.

The trigger rates presented in Table 1 are based on a test run with 25 ns bunch spacing taken in 2012 at 8 TeV and supersede those given in Ref. [2.5]. The rates without the upgrade depend heavily on beam backgrounds and therefore have large uncertainties.

⁵Even more rejection could be achieved by installing additional chambers to be included in the trigger coincidence.

2.2.4 Missing transverse momentum rate at high pile-up

The rate of E_T^{miss} triggers is a very strong function of pile-up. The effect of pile-up crucially depends on the thresholds used to determine which energy depositions are included in the energy sums. The application of these thresholds is distorted by the structure of the out-of-time pile-up that is induced by the gaps between the LHC bunch trains and the response function of the ATLAS calorimeter electronics. Corrections for these effects have been applied for offline ATLAS analysis and they will be implemented in both Run 2 and Run 3 Level-1 triggers. In Run 2 these corrections will be implemented in the nMCM (see Section 3.3.2) module and will reduce the trigger rate for E_T^{miss} and jet triggers by more than a factor of five. The projected rate with the modified nMCM in Run 2 is approximately 9 kHz for a threshold of 90 GeV.

After the Phase-I upgrade, the E_T^{miss} calculation can benefit from the optimisation of the cells included in the calculation of the calorimeter cells in the LAr system as well as jet-area corrections which are described in the section below. The impact of these improvements have been studied in Ref. [2.4] and it was found that the online threshold could be lowered to 70 GeV with a rate of approximately 10 kHz.

2.2.5 Jet rate at high pile-up

The rate of low- p_T and multi-jet triggers is strongly affected by pile-up. If the first-order pedestal shifts are controlled as described in the previous section, there will still be a strong residual effect due to pile-up from algorithms that use clustering or ignore negative fluctuations. This is illustrated in Figure 4a which shows the average number of pile-up jets with $p_T > 20$ and > 40 GeV in full simulation with $\langle\mu\rangle = 140$; there are > 20 pile-up jets per event with $p_T > 20$ GeV. Pile-up also degrades the jet resolution and the response of the calorimeter. The method adopted offline for countering these effects is to use a measurement of the p_T density within each event that excludes the hard physics [2.6]. This provides a measure of the total pile-up activity which can then be removed event-by-event and jet-by-jet by subtracting the density ρ times the area of a given jet. The method results not only in a reduction of the jet multiplicity to more manageable levels even at high $\langle\mu\rangle$ (Figure 4a), but also improves the jet resolution and restores the calorimeter response. Expected values of ρ , measured using the full calorimeter within $|\eta| < 2$, are shown for various $\langle\mu\rangle$ in Figure 4b; the RMS of the ρ distribution increases with increasing $\langle\mu\rangle$. There are multiple ways to calculate an appropriate ρ for a given event. The offline software can be directly implemented in the HLT for jet, tau, and E_T^{miss} triggers, however, the basic scheme requires some adjustment for the coarser Level-1 segmentation and resolution. For example, a measure of the total number of Level-1 towers or segments with energy above threshold should be proportional to ρ . While the ρ metric captures the global fluctuations within an event, local fluctuations also result in degraded resolution. A solution can be adopted that employs $2\pi\phi$ rings for limited η ranges to correct for local backgrounds due to pile-up. This is similar to the scheme implemented for the Heavy Ion program and results in lower rates for jet, tau, and E_T^{miss} triggers. This methodology was deployed for the Run 1 Heavy Ion HLT jet triggers and is expected to result in a sharper jet threshold [2.4] allowing the triggering on lower- p_T jets or a decrease in rates; it can be added at Level-1 given the appropriate architecture.

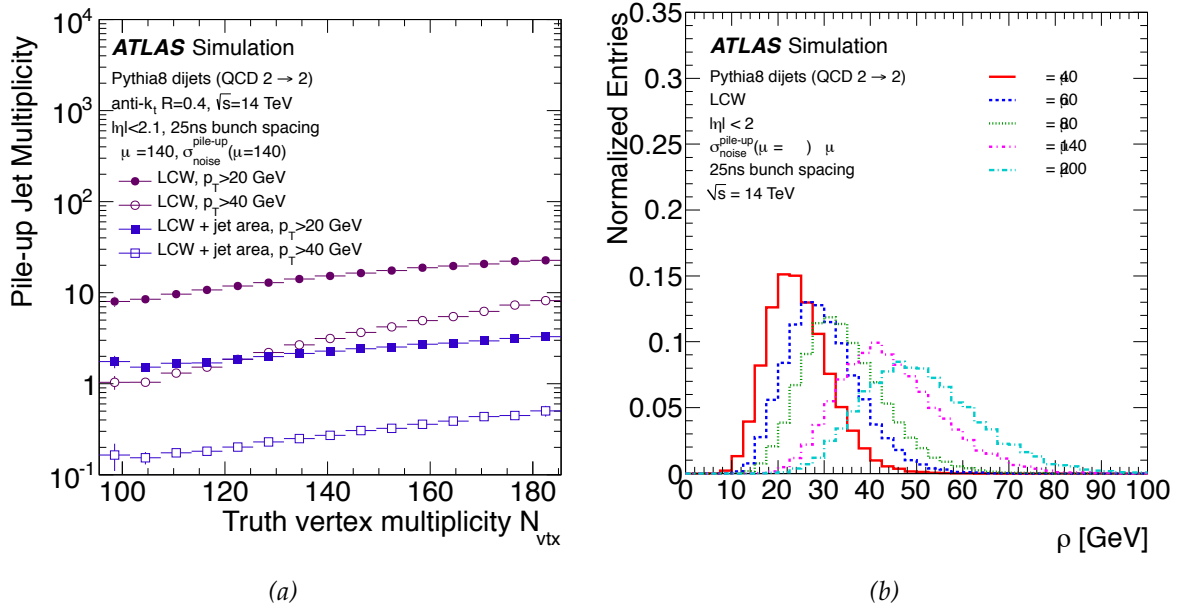


Figure 4: (a) The mean pile-up jet multiplicity as a function of the number of vertices in events with $\langle \mu \rangle = 140$. The closed circles (squares) represent the mean number of pile-up jets with $p_T > 20$ GeV (solid markers) and 40 GeV (open markers) before and after pile-up subtraction using the event-by-event median p_T density, ρ . The jet p_T have been calibrated using Locally Cluster Weighted (LCW). (b) Distributions of ρ for different values of $\langle \mu \rangle$. An increase in the RMS of the ρ distribution is observed with higher values of $\langle \mu \rangle$.

2.2.6 Topological triggering

The Level-1 topological processor described in Section 5.3 will allow the combination of electron/photon, muon, tau and jet objects, as well as event-level quantities, such as E_T^{miss} and its direction. The topological algorithms to be used are still evolving, but those most likely to be used include restriction on angles between objects, as well as the reconstruction of invariant masses of pairs of objects, for example, the invariant mass of forward-jets associated with low- p_T lepton pairs. With the exception of beam-induced backgrounds in the muon system, the dominant backgrounds are from di-jet events which can be reduced with angular cuts. Typically, cuts will be made on the difference in pseudo-rapidity, $\Delta\eta$, the difference in azimuthal angle, $\Delta\phi$, and the variable $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$.

The production, commissioning and installation of the topological processor has been accelerated, and together with the CMX module (replacement for the current Cluster Merger Module, CMM) (see Section 3.3.3), it will be deployed for use in the present L1Calo trigger for Run 2. The use of the topological processor can result in improvements of 30% or more, simply by applying requirements which would be made anyway in the physics analyses, already at Level-1. In almost all cases, inclusive triggers, without topological cuts but with higher thresholds, will be retained in addition to the topological triggers.

Before and after the Phase-I upgrade, the largest impact of the topological trigger will be on di-tau final states (primarily for Higgs studies) with one or both taus decaying hadronically. The rate of most di-tau triggers can be reduced to a tolerable level by imposition of $\Delta\eta$ or $\Delta\phi$ and ΔR cuts with modest loss in efficiency for signal events. This reduction is described below in Section 2.5.1.

The topological trigger is also useful for triggering on the process $ZH \rightarrow \nu\bar{\nu}b\bar{b}$. In this process, the missing energy is likely to point far from any jet, giving good discrimination against the QCD background as is demonstrated in Section 2.5.1.

The topological trigger offers the possibility to use soft muons to tag b -jet candidates at Level-1. As an example, $b(\bar{b})H/A \rightarrow b(\bar{b})b\bar{b}$ benefits from a cut on ΔR between a muon and a jet: compared to a baseline trigger that requires a leading jet with $p_T > 30$ GeV and a sub-leading jet with $p_T > 20$ GeV in coincidence with a $p_T > 4$ GeV muon, the rate decreases by a factor of 5 at the cost of 33% efficiency decrease, doubling the expected significance for this process.

2.3 Level-1 Trigger Menus

Example Level-1 trigger menus before (Run 2) and after the Phase-I upgrade (Run 3) are given in Table 2, assuming a luminosity of $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $\sqrt{s} = 14$ TeV. Both the Run 2 and Run 3 menus are expected to evolve as more sophisticated algorithms are developed and there is substantial scope for lowering Level-1 thresholds, especially in Run 3, where studies on exploiting the Phase-I upgrade are ongoing, as is documented in the companion LAr Calorimeter TDR [2.4]. Also presented in the table are the rates from the Run 1 menu used at the highest luminosities in 2012, scaled to $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $\sqrt{s} = 14$ TeV. The total rate of the Run 1 menu used in these conditions exceeds the maximum allowed rate of 100 kHz by a factor of eight.⁶

In Run 2, prior to the Phase-I upgrade, rates of electromagnetic triggers (electrons and photons, denoted by the prefix “EM”), can be controlled by raising thresholds, adding hadronic isolation (denoted by “H”), or adding electromagnetic isolation (denoted by “I”) and varying thresholds with η to account for energy loss (denoted by “V”). The number following the prefix corresponds to the p_T or E_T threshold of the trigger.

After the Phase-I upgrade in Run 3, with the additional capabilities of the L1Calo upgrade, the trigger rate of EM objects can be reduced with the R_η cut (denoted by “R”) discussed above in Section 2.2.1, giving a factor of more than two and this is exploited in the Phase-I menu. This allows the offline threshold to be lowered by 6 GeV from 38 GeV to 32 GeV for single electrons. As already noted in Section 2.2.1 cuts on additional shower shape quantities as well as improvements in energy resolution from the use of optimal weights when combining LAr Calorimeter layers will result in an additional reduction of the offline threshold by several GeV [2.4].

Tau triggers in the menu are denoted by the prefix “TAU”, with electromagnetic isolation denoted by “I”. After the Phase-I upgrade, the rate of these triggers can be reduced by $\sim 30\%$ with a looser cut on the R_η quantity (also denoted by “R”). These improvements crucially allow a ~ 10 GeV reduction in offline tau thresholds which will significantly improve the reach of the $H \rightarrow \tau\tau$ analysis described in Section 2.5.1.

Muon triggers are denoted by the prefix “MU”. During Run 2, before the introduction of the NSW, one may be able to maintain a trigger threshold of 20 GeV, but only by significantly cutting bandwidth from other triggers, such as those in the topological processor. After the Phase-I upgrade, there will be significant bandwidth available to design many specialised topological triggers.

⁶ The Run 1 configuration assumes the first 12 bunches in each bunch train would be vetoed giving an inefficiency for these triggers of 15%. This avoids the bunch-train dependent effects that will be mitigated by the installation of the nMCM in Run 2 (see Section 3.3.2).

Jet triggers are denoted by the prefix “J”, and E_T^{miss} triggers are denoted by “XE”. In Table 2 it has been conservatively assumed that the jet thresholds and rates will be unchanged, but it is expected that after the Phase-I improvements, the offline E_T^{miss} cut can be lowered from 250 GeV to 200 GeV [2.4].

In many cases, a single energy deposition can give rise to several triggers that overlap in the same region of ϕ and η , for example, deposition of energy in the EM layer of the calorimeter can cause overlapping EM, tau and jet triggers. There is no overlap removal in the Level-1 trigger between multi-objects, except for some specialised instances in the topological processor. For example, the trigger 2TAU20VI_EM20VHI_3J20 is designed to trigger events with one jet, one electron and one tau. (The tau also satisfies the jet trigger, and the electron satisfies all three, the electron, jet and tau triggers.)

Run 1			Run 2			Run 3		
	Offline p_T Threshold [GeV]	Rate [kHz]		Offline p_T Threshold [GeV]	Rate [kHz]		Offline p_T Threshold [GeV]	Rate [kHz]
EM18VH	25	130	EM30VHI	38	14	EM25VHR	32	14
EM30	37	61	EM80	100	2.5	EM80	100	2.5
2EM10	2x17	168	2EM15VHI	2x22	2.9	2EM12VHR	2x19	5.0
EM total		270			18			20
MU15	25	150	MU20	25	28	MU20	25	15
2MU10	2x12	14	2MU11	2x12	4.0	2MU11	2x12	4.0
Muon total		164			32			19
EM10VH_MU6	17,6	22	EM15VH_MU10	22,12	3.0	EM10VHR_MU10	17,12	3.0
			EM10H_2MU6	17,2x6	2.5	EM10HR_2MU6	17,2x6	1.0
TAU40	100	52	TAU80V	180	4.7	TAU80VR	180	3.2
2TAU11I_TAU15	30,40	147	2TAU50V	2x110	3.8	2TAU40VR	2x100	3.9
2TAU11I_EM14VH	30,21	60	2TAU20VI_3J20	2x50,60	5.2	2TAU15VR_3J15	2x40,50	8.1
			2TAU20VI_			2TAU15VR_		
			EM18VHI_3J18	50,25,60	2.8	EM13HR_3J13	40,20,50	3.3
			TAU15VI_MU15	40,20	3.8	TAU11VR_MU11	35,12	6.4
TAU15_XE35	40,80	63	TAU20VI_			TAU15VR_		
			XE40_3J20	50,90,60	4.4	XE40_3J15	40,90,50	5.0
Tau total		238			20			25
J75	200	34	J100	200	7.0	J100	200	7.0
4J15	4x55	87	4J25	4x60	3.3	4J25	4x60	3.3
			J75_XE40	150,150	8.3	J75_XE40	150,150	8.3
XE40	120	157	XE90	250	10	XE70	200	13
Jet/E_T^{miss} total^a		306			25			25
Topological triggers		-			~5			~20
Total		~800			~100			~100

Table 2: Level-1 Trigger menus for various configurations for a luminosity of $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The columns labelled Run 1 correspond to the menu used for the running at the end of 2012 at $\sqrt{s} = 8 \text{ TeV}$. The columns labelled Run 2 correspond to an example menu for $\sqrt{s} \sim 14 \text{ TeV}$ after LS1. The columns labelled Run 3 correspond to an example menu for $\sqrt{s} \sim 14 \text{ TeV}$ after the Phase-I upgrade is completed at the end of LS2. The offline thresholds typically correspond to the point at which the trigger turn-on curve reaches 90–95% of its plateau value. The items listed in this table assume no overlap removal. For example, in the item 2TAU20I_3J20, for two taus and one jet, it is assumed that both tau candidates will also cause jet triggers. In this case, the tau candidates must pass the 20 GeV tau p_T cut and, additionally, the 20 GeV jet p_T cut.

^aJet/ E_T^{miss} items in the Run 1 menu are assumed to be vetoed on the first 12 bunches to avoid huge bunch train effects giving a 15% inefficiency. Also for Run 1 some of the offline jet thresholds were set at the point where the efficiency reached 99% of its plateau value.

2.4 HLT at Phase-I

This section describes the expected performance of the ATLAS High-Level Trigger (HLT) after the Phase-I upgrade.

The ATLAS High-Level Trigger was originally designed to reduce the 100 kHz Level-1 trigger rate by a factor of approximately 500 to around 200 Hz, but in practice the HLT output rate has been considerably higher. The output rate of the HLT is not limited by the speed at which data can be written to disk, but instead by the total amount of offline computing resources, for example disk storage space, available to the ATLAS collaboration in their grid-based distributed data analysis system. During the most recent LHC pp run in 2012, reorganisation of the offline resources allowed approximately 2.3×10^9 events to be recorded to the data streams that were promptly reconstructed. This corresponded to an average rate during periods with colliding beams of about 350 Hz. In addition to promptly reconstructed events, approximately 0.7×10^9 events were written to tape in the “delayed stream” and reconstructed after the 2012 pp run was finished.

The increased average output rate in 2012 resulted in a considerable improvement in the physics performance of the ATLAS experiment. The low- p_T thresholds for single leptons, di-tau events and E_T^{miss} allowed increases in acceptances for many physics processes. In addition, online selection criteria were set loose enough to allow some cuts to be reversed to obtain data-driven backgrounds for most physics analyses.

In Run 2, ATLAS plans further rationalisation of computing resources to allow running with an HLT output rate of 1 kHz. There is a contingency plan, involving considerable loss in physics performance, to run at 0.5 kHz should this be needed. The average output rate in Run 3 is likely to remain the same as in Run 2, as no major increase in grid resources is expected, even though the average rate of interesting physics events could increase by a factor of 10 or more over the values in Run 1.

The migration of much of the rejection from the current HLT menu to the Level-1 menu in the Phase-I upgrade, combined with higher thresholds giving a sample more enriched in physics events, requires a considerable improvement in the HLT performance to meet rejection targets for the Phase-I upgrade. Those rejection targets will be met in part by exploiting the capabilities of the FTK tracking system. FTK will provide full silicon tracking information to the HLT within $100 \mu\text{s}$ of every event that is accepted by the ATLAS Level-1 Trigger at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The availability of this information will have a significant effect on the HLT processing structure. First, event-level quantities such as the primary vertex position and the number of primary vertices will be available for all events in the HLT. This information can be used to determine if a triggered object originated from the same vertex that was identified as the primary vertex of the hard interaction, as well as to apply pile-up dependent corrections. Second, track information can be integrated into the very start of HLT processing which is not possible in the pre-Phase-I architecture for most trigger chains. Typically, calorimeter or muon system based selections are applied first to reduce the rate to a level where the tracking can be run. The availability of early tracking information allows for an event and object selection which is more closely matched to offline selections, improving the physics performance of the HLT. To obtain a similar performance without the FTK would require prohibitively expensive expansion of the HLT CPU farm as discussed in Section 7.3. Details of how the FTK will be integrated into the HLT processing framework is discussed in Section 6.5.1.

2.4.1 HLT selections based on single electron and muon triggers

For the single-lepton triggers, the production of W and Z bosons provides an important constraint on the HLT trigger menu. The goal of the Run 3 physics programme is not necessarily to collect as many single W and Z bosons as possible, but rather to use them to trigger physics processes such as associated Higgs production (WH and ZH) with minimal bias. This is achieved by postponing additional trigger selections, except for those on the single lepton, until the full detector information is available at the HLT.

At $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the total rate of electrons and muons from W boson decay is approximately 1.2 kHz, which by itself would saturate the allowed HLT output rate if it were possible to trigger on all of these. About 50% of these electrons and muons are inside the trigger acceptance, but even electron and muon selections with very tight isolation requirements at the HLT have considerable background from heavy quark decays and fake jets, raising the rate by a counterbalancing factor of two. To obtain reasonable rates, it may be necessary to move some of the physics analyses requirements to the HLT.

The requirements imposed by the HLT can be based on objects with near offline reconstruction quality. Objects that will be considered include muons and tau leptons (without Level-1 trigger seeds) as well as jets, with b -tagging where needed. It will also be possible to use event-level quantities such as E_T^{miss} . This will require full-event reconstruction in the calorimeters and tracking for most events that pass the electron and muon selections. Given the lepton-only rejection rates in Table 3, the rate of full reconstruction for electron and muon events is expected to be approximately 2.5 kHz.

The reconstructed objects will have performance very close to those of the offline reconstructed quantities to avoid unnecessary inefficiencies. Using reconstruction techniques similar to those currently used in the present HLT, but improved to incorporate offline-style pile-up corrections, it will be possible to identify jets down to 20 GeV, far below what will be possible at Level-1 after the Phase-I upgrade.

The requirements moved from the physics analysis into the HLT will depend on the physics process of interest. Examples of cuts that might be applied at the HLT would be an E_T^{miss} cut of 30 GeV or the existence of a jet, in the tracking acceptance, associated to the same primary vertex as the lepton, with a p_T of at least 20 GeV, possibly with a loose b -tag. These HLT selections would be designed to reduce the rate of events seeded from Level-1 single-electron and muon triggers to less than 400 Hz.

2.4.2 HLT selections based on hadronically decaying tau leptons

The 20 kHz Level-1 tau rate puts a considerable demand on the Level-2 ROI processing budget. In Run 1, this issue was solved by requiring a first selection based only on coarse calorimeter information which reduced the rate to a level acceptable for ROI track finding. Then further selection was made using track and more refined calorimeter quantities. Because of the relative coarseness of the calorimeter information used in the first stage, this selection resulted in a 10–15% inefficiency even for high- p_T taus. In [2.7] it was shown that replacing the existing Level-2 algorithm with a simple selection based on the numbers of FTK tracks in a core and an isolation cone, defined with respect to the highest- p_T track in the tau ROI, recovered a significant portion of the inefficiency. When applied to a $H \rightarrow \tau\tau$ final state it resulted in a 22–28% increase in the event yield. It is expected that the purity of the tau objects in Run 3 will be higher than in Run 1, however, an algorithm similar to what was developed in Ref. [2.7] will still be needed to do a fast rejection before more sophisticated algorithms are run. These

Signature	2012 selection HLT rejection	2012 selection used at Phase-I HLT rejection	2012 fraction selected offline	Phase-I (stipulated) HLT rejection
single- e	180	60	$\sim 80\%$	60
single- μ	80	25	$\sim 95\%$	25
di- τ	500	250	$\sim 70\%$	300

Table 3: Examples of HLT lepton selections. The first column gives the rejection of the HLT achieved in 2012. The second column gives expected rejection if the 2012 selection were run on the Phase-I Level-1 output. The third column gives the fraction of 2012 events which are selected by the corresponding offline selection. The last column gives the desired rejection factor of the HLT in Phase-I. Because of the high purity of the single-object triggers, additional rejection must be obtained using event properties such as jets reconstructed in the HLT.

algorithms would use topo-clusters and tracks, and mimic the offline multi-variate selection. Additionally, FTK tracks can be used to do full-scan tau finding to recover inefficiencies due to the relatively high tau p_T threshold at Level-1.

Final states involving hadronically decaying tau leptons predominantly have at least two trigger objects in the final state, allowing for very large rejection factors, as can be seen in Tables 2 and 3. Therefore, it will generally be unnecessary to add objects beyond those used to seed the triggers at Level-1 in order to obtain low HLT output rates. Similarly multi-object tau triggers with electrons and muons will have high rejection rates and will not require the identification of additional objects in the HLT.

2.4.3 HLT selections based on hadronic triggers

The third category of final states given in Table 2 involves jets and E_T^{miss} , and accounts for approximately 20% of the Level-1 bandwidth.

The most important of these hadronic triggers in the initial running at the LHC has been the E_T^{miss} trigger. The E_T^{miss} trigger has figured prominently in the Higgs analyses (e.g. $ZH \rightarrow \nu\bar{\nu}b\bar{b}$ as discussed in Section 2.5.1), as well as in the majority of the SUSY and Exotics analyses with hadronic final states. In Run 2 and Run 3 the expected increase in pile-up will inevitably degrade E_T^{miss} performance and also lead to a large increase in soft jets from pile-up. Sophisticated pile-up-suppression techniques will be used at the HLT, and FTK tracks can be used to provide offline-style E_T^{miss} corrections by correcting the energy of soft clusters in the event. Soft clusters are defined as calorimeter clusters which have $E_T < 20$ GeV. With $\langle\mu\rangle = 40$ at $\sqrt{s} = 8$ TeV this method has been shown to improve the rejection of di-jet events by 30–40% for a fixed efficiency for $ZH \rightarrow \nu\bar{\nu}b\bar{b}$.

For generic E_T^{miss} triggers, and for combined jet E_T^{miss} triggers, the HLT must be closely matched to the offline algorithm to maximise the overlap of events selected by the two algorithms. The E_T^{miss} threshold for a generic E_T^{miss} algorithm will have to be approximately a factor of 2 or 2.5 higher than the Level-1 threshold to obtain a factor of 100 rejection. It is anticipated that there will be a large number of specialised HLT algorithms seeded from Level-1 E_T^{miss} , which are based on jets (with and without b -tags) that can have E_T^{miss} thresholds closer to the Level-1 thresholds. For example, offline E_T^{miss} -based algorithms typically apply cuts on the direction of the E_T^{miss} relative to hard jets in the event, that can be implemented at the HLT.

Finally, for generic single-jet triggers and multi-jet triggers, full calorimeter read-out allows more accurate determination of jet energies using the full offline jet calibration procedure, including pile-up suppression and correction. This allows the HLT thresholds to be placed very close to the offline ones. The full read-out also allows iterative event-level jet-finding algorithms, such as the anti- k_T algorithm, to be used.

Maximal efficiency of the HLT is obtained if the full calorimeter is read out and jet finding and the HLT E_T^{miss} determination are made for each of the hadronically-triggered events. This requires that the HLT supports a full calorimeter read-out rate in excess of 20 kHz (see Section 7.1.1) and that sufficient CPU resources are available for the topological clustering of the calorimeter information, as well as calibration close to what is used offline (see Section 7.3). Contingency for uncertainties in the calculation of offline resources is provided by algorithms that were used in the 2012 run at Level-2 in the HLT. For the E_T^{miss} calculation, fast sums from the full-calorimeter read-out can be used. In 2012, these sums could be used to reduce the E_T^{miss} rate by a factor of four with no loss in efficiency. However, after the Phase-I upgrade, the Level-1 E_T^{miss} determination will be improved and it is expected that this technique will be less effective, reducing the gain by a factor of two. In addition, the fast sums can not easily be corrected for artefacts in the out-of-time pile-up from bunch trains. At Phase-I ATLAS will also have much improved Level-1 trigger information from the LDPS [2.4] system that will allow full anti- k_T jet finding, resulting in a similar gain.

b -tagging is a powerful tool for background rejection which is only available at the HLT level. However, in Run 1, b -tagging algorithms were a large consumer of HLT CPU time. To maintain the same relative CPU usage as pile-up increases, the p_T threshold of b -tagged jets must be raised to compensate for the increase in the number of jets present in the events and the more complicated pattern recognition task. By removing the pattern recognition and track fitting step, FTK allows the HLT to process many more ROIs for b -tagging. FTK tracks can be used directly by the tagger, or they can first be refit with the HLT fitting algorithms to improve the track parameter resolutions. In [2.7] it was shown that simple, pile-up robust FTK b -tagging algorithms could be developed with only modest decreases in efficiency with respect to offline taggers. It should be noted that b -tagging can be applied to soft jets that are invisible at Level-1, but found from either full calorimeter reconstruction or from jet finding based on the results of tracks from the FTK processor.

2.5 Physics Studies

This section concentrates on the role of the Phase-I upgrades for Higgs physics and boosted objects. As already pointed out in Section 2.1, the low single lepton and di-lepton thresholds, as well as the low E_T^{miss} thresholds that are needed for the study of Higgs physics covers almost the entire phase space needed for the study of Standard Model, SUSY, Top and Exotic processes. The Physics studies below concentrate on the role of the topological trigger that is already relevant in Run 2. Additional physics studies can be found in the companion TDRs for the LAr Calorimeter [2.4] and the NSW [2.5].

2.5.1 Higgs couplings and properties

The precise determination of Higgs branching ratios and properties requires samples with large statistics and small (or well understood) trigger biases. For the recently discovered Higgs boson with a mass near 125 GeV, the detectable final states include pairs of vector bosons, ZZ , WW and $\gamma\gamma$ and pairs of fermions, that include $\tau\tau$ and bb . The coupling of

the Higgs boson to top quarks can be probed using the $t\bar{t}H$ final state. The couplings of the Higgs bosons can be further constrained by measuring its associated production with W and Z bosons and its production in Vector Boson Fusion (VBF). The constraints from associated production and VBF avoid the theoretical uncertainties in gluon fusion production of Higgs. The Higgs spin and parity are constrained by measurements on the decaying products of the boson and fermions. Both the branching ratios and spin-parity measurements are aided by having low p_T thresholds on electrons, muons and tau leptons.

Higgs decays to ZZ , WW and $\gamma\gamma$ The hermetic ATLAS electromagnetic calorimeter and associated trigger system ensures that Higgs decays to at least one pair of EM objects ($H \rightarrow ZZ \rightarrow e^+e^-\ell^+\ell^-$, $H \rightarrow \gamma\gamma$ and $H \rightarrow WW \rightarrow e^+\nu_e e^-\bar{\nu}_e$) are efficiently triggered for p_T of 25 GeV or greater, even without the Phase-I upgrade. The largest gains from the Phase-I upgrade, for these processes, come in channels with muons. Due to the limited geometric muon trigger efficiency of 72% per muon in the ATLAS barrel muon spectrometer, events of interest are inevitably lost if di-object triggers are used. The p_T of the leading and sub-leading leptons for $H \rightarrow WW \rightarrow e^+\nu_e \mu^-\bar{\nu}_\mu$ ⁷ are shown in Figure 5. For example, if the leading lepton is an electron and the muon is lost (due to the limited geometric coverage) there is an efficiency loss of approximately 25% if the offline threshold must be raised from 32 GeV to 38 GeV. Without the Phase-I upgrade the overall efficiency loss for the $H \rightarrow WW \rightarrow e^+\nu_e \mu^-\bar{\nu}_\mu$ 0-jet analysis, as documented in reference [2.4], is 12%.

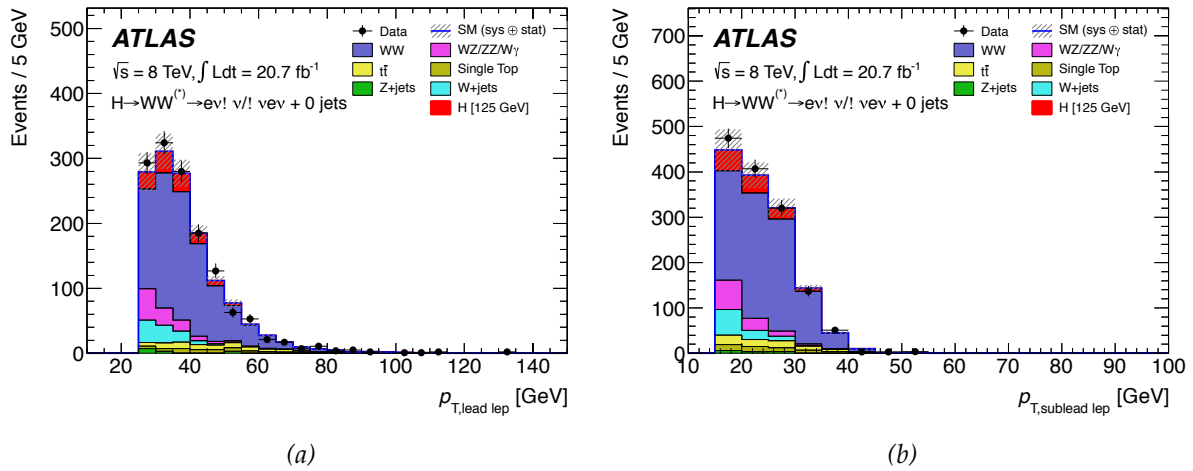


Figure 5: Leading (a) and sub-leading (b) p_T of leptons from $H \rightarrow WW \rightarrow e^+\nu_e \mu^-\bar{\nu}_\mu$

$H \rightarrow \tau\tau$ The analysis of $H \rightarrow \tau\tau$ is based on three channels depending on the decay products of the τ -lepton, $H \rightarrow \tau\tau \rightarrow \ell\ell$, $\ell\tau_h$ and $\tau_h\tau_h$, where ℓ represents e or μ and τ_h represents the hadronic decays. The τ -lepton decays to e ($\tau \rightarrow e\bar{\nu}_e\nu_\tau$) or μ ($\tau \rightarrow \mu\bar{\nu}_\mu\nu_\tau$) have a combined branching fraction of 35%. This leads to final states composed of 46% $\ell\tau_h$, 42% $\tau_h\tau_h$ and 12% $\ell\ell$. Since a large fraction of the final states contain at least one hadronic decay, efficient triggers are needed to identify the hadronic τ decays.

In the $\ell\ell$ channel, di-lepton (ee , $\mu\mu$ and $e\mu$) triggers are used as well as the single-lepton triggers (e and μ). Inevitably, the lepton momentum from subsequent decay are softer (see Figure 6), so that low p_T trigger item combined with other trigger objects will be more efficient

⁷Charge conjugate states are implied throughout the text unless explicitly stated otherwise.

than the single-lepton triggers with relatively higher p_T threshold. Single-electron triggers are also important to recover events that are lost due to muon inefficiencies in the barrel muon spectrometer, as mentioned above for the $H \rightarrow WW$ channel.

Since the di-lepton trigger rates given in Table 2 are manageable (see 2EM15VHI, 2MU11 and EM15VH_MU10), the trigger menu does not rely on further additional trigger object combination or topological selection. On the other hand, in the $\ell\tau_h$ and $\tau_h\tau_h$ channels, the high rates force high thresholds compromising the precision with which the branching ratio of $H \rightarrow \tau\tau$ can be measured. In addition, measuring as much of the p_T spectra of the hadronic decays as possible is useful in constraining the parity of either SM or BSM Higgs particles.

The trigger menu is designed to identify the two accessible event topologies that are currently used in the offline analysis: VBF and boosted events from gluon fusion. The boosted category selects those gluon fusion events in which the Higgs boson is highly boosted in the transverse plane and is associated with a high p_T jet in the opposite hemisphere. In both cases, events have at least one additional jet. Therefore, the trigger is designed to have $e + \tau_h + \text{jet}$ (2TAU20VI_EM20VHI_3J20), $\mu + \tau_h (+ \text{jet})$ (TAU15VI_MU15) and $2 \tau_h + \text{jet}$ (2TAU20VI_3J20). A further rate reduction can be expected when the additional jet requirement is also applied in the $\mu + \tau_h$ channel. The rate reduction by the addition of the one jet requirement (J20) is about 75%, while signal acceptances in the VBF and boosted categories are under control, with more than 95% of the events passing the offline jet requirement retained with the online requirement. As can be seen in Table 2, the Phase-I upgrade allows the offline thresholds for hadronic taus to be reduced from 50 to 40 GeV which is at the peak of the true p_T spectrum shown in Figure 6. In the $e - \tau_h$ final state, without the Phase-I upgrade, there is a 37% loss of trigger acceptance [2.4]. In the $\mu - \tau_h$ final state, the loss in trigger acceptance without the Phase-I upgrade is approximately 32%. Without the Phase-I upgrade an even larger loss is expected for $\tau_h\tau_h$.

To further extend the acceptance, additional topological selections at Level-1 can be employed with lower p_T cuts. An approach, unbiased with respect to production mechanism, is to apply the $\Delta\eta$ cut to the two τ objects. With the topological selection of $\Delta\eta < 2.0$, the rate is further reduced by 25% without losing any signal events, as shown in Figure 6. A more aggressive approach, which primarily increases the acceptance for the boosted events, is to use the ΔR and $\Delta\phi$ topological selections that have a stronger rate suppression than the $\Delta\eta$ selection. The topological selection by $\Delta R < 2.8$ can achieve 50% rate reduction with respect to the one without any topological selection. As an example, in Run 2 the sub-leading tau lepton trigger threshold could be lowered by 5 GeV with the combination of $\Delta R < 2.8$ and the requirement of an additional jet, while adding a negligible rate of only 1.3 kHz.

$H \rightarrow b\bar{b}$ Detection of the process $H \rightarrow b\bar{b}$ is more difficult than $H \rightarrow \tau\tau$. The current most sensitive ATLAS analyses are based on the associated production of a W or Z boson and Higgs boson, referred to as the VH process. In the case of decays of the W boson to e or μ , the single-lepton triggers are efficient for both the Run 2 and Run 3 menu. The largest gains from the Phase-I upgrade come from the added rejection of the topological trigger for the $ZH \rightarrow \nu\bar{\nu}b\bar{b}$ processes, which is considered in some detail below.

$ZH \rightarrow \nu\bar{\nu}b\bar{b}$ The $ZH \rightarrow \nu\bar{\nu}b\bar{b}$ decay was selected in 2012 with the E_T^{miss} trigger. The Level-1 threshold was set to 40 GeV (i.e. XE40) corresponding to 120 GeV offline. With the expected increase of pile-up and centre-of-mass energy for Run 2, the XE40 rate will be above the *total* allowed Level-1 rate (see Table 2). Increasing the E_T^{miss} threshold at Level-1 rapidly suppresses

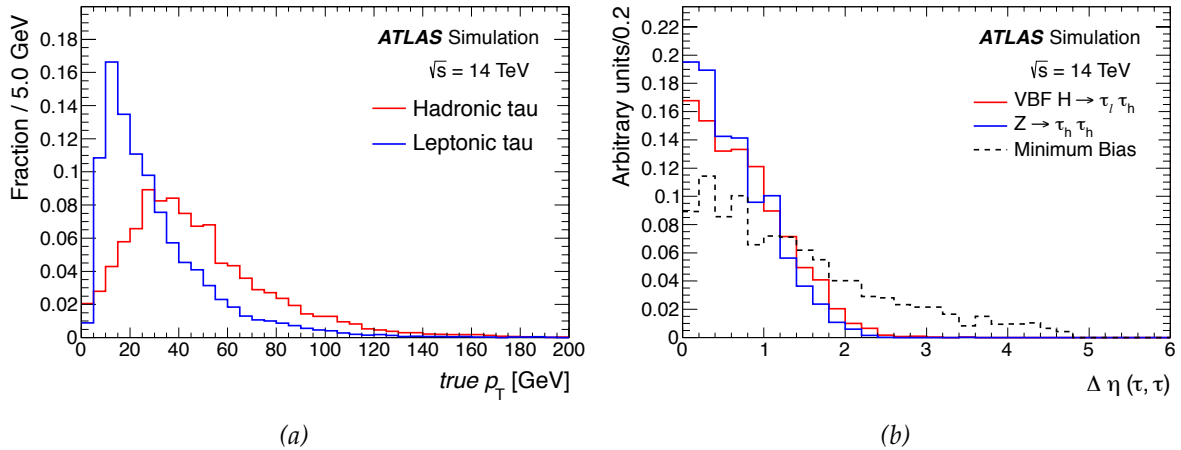


Figure 6: (a) The p_T spectrum of leptons and hadrons from the decay $H \rightarrow \tau\tau$ for $m_H = 125$ GeV Higgs produced in the VBF process. (b) The distribution of $\Delta\eta$ for tau candidates with a leading tau exceeding an online threshold of 20 GeV and the sub-leading tau exceeding an online threshold of 12 GeV.

signal acceptance. Both combinations of E_T^{miss} and jet requirements as well as topological selections have been studied using Run 2 Monte Carlo samples. When combining E_T^{miss} above 50 GeV with an inclusive jet trigger with $p_T > 40$ GeV (i.e. XE50_J40), the Level-1 rate is expected to decrease to approximately 10 kHz, which is still too large to include in the trigger menu. Offline selections to reduce multi-jet background include the requirement of exactly two or three central jets, ($p_T > 20$ GeV and $|\eta| < 2.5$) and a minimum value of $|\Delta\phi(E_T^{\text{miss}}, \text{jets})|$ (the azimuthal angular distance between Level-1 E_T^{miss} and central jets) as well as a maximum value of $\Delta R(\text{jets})$ (minimum radial distance between the jets). These quantities are shown in Figure 7. A loose requirement of the minimum $|\Delta\phi(E_T^{\text{miss}}, \text{jets})| > 1$ in addition to XE50_J40 provides very good signal acceptance for an expected Level-1 rate of ~ 5 kHz. The unique rate is expected to be a factor of three smaller and can be accommodated in the Run 2 and Run 3 trigger menus.

For Run 3, after all of the of the Phase-I upgrade is in place, it should be possible to reduce the Level-1 E_T^{miss} threshold to approximately 70 GeV as described in LAr Calorimeter TDR [2.4] and a trigger strategy more similar to the Run 1 one could be employed. Ultimately both approaches, standalone E_T^{miss} triggers and combined topological triggers, as described above, are likely to be used.

$t\bar{t}H$ The newly discovered Higgs boson with mass of 125 GeV is too light to decay into pairs of top quarks, but the process $t\bar{t}H$ can be used to probe its coupling to top quarks. Initially, the decays of the Higgs to WW , bb , $\tau\tau$ and $\gamma\gamma$ [2.8, 9, 10, 11, 12] will be exploited.

In Run 1, analyses of most final states including $t\bar{t}$ have used triggers based on the leptons from top decay, with p_T thresholds of 25 GeV. For muons, this threshold can be maintained before and after the Phase-I upgrade. For electrons, the Run 2 trigger with an offline threshold of $p_T > 38$ GeV will have a significant inefficiency. This can be addressed by using a trigger that adds an additional jet. For example, EM18VH_3J20 will have an additional inefficiency for $t\bar{t} bb$ signal of about 2.5% when compared to the Run 1 trigger and a unique rate of about 2 kHz. In the WW and $\tau\tau$ final states the situation is even better.

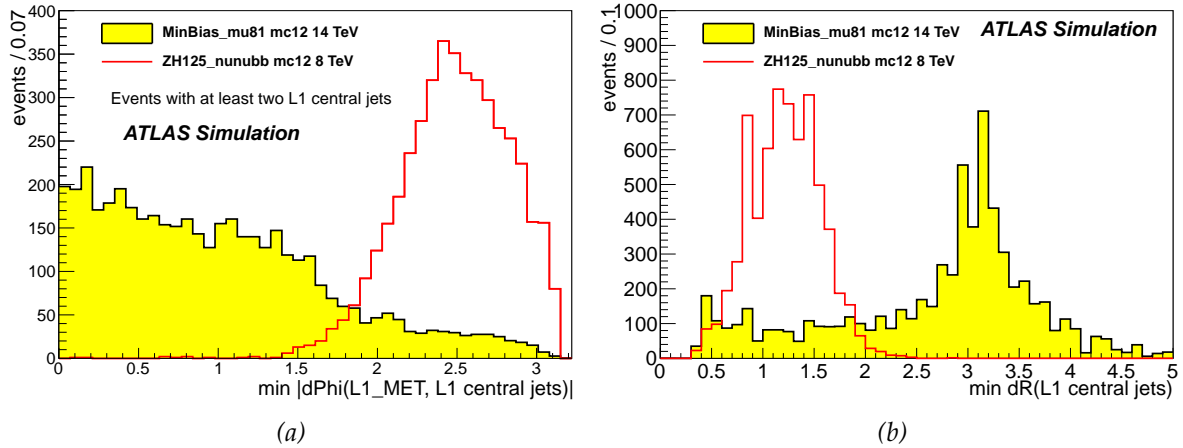


Figure 7: (a) Minimum azimuthal angular distance between Level-1 E_T^{miss} and Level-1 central jets with $p_T > 20 \text{ GeV}$ and $|\eta| < 2.5$ for minimum bias (filled histogram) and ZH (open red histogram) events with at least two central jets. (b) Radial distance (ΔR) between the Level-1 central jets. The distributions are normalised to the same area. (The ZH signal was simulated at 8 TeV, however, the distributions of the 14 TeV signal are expected to be very similar.)

2.5.2 Boosted objects

The study of Lorentz-boosted high- p_T objects such as W and Z bosons and top quarks has been a major ATLAS Run 1 priority in both Standard Model measurements and in searches for new physical phenomena. The increased energy of Run 2 and beyond will increase the importance of these boosted modes. The study of high- p_T Higgs (now relevant for BSM scenarios) to $b\bar{b}$ has long been advocated as a unique opportunity to identify Higgs bosons in this decay channel [2.13]. Many special techniques have been developed to improve the separation between normal QCD jets and those containing almost all of the decay products of heavy particles such as W bosons, referred to here as EW jets. One defining characteristic of these EW jets is that the energy distribution within the jet is spread considerably wider (contained within typically $R = 1.0$ or larger) than for QCD jets ($R = 0.4$ or 0.6) and that they have multiple hard cores reflecting the partons produced in the EW decay. Jet triggers designed to be efficient for narrow QCD jets will necessarily be inefficient for these wider EW jets. Additionally, lepton triggers do not capture this physics as leptons produced in the decays of these high- p_T particles merge with neighbouring jets and are not isolated. This is illustrated by Figure 8a which shows the acceptance for $Z' \rightarrow t\bar{t}$ in a Run 1 search [2.14]. At large $m_{t\bar{t}}$, events selected with boosted objects and large-jet triggers dominate the overall acceptance; the e +jets channel acceptance decreases with increasing $m_{t\bar{t}}$ as the electrons merge into jets. While the methods for efficiently identifying EW jets can be implemented in the HLT using the standard offline software, these highly-iterative complex algorithms are not appropriate for the Level-1 environment. It is crucial to preserve the acceptance at Level-1 for these critical objects that can span more than two units in η as illustrated in Figure 8b which shows the separation between the quarks resulting from the hadronic decays of high- p_T W bosons [2.15]. Typical p_T in physics analyses are $> 200 \text{ GeV}$ for bosons and $> 300 \text{ GeV}$ for top quarks when not constrained by the trigger acceptance.




Figure 8: (a) The selection efficiency as a function of the true $m_{t\bar{t}}$ for simulated Z' resonances at various mass points. The μ +jets channel is shown with grey lines and the e +jets channel with black lines. Dashed lines show the boosted selection and solid lines the total selection efficiency. (b) The angular distance between the light quark and anti-quark from $t \rightarrow Wb$ decays as a function of the p_T of the W boson. The distribution is at the generator level and does not include effects due to initial and final-state radiation, or the underlying event.

References

- [2.1] T. Sjostrand, S. Mrenna, and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput.Phys.Commun.* **178** (2008) 852–867, [arXiv:0710.3820](https://arxiv.org/abs/0710.3820) [hep-ph].
- [2.2] A. Martin, W. Stirling, R. Thorne, and G. Watt, *Parton distributions for the LHC*, *Eur.Phys.J.* **C63** (2009) 189–285, [arXiv:0901.0002](https://arxiv.org/abs/0901.0002) [hep-ph].
- [2.3] ATLAS Collaboration, *Summary of ATLAS Pythia 8 tunes*, Tech. Rep. ATL-PHYS-PUB-2012-003, CERN, Geneva, Aug, 2012. <https://cds.cern.ch/record/1474107>.
- [2.4] ATLAS Liquid Argon Calorimeter Group, *ATLAS Liquid Argon Calorimeter Technical Design Report: Trigger Electronics for the LHC Phase-I Upgrade*, Tech. Rep. CERN-LHCC-2013-???. ATLAS-TDR-??, CERN, Geneva, (in preparation). <https://twiki.cern.ch/twiki/bin/view/LAr/LArPhaseITDR>.
- [2.5] ATLAS Collaboration, *New Small Wheel Technical Design Report*, Tech. Rep. CERN-LHCC-2013-006. ATLAS-TDR-020, CERN, Geneva, Jun, 2013. <https://cds.cern.ch/record/1552862>.
- [2.6] M. Cacciari, G. P. Salam, and G. Soyez, *The Catchment Area of Jets*, *JHEP* **0804** (2008) 005, [arXiv:0802.1188](https://arxiv.org/abs/0802.1188) [hep-ph].
- [2.7] ATLAS Collaboration, *Fast TracKer (FTK) Technical Design Report*, Tech. Rep. CERN-LHCC-2013-007. ATLAS-TDR-021, CERN, Geneva, Jun, 2013. <https://cds.cern.ch/record/1552953>.
- [2.8] ATLAS Collaboration, *Search for the Standard Model Higgs boson produced in association with top quarks in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, Tech.

- Rep. ATLAS-CONF-2012-135, CERN, Geneva, Sep, 2012.
<https://cds.cern.ch/record/1478423>.
- [2.9] ATLAS Collaboration, *Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics*, [arXiv:0901.0512](https://arxiv.org/abs/0901.0512) [hep-ex].
- [2.10] C. Boddy, S. Farrington, and C. Hays, *Higgs boson coupling sensitivity at the LHC using $H \rightarrow \tau \tau$ decays*, *Phys.Rev.* **D86** (2012) 073009, [arXiv:1208.0769](https://arxiv.org/abs/1208.0769) [hep-ph].
- [2.11] ATLAS Collaboration, *Search for $t\bar{t}H$ production in the $H \rightarrow \gamma\gamma$ channel at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2013-080, CERN, Geneva, Jul, 2013.
<https://cds.cern.ch/record/1564319>.
- [2.12] A. Liss and J. Nielsen, *Physics at a High-Luminosity LHC with ATLAS*, Tech. Rep. ATL-PHYS-PUB-2013-007, CERN, Geneva, Jul, 2013.
<https://cds.cern.ch/record/1564937>.
- [2.13] J. M. Butterworth, A. R. Davison, M. Rubin, and G. P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys.Rev.Lett.* **100** (2008) 242001, [arXiv:0802.2470](https://arxiv.org/abs/0802.2470) [hep-ph].
- [2.14] ATLAS Collaboration, *A search for $t\bar{t}$ resonances in the lepton plus jets final state with ATLAS using 4.7 fb^{-1} of pp collisions at $\sqrt{s} = 7$ TeV*, *Phys.Rev.* **D88** (2013) 012004, [arXiv:1305.2756](https://arxiv.org/abs/1305.2756) [hep-ex].
- [2.15] ATLAS Collaboration, *Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, [arXiv:1306.4945](https://arxiv.org/abs/1306.4945) [hep-ex].

3 Level-1 Calorimeter Trigger

3.1 Introduction

The Level-1 calorimeter trigger (L1Calo) is a pipelined system processing signals from the electromagnetic and hadronic calorimeters in real-time to produce trigger signals to the Central Trigger Processor (CTP). The initial system produces object counts above threshold for e/γ , τ , jet, ΣE_T , E_T^{miss} and missing E_T significance (XS) using custom hardware processors based on FPGA technology. In the future, the LHC luminosity will increase beyond $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and the associated pile-up will also increase. To manage increasing rates in this regime, whilst preserving an effective and efficient calorimeter trigger, significant upgrades will be necessary. Specifically, more refined processing of electromagnetic calorimeter information, at higher granularity, will enable better discrimination between photons, electrons, taus and jets and hence provide efficient single object triggers for electroweak-scale physics. Improved processing, also at higher granularity, will enable better identification of jets and, using a wider search window, will open up the possibility of identifying broader jets. In addition to this, processing of event topology information (based on position and transverse energy estimated within the resolution available in the trigger) will enable more exclusive triggering targeted at specific physics processes.

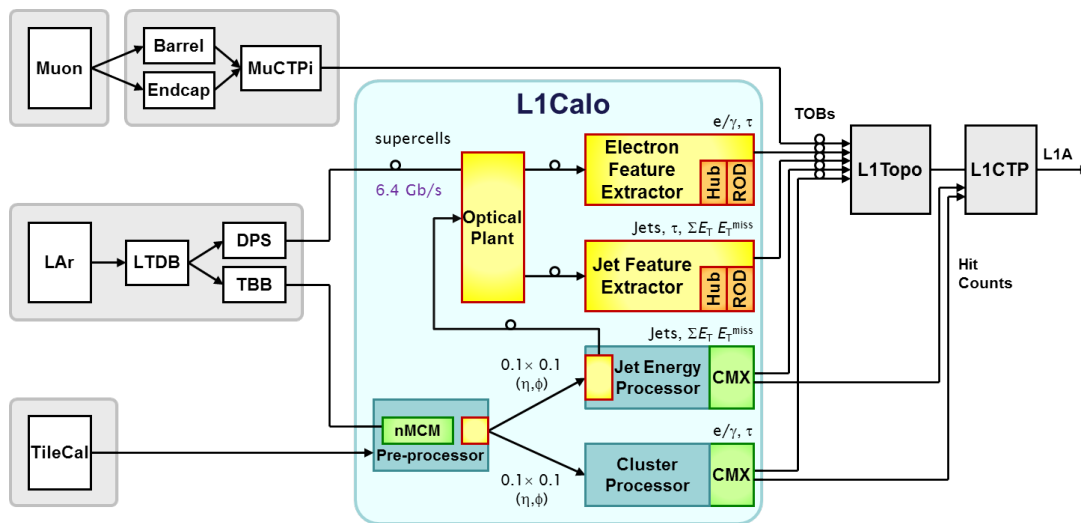


Figure 9: Overview of the Level-1 trigger system planned for Run 3

Figure 9 gives an overview of the Level-1 trigger system for Run 3, with L1Calo set in the context of other system components. New signal digitisation in the current Pre-Processor system will improve bunch crossing identification and E_T estimation in the face of increased pile-up and will be based on a replacement of the Multi-Chip Modules (MCM). The existing Common Merger Modules (CMM) in the Cluster Processor (CP) and Jet/Energy Processor (JEP) systems will be replaced with new Common Merger Modules (CMX) to enable transmission of position and transverse energy estimates to a new Topological Processor (L1Topo), in addition to existing count-above-threshold information sent to the CTP (see Section 5.3). L1Topo will run algorithms to flag interesting event topologies to the CTP. Algorithms will rely on event features such as opening angles between Trigger Objects (TOBs), effective masses of TOB pairs, etc. These upgrades will be installed during LS1 to meet anticipated triggering challenges in the subsequent run.

For Run 3 the CP and JEP systems will be replaced with new hardware and algorithms working with finer granularity information. Liquid Argon Calorimeter (LAr) signals will be digitised by a new Digital Processing System (DPS) to produce E_T measurements at a finer granularity (illustrated in Figure 20 of Section 3.3.1). Information from the Tile Calorimeter will be derived from the existing digitisation system (within L1Calo) at maximum granularity available there. These data will be sent to the electromagnetic and jet feature extractor subsystems (eFEX and jFEX respectively) over multi-gigabit optical fibre links. The eFEX subsystem will employ new cluster-finding algorithms on the higher granularity data to produce more refined TOBs. The jFEX subsystem will use jet-finding algorithms on higher granularity data.

This chapter describes components to be developed for installation during LS1 and LS2. Section 3.2 discusses the algorithms and anticipated performance of the system, whilst Section 3.3 describes the system evolution from the present architecture. The full Run 3 architecture is presented in Section 3.4, with subsequent sections covering firmware, software, testing, installation, commissioning and evaluation of resultant performance. Section 3.10.1 introduces the gFEX which is an option being considered for global event processing. (Since, at present, it is only an option under study, it has not been referenced in other sections or diagrams.) The relationship to Phase-II upgrades is discussed in Section 3.11, in the context of forward compatibility, and the final section introduces the project organisation.

3.2 Algorithms & Performance

3.2.1 Input data

A key feature of the DPS is that it will provide finer granularity inputs to the L1Calo algorithms. Currently these are “trigger towers” of $\sim 0.1 \times \pi/32$, formed by analogue summation of calorimeter cells. In the proposed design, the DPS will provide information from up to ten “SuperCells” within each electromagnetic trigger tower, where a SuperCell is a sum of four or eight calorimeter cells. Each tower (see Figure 20) will contain one SuperCell from layers zero (Pre-Sampler) and three (back sample) of the EM calorimeter, and four SuperCells of finer granularity in layers one and two. Details of the signal E_T scales are still being optimised, but in the studies below a baseline E_T scale of 250 MeV/count has been used. The full granularity will be available to the e/γ and τ triggers in the eFEX, while the jet and E_T triggers will use a granularity comparable to the current towers (described below).

3.2.2 Electron/photon trigger

The electron/photon trigger is one of the key physics triggers. Rising luminosity not only increases the rates of both signal and background, but the increased pile-up levels reduce the effectiveness of isolation in rejecting background unless lower efficiencies are accepted. However, much of the physics motivation of ATLAS remains in studying electroweak-scale objects, including Higgs bosons, so increasing thresholds reduces acceptance. One of the key motivations of the upgrade is therefore to improve the selectivity of the electromagnetic trigger through the use of finer granularity information.

There are two key elements to an e/γ trigger. The first is forming a cluster which is large enough in area to contain the energy of an electromagnetic shower, providing a sharp turn-on, while being small compared with a typical jet, in order to reduce the background rate. In LHC conditions this is not, by itself, sufficient except at high E_T , and so the second element is to use information, such as the shape of the cluster or the activity around it, to provide

additional rejection against jets. The improved segmentation of the calorimeter information available after LS2 allows both of these elements to be refined relative to the existing trigger.

The current Level-1 EM trigger [3.1] is based on towers of $\sim 0.1 \times 0.1$, with no depth segmentation in EM or hadronic calorimeters. The E_T of the candidate is estimated using the sum of two electromagnetic towers, adjacent either in η or ϕ , while the sums of E_T in 2×2 hadronic towers behind the cluster and the surrounding 12 towers in the electromagnetic calorimeter are used for isolation.

In the Run 3 algorithms, electromagnetic clusters are built around a “seed” - a SuperCell in layer two of the electromagnetic calorimeter, where most of the shower energy is deposited and which is more energetic than any of the surrounding SuperCells. A cluster is formed by summing this cell with the more energetic of its neighbours in the ϕ direction, then summing both neighbours in the η direction, to form a cluster of 3×2 SuperCells. The corresponding cells from layer one are added, plus the pair of cells in layers zero and three corresponding to the towers containing the seed and its neighbour in ϕ . Given that the full calorimeter granularity is not available to the Level-1 trigger, this approximates the clustering used in the ATLAS High-Level Trigger and offline. The main part of the cluster slides in steps of 0.025 in η , providing good containment of the showers wherever they are within a tower, and resulting in a turn-on curve with less dependence on the position of the shower within the tower than the current trigger. The efficiency curve for this selection in the barrel region of the detector has already been shown in Figure 3a. With proper optimisation and calibration it is expected that similar results will be obtained for the remaining η regions.

Figure 10, shows the rate vs E_T for the EM cluster alone in Run 3, compared to the Run 1 trigger algorithm, in simulated Run 3 conditions. The rates are plotted against the offline electron p_T for which the trigger is 95% efficient. This simplifies comparison between the two systems, which have different effective E_T scales, and also makes it possible to plot rates for the η -dependent thresholds described above. Without any additional vetoes, the proposed Run 3 system produces a $\sim 20\%$ lower rate for the same efficiency. The plot also compares rates using a fixed, η -independent threshold (giving an η -averaged 95% efficiency) with the η -dependent thresholds. The gain in rate from the η variation is modest, but the more uniform turn-on is a significant benefit.

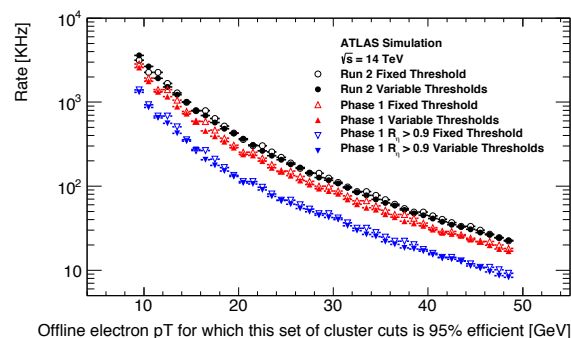


Figure 10: EM trigger rates for Run 2 and Run 3 systems. The E_T scale is the offline electron p_T for which the E_T thresholds used would give 95% efficiency.

The greatest gain in jet rejection comes from use of the finer granularity of the EM calorimeters in layers 1 and 2 to build more powerful jet vetoes. The example that has proved most successful is an “ R_{core} ” variable, defined as the ratio of the E_T in an inner cluster to the total in a larger region which includes the core. A narrow EM object should be mostly con-

tained within the inner cluster, resulting in higher values of R_{core} than are typically produced by more extended QCD jets. One specific example of this is the R_η variable used in offline electron identification, which defines, in layer 2, the ratio of a cluster of 3×7 cells to a 7×7 cell window centred on the cluster. The nearest analogue of this in the Run 3 trigger is $3 \times 2 / 7 \times 2$ SuperCells. Cuts in R_{core} using both larger and smaller environments have been studied. The largest outer region which can be formed within the baseline 3×3 trigger tower environment is 9×3 SuperCells, and studies have not found any gain in performance from using regions larger than this.

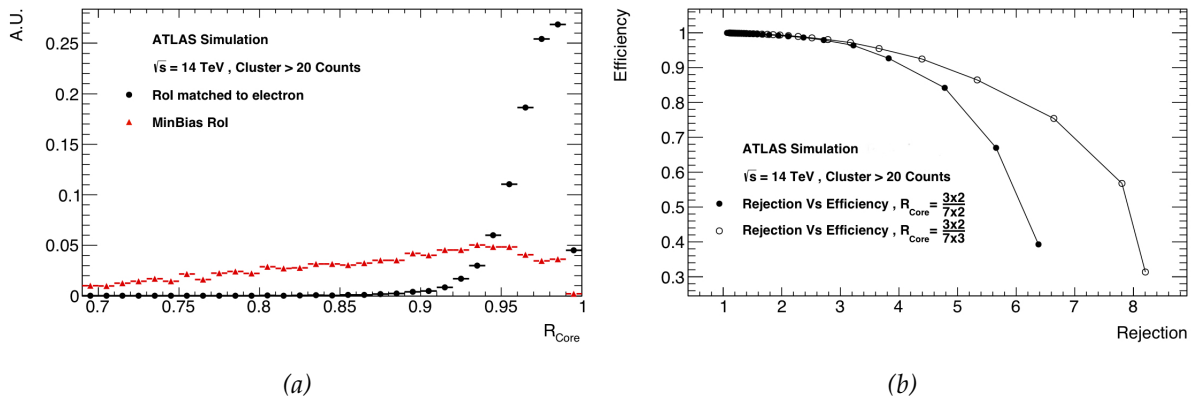


Figure 11: (a) R_{core} distributions for clusters matched to reconstructed electrons and for clusters found in minimum-bias QCD events. (b) Electron efficiency vs rejection of clusters in minimum bias events for two different R_{core} variables. All clusters were required to pass a 20 GeV threshold.

Figure 11a shows R_{core} distributions for electrons from $Z \rightarrow e^+e^-$ events and for jets passing a cluster E_T threshold of 20 GeV. Here, an inner core of 3×2 SuperCells is compared to a larger environment of 7×3 SuperCells. Figure 11b shows electron efficiency vs the rejection factor for clusters in minimum-bias events for this R_{core} variable and the offline-like $3 \times 2 / 7 \times 2$ quantity.

Similar variables may also be formed in layer 1, but as a smaller fraction of the e/γ E_T is deposited in this layer, they give a less clean separation between e/γ and jets, and are strongly correlated with the layer 2 variables. Some studies have shown modest improvements in rejection when using layer 1 in combination with layer 2, but the effect is small and not seen in all studies, so further work is needed to optimise this.




Figure 12: EM trigger rates for Run 2 and Run 3 systems, including R_{core} cuts and hadronic isolation. The E_T scale is the offline electron p_T for which the E_T thresholds used would give 95% efficiency.

In addition to the R_{core} variable in the electromagnetic calorimeter, the hadronic towers in the same window can be used to provide additional jet rejection. Figure 12 shows the rates vs electron p_T predicted for the cluster alone, cluster with an R_{core} cut and cluster with R_{core} cut and hadronic isolation, where the sum of p_T in the 2×2 tower hadronic region behind the EM cluster is used for isolation. Also shown for comparison is the prediction for the Run 2 system with and without isolation. It can be seen that the expected rate for a given electron p_T is 50% of that for the Run 2 system, or equivalently for the same rate it is possible to trigger on electrons of approximately 10% lower p_T .

3.2.3 Tau trigger

Hadronic tau triggering essentially involves the selection of narrow mixed EM and hadronic clusters and separating these from the jet background. It is particularly challenging since hadronic tau decays include a range of different states, are less distinctive than electrons, and also typically have a lower observable E_T due to the neutrinos present in the decays. The approach proposed for tau triggering is to use fine EM granularity information in forming clusters and jet discriminants in the eFEX. The jFEX could also be used to search for taus using a larger area in the algorithms, but without the fine granularity.




Figure 13: (a) Efficiency turn-on for different tau clusters compared with the current trigger. Efficiency is measured relative to truth-matched reconstructed hadronically-decaying taus. (b) Predicted tau trigger rates, E_T cluster only, vs the p_T of the tau at which the efficiency reaches 95%.

The E_T cluster for the tau trigger needs to be sufficiently large to provide a sharp turn-on, but as small as possible to minimise the jet background rate. Figure 13a shows efficiency curves for hadronic taus, comparing the turn-on for the current Level-1 trigger algorithm with a cluster formed from the proposed SuperCells. The new cluster uses a region of 5×2 SuperCells in layer 2 (0.125×0.2), 3×2 in layer 1, 0.1×0.2 in Layer 0 plus 0.2×0.2 clusters in layer 3 and the hadronic calorimeter. In addition, a 3×3 tower cluster (0.3×0.3 , EM+Had) has been studied, and the current algorithm using towers built from SuperCells. For efficiency studies, reconstructed tau candidates matched to a true hadronically decaying tau were used, and the trigger tau was required to match these within $\Delta(R) < 0.2$. The new SuperCell-based cluster and the large cluster produce turn-ons that are slightly improved compared to the current system. Figure 13b shows rates for the different cluster definitions vs the tau E_T for which the trigger is 95% efficient. Here a smaller cluster performs better than a larger one.

It can also be seen that all of the SuperCell-based triggers produce higher rates, even where identical algorithms are used. This is under investigation, but there are indications that the SuperCells are more sensitive to out-of-time signals, which may be due to the incomplete DPS simulation used.

Jet rejection relies on the narrower width of the tau cluster compared to a jet. This can be exploited via isolation, such as the R_{core} variables used in the EM trigger, or through a cluster width variable in the EM calorimeter, which has the fine granularity in η . Preliminary studies indicate that greater rejection can be achieved using the R_{core} isolation than a width variable, so that is used here to estimate performance. A number of different cluster/environment size combinations have been studied, with best results coming from the same inner cluster as used in the EM trigger, 3×2 SuperCells, and a slightly larger outer environment, with 9×3 SuperCells. No advantage has been found in using larger environments, but the connectivity of the eFEX allows expansion of the environment in the η direction if necessary. Figure 14 shows the R_{core} distributions for clusters matched to reconstructed taus and for clusters in minimum bias events.

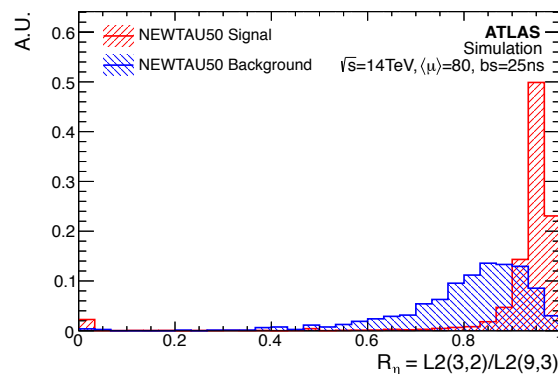


Figure 14: R_{core} distributions for truth-matched tau clusters and clusters from QCD background events

3.2.4 Jet trigger

In the jFEX the finer granularity of the SuperCells is not available. Instead the inputs are towers of $\sim 0.1 \times 0.1$ in the region $|\eta| < 2.5$ and 0.2×0.2 between $2.5 < |\eta| < 3.2$. In the FCAL region the ϕ granularity is coarser again, ~ 0.4 , while the η granularity varies between 0.1 and 0.4. Compared with the current system, this is an improvement in spatial granularity (the current trigger uses 0.2×0.2 for $|\eta| < 3.2$ and has no η segmentation in the FCAL) and also in least count, with a digit scale of 250 MeV/count rather than 1 GeV, and more importantly an improved ability to resolve small signals ($< \sim 2$ GeV) from noise.

The nature of the Level-1 trigger means that offline-style jet-finders, which sequentially search for jets, removing elements used in each jet before searching for the next, are not suitable at Level-1, which must identify all possible jets in a limited and fixed latency. Level-1 jet algorithms therefore must be executed in parallel within a set of overlapping window environments, and cannot use information from outside those environments. The size of the environment is limited by the available bandwidth and the necessary interconnectivity to share data between FPGAs, modules and crates. In the baseline design the environment is 9×9 towers, though options that would allow this to be increased are being investigated.

A 9×9 tower window has an area slightly larger than a $\Delta R = 0.4$ cone. Since the windows overlap, it is necessary to identify which windows should be considered as jet candidates, and

unless the cluster is significantly smaller than the window this cannot be done by requiring that the jet cluster itself is larger than overlapping neighbours. A technique that has been used in the current trigger is to seed the jet using a smaller cluster, which can be identified as a maximum using the information in the window, and then to build a larger cluster around this to estimate the jet E_T . One simple implementation of that would be to form a 3×3 tower cluster in the centre of the window and compare this to nearest and next-to-nearest neighbours (i.e. overlapping cores) in order to determine whether it is a local maximum.

Another approach is based on Gaussian filtering [3.2]. Here towers are summed with a weight which decreases as a Gaussian function of the distance from the centre of the cluster. The motivation is that more distant towers are more likely to have a larger (fractional) contribution from pile-up or other jets, so are given a lower weighting. The key parameter in determining performance is the value used for σ (the Gaussian width) – large values (0.3–0.4) give good inclusive jet performance, while small values (~ 0.1) are good for identifying nearby jets, and perhaps for jet substructure. Within the context of the Level-1 jFEX, a $\sigma = 0.1$ Gaussian cluster can be used to provide a seed, and a wider cluster (Gaussian or otherwise) used for jet E_T measurement. Because of the narrow nature of the $\sigma = 0.1$ clusters it is sufficient to compare only with nearest neighbours. It is also possible to provide more than one jet cluster definition seeded on this, allowing some optimisation for different purposes.

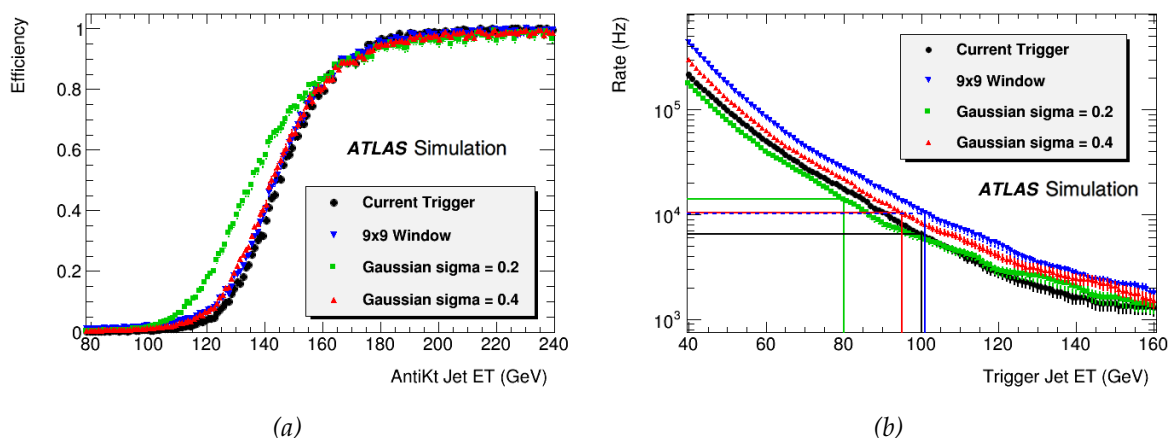


Figure 15: (a) Single jet turn-on curves for different jet algorithms. In each case an offline anti- k_T jet was considered to have been found if there was a trigger jet within $\Delta(R) < 0.4$. The thresholds for the different algorithms were chosen so that the 90% efficiency points would correspond to that of the current algorithm with a 100 GeV threshold. (b) Rates from minimum bias Monte Carlo study. Since the effective E_T scales of the algorithms differ, the lines indicate the threshold values for each algorithm corresponding to 100 GeV in the current trigger, and the corresponding inclusive jet rates.

Figure 15a shows single jet turn-on curves for a simple sliding window algorithm and two Gaussian jets, with $\sigma = 0.2$ and 0.4 , both seeded off a $\sigma = 0.1$ cluster. All are implemented within a 9×9 window. It can be seen that most of the options have similar turn-ons, with the smaller Gaussian being slightly slower to reach full efficiency and hence requiring a lower threshold. Only modest improvements in turn-on come from using a larger environment for the Gaussian algorithm. Figure 15b shows the corresponding rates obtained from minimum bias Monte Carlo at a simulated luminosity of $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. As with the other triggers, it is important to note that different threshold values are needed to achieve the same efficiency with different algorithms. Indicated on the plot are the rates for a 100 GeV threshold with the

current trigger, and the equivalent thresholds for the other triggers discussed. Equivalence is here defined by choosing thresholds which reach 90% efficiency for the same offline jet p_T . The current trigger is much more sensitive to pile-up than the proposed upgrade, and measures which will be implemented for Run 2 to control pile-up effects at the start of the bunch train are modelled here. Without these the rates for the current trigger would be considerably higher. With this, the proposed upgrade triggers produce slightly higher rates than the current system. There is a slight gain from using a Gaussian weight with a $\sigma \approx 0.4$ over either a very narrow σ or a very broad one (represented by the unweighted 9×9 tower cluster). However, calibrations and pile-up control for the upgrade algorithms have not yet been optimised, so further improvements may be possible.




Figure 16: (a) Turn-on curves for a 4-jet trigger in $t\bar{t}$ events. Thresholds were chosen to match the 90% efficiency point for a 4×25 GeV trigger using the current system. (b) Predicted 4-jet trigger rates from minimum-bias Monte Carlo. The lines indicate thresholds giving equivalent 90% efficiency points and the corresponding rates. Note that for many processes the plateau efficiency for the upgrade trigger is higher than is possible with the current system.

Figure 16a shows the efficiency for a four-jet trigger for $t\bar{t}$ events selected as containing 4 jets using an $R = 0.4$ anti- k_T jet-finder. Here the slopes of the turn-ons are similar, but the algorithm based on the current inputs is not fully efficient on the plateau. This is because the coarser granularity limits the ability to resolve nearby jets, resulting in a lower efficiency for multi-jet triggers. The unweighted 9×9 window sum shows some excess efficiency at low E_T , an effect that is also seen when larger σ values are used in the Gaussian algorithm. This is probably partially due to overlap of nearby trigger jets resulting in some double-counting of E_T , but perhaps also because of the greater sensitivity of the inputs to small signals than the current trigger towers have. It becomes problematic if it leads to an increased rate. Using a relatively narrow $\sigma = 0.2$ for multi-jet triggers largely eliminates this effect while still maintaining a good sharpness and plateau efficiency. Figure 16b shows the corresponding 4-jet trigger rates, with thresholds corresponding in efficiency to a 25 GeV 4-jet trigger in the current system. Here it can be seen that the narrow Gaussian trigger, even at this early stage of optimisation, offers a small improvement in multi-jet trigger rates as well as an improved plateau efficiency. The unweighted 9×9 window illustrates that without some measures to control overlap or pile-up effects, multi-jet rates in this high luminosity environment grow rapidly.

The proposed jFEX will provide E_T measurements based on different environments or radial weighting, therefore allowing the efficiency and rate to be optimised separately for inclusive and multi-jet triggers.

3.3 System Evolution

The following sections describe the evolution of the L1Calo architecture from the present day system. The biggest changes take place during LS2, but some upgrades will be installed in LS1 to enhance the trigger selectivity in the intervening years, and new components used to realise these improvements are also described.

System requirements for Run 3 will be included in the existing User Requirements Document.

3.3.1 Architecture

Initial system architecture - Run 1 The current L1Calo system (Figure 17) receives analogue signals for trigger towers of 0.1×0.1 ($\eta \times \phi$) from the electromagnetic and hadronic calorimeters. These signals are sampled in the Pre-Processor Modules (PPMs) at 40 MHz and the resultant digital data are sent to the CP and JEP subsystems.

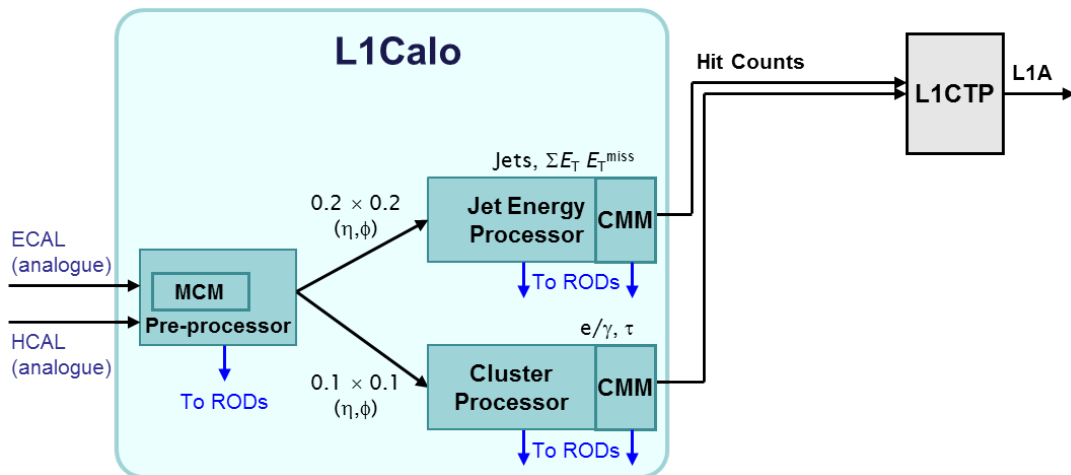


Figure 17: System Architecture during Run 1. See the text for numbers of modules.

The CP subsystem contains 56 Cluster Processor Modules (CPMs) in four VMEbus crates. Each CPM processes an area of calorimeter data to identify and count energy deposits indicative of isolated e/γ and τ particles. The JEP subsystem contains 32 Jet/Energy Modules (JEMs) in two VME crates. Each JEM identifies and counts energetic jet candidates and computes transverse energy quantities for the area of calorimeter examined.

The results from the CPMs and JEMs are transmitted over crate backplanes for summation in the CMMs, before being transmitted to the CTP. The 8 CMMs in the CP subsystem and 4 in the JEP subsystem use the same hardware but different firmware. The outputs from the CMMs to the CTP comprise counts of e/γ , τ and jets that pass programmable thresholds, and hit flags indicating which thresholds have been passed for transverse energy quantities.

On receipt of an L1A, all the L1Calo modules provide the DAQ system with read-out data. The CPMs, JEMs, and some CMMs also provide Region-of-Interest (ROI) data for the High-Level Trigger (HLT). Both the read-out and ROI data are transmitted via Read-Out Driver

(ROD) modules. These aggregate the data from multiple processor modules and provide buffering and flow control before building and transmitting event packets.

System architecture during Run 2 During LS1, a number of upgrades will be made to L1Calo. On the PPMs the MCMs, which digitise, calibrate and filter the calorimeter signals, will be replaced with newer modules (nMCMs) that provide 80 MHz digitisation, lower noise and greater flexibility to handle pile-up. Downstream of this, L1Calo will be enhanced to allow topological triggers; the firmware on the CPMs and JEMs will be modified and the merger modules will be replaced. This is shown in Figure 18.

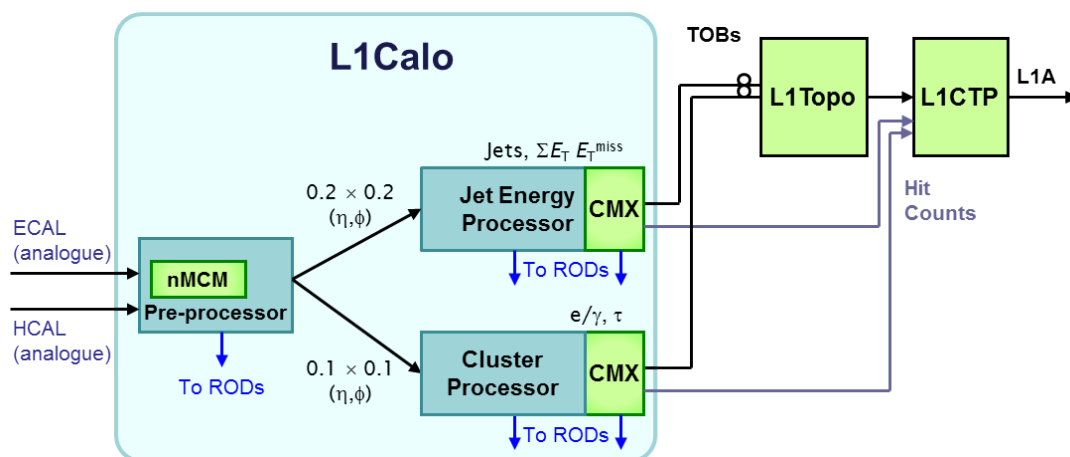


Figure 18: System Architecture during Run 2. New components are shown in green.

During Run 2, the CPMs and JEMs no longer output hit counts for e/γ , τ (CPM) and jets (JEM). Instead, they output TOBs, which comprise the location, energy and type of object identified and also sums for E_T , E_T^{miss} and XS. As this requires extra bandwidth, the data are driven on to the crate backplane at 160 MHz (compared to 40 MHz in the initial system). The CMMs are each replaced by an enhanced version of that module, the CMX, capable of receiving and processing the additional data.

The CMX, like the CMM, sends counts of objects over threshold to the CTP. TOBs for the whole $\eta - \phi$ range processed by L1Calo are transmitted optically by the CMX to the new L1Topo, which also receives data from L1Muon (see Section 5.3). L1Topo forms combined trigger objects, based on the full event topology, and transmits them to the CTP.

In addition to real-time trigger processing, as with the previous modules, the CMX provides ROI and read-out data to the HLT and DAQ systems. These data are transmitted via existing RODs with updated firmware.

System architecture during Run 3 During LS2 there will be further upgrades, both to L1Calo and the upstream calorimeter electronics. Two new subsystems will be added to L1Calo: the electromagnetic and jet feature extractors, see Figure 19. To achieve the increased discriminatory power necessary to handle the LHC luminosities planned beyond LS2, these will process calorimeter data at a finer granularity than the pre-LS2 L1Calo system. Initially at least, they will augment the existing L1Calo electronics, operating in parallel with the CP and JEP systems. Once the outputs of the eFEX and jFEX have been validated, removing the CP and JEP systems from the L1Calo processing chain will be an option (except for that section of the JEP used to provide hadronic data, as described below).

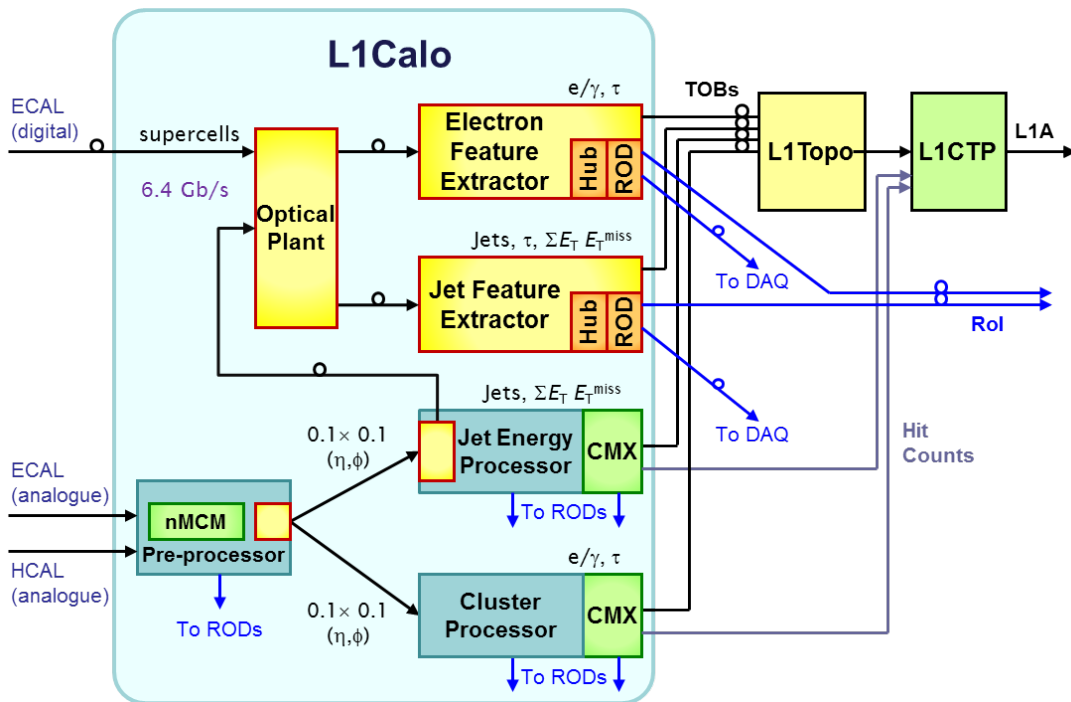


Figure 19: System Architecture during Run 3. New components are shown in yellow / orange.

Prior to LS2, the LAr and Tile calorimeter electronics provide analogue data which are digitised in the PPMs. During Run 3 the digitised LAr calorimeter data are provided to L1Calo by the DPS. For each tower of 0.1×0.1 , L1Calo receives ten samples derived from longitudinal segments and transverse sums of groups of calorimeter cells, each sum forming a SuperCell (see Figure 20). Analogue Tile calorimeter data are digitised in an upgraded PPM, using the nMCM fitted in LS1. A new PPM LVDS Cable Driver (LCD) daughter card outputs these data at 960 Mbaud (800 Mb/s payload), twice the current rate. Since there is now no summation of digital tower data, Bunch-Crossing Multiplexing (BCMux) can be used to halve the required transmission bandwidth. Together, these allow an increase in the granularity of the hadronic data transmitted, to towers of 0.1×0.1 . These hadronic data are received on the JEMs, by new daughter cards which output copies of the data optically, for the eFEX and jFEX subsystems.




Figure 20: The trigger granularity from each 0.1×0.1 trigger tower after the upgrade of the LAr Calorimeter electronics. Ten E_T values are provided from “1-4-4-1” longitudinal/transverse samples, each forming a SuperCell.

The FEX subsystems require a large proportion of their input to be duplicated. Most of the duplication is performed by the DPS before transmission to the FEXs via an optical plant. The

optical plant breaks apart the fibre bundles and reforms them, changing the mapping from that employed by the calorimeter and JEP electronics to that required by the FEX subsystems. It also implements additional duplication of signals.

The eFEX subsystem identifies isolated e/γ and τ candidates. The jFEX subsystem identifies energetic jet candidates and also calculates the ΣE_T , E_T^{miss} and XS for each bunch crossing. The results comprise TOBs, which are transmitted to L1Topo over optical fibres. In L1Topo, results from all FEXs are used as input to topological algorithms and the results are transmitted to the CTP.

The eFEX and jFEX modules are both implemented following the ATCA standard. The eFEX comprises two shelves of 12 eFEX modules each and the jFEX comprises a single shelf of eight jFEX modules.

On receipt of an L1A the FEX modules provide ROI and read-out data to Level-2 and the DAQ system respectively. Each FEX module outputs these data to the shelf backplane. Two Hub modules in each shelf aggregate the data and implement the required ROD functionality on daughter boards. Additionally, the Hub modules provide nodes on the TTC, control and monitoring networks. The hardware of the Hub modules is common to the eFEX and jFEX subsystems, but the read-out firmware differs in order to handle different data.

System architecture during Run 4 A further upgrade of L1Calo is planned during LS3 to handle even higher luminosities. (This is not the subject of this document, but it does have implications for forward compatibility of hardware.) To meet this demand, the calorimeter trigger will be split into a Level-0 part (L0Calo) and a new Level-1 part (L1Calo). The analogue calorimeter signals and remaining PPM and JEP subsystems become obsolete and will be removed. All calorimeters (both electromagnetic and hadronic) will perform on-detector digitisation. The eFEX and jFEX subsystems will then be modified to constitute L0Calo. Further details can be found in Section 3.11.

3.3.2 nMCM

Most of the signal processing in the initial Pre-Processor subsystem is performed in MCMs. Each MCM receives and digitises four analogue calorimeter trigger signals, extracts the E_T for each signal, assigns the correct bunch crossing, and sends the results serially to the CP and JEP subsystems [3.1].

During LS1, the MCM will be replaced with a new device, the new MCM (nMCM). In this plug-compatible device, shown in Figure 21, the ASIC is replaced by an FPGA. The four inputs from the motherboard are digitised to ten-bit precision at 80 MHz in two dual-channel ADCs. The input voltage range corresponds to an E_T deposition range of 0-256 GeV with ~ 0.25 GeV E_T least count. The phase of the ADC strobes can be adjusted in steps of about one nanosecond over the LHC bunch-crossing period of 25 ns, to position the sampling point at the maximum of the analogue pulse.

The resulting 10-bit digital data are routed to the CALIPPR FPGA, which performs bunch-crossing identification, E_T extraction and final calibration, noise suppression, and pedestal subtraction. In general this signal processing runs at the LHC clock frequency of 40 MHz, discarding every other 10-bit data word. However, saturated-pulse processing in the CALIPPR device uses the full 80 MHz samples. The results of the digital signal processing are transmitted serially on LVDS links to the CP and JEP subsystems at 480 Mbaud (400 Mb/s payload). Data is sent to the CP for individual 0.1×0.1 trigger towers, while the JEP is sent data pre-




Figure 21: The new Pre-Processor Multi-Chip Module

summed into 0.2×0.2 jet elements. For Run 3 operation, the serial link to the JEP system will be upgraded to 960 Mbaud (800 Mb/s payload), to carry 0.1×0.1 hadronic E_T values for onward transmission to eFEXs and jFEXs. These data will be further summed to provide the 0.2×0.2 jet elements for the JEP itself.

The CALIPPR FPGA reads out event data to DAQ on receipt of a L1A signal. It can also accumulate energy deposition rates and distributions of deposited E_T above programmable thresholds for each channel, for read-out to a separate monitoring system.

The nMCM includes a programmable analogue signal generator which may be used for standalone tests. The amplitude, width and pedestal level of the generated analogue signals are controlled from the CALIPPR FPGA.

Use of an FPGA gives flexibility to the system, for example in the choice and sign of digital filter coefficients, and the possibility of using different types of filter. One type under study is the auto-correlation filters, based on the correlation between ADC samples. These filters may give a better performance in high pile-up conditions. Figure 22 compares the performance of the current (matched) and autocorrelation filters in determining E_T at the start of an LHC bunch train and in the ensuing train bulk.

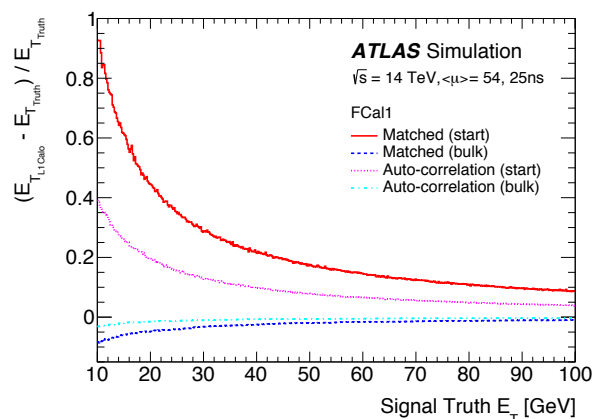


Figure 22: Performance of digital filters

In the Run 1 Pre-Processor, the same E_T values are sent to the CP and JEP subsystems. For each tower, a single look-up table is used for calibration, the value being optimised for the electromagnetic energy scale. Including a second look-up table would allow provision of separate, hadronically calibrated E_T values to the JEP, and this option is under study. When a calorimeter signal exceeds the dynamic range of the ADCs, timing information must be obtained from the leading edge of the pulse rather than the (unknown) position of the pulse

peak. This “saturated” BCID algorithm can be improved using the 80 MHz samples from the nMCM, to ensure the correct BCID up to the highest possible energies (see Figure 23).

Finally, the nMCM offers the possibility to correct for pile-up induced fluctuations of the pedestal. Such a correction would greatly increase the performance of all triggers which are affected by pile-up, particularly E_T^{miss} triggers.

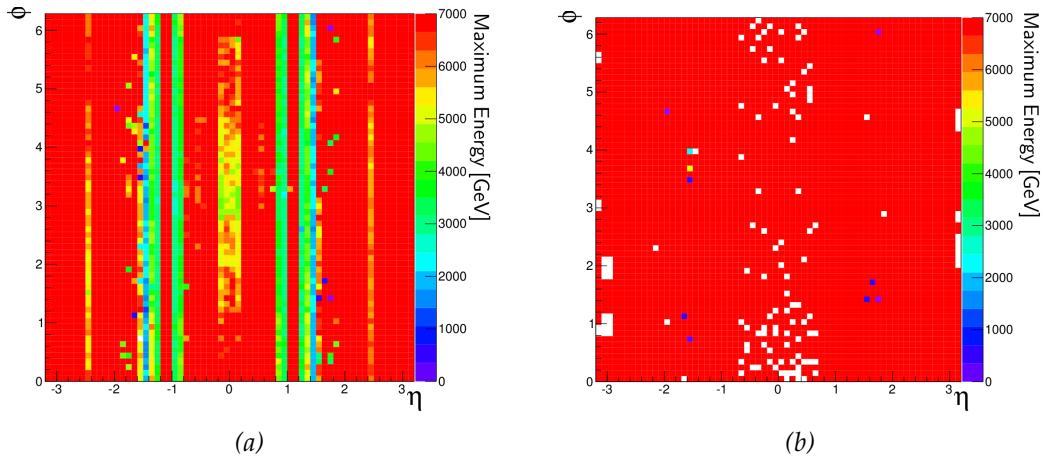


Figure 23: Maximum energy with full BCID efficiency for (a) the current and (b) the improved BCID algorithm for saturated pulses

3.3.3 CMX modules

The Common Merger module extended (CMX) is a replacement for the CMM and supports all Run 1 CMM functionality. This is enhanced in two important ways, with faster communication over existing data paths and provision of inputs to L1Topo. Figure 24 shows the block diagram of the CMX module.

The CMX provides all the multiplicity-based L1Calo trigger algorithms by counting the number of TOBs above threshold. Input data from the CPMs and JEMs arrive at 160 Mb/s (instead of 40 Mb/s), sufficient bandwidth to permit TOB information to be sent in replacement of counts over threshold. CMM functionality is retained on the CMX boards to gather and merge sums as preparation for triggering on multiple TOBs. Threshold-dependent isolation criteria are applied to data from the CPMs. At each jet threshold, E_T calculations of two different jet sizes are performed on data from the JEMs. E_T thresholds (possibly varying with η) are defined for particular TOB types, along with other criteria which may be applied to the TOBs. Firmware upgrades in the CPMs and JEMs are necessary to support these changes.

The total multiplicity for each TOB type and threshold is calculated by producing crate-level counts, in crate CMXs. These are sent over LVDS links to one of the CMXs to form a system-wide sum which is then passed to the CTP. The large logic capacity of the CMX FPGAs permits approximately twice as many selection and isolation thresholds for electromagnetic and tau clusters and similarly doubles the number of jet thresholds when compared to the CMM. The LVDS communication between CMXs and to the CTP is capable of running at 160 Mb/s (instead of 40 Mb/s between CMMs), though it is anticipated that operation will take place at 80 Mb/s between CMXs and 40 Mb/s to CTP. The directionality of these LVDS connections is selectable to facilitate initial testing.

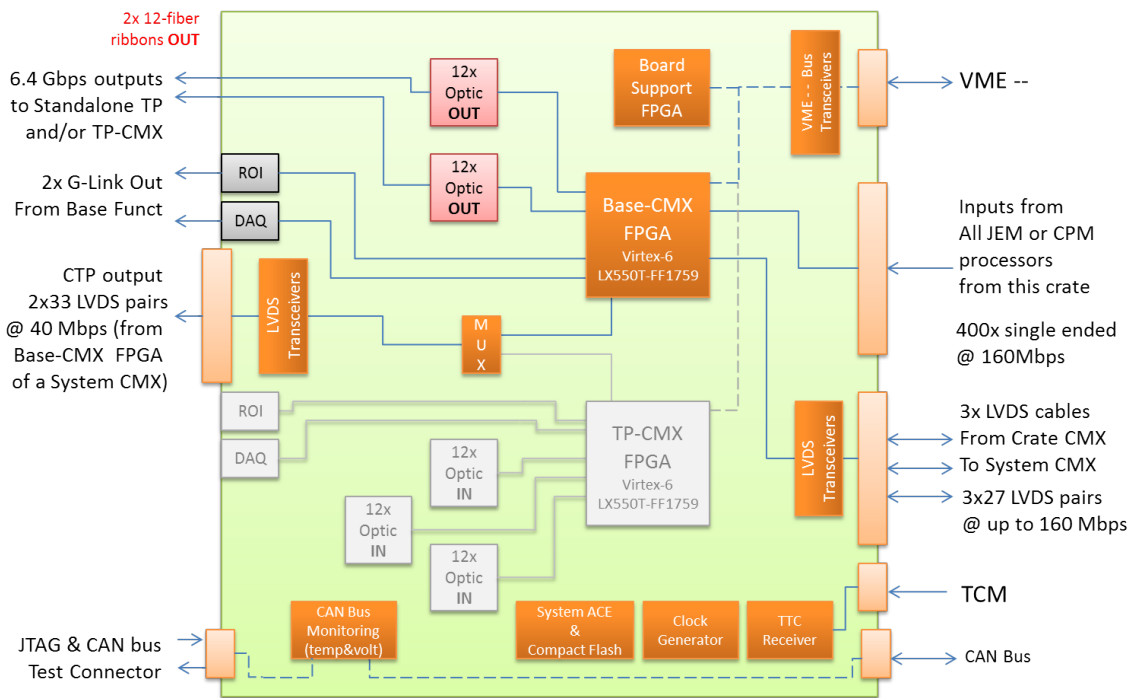


Figure 24: Block diagram of a CMX module showing inputs, outputs and principal components

In addition to the multiplicity analysis of TOBs, the CMX also collects and manipulates the TOB information and sends it on 6.4 Gb/s fibres to the L1Topo processor. Zero suppression of candidates will be performed in each CMX in parallel, before grouping the candidates by output fibre. This simplifies the first stages of input handling in L1Topo.

Two copies of the resulting output lists will be sent on two separate fibre groups (see Figure 24, outputs to standalone TP); this supports two L1Topo modules, with identical sets of inputs. Each set of inputs could contain up to 12 fibres, but studies indicate that fewer than this number will be required to transmit the actual number of TOBs. More elaborate subdivision and duplication schemes using the 24 output fibres are supported by the CMX hardware, but are not currently planned in the firmware. These will be needed by the additional L1Topo modules planned for installation during LS2.

ROI and DAQ read-out of the CMX is supported over fibres running the G-link protocol and using existing RODs. Board timing, monitoring and initialisation are supported in a similar manner to the CMM,

As shown in Figure 24, the board layout supports addition of a second FPGA, connected to 36 input fibres and to a second G-link output path. In the first instance, this second FPGA and its read-out links will be fitted on only one prototype of the board to facilitate testing. The second FPGA could provide limited topological processing. However, the 36 input fibres are insufficient to receive full outputs from all other CMX boards, and the second FPGA is not an alternative to the full L1Topo system. If further local processing is required in the CMX (for example for one type of TOB), a CMX could be retrofitted with a second FPGA.

Parameter	eFEX subsystem	jFEX subsystem
Calorimeter coverage in η	$ \eta < 2.5$	$ \eta < \sim 4.9$
Algorithm window size	0.3×0.3	0.9×0.9
Core area processed per module	1.7×0.8	$\sim 9.8 \times 0.8$
Environment processed per module	1.8×1.0	$\sim 9.8 \times 1.6$
Link speed (Gb/s)	6.4	6.4
<i>Hardware per FEX Module</i>		
Input fibres	144	288
Output fibres	36	24
Processor FPGAs	4	6
Backplane interface FPGAs	1	0
Control FPGAs	1	1
Configuration FPGAs	1	1
Total FPGAs per FEX	7	8
FPGAs per Hub module	3	3
FPGAs per ROD daughter card	1	1
<i>Hardware per subsystem</i>		
FEX modules	24	8
Hub modules	4	2
ROD daughter cards	8	4
ATCA shelves	2	1
Total input fibres	3456	2304
Total output fibres	864	192
Total FPGAs	188	74

Table 4: Run 3 principal system parameters for the baseline design

3.4 Phase-I Architecture

The following sections describe the architecture of the system to be built, installed and commissioned during LS2 to meet the challenges of running at luminosities of up to $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $\sqrt{s} = 14 \text{ TeV}$. Table 4 lists the principal system parameters.

3.4.1 Input interfaces

This section describes the inputs to the eFEX and jFEX systems in the baseline design, which assumes 6.4 Gb/s links. Higher link speeds, such as 10 Gb/s, are under investigation and will be used if tests prove successful. The impact of this improvement on the baseline design is also outlined.

Input sources Both the eFEX and jFEX systems receive digitised E_T values on optical fibres. The EM, hadronic endcap (HEC) and forward calorimeter (FCAL) inputs are all provided by the LAr DPS [3.3]. However, the TileCal will continue to provide analogue trigger signals, which are digitised in the PPMs and transmitted to the JEMs by new link daughter cards. New JEM input daughter cards will be built to provide optical signals to the FEXs, necessitating a corresponding increase in the PPM to JEM link transmission speed. This is shown in figure 19.

Optical link speed and content A single 6.4 Gb/s link can carry 128 bits of data per bunch crossing using 8b/10b encoding. Eight bits are used to send a cyclic redundancy check (CRC), to detect transmission errors. This leaves 120 bits per fibre for trigger information.

For the EM layer, there are two distinct sets of links with different content for the two L1Calo FEX subsystems. The eFEX needs the finest granularity SuperCell information available from the DPS, which comprises 10 SuperCells per 0.1×0.1 tower as shown in Figure 20. In contrast, the jFEX needs only the E_T summed over each 0.1×0.1 tower. In the hadronic layer, both eFEX and jFEX receive the same 0.1×0.1 granularity towers.

Each EM fibre to the eFEX carries the data from the 20 SuperCells in two adjacent 0.1×0.1 trigger towers. A dynamic range of 10 bits E_T per SuperCell requires 200 bits to be transmitted. Since this exceeds the available bandwidth, the baseline transmission scheme assumes the BCMuX protocol as used in the existing L1Calo system (c.f. Section 4.5 in Ref. [3.1]). Using this protocol, data for 20 SuperCells can be transmitted in a total of 110 bits per bunch crossing. This leaves 10 unused bits which could be used to extend the dynamic range to 11 bits per SuperCell. The best use of the dynamic range is still under study but the value of the least significant bit for each layer is expected to be in the range 64–256 MeV.

The jFEX fibres for the EM layer will carry eight tower sums covering a 0.4×0.2 region of $\eta \times \phi$ space. The BCMuX protocol can not be used with summed signals. After reserving eight bits for CRC data, an option would be to send 13 bits per trigger tower leaving 14 bits to send the total energy (not E_T) for the entire 0.4×0.2 region. This might be of benefit to E_T^{miss} trigger algorithms. The least significant bit for the trigger tower data is expected to be either 128 or 256 MeV.

In the hadronic layer, there is no fine-grain information available from the calorimeters. Although they are segmented in depth, the baseline is that the hadronic layer fibres to both eFEX and jFEX will carry data for eight 0.1×0.1 towers (summed in depth) from a 0.4×0.2 region, exactly as for the EM layer fibres to the jFEX. This mapping is expected to apply to both the LAr HEC and the TileCal.

The FCAL detector has three layers, each with a different granularity, and it is planned to send finer-granularity data to the jFEX than is used in the initial L1Calo system. A set of three fibres will carry the twelve FCAL1, eight FCAL2 and four FCAL3 SuperCells per 0.4 slice in ϕ (i.e. eight values per fibre) as for the other hadronic fibres.

Studies are under way to determine if the link speed can be raised above the baseline 6.4 Gb/s. If this were possible the additional bandwidth might be used to send data for a larger number of towers per fibre, to increase the dynamic range of transmitted data, or to remove the BCMuX protocol. For the jFEX this would permit an increase in the size of the jet window processed up to 1.7×1.7 , and for eFEX and jFEX give flexibility to send additional data in Run 4. For the eFEX the dynamic range could be increased to allow the full precision of data from the DPS to be transmitted. Some studies show that the use of BCMuX might result in a small loss of input data at high luminosities. Further investigation is needed to determine the best use of any higher bandwidth.

Link fan-out Each FEX module processes a fixed core area in $\eta - \phi$ space. To do this, it also needs copies of data for the surrounding SuperCells or towers required by the sliding window algorithms. Electrical re-transmission of multi-gigabit signals between modules over the backplane is precluded by latency and board layout considerations, so the environment data need to be duplicated at source and sent to each module. This will mainly be done by driving two separate copies of the same data on separate fibres from source FPGAs. There

may be some difficult regions – such as corners where four FEX modules meet in $\eta \times \phi$ space – where extra copies are needed. This will depend on the detailed mapping between DPS and FEXs, which is still being studied. It is also expected that hadronic endcap data will require additional duplication, so that two identical copies may be sent both to eFEX and jFEX subsystems.

Since the organisation of fibres in ribbons from the DPS will in general be different from that needed by each eFEX or jFEX module, an intermediate optical patch panel is needed between the DPS outputs and the FEX inputs. In most cases, where no optical splitting is required, this will be implemented by special fibre bundles, which re-group fibres between multi-ribbon connectors.

Where additional copies of data are required, beyond the capacity of the DPS, two-fold passive optical splitting will be used. Studies show that the received optical power is adequate. If larger numbers of copies are required, active optical splitting may be required. This option is under study.

3.4.2 The Electron Feature Extractor (eFEX) module

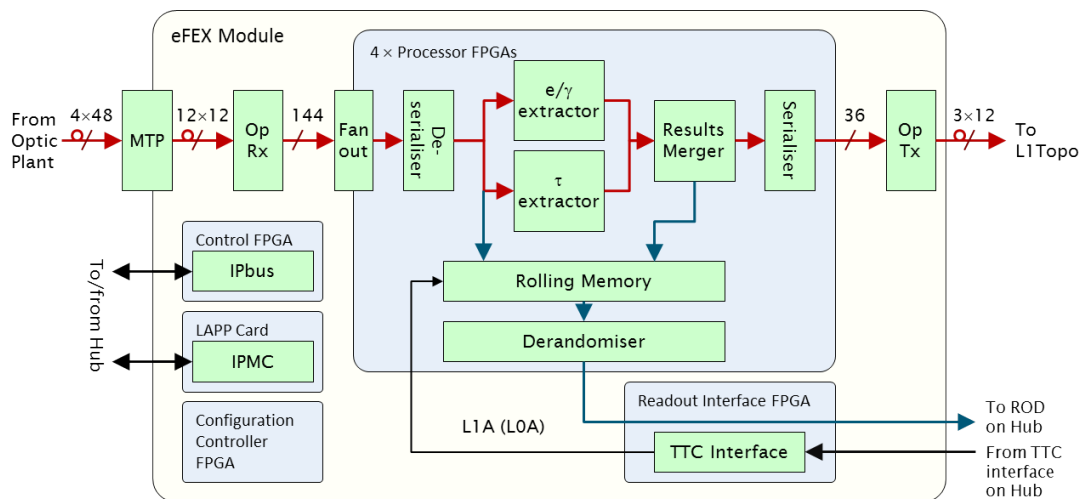


Figure 25: Block diagram of the eFEX module showing the real-time and read-out data paths. Control and monitoring signals are omitted except for the L1A.

Overview The function of the eFEX module is to identify isolated energy deposits indicative of electrons, photons and tau particles. As shown in Figure 25, this is done in a bank of four Processor FPGAs, and the results transmitted optically to L1Topo. On receipt of an L1A, data are copied from the real-time path to the read-out path and from there to a ROD daughter card mounted on a Hub module (see Section 3.4.5). The eFEX is an ATCA module, conforming to the PICMG 3.0 Revision 3.0 specification. The functional elements of the eFEX design are described in the following sections.

Processing area The algorithms used by the eFEX are described in Section 3.2. The algorithm defines a core area in which energy deposits are found, surrounded by an environment area used to establish isolation. Each eFEX processes a core area of up to 1.7×0.8 of calorimeter data, for which it receives environment data covering 1.8×1.0 (see Figure 26). Depending on their location in the subsystem some eFEXs process slightly fewer core inputs than this, but

the hardware and firmware of all the modules are the same. The small differences in the core areas are implemented via programmable parameters. A total of 24 eFEX modules are required to process all of the data from the calorimeter within the range $\eta < |2.5|$.




Figure 26: eFEX module processing window showing core (red) and environment areas (in units of 0.1×0.1), for one of the three possible module η positions. The Algorithm environment area is shown in yellow and data carried by EM and hadronic fibres, but not used in this module, in green and blue respectively.

Input data format The eFEX modules receive data from the electromagnetic and hadronic calorimeters on optical fibres. From the electromagnetic calorimeters, each fibre carries the data for two adjacent triggers towers, i.e. an area of 0.2×0.1 . From the hadronic calorimeters each fibre carries the data for an area of trigger towers of 0.4×0.2 . Further details are given in Section 3.4.1.

Data reception and fan-out Given the input data format described above, the environment processed by a module does not necessarily correspond to an integral number of fibres. To process an environment of 1.8×1.0 , each eFEX module must receive electromagnetic calorimeter data on at least 90 fibres and hadronic calorimeter data on at least 25 fibres. The baseline design uses 100 and 36 fibres respectively to allow for the partial use of the data on fibres at module boundaries.

The optical connection to the eFEX is made via four 48-way MTP connectors (not fully populated) mounted in Zone 3 of the ATCA backplane. The eFEX input fibres are organised into twelve ribbons of 12 fibres, mechanically supported by a rear transition module. On the eFEX side of the MTP connectors, ribbons of 12 fibres carry the signals to 12 Avago MiniPOD™ [3.4] transceivers, mounted in-board, which perform optical to electric conversion. For those parts of the calorimeter where the mapping is straightforward, only 136 of these fibres are required to carry the data. The eight additional inputs are required in regions where mapping is difficult.

A total of 136 multi-Gb/s signals are received from the calorimeters (excluding spares). Of these, 52 are transmitted from the MiniPODs to a single Processor FPGA, 72 are transmitted to two Processor FPGAs (required by the overlapping nature of the eFEX algorithms) and 12 are transmitted to three FPGAs. The eFEX must therefore handle upwards of 230 multi-Gb/s signals.

Three-way fan out, where required, is implemented using discrete, high-speed electrical buffers. While this is also the preferred method of implementing two-way fan out, many additional components are required, and detailed layout studies may preclude this possibility. The fall-back solution for signals requiring two-way fan out is to loop them through the

multi-gigabit receiver/transmitter pairs in the Processor FPGAs, a technique known as PMA loopback. In this scheme, a multi-Gb/s signal is transmitted from a MiniPOD to a single FPGA. Once the signal has been received by the FPGA and equalisation has been performed, but before it has been decoded, it is routed from the high-speed receiver to the paired high-speed transmitter, and thence to a second FPGA. There is a latency penalty of ~ 25 ns and some degradation of signal quality associated with this method.

Processor FPGAs Most of the eFEX functionality, including all the real-time functionality, is implemented in the four Xilinx Processor FPGAs. Each one processes a 0.6×1.0 region, receiving calorimeter data on 58 serial inputs at 6.4Gb/s and identifying e/γ and τ objects using the algorithms described in Section 3.2. For each object found, a Trigger Object (TOB) datum is generated, comprising the object type (e/γ or τ), the energy measured and the η or ϕ co-ordinates. The FPGA also records the input and output data in scrolling memories and, on receipt of a L1A signal, transfers the data for the triggered bunch crossing to a de-randomising buffer. Here a data packet is built (including the bunch number) and transmitted it to the Read-out Interface FPGA. A slow-control interface is provided for all programmable parameters within the FPGA via IPbus .

One of the four Processing FPGAs is responsible for collecting all the TOBs, forming a single output packet, and transmitting copies serially at 6.4Gb/s via electrical-optical transmitters to up to six L1Topo modules. The limiting resource in the Processor FPGAs is the number of multi-gigabit receivers available. Of the 80 available in the proposed device, 60 are used, while the combinatorial logic is estimated to need less than 20% of internal resources. Should the specification of the eFEX inputs change, the specification of the real-time outputs will be updated to match, so that outputs can be connected back to the inputs for diagnostic purposes.

Read-out Interface FPGA This FPGA implements the following two functions. By locating these on the eFEX, in an FPGA, the flexibility to develop the system for Phase-II is maximised.

1. **Read-out interface to the backplane:** The Read-out Interface FPGA receives read-out data from the Processor FPGAs following each L1A. From these it builds a single packet per event, which it transmits to the Hub module via the ATCA backplane.
2. **TTC interface:** This FPGA receives all TTC information and commands transmitted to the host eFEX crate, via the crate backplane. All signals and information necessary to initiate read-out and build event packets are transmitted to the Processor FPGAs. A control and status interface to the Control FPGA is also provided. In Run 4 this device will interface to the potentially changed timing and control signals.

Slow control and environment monitoring The Control FPGA provides the interface between the eFEX and the IPbus network on the ATCA backplane, which is used for functional control of the eFEX system. It contains those control and status registers that concern the operation of the eFEX as whole (i.e. all registers not specific to one particular Processor FPGA). This includes status and control registers for the Read-out Interface and Configuration Controller FPGAs. The Control FPGA also contains a switch to route IPbus traffic to/from those registers that are implemented in the Processor FPGAs.

Environment monitoring on the eFEX is implemented using a daughter board designed by LAPP (Annecy), which collects voltage and temperature data from all FPGAs and from

dedicated sensors on the module. These data are transmitted to the Control FPGA, which transmits them to the Hub module across the ATCA backplane. From there they are transmitted to the ATLAS DCS system. See Section 3.4.5 for more details.

Layout MTP connectors, which receive the fibres from the optical plant, are mounted on the rear of the board to facilitate easy insertion and extraction of the module. The smaller number of optical transmitters, that transmit the real-time output data, are mounted on the front panel, owing to space restrictions at the rear. Much of the glue logic is placed on the underside of the module. Routing of over 200 multi-Gb/s signals presents a significant challenge for the design of the eFEX PCB. Track lengths and routing complexity is minimised by placing Avago MiniPOD receivers close to the Processor FPGAs that they serve. Routing of the fibre-optic ribbons carrying data from the MTP connectors to the MiniPOD receivers is constrained such that they need to twist and curve to bypass the large heat sinks on the processing FPGAs.

FPGA configuration logic The configuration of the Processor, Control and Read-out Interface FPGAs is controlled by a seventh FPGA, the System ACE FPGA, containing an IP core supplied by Xilinx. The configuration data for the Processor, Control and Read-out Interface FPGAs are stored in an SD memory card mounted on the eFEX. (Configuration data for the removable device is stored in its own internal flash RAM.)

Power The eFEX conforms to the full ATCA PICMG 3.0 Revision 3.0 specification regarding power and power management (including the facility for hot swapping, although this is not expected to be used in the trigger system). The power management functions are implemented via the LAPP daughter card. For those power supplies feeding multi-gigabit receivers and transmitters (including those in the FPGAs) a point-of-load architecture is used. The total power consumption of the eFEX is estimated to be 190 W per module.

3.4.3 The Jet Feature Extractor (jFEX) module

Overview The jFEX processor system performs jet, large-area tau, E_T^{miss} and total- E_T trigger algorithms on 0.1×0.1 electromagnetic and hadronic towers, as described above. The jFEX system (Figure 27) consists of a single ATCA shelf, equipped with several module types and fed by an optical fibre plant. The fibres are routed via rear transition modules into jFEX modules, where the physics algorithms are run. Timing distribution, read-out and control functionality are implemented on two Hub modules located in the fabric slots of the shelf.

Each jFEX module processes calorimeter data from a ϕ -octant, covering the entire η range, including barrel, endcap and forward calorimeters. A total of eight jFEX modules is required for full coverage of the calorimeter solid angle.

Jets and taus are identified on the jFEX modules as local maxima of energy deposits in a sliding window of 0.9×0.9 . This requires that each jFEX module is supplied with a copy of the data from neighbouring ϕ -octants. The data duplication is performed at the data source, the Digital Processor System. A diagram of the jFEX module is shown in Figure 28. The design foresees a flexible definition of jet algorithms implemented within the 0.9×0.9 window (e.g. in the shape of the region for the energy summing or by using radial weighting), allowing for several definitions to be used concurrently. The information transmitted from the jFEX system to L1Topo (jet energies and total energy) allows area-based pile-up correction of the jet energies.

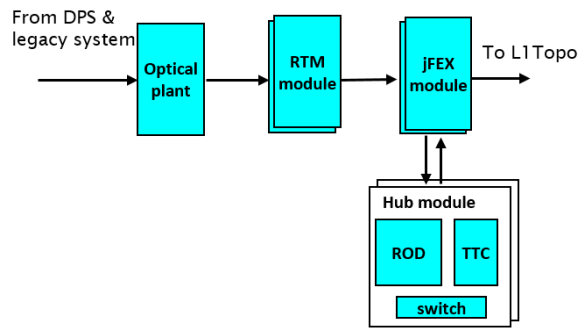


Figure 27: The jFEX processor system

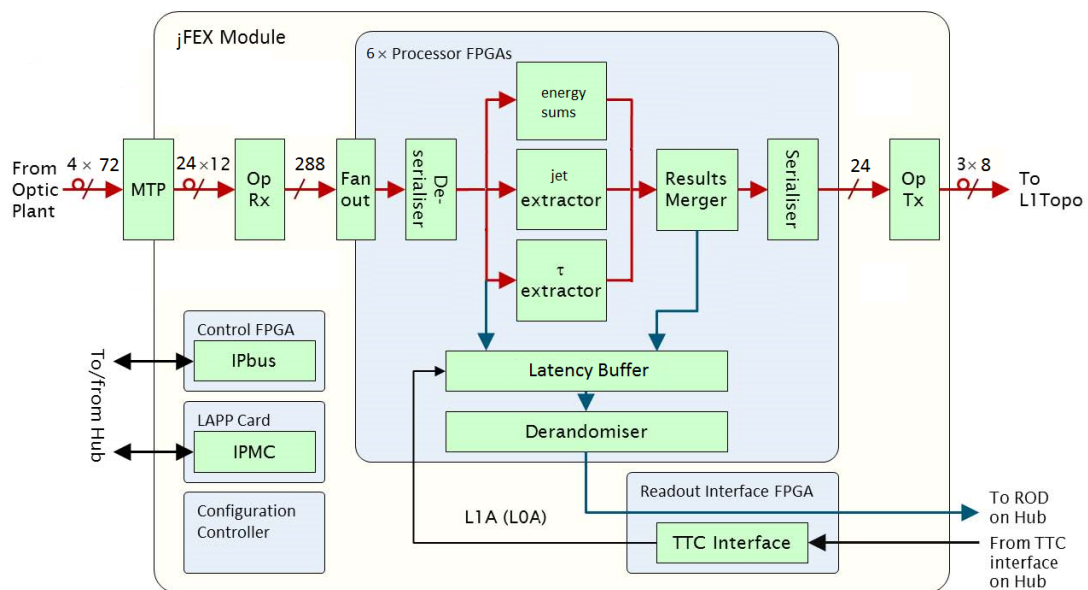


Figure 28: Block diagram of the jFEX module

Real-time data path Real-time input data from the DPS (including LAr, HEC & FCAL) and input from the JEM daughter board (TileCal) enter the jFEX module, via the rear transition module. This provides the mechanics for feeding fibres into the jFEX module on four 72-way MTP/MPO connectors in ATCA zone 3. This allows a total of 288 input fibres per jFEX.

On the jFEX module, the four fibre bundles are each split into six twelve-way ribbons that are received in MicroPOD™ [3.4] electro-optical devices. These devices support data rates greater than 10 Gb/s per fibre. They require custom-built mechanics to provide stability and to dissipate the heat generated. The MicroPOD outputs are routed, as differential signals, to the GTH high-speed receivers in the Processor FGAs.

The Processor FGAs identify and report TOBs (i.e. jet and tau candidates) and calculate global quantities (E_T^{miss}). This is done separately in segments of η , so as to allow for downstream pile-up corrections. All six Processor FGAs perform these calculations. In addition one Processor FPGA consolidates the results before they are sent, optically via MicroPOD devices, to L1Topo. So as to be able to duplicate data into more than one L1Topo module, two MicroPOD devices (i.e. 24 optical fibres) are allocated for that purpose. The fibre bundles are routed through the front panel of the jFEX on standard MTP/MPO connectors.

Signal duplication The core tower of the jet algorithm is defined as the central 0.1×0.1 tower of the 0.9×0.9 window, to which any TOB generated in that window is assigned. Each Processor FPGA is responsible for a core area of 1.2×0.8 , requiring data from an environment of 2.0×1.6 (see Figure 29). As neighbouring Processor FPGAs (and jFEX modules) cover contiguous core areas, and the data environments required extend beyond the core areas, a significant amount of input data need to be duplicated between FPGAs and between jFEX modules.

Data sharing between jFEX modules is accomplished by sending two copies of all data from the DPS and JEM daughter cards, with no data transmitted between jFEX modules. On the jFEX module, the data is partitioned between six Processor FPGAs. Each of the FPGAs needs access to the core towers it processes, and a surrounding environment of four towers in each direction. Therefore a large amount of data needs to be duplicated between neighbouring FPGAs. For inter-FPGA data duplication the baseline design makes use of the PMA loopback scheme, as described in Section 3.4.2.




Figure 29: jFEX FPGA processing area of 12×8 core towers (green) and the environment received from DPS (purple) and duplicated on the jFEX module (orange)

The data are transmitted with 8b/10b encoding at 6.4 Gb/s line rate. As outlined in Section 3.4.4, the net data volume of 128 bits per bunch crossing per fibre allows the transmission of eight energy samples on each FPGA input signal with a checksum. Each signal carries data from an area of 0.4×0.2 . Two samples, electromagnetic and hadronic energy, are required for each trigger tower.

A single large FPGA with 80 inputs can thus receive data for a maximum area of 20×16 trigger towers. This is 12×8 core towers plus environment data. Irregularities in the calorimeter overlaps and in the FCAL require two of the jFEX Processor FPGAs to be larger devices with 96 high-speed inputs.

Timing All real-time data transmission and processing is performed synchronously at multiples of the LHC bunch clock. The clock as well as TTC control data are received in backplane zone two on two signal pairs each from the hub modules. Either hub can be selected to act as a clock source. After signal conditioning (jitter reduction) and frequency multiplication, the clock is routed to the high-speed link circuitry and the FPGA.

Read-out Read-out to the DAQ and HLT is implemented as described in the Section 3.4.2. Processor data are read out into two Hub modules on six high-speed links each (one link per

Processor FPGA), in Zone two of the ATCA backplane. The backplane bandwidth is sufficient for both incoming data and intermediate and final results. In normal running, the input data are expected to be read out only in the unlikely event of checksum mismatches in the data received.

Configuration, control, monitoring and diagnostics The real-time processors are configured on power-up through the Control FPGA. The configuration data are stored locally on an SD card. The Control FPGA itself is configured via an SPI flash device and the control firmware includes the algorithms to read and write the SD card and sequence the configuration data into the processor FPGAs.

IPbus is used to control and monitor the operation of all FPGAs, allowing parameters for the algorithms to be set, high-speed links to be controlled, playback and spy memories attached to the real-time path to be accessed, and the environment (temperatures and voltages) to be monitored. Low-level, ATCA-compliant control is achieved by an ATLAS standard IPMC module.

Options The baseline design described above allows for the detection of jets and large-area taus in an environment of 0.9×0.9 . There is a strong interest in considerably higher input data rates (9.6 Gb/s or above), which would support significantly larger jet sizes. Design studies and R&D are under way to study use of higher data rates, higher signal counts and data duplication schemes in dense electro-optical systems. Initial studies indicate that the jFEX system could utilise an increased transmission speed to reach a jet size of 1.7×1.7 , corresponding to a jet radius parameter of 0.85. A remapping of the calorimeter signals would provide each jFEX module with a slice in η (full ϕ coverage), allowing each jFEX to perform event-by-event pile-up subtraction in slices of η to improve E_T^{miss} determination, and to correct the jet energies measured within the jFEX system.

Phase-II The jFEX system will be designed to operated as part of the proposed L0-trigger in Run 4. Apart from the change to on-detector digitisation of TileCal data, the real-time data path and the processing are expected to run unchanged, though the sliding window algorithms might be upgraded. The FPGAs will be chosen to provide enough logic resources for algorithmic upgrades.

The distribution of clock and timing signals and read-out are expected to change in Run 4. These functions are handled on the Hub module, and any changes to the jFEX would be limited to firmware.

3.4.4 Output interfaces

Real-time outputs from L1Topo to CTP As noted in Section 5.3, output data from L1Topo will be transmitted to the CTP subsystem either electrically or optically. The choice may be made separately for individual L1Topo modules, depending on the relative latency of the algorithms providing the signals. The decision on which topological algorithms to run has not yet been finalised, and will change with time. However, the results transmitted to the CTP will take the following two forms:

Trigger Object Counts These are the numbers of TOBs above thresholds, which the CTP will continue to use as a primary input. In Run 1, these were computed in individual CPM and JEM modules, summed in Common Merger Modules, and sent to the CTP; after

LS1, the CPM and JEM modules provide TOB information to CMX modules where the hit counts will be computed and sent to the CTP; and for Run 3 use, the counts are computed in one L1Topo module from TOB information provided by eFEX and jFEX modules. The CTP does not have sufficient inputs to accept all the new signals as well as all the original multiplicity signals, so a transition in stages will be needed. When the hit counts are available from L1Topo in Run 3, some parts of the previous (analogue-based) trigger will be retired.

Topological flags The hit flags above use the majority of the available CTP input bits. CTP bits added in LS1 will be available to receive L1Topo output bits, probably indicating that individual topological algorithms have been satisfied, or that topological cut values have been met.

ROI outputs to the ROIB The eFEX and jFEX Hub modules will collect ROI data from individual eFEX and jFEX modules, and provide these to be used as ROI seeds in the High-Level Trigger. Depending on future evolution of the ROI Builder, the information may be sent either via S-Link or via Ethernet. The information will include the η and ϕ coordinates and the energy for each trigger object found. As an indication of bandwidth requirements, the ROI data sent by S-Link uses about 5% of the available bandwidth when the present system is running at 100 kHz. Most of this is due to the standard per-event packet structure, so increasing the number of ROIs tenfold would raise the link occupancy only to about 10%.

The subsystem generating the largest possible number of ROIs is the new eFEX, whose electron-finding algorithm may potentially generate up to one ROI per pair of trigger towers, giving 72 ROIs per eFEX or a system limit of 1600 eFEX ROIs. This is far beyond the point at which ROIs are useful to assist HLT processing, and a cut will be imposed on each eFEX and jFEX module to limit the number of ROIs sent. A further cut may be imposed at the ROD in the Hub module to limit the number of ROIs per eFEX and jFEX crate. In both cases, the ROI data will include flags if ROI numbers have been limited.

DAQ read-out from eFEX, jFEX, and L1Topo subsystems It is desirable to record all inputs and outputs from the trigger system, so that trigger operation may be checked event-by-event. In the trigger prior to LS1, all inputs and outputs are read out from each individual module. After LS2, much larger data volumes are involved, and more optimisation is necessary. The underlying idea is that a checksum is sent on each link (or for a small number of links working together), so that correct reception of data received in L1Topo, eFEX and jFEX modules can be confirmed for each bunch crossing. The checksum is computed at the transmitting end, and recomputed and verified at the receiving end. A copy of the transmitted data is always read out to DAQ. At the receiving end, the input data is only read out if the checksum does not match. The read-out bandwidths are chosen to meet this lower requirement.

Links will be developed to run reliably. If a link begins to fail, a large number of checksum mismatches may occur, potentially swamping the read-out system. To avoid this, all failing events will be flagged as faulty, but a limited number of failing events will be read out (for example, a maximum rate per LHC orbit). The cut will need to be imposed on individual eFEX, jFEX and L1Topo modules.

Apart from this handling of input data, read-out to DAQ will consist of a programmable number of time-slices of all input and output data from all L1Calo modules.

Phase-II: output from L0Calo to L1Calo and L1Track In Phase-II operation, it is expected that the eFEX, jFEX and L1Topo modules will be relabelled to become parts of the Level-0 trigger in a two-level (Level-0 and Level-1) hardware trigger system, which will also require a split-level CTP (i.e. CTP0 and CTP1). The design allows for these changes, as follows:

- DAQ read-out will occur following a Level-1 Accept (i.e. not a Level-0 accept).
- ROI read-out to HLT, if required, will also occur after a Level-1 Accept.
- Level-0 ROIs (L0ROI) will be sent from the Level-0 trigger to the Level-1 trigger at the L0A rate of up to 500 kHz.
- L1Track Regional Read-out Request (R3) data will be obtained by processing a further copy of L0ROI data, also at up to 500 kHz.

3.4.5 Common modules

The requirements of the eFEX and jFEX subsystems have been studied in detail to see to what extent the systems can use the same hardware. It has been concluded that a common design can be used for system infrastructure (Hub, ROD and TTC modules and backplane use). However, there are sufficient differences between the processor modules themselves to require separate hardware designs. One area of particular concern is the high density of high speed PCB signal tracks, which would be strongly impacted in any attempt at a unified design. Both modules are near the technical limits of what can reasonably be built and a single module to provide eFEX and jFEX functions would have to meet a superset of requirements. This would greatly increase the technical challenges, risks and time scales in what is already a demanding construction schedule.

The eFEX and jFEX subsystems have been designed to allow those functions not concerned with algorithmic trigger processing to be implemented on hardware modules that are common to both subsystems. The functions implemented on such common hardware include low level ATCA module control, read-out, monitoring and the interface to the TTC system. There is also a common module for the testing and commissioning of the FEX modules.

The following sections describe the hardware that is common to the eFEX and jFEX subsystems. It should be noted, however, that the common elements in these subsystems are not limited to this hardware. Even where dedicated hardware is used the subsystems employ common design standards, principles and, where feasible, firmware. One example of this is the IPbus firmware and associated software, which was originally developed for a different experiment. This is used throughout the L1Calo system to implement high-level module control and the interface to ROD Crate DAQ.

The common modules described in this section consist of the IPMC mezzanine on each module, the ATCA Hub module and its two main mezzanines, the TTC interface module and the ROD module, and Test Module.

IPMC mezzanine For the purposes of monitoring and controlling the power, cooling and interconnections of modules, the ATCA specification defines a low-level hardware management service based on the Intelligent Platform Management Interface standard (IPMI). Every ATCA module must have an Intelligent Platform Management Controller (IPMC), to provide an interface to the shelf manager via the IPMI bus. Within L1Calo, all ATCA modules use the IPMC mezzanine produced by LAPP for this function, as recommended by the ATCA in ATLAS working group. This mezzanine also supports communication over Ethernet.

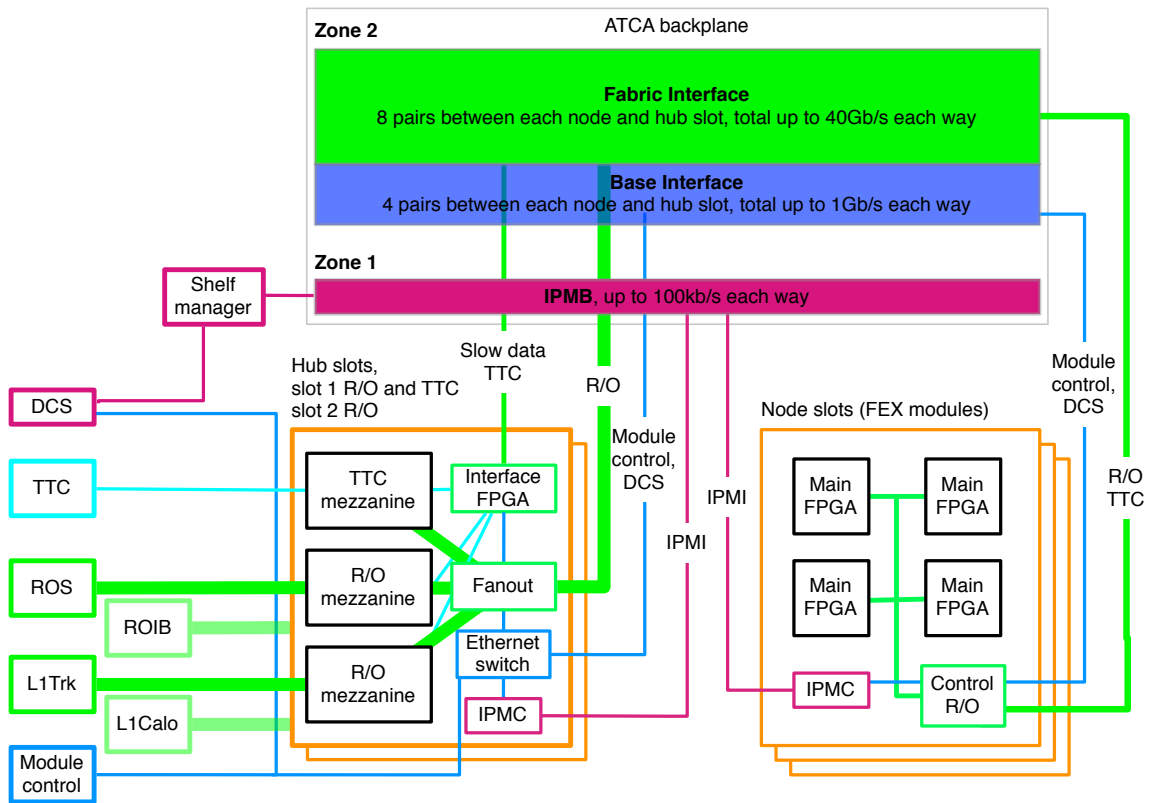


Figure 30: Block diagram of the ATCA Hub Module

ATCA hub module The ATCA Hub Module (henceforth “Hub”) provides a range of functions to support the FEX subsystems. Each L1Calo ATCA shelf will host two Hubs, in logical slots 1 and 2, which are hubs in the ATCA Base and Fabric Interface networks.

The Hub hosts several mezzanine boards and provides on-board logic to perform the following functions: route data to and from the FEX nodes in the shelf, aggregate FEX status information, route FEX control commands, and facilitate ATCA system monitoring. The design of the module provides maximum flexibility for both Run 3 and Run 4 operations. Thus, much of the Hub functionality resides on mezzanine cards that can be replaced or re-configured. For example, both the TTC and read-out requirements will change during the Phase-II upgrade, so these functions are housed on mezzanines that can be upgraded without redesigning the Hub.

In the L1Calo ATCA shelves, an Ethernet network on the Base Interface connects each Hub to every other module in the shelf (including the other Hub) in a dual-star layout. The use to which this network is put differs between the two Hubs in the shelf. For the Hub in Logical Slot 1, this network is used for module control (the IPbus network). For that in Logical Slot 2, it is used to collect DCS traffic from the FEX modules.

The Fabric Interface of each L1Calo shelf also forms a dual-star network, in which each Hub module is connected to every other module in the shelf via eight differential signal pairs. Each pair can carry a signal of up to 6.4Gb/s. The uses to which they are put are as follows. One pair is used to transmit the recovered TTC clock from the Hub to each module in the shelf, as described in Ref. [3.5]. A second pair is used to transmit the decoded TTC signals and near-real-time signals such as ROD busy. The remaining six pairs can be used by the Hub to receive read-out data. In the eFEX system each Hub receives read-out data on four signal

pairs from each of the 12 eFEX modules in the shelf. In the jFEX system each Hub receives read-out data on six signal pairs from each of eight jFEX modules. In each of these cases the Hub receives read-out data on a total of 48 signal pairs, but the signal pairs populated in the two cases are not identical. They are different subsets of a total of 72 hardware pairs used for this purpose.

The Hub module has sites for three mezzanines, two of the sites adjoin the front of the module, whilst the third faces towards Zone 3 of the backplane. The read-out data received by the Hub are fanned out to all three of these sites and read-out control signals from the sites are merged in a Xilinx Virtex 7 FPGA (the Interface FPGA). In addition to this, one of the mezzanine sites has connectivity for TTC signals, to and from the Hub. Thus, a Hub module may host a maximum of three ROD mezzanines, or two ROD mezzanines plus one TTC mezzanine.

In Run 3, the functions of the two Hub modules in an ATCA shelf are partitioned as follows (see Figure 30). The Hub module in slot 1 houses one ROD mezzanine card and a TTC mezzanine. It provides the TTC and module-control interface for the whole shelf, and transmits read-out data to the ROS. The Hub module in slot 2 houses one ROD mezzanine, which transmits read-out data to the ROIB. It also provides the DCS interface for the shelf.

In Run 4, each Hub module carries two ROD mezzanines, with the extra hardware being used to transmit read-out data to L1Track and L1Calo. One ROD mezzanine (on the slot 2 Hub) is spare, to support unforeseen requirements at Run 4.

As noted above, the two Hub modules in each ATCA shelf will provide slightly different functionality using the same base hardware design. The Hub module in slot 1 (Figure 30 with mezzanines supporting TTC and read-out to ROS and L1Track) will provide read-out for the ROS, TTC decoding, module control via Ethernet switch and IPMC functionality. The Hub module in slot 2 (Figure 30 with mezzanines supporting the ROIB and L1Calo read-out) will provide read-out for the ROIB, DCS interface via Ethernet switch and IPMC functionality. These differences may expand in Phase-II when additional Hub demands are likely to arise.

The Read-Out Driver (ROD) module In Run 3, when the eFEX and jFEX subsystems receive an L1A signal, they transmit the read-out data associated with that trigger decision to the ROS and ROIB systems. The key element in this process are the ROD mezzanine cards housed on the Hubs. Each of these receives the read-out data from all of the FEX modules in that shelf. The RODs check, re-organise and buffer the data before transmitting them downstream. For each shelf two RODs are required. The slot 1 Hub houses a ROD that supplies data to the ROD, and the slot 2 Hub houses a ROD that supplies data to the ROIB.

In Run 4, the RODs are required to implement a different read-out system. Here the FEX subsystems form part of the L0Calo system and the transfer of data from the FEX modules to the RODs is initiated on receipt of an L0A signal. In addition to checking and reorganising the data, the RODs hold them pending the arrival of an L1A signal. If this signal is received, the RODs transmit the corresponding data to the DAQ System, to HLT, to L1Track and to L1Calo. The data for each of these destinations is handled by a different ROD.

A single hardware design is used for all of the RODs in the FEX subsystems, in both Run 3 and Run 4. However, multiple firmware designs are required to handle the data from the different types of source (eFEX/jFEX), for the different types of destinations and for each run. The form of the ROD is based on the FMC standard, but with an extended area.

ROD I/O Read-out data reach the ROD via a 400-pin FMC connection to the Hub. They are received as differential signals of up to 6.4Gb/s. Each ROD in the eFEX subsystem receives four inbound signals from each of 12 eFEXs in the shelf, whereas each ROD in the jFEX subsystem receives six inbound signals from each of eight jFEXs in the shelf. Thus, each ROD receives data on 48 input signals. However, the pins on which these data arrive differ from the eFEX to the jFEX subsystems and a total of 72 physical pairs are required to satisfy both cases. There is also a signal path from the ROD mezzanine, via the Hub and backplane fabric, to each eFEX module, enabling flow-control signals to be sent. This path is shared by all of the ROD mezzanines on a Hub and also with TTC commands. The Hub manages the traffic.

The read-out data are transmitted from the ROD optically. The precise requirements of the downstream systems have yet to be determined, but the baseline implementation of the ROD has one Avago MiniPOD transmitter, providing a total bandwidth of up to 120 Gb/s for read-out data. It also has one MiniPOD receiver to implement the return links required by such protocols as S-Link. This hardware is fully compatible with the proposed detector read-out upgrades described in section 7.2.1 and adoption of such a scheme, as described for the Liquid Argon detector's front-end electronics, is under active consideration.

IPbus is routed to the ROD mezzanine via the FMC connector, for module control.

ROD functionality The data received at the ROD are 8b/10b encoded and accompanied by a checksum. They are decoded and the checksum evaluated. The latency constraints on the trigger do not allow for error correction, but errors are flagged and link performance is monitored. The data from the 48 links are merged to produce a single packet. This is formatted with a 10-bit CRC, 8b/10b encoded, and placed in an output buffer before being transmitted optically downstream. Flow control handles back pressure and if the buffers reach a programmable threshold a BUSY signal is passed to the CTP (via the Hub).

In Run 4, additional functionality is implemented: data are buffered until an accept signal (the L1A) is received. Those data accepted are transmitted downstream, whereas data not accepted are discarded. (That data have not been accepted is established by receipt of an L1A corresponding to subsequent data.)

Slow control of the ROD functions, and reading the ROD status, is achieved using IPbus. The occupancy of the ROD data buffers is monitored and this information is available via the IPbus, as are samples of the read-out data.

ROD FPGA The core functionality of the ROD is implemented in a large, Xilinx Virtex 7 FPGA. The two key requirements on this device are the amount of high-speed transceivers and the amount of internal memory available. Given the large number of high-speed signals received, a high-end FPGA device is required, with 80 multi-gigabit transceivers.

The amount of memory required on the FPGA is determined by the mode of operation in Run 4, where data are held until an L1A signal is received. For the jFEX, it is estimated that the maximum volume of data that can be sent to the ROD per event is 4000 bits. For the eFEX it is half this, so it is the jFEX subsystem that determines the requirements. Given a theoretical maximum L0A rate of 20 MHz, data arrive at the ROD from each jFEX at 80 Gb/s. Aggregated over the eight eFEXs, this gives a received data rate of 640 Gb/s per ROD. 13 Mb of memory is required to buffer data at this rate for the baseline L1A latency of 20 μ s. The Virtex 7 FPGAs under consideration for the ROD can provide more than twice this.

The configuration data of the ROD FPGA is held locally in flash RAM. Access to the RAM is provided via IPbus, allowing it to be read and modified remotely.

DCS As stated above, the ATCA specification defines a complete hardware platform management layer based on IPMI, with an IPMC on each module and overall control residing locally in the ATCA Shelf Manager in each shelf and globally in the System Manager.

There are no direct interfaces between this IPMI layer and WinCC, the system used by DCS, so a dual pronged approach will be adopted. All parameters requiring monitoring in the system will be reported directly to DCS from each module over the module control/DCS Ethernet network and the subset that are necessary for safe operation of the system will also be reported to the local Shelf Manager over IPMI. The System Manager and Shelf Managers will autonomously control the local powering and cooling of the ATCA shelves, subject to overall control coming from DCS.

TTC In order to accommodate changes in the way trigger information is distributed in Run 4, the TTC-L1Calo interface is implemented on a mezzanine card, which can be replaced with minimal expense. This card is mounted on the Hub. Its form is based on the FMC standard, but it will have an extended area if necessary.

The TTC mezzanine receives signals from the TTC optically. It recovers the clock and transmits this, via the Hub and ATCA backplane, to all modules in a shelf. It also decodes the control information (event counter reset, the bunch counter reset and L1A). These are passed down to the Hub module, where they are re-coded to an internal L1Calo format and transmitted, via the backplane, to all of the modules in the shelf. This decoding and re-coding of the TTC commands ensures that in Run 4, when the TTC system changes, only the front end of the L1Calo interface needs to be modified. The internal protocol used between the Hub and FEX modules is unchanged.

Test module The main purpose of the Test Module is to certify the many multi-gigabit data paths on the eFEX/jFEX modules during prototyping and prior to installation in ATLAS. However, since this module will be built first, it will also be used to confirm the high speed signalling and technology choices intended for the FEX modules.

The Test Module generates test vectors for transmission to the FEX modules and records the results from their outputs. The baseline signalling bandwidth is 6.4Gb/s, but, for the reasons described above, the module should be capable of generating and receiving data in the 6.4Gb/s to 10Gb/s range.

The majority of Test Module functions are provided by two large FPGA devices, the first holding data for transmission to the FEX modules, while the second provides memory to record returning output data. The memory depth is sufficient to source and sink test data for a full a LHC orbit period (3564 bunch crossings). The data received from the FEX modules may be read out for offline analysis or checked in real-time against an expected pattern. Long-term testing, for example to measure bit error rates, is achieved by iterating over the test data.

High-speed signal ports on the main FPGAs are connected via Avago MiniPODs to 12-way optical fibre ribbons, each module providing 96 optical outputs and 24 optical inputs. Two Test Modules are needed to source and sink all the inputs and outputs for one eFEX module, and three to fully populate a jFEX module.

Two smaller FPGA devices, the Control/TTC and System ACE FPGAs, provide control and timing functions. The Test Module can use the TTC clock signal, or provide a local clock and real-time control signals over the shelf backplane. The Control/TTC FPGA also provides the interface to the IPbus network, which is used for the functional control of the module and environmental monitoring.

The module is configured by the System ACE FPGA, using configuration data held in an SD memory card or loaded via IPbus. The module conforms to the ATCA standard and has an estimated power consumption of under 100 W.

3.5 Latency Considerations

Latency is discussed in detail in Section 5.4. A detailed estimate of the system latency has been prepared [3.6], showing a cost of 13.5 Bunch Crossings for the eFEX and 14 Bunch Crossings for the jFEX. In each case, most of the time is consumed by data conversion from high-speed serial to parallel formats, and the algorithmic processing requires only 5 Bunch Crossing periods. The latency estimate will be refined when the final choice of algorithms is complete, and measurements have been made on working hardware.

3.6 Firmware

All of the core algorithmic functionality of the L1Calo Phase-I system, and much of the supporting functionality such as slow control, is implemented in firmware. The hardware in which this functionality is implemented is described in Section 3.4. It is also useful to look at a functional breakdown of the firmware required. Such a breakdown does not map transparently onto the hardware, as some FPGAs house more than one functional element, some have multiple firmware loads (for example, for diagnostic purposes) and some functional elements, such as control functions, are common to multiple hardware designs. The high-level functional breakdown of the firmware required for the FEX subsystems is as follows.

Trigger algorithms

- eFEX/jFEX feature-extraction algorithm
- eFEX/jFEX board-level merging algorithm

Read-out and data monitoring

- Upstream read-out logic on processor modules (versions for eFEX/jFEX)
- Downstream read-out logic on ROD (versions for eFEX/jFEX and DAQ/ROI)
- eFEX/jFEX rate-monitoring logic

Infrastructure

- High-level ATCA control logic (IPbus)
- Low-level ATCA control and monitoring logic
- TTC interface
- High-level ATCA hub logic
- High-speed transceiver logic
- Test Module data source/sink logic (multiple versions)
- eFEX/jFEX commissioning/diagnostic logic

Firmware for a single FPGA device is typically developed by a single engineer. However, owing to the large amount of firmware required, the work will be distributed amongst participating institutes. The firmware is divided into functional blocks, the divisions between which are chosen to minimise interfaces. For example, the real-time algorithmic functionality of an FPGA might be split from the slow control functionality.

To facilitate the sharing of knowledge and design files amongst the participating institutes a common set of tools are used. Due to the expertise in Xilinx devices that resides within the collaboration these will be used to implement the bulk, and potentially all of the field-programmable logic. Firmware is developed using VHDL entry into a Xilinx-only tool flow and SVN is used to maintain a common archive of the source code. A house style guide for VHDL has been developed to facilitate the writing of code that can be understood and transported easily between designers and institutes.

3.7 Software

3.7.1 Online software

The new electronics for Run 3 (and Run 2) will be integrated into the existing online software. The aspects to be covered include module configuration under the TDAQ run control, the use of configuration databases to store the configuration and conditions data for the modules, the detailed bit-level simulation of the modules to enable comprehensive testing, the low level control interface to the modules, and the operational monitoring when ATLAS is running.

Most of the upgrades prior to Run 3 are to existing VMEbus modules or new VMEbus modules (e.g. CMX). For these the associated software work will be mainly updates to existing code or provision of new versions of packages based heavily on current packages. However, for the new ATCA modules [3.7] such as the Run 2 L1Topo and the Run 3 eFEX and jFEX systems, more significant software developments are required.

IPbus [3.8] will be used for the underlying access software for configuring ATCA modules. A custom configuration package for each module type will be developed on top of IPbus. Each package will conform to an existing interface for current VMEbus-based L1Calo modules, so that the same run control software can be used for the whole L1Calo system, presenting a uniform interface to TDAQ. The module configuration packages will take their data from new folders to be added to the L1Calo COOL database and from the extended trigger menu database.

The existing bit-level simulation package will be extended to include each new module to enable testing of prototype modules and the final installation. The package can generate arbitrary test vectors and patterns of triggers to be run through the system and compared with results from actual hardware. These simulation packages are more detailed than those presently used in offline analysis and, in time, are expected to replace them.

A suite of test and low level diagnostic software will be required. For the Run 2 VMEbus-based modules this will be via an extension of existing diagnostic software. For the new ATCA modules, test software will be layered above the IPbus software (which includes some diagnostic tools).

Operations monitoring software for the new modules will be developed and include monitoring of low level trigger rates and error counters. It will be provided as part of the online software under TDAQ run control and published to the TDAQ information service. Higher level monitoring will be written within the offline framework.

3.7.2 Offline software

Detailed performance studies are required to select the best algorithms to be executed in the processors, and also to determine the optimal values of parameters such as the digitisation

scale. Results from such studies are presented in Section 3.2. The studies will be continued to optimise performance further.

Re-formatting of L1Calo readout data during offline event reconstruction is performed to enable studies of data quality and trigger performance, or as input to physics analyses.

Offline monitoring and diagnostics software compare the full data read out from each module with Monte Carlo emulation of each module to validate the flow of data through the system and to localise any hardware problems.

3.8 Testing, Installation & Commissioning

3.8.1 Testing

Module tests are used to verify that production modules, deployed in the trigger, are known to be fully working.

There are three main individual module tests:

1. On receipt from manufacture, and after initial power tests, all modules undergo JTAG continuity checks, to establish that FPGA connections are correct. Dummy connectors are used to include backplane and other electrical connections in this test. Correct behaviour of module power controls is verified, and module power consumption is checked.
2. FPGA devices are loaded with firmware, and a series of register-level tests is run to establish access to all module parameters.
3. Diagnostic firmware is loaded, and BERs are measured for all high-speed links using random data patterns. The data eye opening is also measured. These tests establish that board traces have been manufactured with the correct impedance, and that the optical components are functioning to specification.

At this point, system-level tests are performed in a test rig environment comprising of one or more FEX modules, Hub modules with RODs, computer and read-out hardware, TTC timing infrastructure, and a DCS subsystem. A custom test module, (see Section 3.4.5) will be preloaded with test vectors, which are then driven into all the optical link inputs of the module(s) under test. A pre-defined pattern of L1A signals will be generated, causing read-out to be initiated for the corresponding bunch crossings. A detailed software simulation of the trigger algorithm will be run using the same algorithm settings as those loaded into the module hardware. Knowing the timing of the L1A signals relative to the test vectors, the simulation will calculate the expected content of the read-out data for comparison with the actual read-out data. The test will be run at the maximum expected average and burst L1A rates, and will use a variety of test vectors designed to reach the boundaries of the algorithm.

Each module under test routinely self-monitors all input links for errors. As the tests progress, additional modules are added to the test-rig, so that realistic operating conditions are incrementally built up (including representative loads on power supplies and backplane signal occupancies). The system will be power cycled many times, and long-duration tests will be run to verify system stability.

These tests will be run at both development centres and at CERN prior to installation in ATLAS.

3.8.2 Installation

The eFEX and jFEX subsystems will comprise three ATCA shelves, two for eFEX and one for jFEX. These will be installed in USA15, adjacent to optical patch panels and another three ATCA shelves for the LAr DPS, in a location with the lowest possible latency path from the detector through to the CTP. The only viable location for these is at the end of the row of racks currently occupied by L1Calo, necessitating relocation of other hardware. The ATCA power and cooling requirements for these four racks will be reviewed and upgraded where necessary.

Four sets of optical fibres will need to be installed for the new system: from the LAr front end crates to the DPS via the holes in the shielding wall (responsibility of the LAr group), from the DPS to the eFEX and jFEX via optical patch panels, from the JEM daughter cards to the eFEX and jFEX via optical patch panels and from the FEXs to the L1Topo crate. Careful planning will be needed to minimise disruption to the existing infrastructure, as all these sets of fibres will share already heavily used underfloor cable trays,

The initial L1Calo system will be kept in operation for the start of Run 3 running. Once the DPS and FEXs are fully commissioned, it would be possible to decommission and remove some of the old L1Calo crates, though the Receiver and Pre-Processor crates for TileCal will remain in use.

3.8.3 Commissioning

The main commissioning phase will commence when all infrastructure (including ATCA shelves, fibre plant, timing (TTC) and read-out fibres) is in place and the production modules begin to arrive. Tests as described in Section 3.8.1 will be repeated on the installed hardware, thus allowing qualification of individual parts of the system in isolation from their clients and data-sources. Tests of individual link performance will be checked in real-time using hardware error counters (e.g. parity) or hardware pattern-comparison at the front-end of each module. Joint tests with signal providers (e.g. the LAr DPS) and system clients (CTP, ROIB, DAQ) will be performed.

Final system qualification will use cosmic ray triggers and high-rate stress tests to probe operational regions which remain after systematic vectors (like ramps and repetitive simple patterns) have been used. Offline data will be compared to simulation based on read-out of many events, in search for all discrepancies. Full qualification of the complete system in the final ATLAS environment is important for detection of problems from powering, cooling, electronic noise and system timing.

3.9 System Performance

3.9.1 Validation

First beam Although the new trigger system will be thoroughly checked during installation, final tuning and verification require LHC beam. For example, while cosmic data taking is good enough to measure individual signal and overall system timings to the bunch-crossing level, synchronised LHC beam data is required to reach the ultimate precision. The optimum settings and calibrations in the new filters in the DPS will require feedback from colliding beam data and experience at the higher luminosities, and the performance of the new trigger will improve as these approach their ultimate resolution. It is during this tuning period that

it is most important to maintain the legacy trigger, providing a stable and understood trigger for ATLAS physics and a baseline for measuring the efficiency of the new trigger.

Much of the eFEX and jFEX timing and debugging can be done with cosmic data. However, the signal patterns for physics data are unlikely to be fully covered in any test, and a very thorough check of the trigger decisions will be needed in early running. One of the most powerful tools is a precise, event-by-event and bit-by-bit, simulation of the trigger logic, comparing the processing of the calorimeter signals through the trigger hardware with the values predicted by simulation. It will not be possible to connect all outputs of the new trigger processors to the CTP while the old trigger is still needed. However, the processing of the non-connected trigger outputs can still be checked from readout of the processor outputs and ROIs.

Detailed performance studies Once post-LS2 data taking has been well established and the fundamental calibrations are becoming more stable, it will be possible to study the performance of the new trigger algorithms. These can be done with standard tag-and-probe methods, and by comparison with the performance with the legacy calorimeter trigger system. The rates and efficiencies of the candidate new trigger items can be studied in the early running data, and the exact parameters of the trigger algorithms can be investigated offline by re-running with different parameters. The best environments and ratios can be found for the electron R_{core} algorithm, and all other trigger objects can be investigated similarly.

The performance of the more delicate trigger algorithms is particularly sensitive to details of pile-up response and choices made in the signal processing. Bunch-by-bunch luminosity and beam backgrounds are not easy to predict, and a reasonable quantity of stable data taking will be needed to study the performance of the more sophisticated triggers. Missing energy, for example, is often the most difficult to control. It is likely that some algorithms, for example electron and τ , are qualified early, and that some triggers will be moved to the new digital processor earlier than others. Only when individual trigger objects are understood would they be used in the Topological trigger. Even when all trigger processing is running on the new hardware, the older processors will remain for some time for checking purposes, and it may be desirable to avoid radical changes to the trigger menu during running periods. However, all elements of the new digital trigger are expected to be fully operational after a few months of post-LS2 running, and will entirely replace their predecessor subsystems during the first year of running.

3.9.2 Calibration and monitoring

Introduction Proper calibration of a calorimeter trigger is necessary for its efficient operation. Good uniformity of energy scale between individual SuperCells ensures sharp turn-on in trigger efficiency, minimising rate from events that are not on the efficiency plateau. Knowledge of absolute energy scale is useful for proper adjustment of Level-1 thresholds with respect to thresholds in HLT and offline reconstruction and analysis. In this paragraph, SuperCell level and cluster (object) level calibrations are briefly discussed. All trigger calibrations are determined with respect to offline energies.

Cell-level calibrations Calibrations are applied to individual SuperCell and trigger tower energies, by the DPS in the case of eFEX LAr inputs and by the Preprocessor for jFEX and

Tile signals. However, FEXs also need to be able to disable individual channels or to apply thresholds on their energies.

Feature Extractors will receive calibrated signals from the DPS that are time-aligned and assigned to an interaction bunch crossing. The signals need to be calibrated to the electromagnetic scale for eFEX, while keeping the possibility of using a different scale for jFEX. Both relative (cell-to-cell) and absolute calibration is expected to be known to percent level.

Cell-level calibrations include corrections for areas with reduced response due to dead or disabled noisy cells, and corrections for reduced high voltage in the LAr.

Out-of-time pile-up results from overlapping LAr signals, causing a baseline shift at the beginning of bunch trains. This is the dominant source of L1Calo trigger rate dependence on bunch number. Corrections for this effect will be applied in the DPS, and any possible remaining bunch number dependencies could also be corrected in the feature extractors.

Cluster and object-level calibrations The eFEX and jFEX module designs are determined primarily by their I/O needs, and only a fraction of the FPGA resources are needed for basic sliding-window algorithms. The extra capacity may be used to apply additional calibrations at cluster (or object) level, which is not possible in the pre-upgrade system. Options under study for the eFEX include cluster leakage corrections, dead material corrections based on energy deposited in the pre-sampler and total cluster energy, and determination of cluster position using weighted mean of energy deposits in individual SuperCells. For the jFEX, options include noise and pile-up suppression by weighting tower energies depending on distances from the jet centre, corrections for jets and global quantities, and underlying event corrections

Monitoring Monitoring runs in real time on a subset of events, and also on all recorded events during Tier-0 reconstruction, providing summary and detailed information for rapid diagnosis of problems. The information available includes distributions of input quantities (e.g. FADC hit-maps), scans for flags indicating hardware problems, and checks for discrepancies between read-out and calculated trigger response.

3.10 Optional Additions

3.10.1 The Global Feature Extractor (gFEX) module

Overview The Global Feature Extractor, gFEX, is currently under study as an optional addition to the L1Calo system. It is intended to identify large-area jets using algorithm window sizes up to 1.8×1.8 and detector data with a granularity of 0.2×0.2 in $\eta - \phi$. The architecture will permit event-by-event local pile-up suppression using η dependent baseline subtraction techniques. The gFEX architecture is suitable for other algorithms with the same granularity, including E_T^{miss} and ΣE_T .

The gFEX system (Figure 31) is implemented in a single ATCA shelf with a dual-star backplane fabric. Two gFEX modules each cover half the calorimeter η range. Two Hub modules located in the fabric slots of the shelf perform timing distribution, read-out and control. Input data are carried to the gFEX modules on optical fibres which are routed via rear transition modules to the main boards, where the signals are converted to electrical signals in receiver modules and distributed to processor FPGAs. Output data to L1Topo are transmitted through the front panel over optical links, while timing distribution, read-out and control functions use the ATCA backplane.

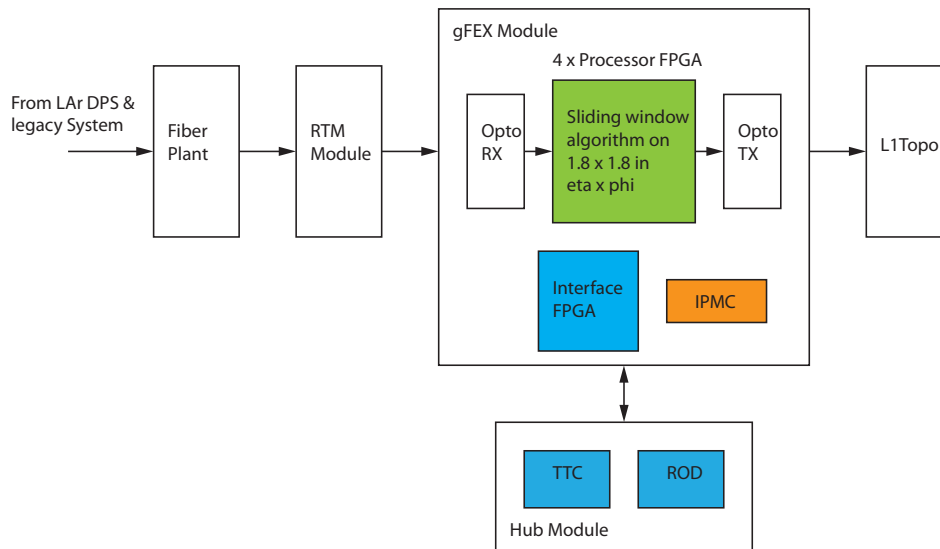


Figure 31: gFEX System

Input data Electromagnetic jet elements are typically produced by four FPGAs on each DPS module, each of which produces one multi-gigabit gFEX output link with a coverage of 0.8×0.4 . A full 0.8 -wide ring in η in the electromagnetic layer therefore requires 16 multi-gigabit links.

For Run 3, the gFEX obtains TileCal data in a similar way to the jFEX, using three upgraded JEM daughter cards per module, two with coverage of 0.8×0.6 and the third with 0.8×0.4 . A 0.8 -wide η ring requires 12 multi-gigabit links.

If a 0.2 in ϕ and full- η geometry is used, the input mapping to gFEX becomes more difficult. One option is to send the hadronic data over higher-speed links in overlapping segments (e.g. 3.2×0.2 at 10 Gb/s).

Calorimeter granularity is coarser at high η , so fewer links will be required to produce full rings at these η values.

Processor FPGAs The gFEX functionality, including the complete jet identification, is implemented in four Processor FPGAs. The η coverage of each processor FPGA includes one “centre” 0.8 -wide η ring that comprises the core coverage, as well as an additional η ring on each side to provide a 1.8×1.8 environment. The η coverage of the four FPGAs on the $\eta > 0$ gFEX is shown in Table 5, the $\eta < 0$ gFEX has a mapping with opposite η signs. It should be noted that in addition to “regular” jet identification up to $\eta = 3.2$, one FPGA also receives the full environment needed to process FCAL jets. To receive and process up to 32 input links per 0.8 η ring requires FPGAs with up to 96 transceivers.

FPGA No	η range	Core η range
1	$-0.8 - 1.6$	$0.0 - 0.8$
2	$0.0 - 2.4$	$0.8 - 1.6$
3	$0.8 - 3.2$	$1.6 - 2.4$
4	$1.6 - \text{FCAL}$	$2.4 - \text{FCAL}$

Table 5: η coverage of the $\eta > 0$ gFEX FPGAs

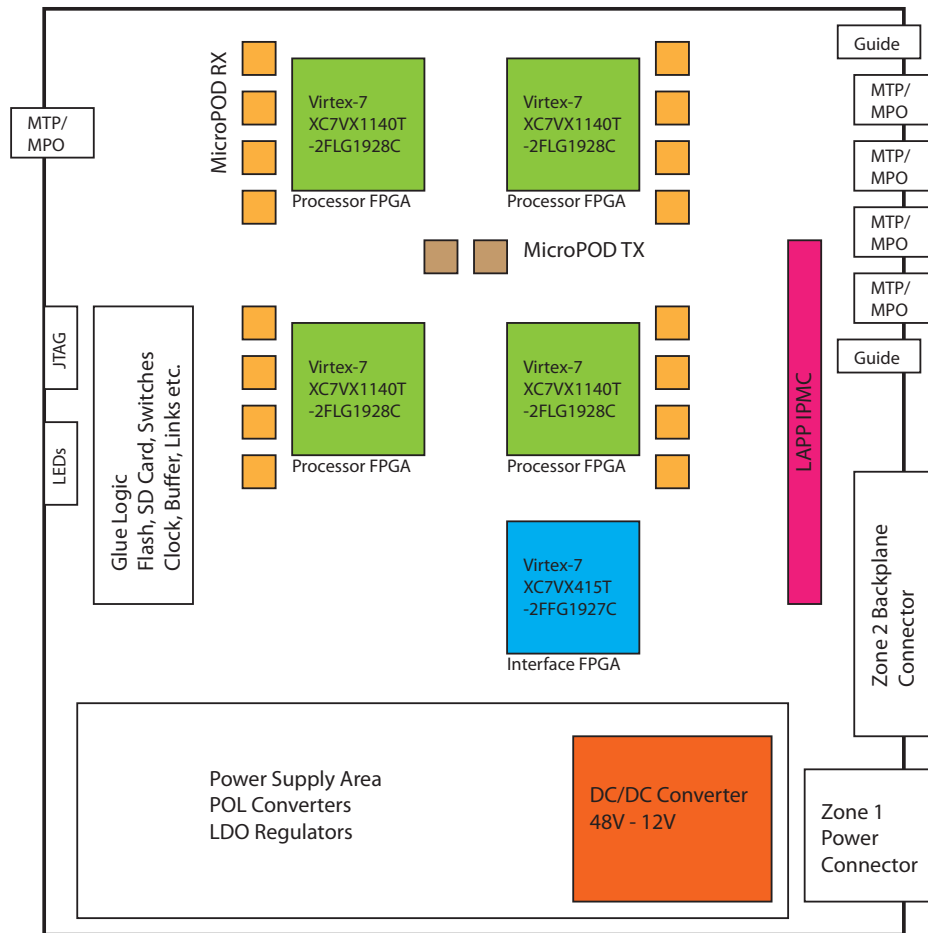


Figure 32: gFEX conceptual board layout

gFEX module design A conceptual layout of the gFEX module is shown in Figure 32. Each gFEX module receives data for six η rings from both calorimeters with up to 32 fibres per ring, giving a total of up to 192 fibres per module. The input fibres are connected to 16 12-channel parallel-optic receivers, from which the data are distributed to the processor FPGAs.

Fan-out of input link signals on a gFEX module is 1:1, 1:2 or 1:3 for different η ranges. This is achieved using a combination of PMA loopback and/or high-speed fan-out buffers. The two η rings on either side of $\eta=0$ are used by both gFEX modules and the data must therefore be duplicated upstream of the gFEX system. Since data from these η rings are used by only a single FPGA on each module, passive optical splitting would be sufficient.

TTC, read-out and configuration An interface FPGA on each module is responsible for TTC and read-out to a Hub module, with a Hub module providing a clock source to each gFEX module. The read-out data from processor FPGAs are transmitted to the Hub module over Zone 2 of the ATCA backplane.

The processor FPGAs will be configured on power-up through the interface FPGA. The configuration bit stream will be stored locally on an SD card. The interface FPGA itself is configured via a QSPI flash device, and the control firmware includes the algorithms to read and write the SD card and sequence the configuration data into the processor FPGAs.

3.11 Relation to the Phase-II Upgrade

It is expected that all calorimeters will convert to on-detector digitisation for Phase-II operation. Hadronic data would be provided in real-time for the trigger, either from new calorimeter RODs or by an extension of the Digital Processing System (DPS). Additional EM information would also become available in calorimeter RODs. The hadronic information, and additional EM information if needed, would enter the eFEX and jFEX using inputs prepared in Phase-I but mostly unused at that stage. The feature extractors would thus have data from all ATLAS calorimetry, and would be able to identify all current calorimeter-based trigger signatures with improved performance. With the loss of the analogue signals, the previous analogue-based L1Calo would be retired.

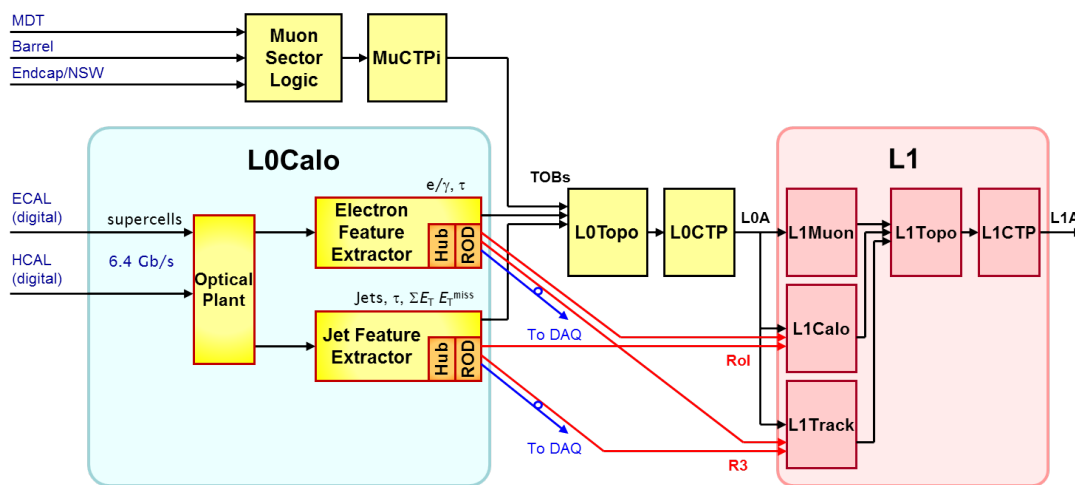


Figure 33: System Architecture during Run 4. The new L1 system is shown in red / pink. Other modules (yellow / orange) are adapted from the previous system to form the new L0Calo.

The hardware trigger system proposed for Phase-II uses a two-stage architecture, as illustrated in Figure 33. To provide the two-stage hardware trigger, the Phase-I trigger described above would be relabelled as the Level-0 trigger, and would be modified to provide ROI seeds with a rate of up to 500 kHz to initiate data capture in L1Track and possibly in L1Muon. The second stage of the Phase-II hardware trigger, i.e. the Level-1 Trigger, would start processing following a Level-0 Accept (L0A) decision. The additional time allowed implies that this stage of trigger could have available not only the Inner Detector tracks (from L1Track), but also more refined fine-grain calorimeter data (obtained by additional calorimeter ROD interrogation following the L0A), and improved muon track information (from the high-precision but slower MDTs). Processing for the Level-1 trigger decision would culminate in global trigger decision logic incorporating both topological and combinatorial functions.

New electronics would be required for the Level-1 trigger, as nothing comparable is planned earlier. The Phase-I eFEX and jFEX subsystems will be prepared in advance as far as possible. However, for some external interfaces neither the data content nor the physical transmission medium can yet be defined. In practice, this means that some external interfaces on the eFEX and jFEX modules will be provided by daughter boards which could be changed at Phase-II. The daughter boards involved are:

- **Timing Distribution:** In Phase-I this is provided via the TTC system. In Phase-II, more signals will be needed, including the Level-0 and Level-1 Accepts and the respective event numbers.
- **Voltage, Current and Temperature Reporting:** The future evolution of the DCS system is unknown at present.
- **Read-out:** In Phase-I, read-out is via S-Link. In Phase-II this is likely to change, and use of GBT links for data transfer off detectors is under discussion. It is not clear what protocols will be used for read-out of trigger subsystems in USA15.
- **ROI Read-out:** In Phase-I, Region-of-Interest data is provided from dedicated RODs via the ROI builder to the High-Level Trigger. This function is similar to the DAQ read-out above.
- **R3 Read-out:** The Phase-II Level-1 Track Trigger will require the coordinates of Level-0 features to initiate sparse read-out of the inner detector, via the Regional Read-out Request (R3) mechanism. These data must be provided from eFEX and jFEX, but the details are not yet established
- **Level-1 seeds Read-out:** The Level-1 trigger is likely to need ROI-like information from the Level-0 system. Requirements are not yet known.

3.12 Project Organisation

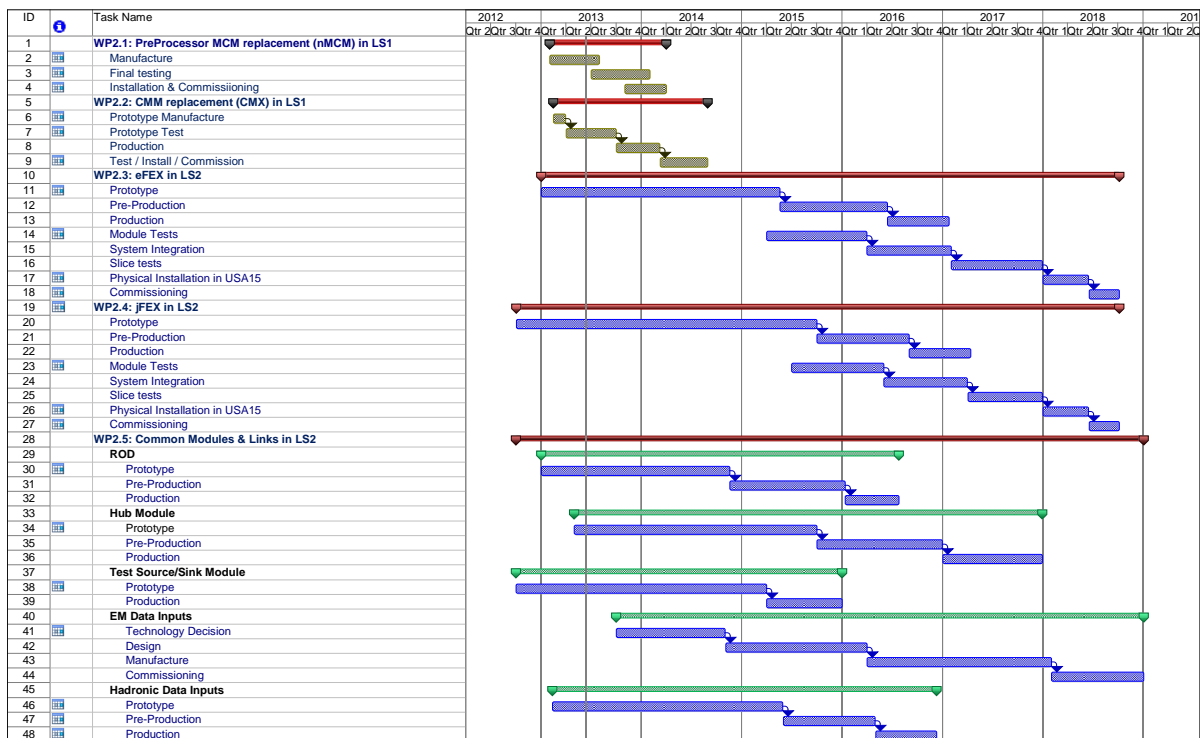


Figure 34: L1Calo upgrade project work breakdown

The project will be organised in a number of work packages, each with a subset of tasks. The time evolution of the principal tasks is shown in Figure 34. The Phase-I accelerated items are shown first, being the nMCM, CMX and L1Topo modules, and are brought to completion by the end of LS1 in December 2014. Associated with the development of each module is a parallel activity in production of FPGA firmware and control, test and monitoring software.

For simplicity, these have been excluded from the chart, but nonetheless form a very significant part of the overall project.

Table 6 shows the project major milestones. For completeness of the technical description L1Topo is included here, but is reported under Central Trigger within the MoU.

Milestone	nMCM	CMX	L1Topo	eFEX	jFEX	Modules & Links
PDR	N/A	Jun-11	Jun-12	Sep-13	Sep-13	Mar-14
IDR	N/A	N/A	N/A	Oct-15	Apr-16	Oct-15
PRR	Dec-12	Jan-14	Oct-13	Jun-16	Dec-16	Oct-16
Tested	Dec-13	May-14	May-14	Dec-17	Dec-17	Dec-17
Installed	Jan-14	Sep-14	Sep-14	Jun-18	Jun-18	Jun-18
Commissioned	Apr-14	Dec-14	Dec-14	Dec-18	Dec-18	Dec-18

Table 6: Principal L1Calo upgrade project milestones. PDR = Preliminary Design Review (review of prototype specification, before further design work proceeds); IDR = Initial Design Review (review of pre-production design (and schematics) before manufacturing that version); FDR = Final Design Review; PRR = Production Readiness Review (final review of production design before starting volume manufacturing).

References

- [3.1] R. Achenbach, P. Adragna, V. Andrei, P. Apostologlou, B. Asman, et al., *The ATLAS Level-1 Calorimeter Trigger*, **JINST 3 (2008) P03001**.
- [3.2] Y.-S. Lai and B. A. Cole, *Jet reconstruction in hadronic collisions by Gaussian filtering*, [arXiv:0806.1499 \[nucl-ex\]](https://arxiv.org/abs/0806.1499).
- [3.3] ATLAS Liquid Argon Calorimeter Group, *ATLAS Liquid Argon Calorimeter Technical Design Report: Trigger Electronics for the LHC Phase-I Upgrade*, Tech. Rep. CERN-LHCC-2013-???. ATLAS-TDR-??, CERN, Geneva, (in preparation). <https://twiki.cern.ch/twiki/bin/view/LAr/LArPhaseITDR>.
- [3.4] Avago, *MicroPODTM and MiniPODTM 120G Transmitters/Receivers*, (online). http://www.avagotech.com/pages/minipod_micropod.
- [3.5] PCI Industrial Computer Manufacturers Group, *Physics Design Guide for Clocks, Gates & Triggers in Instrumentation*, PICMG PDG.0 R1.0, (online), Apr, 2013. http://picmg.org/docs/PDG_0-R1_0-RELEASED-2013-04-23.pdf.
- [3.6] N. Gee, *Phase-I Latency Envelopes for the Level-1 Trigger*, Tech. Rep. ATL-DA-ES-0059 v.1, EDMS Id 1256858 v.1, CERN, Geneva, Dec, 2012. <https://edms.cern.ch/document/1256858>.
- [3.7] P. Farthouat, *ATCA in ATLAS*, Tech. Rep. ATU-GE-ES-0001 v.1, EDMS Id 1304001 v.1, CERN, Geneva, Jul, 2013. <https://edms.cern.ch/document/1304001>.
- [3.8] *CACTUS Project webpage*, (online). <https://svnweb.cern.ch/trac/cactus>.

4 Level-1 Muon Trigger

4.1 Level-1 Muon Endcap Trigger

4.1.1 Introduction

The current Level-1 muon endcap trigger is based on the Thin Gap Chambers (TGCs) of the middle endcap stations, referred to as the Big Wheel (BW), i.e. the BW-TGCs, which are covering an η range of $1.0 < |\eta| < 2.4$. Over the Phase-I shutdown, muon chambers located at the endcap inner station, called the muon Small Wheel, will be replaced by new chambers, sTGC and MicroMegas detectors with high-rate tolerance and improved resolution [4.1]. Together, these are referred to as the New Small Wheel, NSW. The coverage of the NSW is $1.3 < |\eta| < 2.7$ as shown in Figure 35. The primary motivation for introducing the new chambers is to improve the muon tracking performance under the large expected cavern background of neutrons and photons, and at the same time improve the rejection of fake muons in the Level-1 muon trigger by incorporating track-vector information from the NSW. To further improve the rejection of fake triggers in the full η coverage of the BW-TGC, hit information of the EIL4 TGC chambers and energy-loss in the Tile-calorimeter D-layer cells (D5, D6) in the range $1.0 < |\eta| < 1.3$ will be used (see Figure 35). The incorporation of the Tile calorimeter D-layer cells into the Level-1 muon trigger is described in Section 4.3.

Without these proposed improvements, the trigger rate above a p_T threshold of 20 GeV would be 51 kHz at a luminosity of $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, whereas introducing information from the EIL4-TGC and Tile calorimeter D-layer cells in LS1 will reduce the rate to 32 kHz. A further rate reduction to 15 kHz will be achieved in Phase-I, where the NSW is introduced into the Level-1 muon trigger logic as shown in Figure 36. Reducing the fake Level-1 muon rate will permit to keep a low p_T trigger threshold for single muons (20 GeV), thus maintaining the physics acceptance of the ATLAS experiment.

4.1.2 Overview of the new system

Figure 37 shows an overview of the Level-1 muon endcap trigger system. Track vector information from the NSW is combined with results from the current Level-1 muon trigger (based on the BW-TGC [4.2][4.3]) at the new Sector Logic board. The diagram inside the blue box of Figure 37 corresponds to the part of the current trigger system that is based on the BW-TGC and will be used as is after the Phase-I upgrade. The BW-TGC, which covers the range $1.0 < |\eta| < 2.4$, consists of three stations (TGC1, TGC2 and TGC3). The TGC1 station has three layers and the outer two stations (TGC2 and TGC3) have two layers each, for seven layers in total. The TGC3 is referred to as the pivot plane. The trigger algorithm extrapolates pivot-plane hits to the interaction point (IP) to construct roads following the infinite-momentum (straight) path for a track. Deviations (δR and $\delta\phi$) from this path of hits in the trigger planes are related to the momentum of the track. Coincidence signals are generated independently for R and ϕ . A 3-out-of-4 coincidence is required for the doublet planes of TGC2 and TGC3, for both wires and strips; a 2-out-of-3 coincidence for the triplet wire planes; and 1-out-of-2 possible hits for the triplet strip planes. The hit position information with granularity of ROIs and deviations (δR and $\delta\phi$) is sent to the Sector Logic board located in USA15.

The NSW is composed of two detector technologies, sTGC and MicroMegas. Both detectors measure the positions and incidence angles of muon tracks. Information on the candidate muon tracks, which are pointing to the IP within ± 15 mrad deviations, are provided to the




Figure 35: The ATLAS Muon Spectrometer with the view of the New Small Wheel and the Tile calorimeter

Endcap Sector Logic. This information consists of position (R and ϕ) and $\delta\theta$, the deviation of the incidence angle from a straight line to the IP.

The final trigger decision is done by merging the $R - \phi$ coincidence of signals from the BW-TGC and the information from the NSW.

Figure 38 shows the pivot plane formed by the TGC doublet plane (TGC3) furthest from the IP. The pivot plane is divided into two regions, Endcap ($|\eta| < 1.9$) and Forward ($|\eta| > 1.9$). The Endcap region is divided into 48 trigger sectors in ϕ , where a trigger sector is a logical unit that is treated independently in the trigger. Similarly, the Forward region is divided into 24 trigger sectors. The segmentation of trigger sectors projectively corresponds to the layout of the TGCs in the Big Wheels. The red lines in Figure 38 show projective boundaries on the NSW detector, which covers $1.3 < |\eta| < 2.4$ and whose structure has octant symmetry. Each octant has a Large NSW sector and a Small NSW sector. Boundaries of the NSW do not coincide with the segmentation of the trigger sectors. The segmentation of the trigger sectors and granularity of the Region-of-Interest (indicated by a red box labelled ROI in Figure 38) for the Phase-I upgrade are not changed from those of the present system. The sizes of the ROIs are approximately 0.025×0.030 in $\eta - \phi$.

Two types of sector logic boards will be developed, the “Endcap Sector Logic” board and the “Forward Sector Logic” board for the endcap and forward regions respectively. A single Sector Logic board serves two adjacent trigger sectors, therefore 24 Endcap and 12 Forward Sector Logic boards per side are required.




Figure 36: The expected trigger rate of Level-1 muon trigger as a function of p_T threshold. The same trigger condition in 2012 (black), the Pre-Phase-I configuration (blue) and the Phase-I configuration with NSW (red).

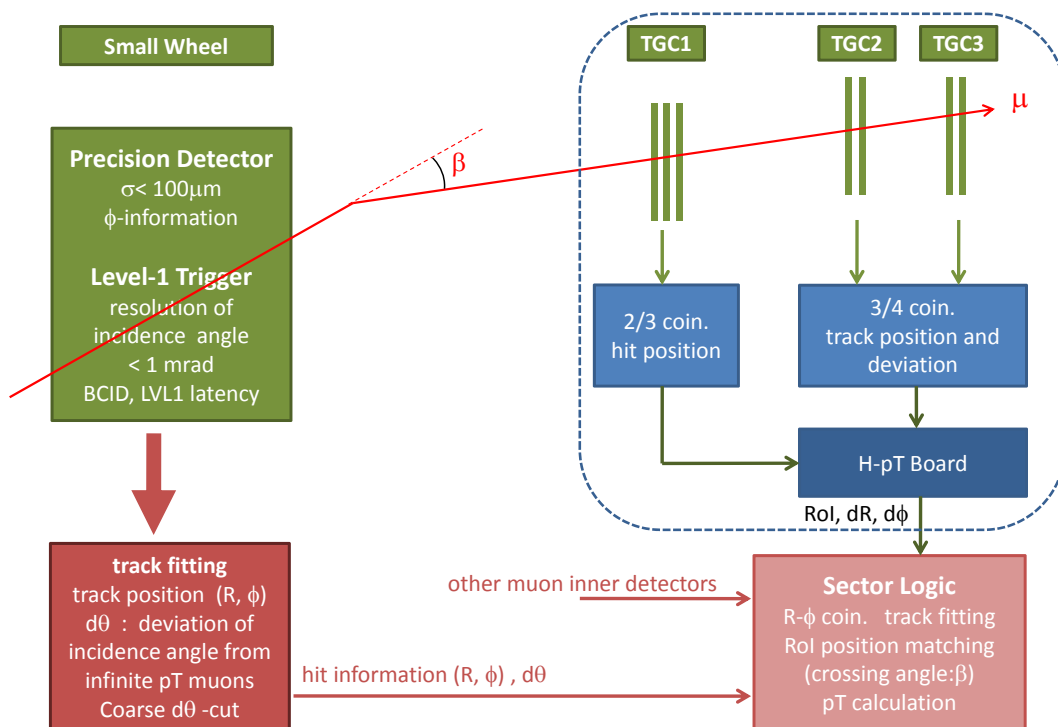


Figure 37: Schematic overview of the system

4.1.3 Interface between the New Small Wheel and the New Sector Logic

Figure 39 shows the signal distribution scheme from the NSW to the Sector Logic boards. Vector data of track segments, which are found in sTGCs and MicroMegs separately, are merged by the NSW trigger electronics. Combined track information is fanned-out and delivered to corresponding Sector Logic boards. The Sector Logic boards covering boundaries of NSW sec-

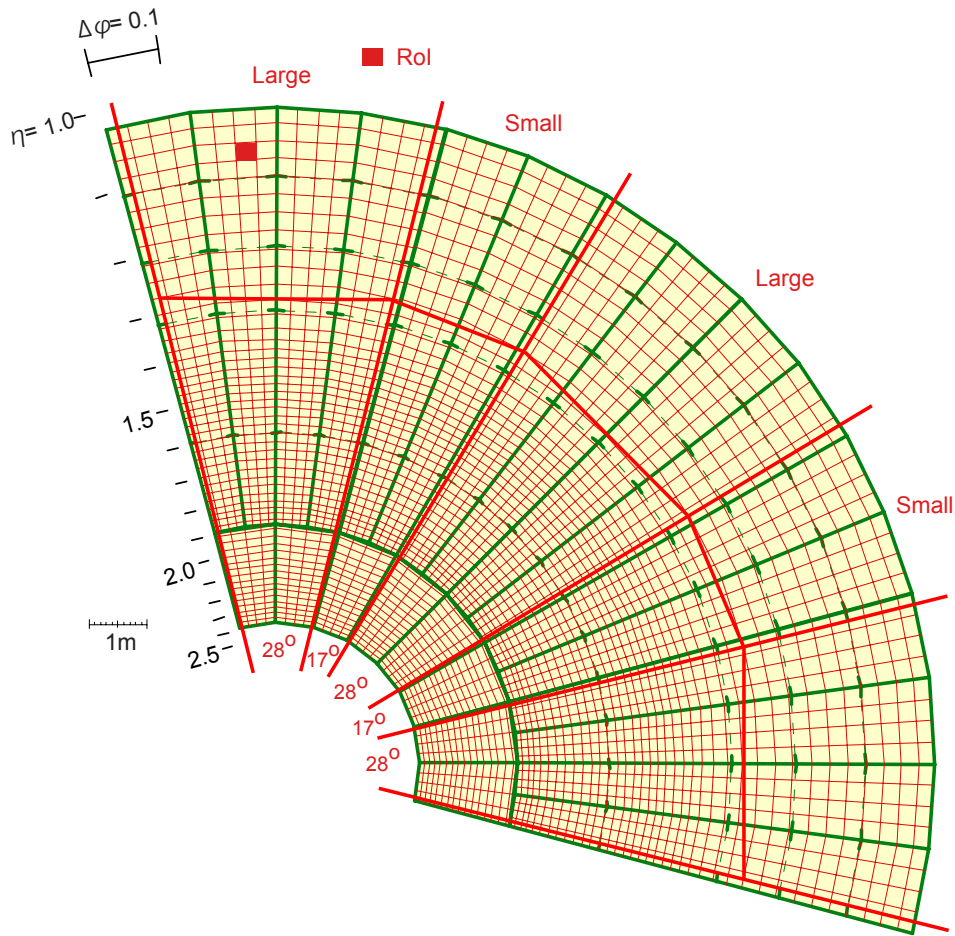


Figure 38: The Trigger Sector Segmentation and mapping

tors have to receive signals from both a Large NSW sector and a Small NSW sector. A Large NSW sector has to deliver track information to six Endcap trigger sectors and four Forward trigger sectors as shown in Figure 38. Similarly, track information of a Small NSW sector are provided to four Endcap trigger sectors and three Forward trigger sectors. The maximum number of signal fan-outs for distribution to Sector Logic boards is 7. (A single Sector Logic board serves two trigger sectors.)




Figure 39: The Interface between the NSW trigger electronics and the New Sector Logic

Table 7 describes the data format from the NSW trigger electronics to the New Sector Logic. One track segment is represented as 24 bits of data, which consist of 2 bits of hit information for each detector (sTGC and MicroMegs), 5 bits for $\delta\theta$, and 6 bits and 8 bits for ϕ and R position information, respectively. Required resolutions (1 digit) are approximately

1 mrad (± 15 mrad full scale) for $\delta\theta$, 10 mrad for ϕ and 0.005 in pseudo-rapidity η . The 2-bit hit information will provide an indication of track quality with details yet to be defined. One NSW sector can transmit data for up to 8 tracks. Table 7 also shows the format of the data transmitted from the NSW trigger electronics to the New Sector Logic. It is comprised of two IDLE codes for alignment purposes, 3 bytes of data for up to four track segments, the NSW Sector ID (4 bits) and a BCID number (12 bits) for every LHC bunch crossing (40 MHz). One 2-byte word is transferred after 8b/10b encoding at 320 MHz. Two optical fibres are assigned to each NSW sector, and the bit rate per link is 6.4 Gb/s. A commercially available standard serial link is used for the application.

Format of a track vector in NSW (24-bit/track)

Field:	sTGC hit	MM hit	$d\theta$ (mrad)	ϕ index	R index	rsv
Num of bits:	2	2	5	6	8	1

Max. number of tracks per NSW sector is 8.

Data Format from NSW trigger electronics to Sector Logic

Words (16-bit)	first byte	second byte
Word-0	comma	comma
Word-1	track-0	
Word-2		
Word-3	track-1	
Word-4	track-2	
Word-5		
Word-6	track-3	
Word-7	ID (4-bit)	BCID (12-bit)

$$8b/10b \text{ encoding} \times 16 \text{ bytes} = 6.4 \text{ Gbps}$$

Table 7: Data format. The comma (,) character is defined as an IDLE code for alignment purposes in the 8b/10b encoding.

4.1.4 The New Sector Logic board

Figure 40 presents a block diagram of the New Sector Logic board. The New Sector Logic board has three groups of input signals. Two groups will be used to receive R and ϕ coordinates (wire and strip data respectively) from the BW-TGC. These input links will implement the G-Link HDMP technology as used by the front-end electronics and trigger processor boards on the BW-TGC that will not be upgraded in Phase-I.

The third group of input signals are used to receive track information from the NSW (R , ϕ , $\delta\theta$), and additional input ports are introduced as spare inputs for future use, for example to receive inputs from the new muon inner detectors (EIL4, EE, BIS) covering the region of $1.0 < |\eta| < 1.3$ that are currently being studied.

Referring to Figure 40, signals from the BW-TGC and NSW are aligned with the BCID clock and decoded at the first stage of the Sector Logic board. The $R - \phi$ coincidence is made in the first Look-Up-Table (LUT), where δR (wire) and $\delta\phi$ (strip) signals from the BW-TGC are merged and p_T values are calculated. The contents of the LUT are fully programmable and can be set for each ROI independently. The outputs from the first LUT are p_T values and positions (ROIs) of track candidates. The inputs to the second LUT are the track vectors measured in the NSW and the output of the first LUT. Position matching between ROIs of the BW-TGC and track vectors in the NSW, as well as angle consistency (δR and $\delta\theta$) between two objects are required. The allowed position deviations between BW ROIs and the NSW are programmable values for each ROI and p_T value. For each ROI and track-vector match, the p_T values of track candidates are re-calculated with ~ 10 sets of p_T thresholds. As baseline implementation, the two highest- p_T muons are selected and their p_T value, ROI and a quality flag of coincidences among muon detector stations are transmitted to the MUCTPI. The hardware resource will allow for the selection of more than two muon candidates if that is required by physics. The amount of data to be transmitted to the MUCTPI is 64 bits per bunch crossing. The hardware specification of a serial link and the information content to be transmitted are described in Section 5.2.4.

All the input and output data sets of Level-1 accepted events are read out through the data acquisition system in parallel with the primary trigger logic.

The trigger processing and read-out functions for a trigger sector are implemented in a signal FPGA. The numbers of optical inputs are 12 from the BW-TGC, 6 from the NSW and 4 from other muon inner detectors per Sector Logic board, which covers two trigger sectors. Other inputs and outputs are 2 optical outputs for the MUCTPI, TTC signals and data read-out to the ROD. Module configurations and monitoring are performed via the backplane of the crate.

The Sector Logic will be implemented in 9U VME64x standard, which has enough performance for this application. The module is designed only for the Phase-I upgrade and is not compatible with the Phase-II upgrade, in which a large number of changes are expected that incorporate information from MDTs into the trigger logic in order to improve the p_T resolution. It is too early to decide on the specification of the sector logic for the Phase-II upgrade.

4.1.5 Latency

The latency numbers are given in Table 8, where numbers of the BW-TGC are measured and all others are estimated. Input signals from both the BW-TGC and inner muon detectors are delivered to the Sector Logic board at the timing of 41 BCs from pp collisions or earlier. The estimated latency for the signal processing of the Sector Logic is 10 BCs except for serialisation/de-serialisation time of the optical links. The delivery time to the MUCTPI is 57 BCs, which is an increase of 3.5 BCs from the present system. The maximum latency and its margin are discussed in Section 5.4.

4.1.6 Project planning

The schedule for the upgrade project is shown in Figure 41. The total number of boards to be constructed is 100 including spares. The milestones for the PDR and PRR of the hardware are Oct 1, 2015 and Nov 1, 2016 respectively. Firmware development and test are the heart of the project to ensure the performance. A wide variety of activities is necessary for software

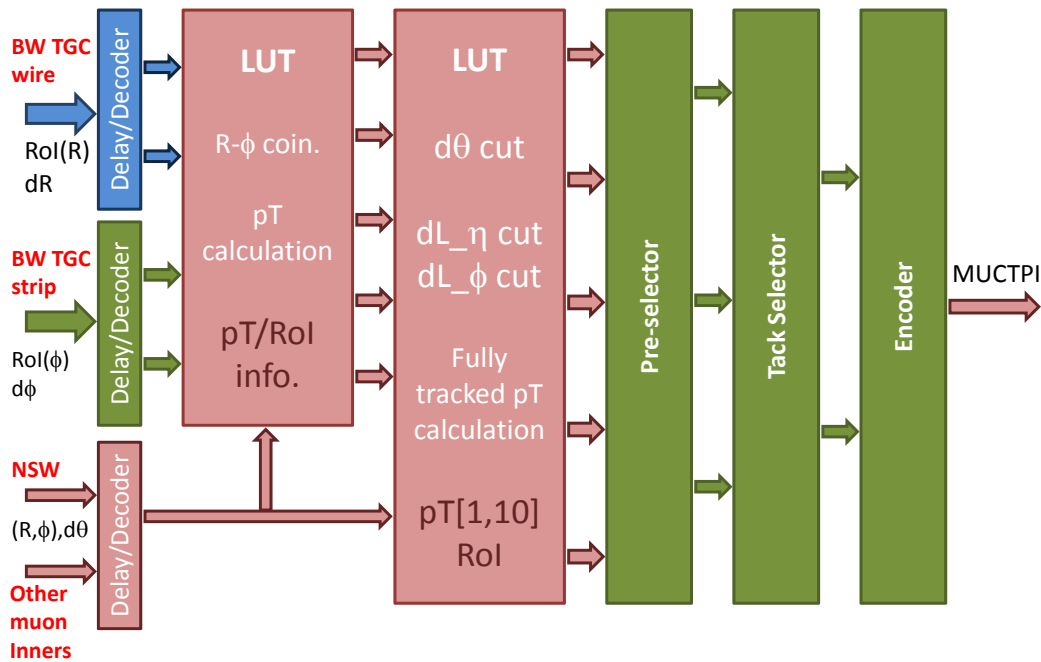


Figure 40: New Sector Logic block diagram

developments. The development of the simulation software is an urgent task, which is followed by trigger performance studies. Early development of the online software is necessary for comprehensive tests of the firmware as well as the hardware. Offline software must also be prepared for monitoring and data quality assurance.

4.2 Level-1 Muon Barrel Trigger

4.2.1 Introduction

During LS1, the Level-1 muon barrel trigger coverage will be extended by about 3% in the toroid feet and elevator regions through the addition of RPC chambers. In LS2, the upgrade of the inputs to the MUCTPI from electrical to optical will need to be matched by the upgrade of the electronics boards used to send trigger data from the muon barrel trigger to the MUCTPI.

Sectors 12, 13 and 14 of the barrel spectrometer have a trigger coverage about 20% lower compared to the other sectors as there are no RPC trigger detectors in the areas occupied by the toroid feet support structure (sectors 12 and 14) and by the elevator shaft (sector 13). The trigger coverage will increase after LS1 with the installation of new RPCs outside the toroid feet support and the elevator structures, placed in the projective area of the trigger holes. The new trigger towers will be integrated in the current barrel Level-1 system and read out by the same sector logic boards currently used for the feet and the elevator sectors.

4.2.2 The new Sector Logic to MUCTPI interface board

The barrel Level-1 muon trigger system uses programmable coincidences of three concentric layers of RPC detectors to detect the muons produced at the interaction point [4.6]. The on-detector Level-1 trigger electronics receives the RPC front-end signals and executes the trigger coincidence algorithm, while the Sector Logic (SL) boards, located in USA15, collect the trigger results coming from the on-detector electronics related to the same trigger sector.

New Small Wheel (sTGC + MM)

	nsec	CLK	Total CLK
TOF from interaction point to NSW (10m)	34	1.5	1.5
Signal Processing on detector		12.5	14
Optical Fibre (90m)		18	32
Signal Processing at USA15 (including sTGC and MM merge board)		6	
Serializer + Optical Tx			38
Optical Fibre (5m)		1	39
Signal fanout		1	40
Optical Fibre to Sector Logic (5m)		1	41

Big Wheel TGC (measurement)

	nsec	CLK	Total CLK
TOF from interaction point to TGC	65	2.5	2.5
Propagation delay on wire/strip	15	1	3.5
TGC response	25	1	4.5
ASD	10	0.5	5
Cable to PS-Board (12.5m max.)		2.5	7.5
Variable Delay, Bunch ID, OR		2	9.5
Variable Delay		1	10.5
3/4 Coincidence Matrix or 2/3 Coin.		2	12.5
LVDS Tx (SN65LV1023)		1	13.5
Cable to H-pT Board (15m max.)		3	16.5
LVDS Rx (SN65LV1224A)		2	18.5
Variable Delay		1	19.5
H-pT Matrix		2	21.5
G-Link Tx (HDMP-1032A) + Optical Tx		1.5	23
Optical Fibre to Sector Logic (90m)		18	41

New Sector Logic

	nsec	CLK	Total CLK
Receive signals from BW and NSW			41
Optical Rx + De-serializer		2	43
TGC R-phi coin. (LUT)		3	46
BW - NSW coin. (LUT)		3	49
Track selection		3	52
pT encoding		1	53
Serializer (64-bit/clk,3.2Gbps) + Optical Tx		2	55
Optical fibre to MUCTPI (10m)		2	57

Table 8: Latency

A schematic view of the off-detector electronics, showing the changes foreseen for Phase-I, is shown in Figure 42. The SL trigger results are sent to the MUCTPI via Interface Boards (IBs). Each IB is installed adjacent to an SL board and receives SL data via the RODbus backplane. To provide the required optical output, the 64 VMEbus-based IBs that currently provide electrical connectivity from the barrel SL to the MUCTPI will be replaced. Each new board will be equipped with one FPGA to perform the local logic and the data serialisation, and one optical transceiver to send the data to the MUCTPI.

Figure 43a shows the I/O signals of the current IB. The board is connected to the adjacent SL via the RODbus backplane [4.7]. This backplane has a 48-bit connection used to send/receive data to/from the SL, and a 14-bit connection available for monitoring and testing. The current front panel of the board has a 32-pin output connector used to send the SL trigger data to the MUCTPI, and a 16-pin optional input connector originally foreseen to receive TileCal trigger data (as an option in case of higher than estimated backgrounds). The 14-bit backplane bus is connected to various connectors on the front panel of the IB providing input/output control and monitoring signals to the SL, i.e. busy, trigger, and TTC.

The format of the 32-bit trigger data, output by the current IB at 40 MHz, is shown in Table 9. A maximum of two trigger candidates per trigger sector are sent from the SL. Each candidate is identified by 5 bits for the ROI, 2 bits for the ϕ overlap information (asserted in case the candidate originated from a region overlapping in ϕ with the adjacent sector) and 3 bits for the p_T threshold. Three additional bits of the BCID trigger word are sent for

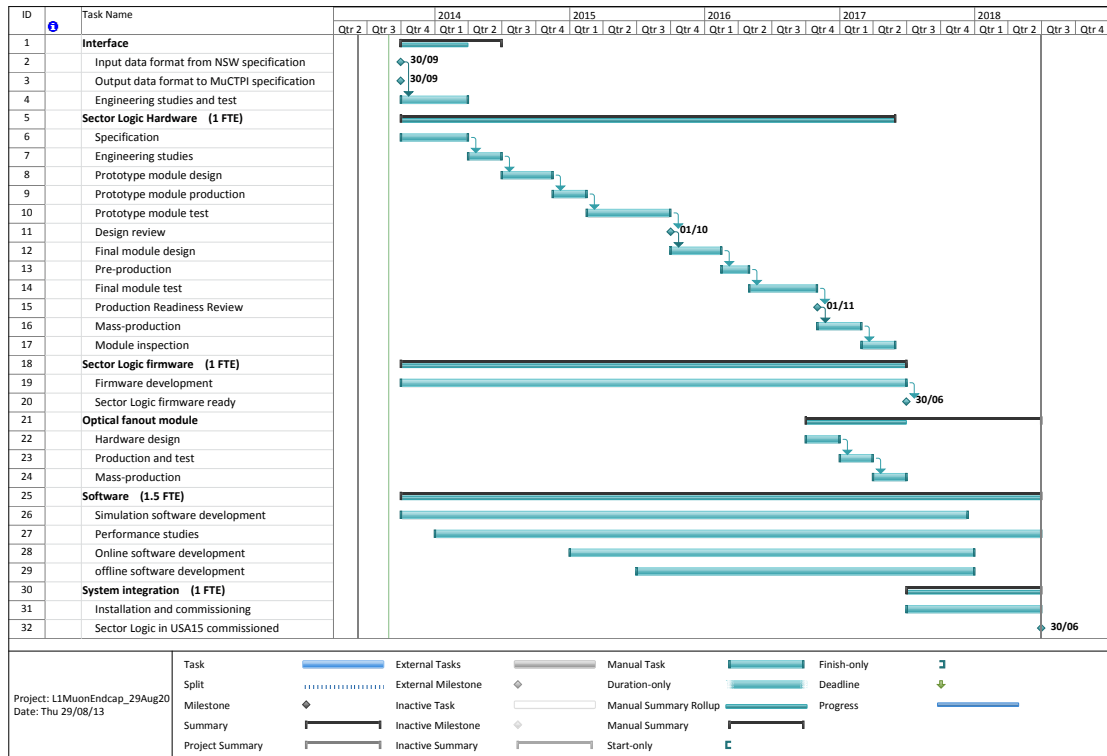


Figure 41: Schedule of New Sector Logic board




Figure 42: Scheme of the barrel Level-1 muon trigger system, showing the parts to be upgraded for Phase-I

synchronisation purposes, and one additional bit is used to indicate more than two trigger candidates found in a sector.

Figure 43b shows a schematic block diagram of the new IB. An FPGA will receive the trigger data coming from the SL (via the 48-bit connection of the backplane RODbus), serialise and send it to an optical transceiver on the board. In order to have a robust and flexible data

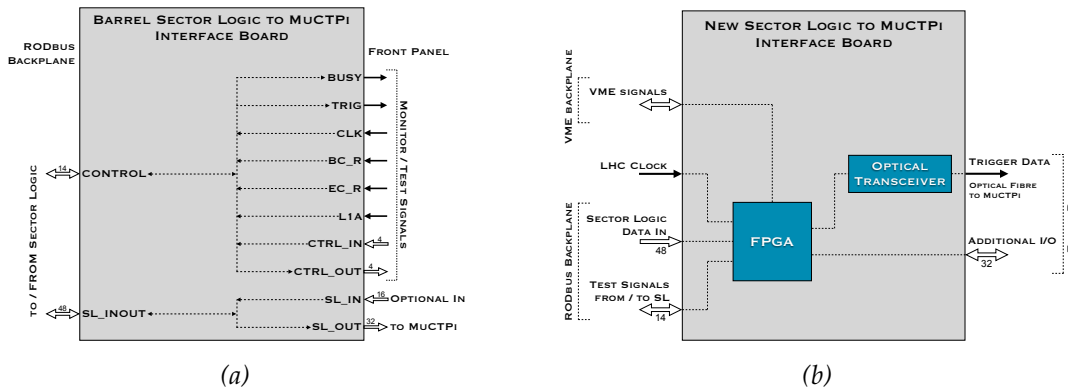


Figure 43: (a) The current SL to MUCTPI Interface Board schema showing its I/O signals. (b) The block diagram of the new Interface Board.

Bit	Usage
0	> 2 candidates in a sector
1-5	1 st candidate ROI
6-7	1 st candidate reserved
8-9	1 st candidate Phi overlap
10-14	2 nd candidate ROI
15-16	2 nd candidate reserved
17-18	2 nd candidate Phi overlap
19-21	1 st candidate p_T
22-24	2 nd candidate p_T
25-26	SL reserved
27-29	BCID
30-31	MUCTPI reserved

Table 9: The current Barrel Interface Board data format

transfer, 64-bit data will be transmitted at 40 MHz to the MUCTPI, using 8b/10b encoding to ensure the DC balance of the link. This will be implemented using a medium-cost FPGA with an internal 6.4 Gb/s serialiser to handle the required data rate, and a commercial 10 Gigabit Ethernet pluggable optical transceiver (standard SFP+) to provide the output. The FPGA logic will be synchronous with the LHC clock, which will be provided either via the RODbus backplane or as an external input coming from the Level-1 barrel TTC VMEbus modules. Thirty-two additional input/output signals will be implemented on the IB front panel for future functional expansion. They could be used to access the backplane 14-bit control signals (routing them to the front panel via the FPGA), to provide optional trigger inputs/outputs, and for testing and monitoring both the IB and the SL. The FPGA will also manage the backplane and front-panel signals, the VMEbus protocol, and the serialisation of the SL trigger data. It will also be used to implement additional features, not foreseen in the baseline IB project, such as the generation of trigger test patterns, monitoring logic, and trigger rate control. Subject to future physics requirements, the FPGA will be used to add one more trigger candidate to the data sent to the MUCTPI.

The additional Level-1 trigger latency needed to sample the SL inputs, serialise the data and transmit them from the optical transceiver is expected to be 3–4 BCs. This additional

latency is documented and tracked in the overall Phase-I Level-1 latency envelope [4.8]. A dedicated software tool will be developed to interact with the IB from the local VMEbus single board computer. This tool will load the FPGA firmware from the Level-1 muon barrel configuration database and control and monitor the board. The tool will be integrated within the Level-1 barrel DAQ framework.

4.2.3 Project planning

The production of the 64 IBs is foreseen for the years 2017–2018, so as to complete installation and commissioning by the end of 2018. A demonstrator, which will allow the assessment of the IB logic and its I/O interfaces, will be designed and implemented during the years 2013–2014. Then a first IB prototype will be built in 2015 and fully tested by the end of 2016. The software required for testing the prototype and the final software to be integrated in the ATLAS DAQ will be developed during the years 2015–2018.

4.3 Level-1 Tile Muon Trigger

The main source of trigger background in the Level-1 muon endcap region are low-momentum protons emerging from the endcap toroid magnets and beam shielding [4.1]. Figure 44 shows the distribution of Level-1 muons as a function of η above an online p_T threshold of 20 GeV (referred to as L1_MU20 in the following). To improve the rejection of fake triggers in the η range $1.0 < |\eta| < 1.3$, i.e. the region of the BW-TGC not covered by the NSW, energy loss in the outermost layer of the Tile hadronic calorimeter (TileCal) [4.9] extended barrel (D-layer cells D5 and D6) will be used in combination with the BW-TGC.




Figure 44: Distribution of Level-1 muons as a function of η and for a p_T threshold of 20 GeV. The combined yellow (shaded) and white (un-shaded) distribution is with the NSW included in the Level-1 muon endcap decision. The yellow (shaded) distribution alone shows the effect of including the outer layer of the TileCal extended barrel in the Level-1 muon endcap decision. (The underlying cyan (dark-shaded) distribution represents offline reconstructed muons after an offline 20 GeV p_T cut.)

The TileCal provides highly segmented energy measurements of incident particles. As shown in the inset of Figure 35, the η range $1.0 < |\eta| < 1.3$ is covered by the TileCal cells D5 and D6, which have a cell granularity of approximately 0.2×0.1 ($\eta \times \phi$).

4.3.1 Tile hadronic calorimeter muon identification

The TileCal signals have two read-out paths, one in which trigger tower energy signals are transmitted to the Level-1 trigger and the other being the standard ATLAS read-out path, via read-out drivers to the data acquisition system. A prototype Level-1 trigger receiver module for the TileCal D-layer muon signals has been developed [4.10] and used to receive data for a small number of D-layer trigger cells during ATLAS running in 2010 and 2011. These data have been used to establish the signal-to-noise ratio (SNR) between the most probable value of the estimated energy distribution of the muon signals and noise, defined as the RMS of energy depositions in cells, using zero-bias data samples.

Figure 45a shows the measured SNR values in cells D5+D6 as a function of the η position of the muons. The variation in the SNR value is due to the different path lengths of muons through these cells (see Figure 35). For the Level-1 read-out path, SNR values of ~ 6 have been measured in the η range of interest. The electronic noise of this read-out path is measured to be ~ 200 MeV and is the largest contribution to the RMS. Also shown in Figure 45a is the SNR obtained by using the standard ATLAS read-out path for analogue signals digitised on the detector. SNR values of ~ 30 have been measured with a total noise of ~ 45 MeV at a μ of 20.




Figure 45: (a) Signal-to-Noise ratios of muon signals of cells D5 and D6 as a function of muon track η , for the standard read-out path (black full circle) and for the Level-1 read-out path (red open circle) as measured during 2011 operation. (b) The noise (see text) in cells D5 and D6 versus μ using the standard read-out path.

Figure 45b shows the measured noise in the standard read-out path, in cells D5 and D6 up to a μ value of 20 using 2012 data. Also shown are the noise values obtained from simulation for larger values of μ . It can be seen that the expected contribution of pile-up to the noise at a μ value of 80 is ~ 105 MeV.

4.3.2 Muon efficiency and fake rate reduction

Each extended barrel in TileCal consists of 64 modules in ϕ , while the Level-1 muon trigger in the endcap region is divided into 48 trigger sectors. For a signal in a muon trigger sector it can be expected that a high- p_T muon will also have traversed one of the two TileCal modules in front of the Level-1 muon endcap trigger sector.

An analysis has been performed in the region $1.0 < |\eta| < 1.3$, on the data of 2012, requiring that for each L1_MU20 at least one of the modules, mapped in ϕ to the associated trigger sector, has a summed energy deposit in the D5 and D6 cells greater than a pre-determined energy threshold.

Figure 46a shows a comparison of the reconstruction efficiency for combined muons, with an offline p_T greater than 25 GeV (open red circles), and for events triggered by L1_MU20 (open blue squares), as a function of the TileCal cell energy threshold. In this analysis, the analogue signals are digitised on the front-end electronics. To emulate the energy resolution measured in the Level-1 trigger read-out path, the energy of each D5 and D6 cell has been smeared by 200 MeV (1σ).

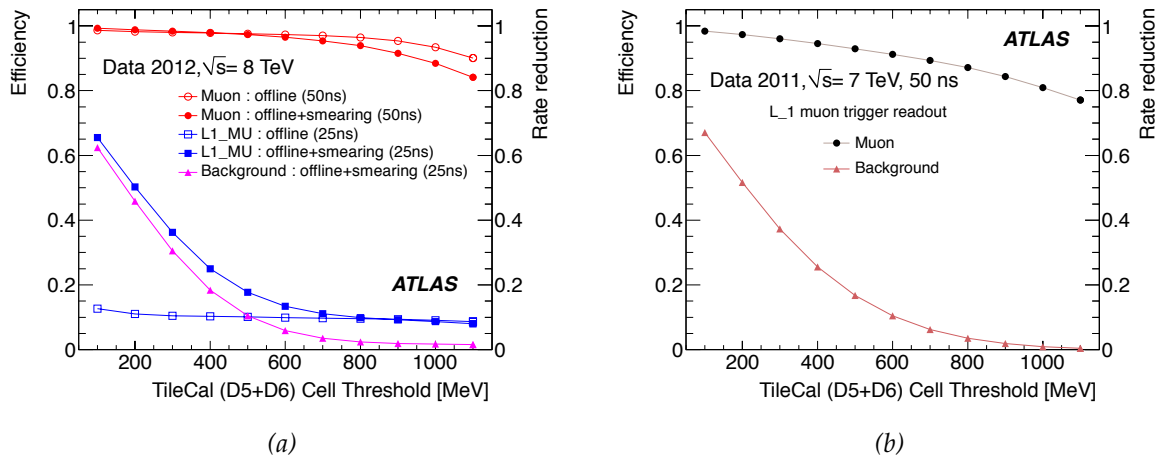


Figure 46: (a) Muon detection efficiency and L1_MU20 (see the text) rate reduction as a function of TileCal cell energy sum threshold. Rate reduction is also shown, the probability that in each bunch crossing the sum of the energies deposited in D5 and D6 will be over threshold, given that there is no offline combined muon. Results were obtained using standard offline read-out data. A smearing of 200 MeV was introduced in the response of each cell to simulate the electronics noise of the Level-1 read-out. (b) Muon detection efficiency and background reduction using a prototype receiver module connected to the Level-1 calorimeter trigger electronics. In 2012 data, 50 ns runs are used to collect enough good muon tracks, 25 ns runs are taken to calculate the fake trigger rate with the slow particles coming from the previous bunch.

The efficiencies obtained with this correction to the data are shown by the full red circles in Figure 46a. Smearing reduces the muon efficiency above a threshold of 500 MeV, and the rejection power below a similar threshold (full blue squares). For a TileCal cell energy threshold of 500 MeV, the efficiency for signal muons is 97%, while the L1_MU20 trigger rate is reduced by 82%. In the same figure is shown (by the solid triangles) the probability that the combined energy deposited in the D5 and D6 cells is above threshold when there is no (offline) combined muon.

These studies assume that unambiguous bunch crossing identification and energy calibration will be performed in the digital filter before summation, similarly to that performed in the Level-1 calorimeter trigger.

Additional studies have been performed that measure the muon efficiency, using the prototype receiver module used in 2010-2011 data taking, and connecting the TileCal D-layer signals to the Level-1 calorimeter trigger [4.11]. In this case, the real D-layer cell noise of ~ 200 MeV was present. As shown in Figure 46b, using the D5+D6 cell information and a

threshold cut of 500 MeV, a muon detection efficiency of 93% is achieved. For the same energy threshold, this analysis confirms the 17% of fake muons found by the Level-1 muon endcap trigger. The measured efficiency is lower than that presented in Figure 46a, which is probably due to the selection being made on the sum of the raw D5 and D6 analogue signals rather than on the sum of individually calibrated cell energies.

4.3.3 System design

To provide the best possible matching in ϕ between the TileCal extended barrel modules and the Level-1 muon endcap sector logic, the receiver board is required to process the D5 and D6 signals from eight TileCal modules and interface with three Level-1 muon endcap sector logic blocks. Covering the entire detector therefore requires 16 Tile Muon Digitiser Boards (TMDB).

The TMDB will be a 9U VMEbus slave module. A block diagram (with the main functional blocks) is shown in Figure 47. The main functions of the TMDB are:

- (a) receive and digitise the TileCal Muon signals (D5 and D6) from eight modules (Analogue Stage);
- (b) provide the correct calibration for each signal (VME FPGA);
- (c) perform signal detection for each cell (Core FPGA);
- (d) provide the BCID information for each cell (Core FPGA using timing information from TTCrx);
- (e) transmit the η , ϕ , and bunch number from the detected cells to the three muon sector logic blocks (GLink);
- (f) connect to neighbour receiver boards to accommodate the non-perfect matching between the eight TileCal modules and the three muon sector logic blocks (LVDS);
- (g) provide ROD data fragments to the DAQ system (ROD).

The 16 boards will be housed in two VMEbus crates in USA15, and the VMEbus will be used for configuration, control, monitoring and the distribution of the LHC clock. TTC signals are received optically by a TTC module in the VMEbus crate and re-transmitted to each TMDB via the VMEbus P0 connector.

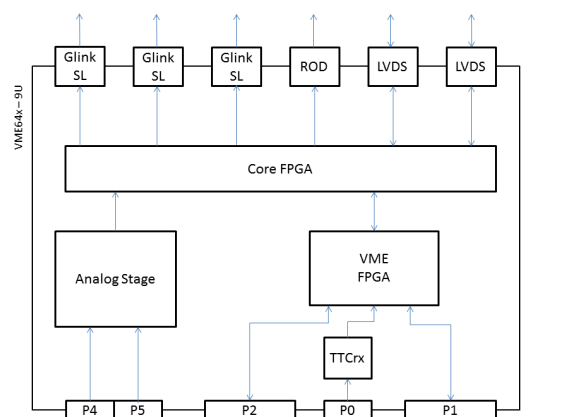


Figure 47: Block diagram of the Tile Muon Digitiser Board

A Merger Module will be designed to receive signals sent to the Sector Logic coming from EI/FI chambers. This module will concentrate signals coming from 4 optical fibres into 3

optical outputs. The merging of these signals will allow to reserve an optical input for the Tile Receiver boards, not available otherwise. This merger module will not be part of the Tile trigger data path, and will therefore not add latency to this subsystem.

4.3.4 Project planning and milestones

As shown in Figure 48, the project schedule aims to deploy the new system before physics data taking in 2015. The project requires the development of new simulation and DAQ software, and new firmware for the Level-1 muon endcap sector logic. The timeline has been divided into three periods, corresponding to the development of the prototype (Q3-Q4 2013), module-0 design and test (Q1-Q2 2014), and full production of the 16 TMDB boards (Q3-Q4 2014). Reviews will take place in each period, following the general ATLAS procedures.

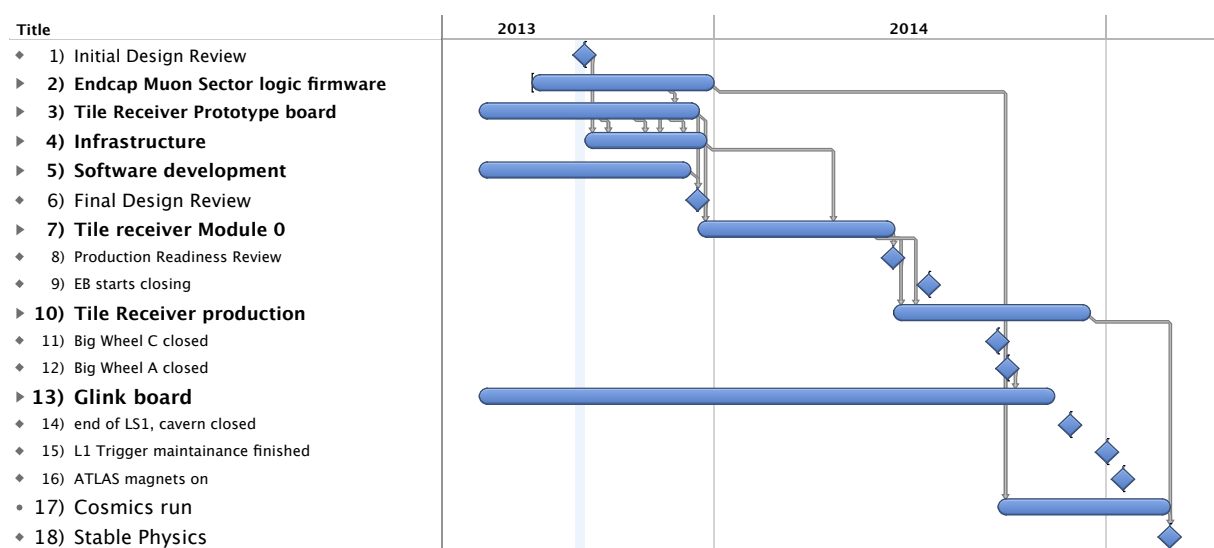


Figure 48: Schedule of the Tile Muon upgrade

References

- [4.1] ATLAS Collaboration, *New Small Wheel Technical Design Report*, Tech. Rep. CERN-LHCC-2013-006. ATLAS-TDR-020, CERN, Geneva, Jun, 2013. <https://cds.cern.ch/record/1552862>.
- [4.2] ATLAS Collaboration, *ATLAS First Level Trigger: Technical Design Report*, Tech. Rep. CERN-LHCC-98-14, ATLAS-TDR-12, CERN, Geneva, Jun, 1998. <https://cds.cern.ch/record/381429>.
- [4.3] ATLAS Collaboration, *ATLAS Muon Endcap Level-1 Trigger: Update of Technical Design Report*, tech. rep., CERN, Geneva, Jun, 2000. https://twiki.cern.ch/twiki/pub/Main/TgcDocument/MuonEndcap_rev01.pdf.
- [4.4] P. Farthouat, *Interfaces and overlaps in the level-1 muon trigger system*, Tech. Rep. ATL-DA-EN-0001 v.5, EDMS Id 114604 v.5, CERN, Geneva, Jan, 2006. <https://edms.cern.ch/document/114604>.

- [4.5] ATLAS Collaboration, *ATLAS High-Level Trigger, Data Acquisition and Controls: Technical Design Report*, Tech. Rep. CERN-LHCC-2003-022, ATLAS-TDR-016, CERN, Geneva, Jun, 2003. <https://cds.cern.ch/record/616089>.
- [4.6] F. Anulli, G. Ciapetti, D. De Pedis, A. Di Girolamo, C. Luci, et al., *The Level-1 Trigger Muon Barrel System of the ATLAS experiment at CERN*, *JINST* **4** (2009) P04010.
- [4.7] M. Della Pietra, A. Aloisio, F. Cevenini, and V. Izzo, *The Read-out Driver for the RPC of the ATLAS Muon Spectrometer*, *IEEE Trans.Nucl.Sci.* **55** (2008) 265–271.
- [4.8] N. Gee, *Phase-I Latency Envelopes for the Level-1 Trigger*, Tech. Rep. ATL-DA-ES-0059 v.1, EDMS Id 1256858 v.1, CERN, Geneva, Dec, 2012. <https://edms.cern.ch/document/1256858>.
- [4.9] ATLAS Collaboration, *ATLAS Tile Calorimeter: Technical Design Report*, Tech. Rep. CERN-LHCC-96-042; ATLAS-TDR-3, CERN, Geneva, Dec, 1996. <https://cds.cern.ch/record/331062>.
- [4.10] Ciodaro, T and others, *Muon Detection Based on a Hadronic Calorimeter*, Tech. Rep. ATL-DAQ-PROC-2010-050, CERN, Geneva, Nov, 2010. <https://cds.cern.ch/record/1308440>.
- [4.11] T. Ciodaro, J. de Seixas, and A. Cerqueira, *Use of Hadronic Calorimetry Information in the ATLAS Level-1 Muon Trigger*, Tech. Rep. ATL-TILECAL-PROC-2013-006, CERN, Geneva, Jun, 2013. <https://cds.cern.ch/record/1553306>.

5 Level-1 Central Trigger System

The Level-1 Central Trigger System consists of the Central Trigger Processor (CTP), the Muon-to-CTP-Interface (MUCTPI), and the Level-1 Topological Processor (L1Topo). While L1Topo is a completely new system that did not exist during the Run 1 phase of the LHC, the original system of CTP and MUCTPI was developed more than 10 years ago and has been successfully operated during the Run 1 phase of the LHC from 2008 to 2013. Its upgrade, whose motivation is given below, retains the basic concepts. The aim is to re-use, wherever possible, the modules from the original system. This is possible to a large extent for the CTP upgrade during the Long Shutdown 1 (LS1). The new CTP will provide more trigger inputs, making room for new trigger sources (e.g. the topological processor), as well as more trigger menu items. For the MUCTPI Phase-I upgrade, however, a new system needs to be developed. The upgrades benefit greatly from new FPGA and optical link technologies which allow more flexibility and better performance. The new MUCTPI will use optical links to receive the sector logic data, and will be able to send detailed Region-of-Interest (ROI) muon information to the topological processor at 40 MHz. All these upgrades allow for considerably more flexibility for physics selections at very high luminosities.

5.1 Central Trigger Processor (CTP)

The Central Trigger Processor (CTP) is the last stage of processing of the Level-1 trigger system. It receives digital trigger signals from L1Calo, L1Topo, the MUCTPI, and various forward detectors and aligns them in time. It further applies the trigger logic configured according to the physics trigger menu, allows for prescaling, and is in charge of managing deadtime.

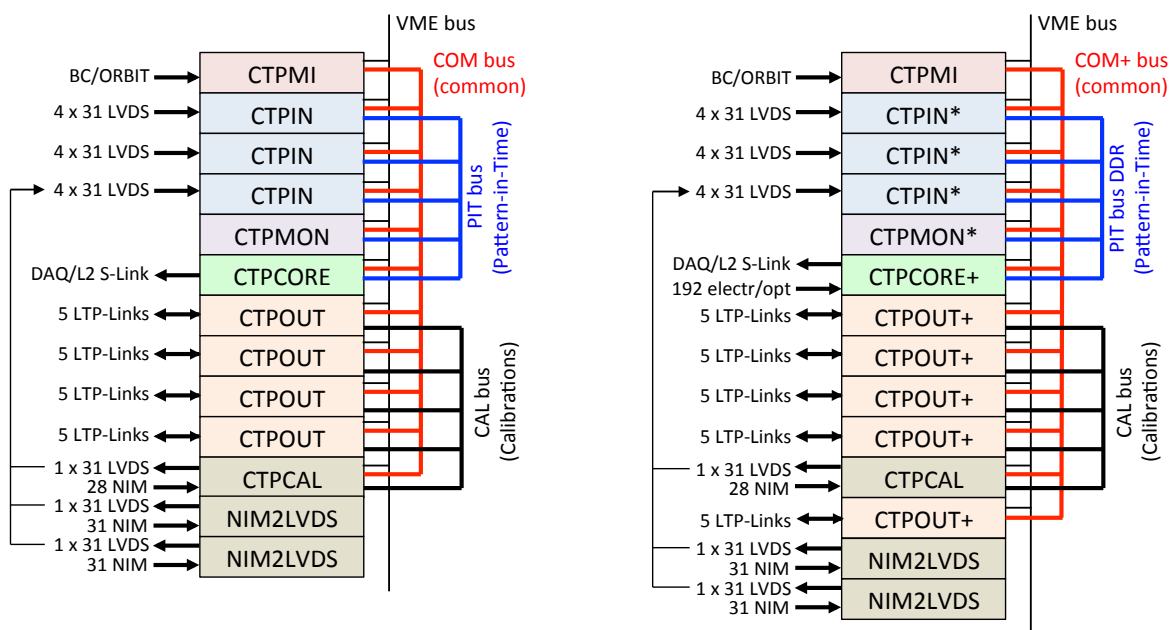


Figure 49: Block diagram of the various components of the CTP. The left picture shows the setup during Run 1, the right picture reflects the changes made during the LS1. Components marked with “+” are newly built, while components marked with “*” use the existing hardware, but with upgraded firmware.

As can be seen in Figure 49 (left), the existing CTP system, implemented as custom-built electronics boards housed in a 9U VMEbus crate, consists of 13 VMEbus boards of 7 different

types, and of 3 custom backplanes. The CTP Machine Interface board (CTPMI) receives the timing signals from the LHC and distributes them to the common bus (COM) on the backplane. The COM backplane also carries the wired-OR busy signals from the sub-detector partitions. The CTP Input Module (CTPIN), of which 3 instances are used, receives trigger input signals from the trigger detectors, synchronises and aligns them, and maps them onto the so-called pattern-in-time (PIT) bus on the backplane. The CTP Monitoring Module (CTPMON) offers per-bunch monitoring of the signals on the PIT bus. The CTP Core Module (CTPCORE) applies a programmable logic to the trigger signals, according to the trigger menu, and forms the Level-1 accept signal that is used to trigger the read-out of all sub-detectors. In addition, the CTPCORE is in charge of sending summary information to the Level-2 trigger and to the data acquisition system. The CTP Output Module (CTPOUT), of which 4 instances are used, sends the trigger, timing and controls signals to the sub-detectors. It can also receive sub-detector calibration request signals and make them available on the so-called calibration bus (CAL) on the backplane. The CTP Calibration Module (CTPCAL) time-multiplexes these calibration requests and makes them available on its front panel, from where they are fed into one of the CTPIN modules. The NIM2LVDS Module collects up to 31 NIM signals from various detectors and funnels them into one LVDS link that can be fed into one of the CTPIN modules.

5.1.1 Motivation

In 2012, the CTP was working near the limits of its capacity, as can be seen from Table 10, which shows the CTP resources used by a recent trigger menu. All of the PIT bus lines, which limit the number of usable trigger inputs, and almost all of the trigger items have been used during Run 1 of the LHC. The main motivation of the upgrade during LS1 is to remove those limitations and increase the number of trigger inputs and the number of trigger items.

In the framework of the trigger upgrade, a new topological processor (L1Topo) will be added to the system, improving the multi-object selection for the increased luminosity (see Section 5.3). L1Topo will receive trigger input from the calorimeter trigger and muon region-of-interest trigger information from the MUCTPI. The CTP will have to be capable of receiving the additional trigger inputs from L1Topo and provide additional trigger items for these inputs. At the same time, the calorimeter trigger merger modules, which send trigger input signals to the CTP, will be replaced with new ones to provide the necessary input to the topological processor. The new L1Calo merger modules (CMX) will also send more trigger signals to the CTP.

The upgrade of the CTP will provide partitioning of the L1A generation for detector commissioning, an improved bunch group masking and bunch-by-bunch trigger item monitoring, and more outputs for sub-detector TTC partitions. In order to reduce the latency for some trigger signals, direct electrical input from the topological processor to the CTPCORE is also foreseen. In addition, optical trigger inputs are available to connect new or upgraded systems, if the overall latency allows doing so.

The CTP upgrade requires a complete re-design of several modules, namely the CTPCORE, the CTPOUT, and the COM backplane. Figure 49 (right) shows the general architecture. The updated version of the modules will in the following be denoted with the suffix "+". In addition, the firmware of the CTPIN and CTPMON modules needs to be modified. It is planned to install and commission the upgraded CTP in USA15 in 2014. The next major

Resource	Used in Run 1	Available in Run 1	Available after LS1
CTPIN input cables (partially used)	9	12	12
CTPIN input signals	212	372	372
CTPIN integrating monitoring counters	138	768	768
PIT signals	160	160	320
CTPCORE trigger items	241	256	512
CTPCORE bunch group masks	8	8	16
CTPCORE max. number of AND terms	6	256	512
CTPCORE max. number of bits in OR terms	6	12	15
CTPCORE per-bunch trigger item counters	12	12	256
CTPOUT cables to TTC partition	20	20	25
CTPMON per-bunch monitoring counters	88	160	160

Table 10: CTP resources used in Run 1, available in Run 1, and available after LS1

upgrade of the CTP is then only foreseen for the Phase-II upgrade, when the Level-1 trigger architecture will change completely and the latency budget is likely to increase.

5.1.2 Trigger inputs

This section addresses the two paths for the CTP to receive trigger input signals: via cables through the CTPIN, and via direct CTPCORE inputs.

CTPIN inputs The CTP can currently only use a subset of 160 of the 372 possible trigger inputs in any given trigger menu configuration. This limitation is due to the custom PIT bus backplane which only connects 160 signals from the input modules (CTPIN) to the core module (CTPCORE) that generates the Level-1 accept (L1A).

It has been demonstrated that via a change in the firmware of the CTPIN, it is possible to increase the number of trigger inputs by doubling the PIT bus transfer rate, i.e. by operating it at 80 MHz using double-data rate (DDR) signalling [5.1]. In this way, the maximum number of trigger inputs to the CTPCORE+ will be increased to 320.

The CTPIN has four identical channels, which receive 31 LVDS trigger input signals at 40 MHz each. After level conversion, an FPGA synchronises the trigger inputs to the internal clock, aligns them with respect to each other using programmable length pipelines, and optionally checks their parity. Finally, a configurable switching matrix implemented in a CPLD is used to select and route the aligned trigger inputs to the PIT bus. A second FPGA is used to monitor the trigger inputs with counters that integrate over all bunches in an LHC turn.

The modified firmware of the synchronisation and alignment FPGA features DDR output registers which drive the 31 LVDS trigger signals onto 16 DDR lines. The output signals of each channel are then sent to the switch matrix CPLD, which selects and routes up to 64 signals from the CTPIN onto any of the 160 PIT bus lines. This implies that there will be a reduction of flexibility in assigning trigger inputs to the LUTs on the CTPCORE+ because the trigger input signals will need to be allocated in pairs to PIT bits by the switch matrix. This reduction is offset by the significant increase in the number of usable trigger inputs.

The implementation of the DDR multiplexing on the CTPIN and de-multiplexing on the CTPCORE+ incurs a latency penalty which is expected to be 2 BC. It is important to note

that the DDR signalling is only internal to the CTP, while the input cables will continue to operate at 40 MHz. Cables that are currently not in use, partially used cables or the additional electrical and optical trigger inputs of the CTPCORE+ make it possible to input additional trigger signals.

Direct CTPCORE inputs The CTPCORE+ features 96 direct electrical inputs from 3 connectors on the front-panel of the module, which are meant for receiving trigger inputs from the topological trigger processor. These inputs are intended for the most latency critical trigger signal path, since they avoid the delay penalty for multiplexing and transmission over the PIT backplane. LVDS signal levels and SCSI VHDCI connectors and cables are used for this interface. Three connectors with 32 differential pairs each are foreseen, enabling the CTP to be connected to up to three L1Topo modules if required. The signals can be operated at 40 MHz (the BC rate), or be over-clocked to twice the BC frequency, possibly even 4-times the BC frequency. The following assumes operation at 80 MHz, giving 192 trigger input signals.

The CTPCORE+ also implements an option to connect trigger inputs optically through 12 serial links, to be used in the Phase-I upgrade. A 12-way ribbon fibre receiver module (Avago MiniPOD™ [5.2]) is foreseen for this purpose. These links will be operated at 6.4 Gb/s with 8b/10b encoding synchronously with the BC frequency. This will enable receiving up to 128 trigger bits per BC on each link, or 1536 trigger inputs in total. It is important to note that depending on the number of trigger bits and the projected use in the trigger menu, complex selection logic may be required in the CTPCORE+ which will have an impact on the overall latency. There is also a latency penalty of about 3 BCs associated with the serialisation and serialisation in the on-chip gigabit transceivers that needs to be considered when using the optical inputs.

5.1.3 Trigger formation

In order to effectively use the additional trigger inputs, the number of trigger items available in the CTPCORE needs to be increased. This number is currently limited to 256 and almost all trigger items have been used in trigger menus during Run 1 operation. The requirement of more trigger items cannot be accommodated by a simple firmware modification and therefore requires the new CTPCORE+ module. The trigger path implemented in the CTPCORE+ module is illustrated in Figure 50.

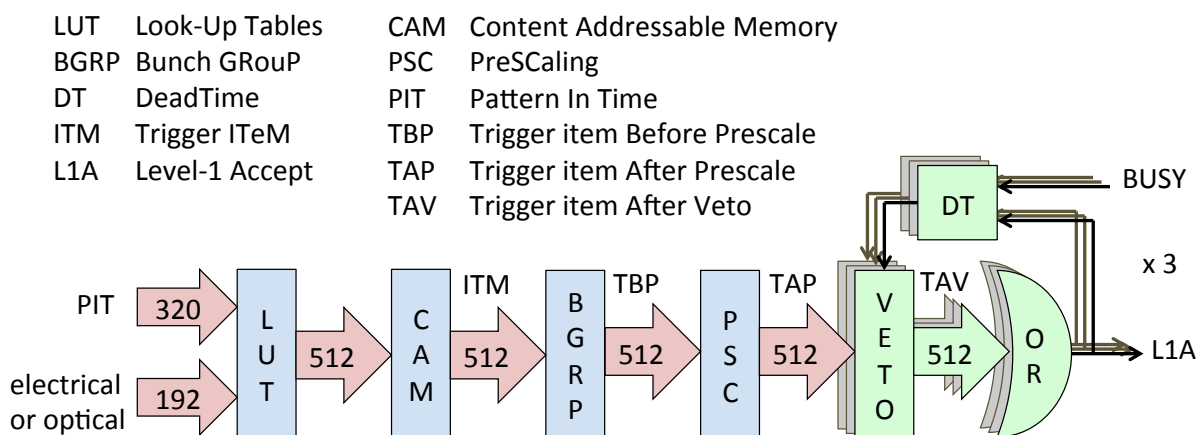


Figure 50: Block diagram of the CTPCORE+ architecture

A set of DDR input registers latches and de-multiplexes the 160 trigger inputs received via the PIT backplane. These 320 trigger signals, together with 192 direct trigger signals, are then combined in an array of Look-Up Tables (LUTs) of up to 15 bits, and large ternary Content-Addressable Memories (CAM) to form 512 trigger items (ITM). This scheme allows for maximum flexibility in the logical combinations of the trigger signals.

In addition to the trigger inputs, there are the internal triggers and bunch groups. In the CTPCORE+, the number of internal triggers (two prescaled clocks and two random triggers in the CTPCORE) is doubled and the granularity of the random trigger frequency is increased. Concerning the bunch groups, the number of 8 bunch groups in the CTPCORE is increased to 16. They are implemented as 16 programmable bunch group masks (BGRP) which are applied to the items at the output of the CAM, indicated as “trigger items before prescale” (TBP) in Figure 50. This allows monitoring of the trigger items before the bunch group masking.

Deadtime generation and busy veto remain unchanged with respect to the existing design, while prescaling was changed from counter-based to random prescaling. There are three independent instances of the deadtime generation and busy veto logic, one per trigger partition.

5.1.4 Outputs to sub-detector partitions

The CTPOUT module will be replaced by the new CTPOUT+ module, which features improved busy monitoring and a built-in pattern generator for testing the TTC distribution network. The number of outputs to the sub-detector TTC partitions will be increased from 20 to 25 by extending the COM+ backplane to provide an additional slot for an additional CTPOUT+ module. The backplane for routing the calibration requests will, however, not be changed, and the additional 5 outputs will therefore not be able to send calibration requests. This is not considered a significant limitation, since it is possible to use one of the other CTPOUT+ slots if calibration request signalling is needed. Only one sub-detector has adopted the calibration request scheme so far.

The CTP upgrade also introduces two additional trigger partitions which allow splitting the L1A generation for detector commissioning and calibration runs. The granularity of the partitions is defined by the number of CTPOUT output cables to the sub-detector TTC partitions, i.e. each output can be part of one of the three trigger partitions. Although the scheme of partitioning the L1A generation in the CTP adds a lot of flexibility, there are some restrictions in the use of the trigger partitions: The first constraint is that all three partitions share the BC and ORBIT signals coming from the CTPMI module, as well as a common trigger menu. Another limitation is that if the same trigger item is required in more than one partition with different prescale factors or bunch group masks, then it needs to be defined twice. Finally, the two secondary partitions do not implement the full flexibility of the trigger type generation: they will only allow the selection of one of 16 pre-defined trigger type values.

Each of the three partitions has its own busy veto logic, deadtime generation, logical OR of the trigger-after-veto (TAV) bits and trigger type logic as shown in Figure 50. All three partitions share the LUT, CAM and prescaler resources. A programmable mask in the veto logic defines which trigger items are included in the L1A generation for a given partition. Only the primary partition sends data to the DAQ system and the Level-2 trigger.

5.1.5 Monitoring

In addition to the existing monitoring capabilities of the CTPCORE, the CTPCORE+ features 256 per-bunch monitoring counters, where each counter builds a histogram of a selected trig-

ger bit. A programmable selection will be implemented to choose which of the approximately 1800 PIT, ITM, TAP and TAV bits to monitor.

Per-bunch monitoring of the PIT bus signals is provided by the CTPMON module, which decodes and selects up to 160 PIT bus signals and builds a histogram with 3564 entries for each of them in order to monitor them on a bunch-by-bunch basis. The firmware of the CTPMON was modified in a proof-of-concept implementation to include DDR input registers, but the memory resources of the histogramming FPGAs were not sufficient to increase the number of signals being monitored to the full number of 320, since this functionality requires a large number of on-chip memory blocks. Instead, a simple mechanism is implemented which allows to select 160 of the 320 trigger inputs for per-bunch monitoring. This is not believed to be a serious limitation, given that in Run 1 only about 50% of the per-bunch monitoring counters in the CTPMON were used.

In the CTPOUT+ module, the monitoring of the busy signal from each of the sub-detector partitions will be improved to allow per-bunch contributions to be measured. This will for instance enable factoring out the busy contributions of some sub-detectors in the LHC abort gap.

5.1.6 Latency

The latency of the CTP used for Run 1 was 5 BC, from the input of the CTPIN to the output of the L1A signal at the CTPOUT. Two of the BCs were spent in the CTPIN and 3 BC in the CTPCORE. The new CTP has 3 paths with different latencies. The path via the CTPIN and PIT bus will have a latency of 8 BC, which is composed of 2 BC for the CTPIN, 2 BC due to the PIT backplane (de-)multiplexing, 3 BC for the CTPCORE+ logic, and 1 BC in the CTPOUT+ for the multiplexing functionality required for the additional trigger partitions. The direct electrical inputs to the CTPCORE+ will take only 5 BC: synchronisation at the input of the CTPCORE will need 1 BC, the CTPCORE+ logic 3 BC, and CTPOUT+ 1 BC. The option of using optical inputs to the CTPCORE+ will use an additional 2 BC for de-serialisation and 1 BC for a switch matrix, i.e. 8 BC of latency in total. Table 11 summarises the latency of the different paths through the CTP.

Trigger Path	Latency
CTP (Run 1)	5 BC
CTP+ via CTPIN	8 BC
CTP+ direct electrical	5 BC
CTP+ direct optical	8 BC

Table 11: Summary of the latency of the CTP used during Run 1 and expectations for the upgraded CTP

5.2 Muon-to-CTP-Interface (MUCTPI)

The ATLAS Muon to Central Trigger Processor Interface (MUCTPI) [5.3, 4] receives muon candidate information from each of the 208 muon trigger sectors, calculates the muon candidate multiplicity for each of the six transverse momentum (p_T) thresholds and sends it to the CTP. To avoid double counting [5.5], the calculation takes into account the potential overlap between candidate muon tracks. For events accepted by the CTP, the MUCTPI outputs a list of muon candidates to the DAQ system, while it sends Region-of-Interest (ROI) data

to the Level-2 trigger to seed the muon processing. In addition, it performs comprehensive monitoring.

5.2.1 Motivation

The existing MUCTPI will have to be upgraded as part of the Phase-I Level-1 trigger upgrade program [5.6]. The main motivations for this upgrade are the following: after the Phase-I upgrade, the MUCTPI is supposed to provide full-granularity muon ROI information at the bunch crossing rate of 40 MHz to the newly introduced topological trigger processor (L1Topo) [5.7]. This requirement calls for a replacement of the existing MUCTPI, because in the existing MUCTPI the full muon ROI information can only be accessed after the Level-1 Accept.

The second main motivation for replacing the MUCTPI is that it needs to be compatible with the new endcap sector logic modules, deployed as part of the New Small Wheel upgrade. The new modules will use serial optical links to transfer the muon candidate information to the MUCTPI, instead of the electrical cables used currently. The optical links will provide a much higher bandwidth which will be used to transfer additional information from the sector logic modules, for example data for more than two muon candidates, additional p_T thresholds, coincidence flags or additional bits indicating the quality of the muon track. The smaller physical dimensions of the optical connections enable increased integration of the upgraded MUCTPI, which in turn allows spare inputs to be added for possible additional muon sectors (e.g. an MDT muon track trigger at Phase-II).

5.2.2 Interfacing of the existing MUCTPI to L1Topo in Run 2

For sending the full ROI granularity at 40 MHz to L1Topo, the MUCTPI has to be redesigned for the Phase-I upgrade. For Run 2, one is limited to the existing system. The 16 octant (MIOCT) modules of the existing MUCTPI each have 2 NIM outputs. With an upgrade of the firmware of the modules the outputs can be used for transferring hit maps or trigger objects to L1Topo. Tests have shown that each output can transfer 320 Mb/s, i.e. 8 bits per LHC clock tick. The 32 outputs available in total can output 256 bits per LHC clock tick. As mentioned in Section 2.2, using muon information in L1Topo has important benefits. It is therefore desirable to provide already during Run 2 muon information to L1Topo. L1Topo does not have NIM inputs, thus, a conversion from NIM to optical serial I/O, running at 6.4 Gb/s is required. This is to be achieved with the MUCTPI-to-Topo interface. This interface, currently being implemented, has 34 NIM inputs and 8 NIM outputs and 8 optical outputs. It is based on a Xilinx VC707 development kit, equipped with two mezzanine boards. One of the mezzanine boards is a custom designed board with LEMO connectors for the NIM inputs and outputs, the other board is a commercially available board with two QSFP transceivers for output to up to 8 fibres. With 6.4 Gb/s per fibre and using 8b/10b coding, two fibres can transfer all data. With the 8 fibres available, up to 4 copies of the muon data can be sent to different parts of L1Topo. The VC707, the two mezzanine cards, and a power supply will be installed in a rack-mountable, pizza-box-like enclosure with provision for vertical air cooling. Installation and commissioning in USA15 will take place in 2014.

5.2.3 Implementation

Figure 51 shows the proposed implementation of the upgraded MUCTPI and its interfaces to other components of the Level-1 trigger system. It is based on FPGAs with a large number of on-chip high-speed serial links as well as high-density parallel fibre optics receiver and transmitter modules.

There are four FPGAs, each of which receives and processes the output from up to 60 sector logic modules, i.e. a quadrant of the muon detector, via 12-channel parallel fibre optic receiver modules. The results of this processing are then sent from each quadrant FPGA, via 12-channel parallel fibre optic transmitter modules, to the L1Topo modules.

A fifth FPGA is used to merge the results from the four quadrants and to send multiplicities for each p_T threshold to the CTP through a single optical link. The latency increase of a few BC for a serial optical link compared to a parallel electrical connection is small and well absorbed in the parallel latency critical path through L1Topo. The merger FPGA also implements the MUCTPI read-out driver functionality and sends the muon ROI information to the Level-2 trigger system for events accepted at Level-1.

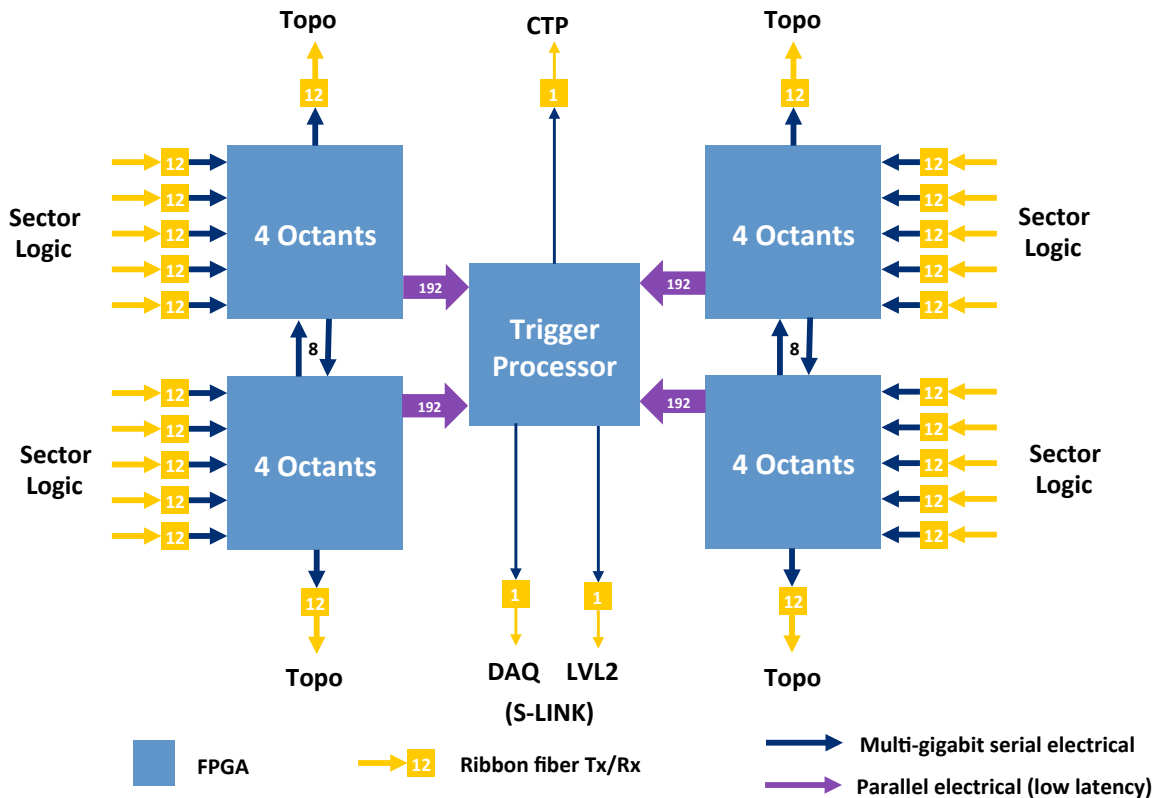


Figure 51: Block diagram of the MUCTPI architecture

The baseline implementation of the MUCTPI will be a single electronics board based on the ATCA standard. The feasibility to integrate the MUCTPI on a single electronics board still needs to be studied in more detail, and it may be necessary to partition it over two modules. The baseline implementation will use FPGAs with on-chip multi-gigabit transceivers/receivers (MGTs), for example an FPGA from the Xilinx Virtex-7 family [5.8] which has 80 on-chip MGTs and a sufficient number of gates and sufficient memory to implement the required logic. The FPGA will be chosen such that a significant fraction of its resources remain free as margin for the implementation of future functionality, for example at Phase-II. As already indicated, 12-

channel fibre-optic transmitter and receiver modules [5.2] will be used for I/O. These modules feature a small footprint and support line rates of up to 10.3 Gb/s. An additional advantage of these modules is that they can be placed mid-board, close to the FPGAs, thereby reducing the required line length of the critical high-speed signals on the PCB and easing signal integrity issues.

5.2.4 Interfaces

This section describes the external interfaces of the MUCTPI: to the sector logic modules, to L1Topo, to the TTC, and to the DAQ and Level-2 trigger systems.

Sector logic As mentioned above, 12-way ribbon fibre receiver modules (Avago MiniPOD) will be used on the MUCTPI for reasons of integration density. On the sector logic modules, standard 10 Gb/s SFP+ fibre optic transceiver modules can be used, since a single output is sufficient. A patch panel will be used to break out the 12-way ribbons (MPO/MTP connector) to individual fibres (LC connectors). The new endcap sector logic will have an optical output, however, for the barrel trigger sectors, the 64 VMEbus interface boards between the Level-1 muon barrel sector logic and the MUCTPI will be replaced with boards that use optical outputs. The optical links from the sector logic modules to the MUCTPI will operate at 6.4 Gb/s with 8b/10b encoding synchronously with the BC frequency. This allows to receive up to 128 trigger bits per BC on each link. For the barrel and new endcap sectors, 32 bits and 64 bits per BC are required respectively. The remaining bits will be used for control characters (IDLE) to maintain and check alignment and synchronisation. Operating the links faster than required also results in a latency saving compared to a slower link speed. The end-to-end latency of such a link has been estimated to be about 4 BC.

For a robust and flexible data transfer, a 64-bit data word per BC will be transmitted between the sector logic and the MUCTPI. The exact sector data format still needs to be decided upon, but the data word will include

- a bit to flag that more than 2 candidates have been found in one sector,
- at least the lower 3 bits of the BCID (possibly more bits),
- data for up to two candidates:
 - ROI number (5 bits for L1Muon barrel, 8 bits for L1Muon endcap),
 - transverse momentum p_T (3 bits for L1Muon barrel, 4 bits for L1Muon endcap),
 - overlap flags,
 - additional bits and flags for track quality, possibly for charge and for flagging more than one candidate in an ROI.

L1Topo Each quadrant of the MUCTPI has one 12-channel ribbon fibre transmitter module to connect to one or more L1Topo processor modules. The optical links to the L1Topo module(s) will operate at 6.4 Gb/s synchronously with the LHC clock and use 8b/10b encoding. The exact information content and data format of the optical links to L1Topo are yet to be defined.

TTC The MUCTPI will also have to interface to the TTC system. Initially an electrical connection as with the current MUCTPI can be used, however, an optical TTC input is also foreseen, most likely on a mezzanine card, in order to be able to adapt to the new TTC distribution system foreseen for the Phase-II upgrade.

DAQ and Level-2 Trigger The Read-Out Links connecting the MUCTPI to the DAQ and Level-2 trigger are based on the 2 Gb/s HOLA S-LINK protocol for compatibility with the existing system. The implementation will use the FPGA on-chip serialiser/serialiser together with a pluggable SFP+ fibre optic transceiver. This approach also enables higher link speeds (up to 10 Gb/s) and/or different protocols to be implemented, e.g. for Phase-II.

5.2.5 Latency

An estimate of the latency for the path from the sector logic through the upgraded MUCTPI to the L1Topo processor is shown in Table 12. There is also a real-time data path from the sector logic through the MUCTPI directly to the CTP; however, this path is not latency critical, since it bypasses L1Topo and the associated delay. The latency penalty associated with the serialisation and deserialisation in the on-chip gigabit transceivers is estimated to be 4 BCs. The serial link latency has been split equally between the transmit and receive ends of the link. It is assumed that the latency for a given subsystem includes the delay for de-serialisation at the input, the serialisation at the output, as well as the delay for the cable connecting to the receiving system.

Description	Latency (BC)
Sector logic serial link input	2
Synchronise to local clock	0.5
Processing	3
L1Topo serial link output	2
Cable to L1Topo (2 m)	0.4
Total	7.9

Table 12: Breakdown of the expected latency in the MUCTPI, for the critical path via L1Topo

Despite the additional processing requirements and the adoption of higher-latency serial optical inputs and outputs, the estimate indicates that the latency of the upgraded MUCTPI system can be kept within the time of eight BCs originally allocated to this subsystem in the overall latency budget for the Level-1 trigger.

5.2.6 Software

For the upgrade of the MUCTPI, new software will be developed, which includes low-level software for interfacing with the MUCTPI hardware, tools for diagnostics and tests of the MUCTPI hardware, online software for configuration, control, and monitoring within the ATLAS run control framework, as well as offline software for simulation, reconstruction, and event monitoring. The exact model of communication still needs to be defined. Since the MUCTPI will make use of the ATCA standard, the previous model of a local area network connection to a single-board computer in the VMEbus crate using a library and a driver to communicate over VMEbus with the hardware modules will have to be replaced. Local area network connection to a work station, a dedicated Ethernet link and the use of IPbus to communicate with the new hardware modules is envisaged as a replacement. In order to keep the effort reasonable and not to develop a private solution, the MUCTPI will make use of a solution provided by ATLAS in the context of the VMEbus replacement working group.

5.2.7 Project planning

A tentative schedule for the Phase-I MUCTPI upgrade project is shown in Figure 52. The project has been split into a number of activities, namely:

- Hardware module development, production and test;
- Firmware development and test;
- Online software development;
- Offline software development;
- Installation and commissioning.

In addition, it may also be necessary to develop a dedicated test module to drive the optical sector-logic inputs of the MUCTPI. The MUCTPI system will be integrated and extensively tested in the lab. This includes operation with realistic inputs and systematic checking of the outputs against the results of the low-level simulation. It is also planned to perform interoperability tests with prototypes of modules from the other subsystems well before the final installation as far as possible in order to check the compatibility and data formats. This concerns in particular L1Topo, the New Small Wheel sector logic, and the muon barrel converter modules.

A preliminary design review (PDR) of the upgraded MUCTPI is foreseen in Q1 2016, a final design and production readiness review (FDR/PRR) in Q1 2017. The upgraded MUCTPI will be installed in USA15 during the Long Shutdown 2 (LS2) and should be ready for integration with other subsystems several months before the end of the shutdown.




Figure 52: Planned schedule of the MUCTPI

5.3 Level-1 Topological Processor

The Topological Processor for the Level-1 trigger (L1Topo) performs real-time event selection based on the geometrical relationship between trigger objects such as muons, jets, E_T^{miss} , electron and tau clusters, and selections based on the variables H_T , M_{eff} , or M_{inv} . A detailed discussion of these selections has been presented in Section 2.2.6. In this section, these objects and variables are referred to as Trigger Objects (TOBs).

The baseline implementation of the L1Topo system is a single ATCA shelf equipped with two or more L1Topo processor modules. They receive data from the Level-1 muon and calo-

rimeter trigger systems via optical fibres. The real-time output from L1Topo is sent to the CTP (see Section 5.1.2). Figure 53 shows the layout of the L1Topo processor module. Event selection is performed by two FPGAs (labelled U1 and U2 in the figure), receiving their input via optical receivers. Module control and read-out functionality are implemented in the separate Control FPGA.

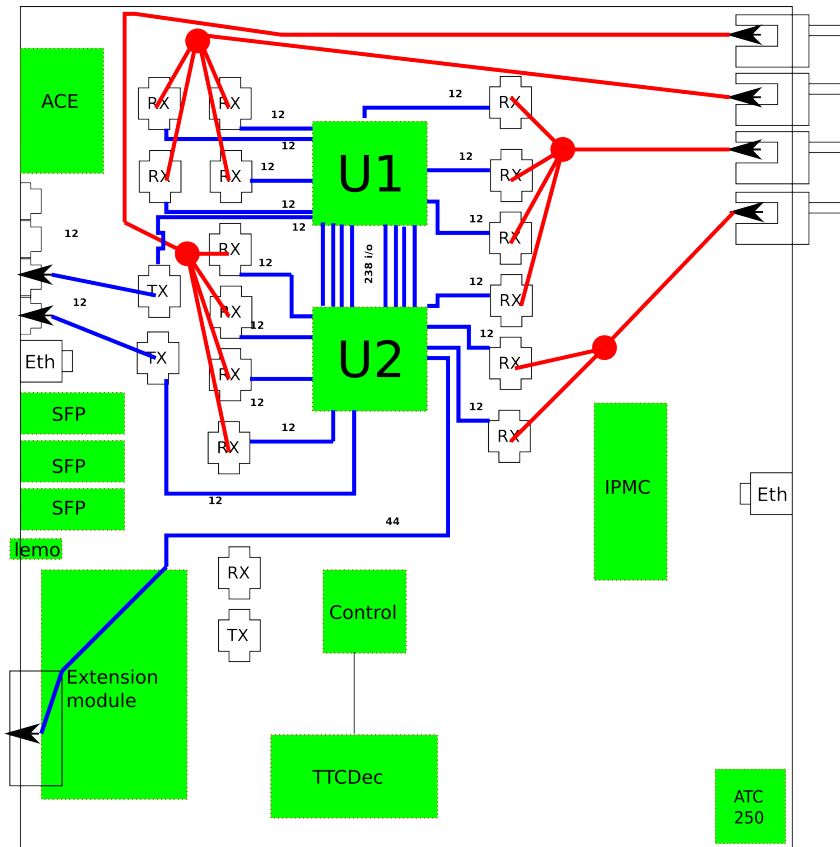


Figure 53: Layout of the L1Topo module, showing the main functional blocks and real-time data path lines

Input and real-time processing The input to an L1Topo module from the Level-1 muon and calorimeter trigger systems is via four 48-way fibre ribbon bundles into four MTP-CPI connectors. From these connectors, the optical signals are routed via “octopus” cable sets splitting the 48-way ribbon fibres into four 12-way ribbon fibres. The latter are routed to the twelve Avago MiniPOD optical receivers (labelled “Rx” in the figure) that perform opto-electrical conversion. The electrical signals are routed into the processing FPGAs via their on-chip multi-gigabit transceivers (MGT) and de-serialised at the LHC clock frequency. The processing FPGAs are Xilinx Virtex-7 devices with 80 MGTs each. In Run 2, the input data to the processors is duplicated at source, and both processor FPGAs are supplied with the same data and operate independently and in parallel. However, a low-latency, parallel real-time communication path of 238 Gb/s aggregate bandwidth is available between these devices. Using the baseline link speed of 6.4 Gb/s with standard 8b/10b encoding provides ~ 5.1 Gb/s of payload for each of the 160 inputs to the two main FPGAs. This is equivalent to a total input of 20480 bits per LHC Bunch Crossing. It should be noted that both the MGTs and optical receivers can operate at up to 13 Gb/s.

Board configuration and control The module initialisation is performed by an IPMC device mounted on the module. The control FPGA is configured by a local memory. The processor FPGAs are configured from an SD flash memory sequenced by the control FPGA. Module control is achieved via IPbus, and IP connectivity is provided by two Ethernet ports, one on the front panel and one via the ATCA backplane.

Clock The L1Topo module is designed to operate at the 40.08 MHz LHC bunch-crossing frequency. For synchronous operation, data transmitters operate at multiples of the LHC bunch clock while the receivers' reference clocks may optionally be derived by local crystal oscillators. The jitter on the MGT clock is reduced by a PLL device. L1Topo receives the LHC bunch clock and associated data through a TTCDec module, based on a TTCRx chip, connected to the control FPGA.

Algorithmic firmware Many topological trigger algorithms are under study, as discussed in section 2.2.6. It is not yet clear which algorithms will be chosen, and the selection may change with luminosity. The firmware is therefore developed in a modular structure, isolating the algorithmic core from the handling of FPGA I/O, and providing generic functions such as sorting of input trigger objects. The algorithms themselves will rely on tools that provide cuts on angles or invariant masses of object pairs, and identification of overlapping objects. The output consists of individual bits indicating the results of each specific topological trigger algorithm. Each module may transmit up to 64 bits to the CTP.

Module read-out In Run 2, the read-out to the ROI builder and DAQ is implemented in the control FPGA. The Read-Out Links are based on the 2 Gb/s HOLA S-LINK protocol for compatibility with the existing system. The implementation will use the FPGA on-chip serialiser/de-serialiser together with a pluggable MiniPOD fibre optic transceiver. An additional read-out path is provided directly from the processor FPGAs onto the backplane. This link is compatible with the Level-1 calorimeter trigger HUB module (see Section 3.4.5), and provides access to higher read-out speeds and additional read-out processing needed after LS3.

Installation The L1Topo ATCA shelf will be installed as close as possible (to minimise latency) to the CTP. The fibres from the L1Calo CMX modules to L1Topo will be routed via the existing "short cut" hole in the floor between USA15 levels 1 and 2.

Phase-II upgrade The L1Topo module is built such that it can be operated as part of the Level-0 trigger after the Phase-II upgrade, although the firmware, content and use of the real-time data paths may change. Processing power can be scaled by the number of modules operating concurrently on the same, or differently selected input data, while input bandwidth can be scaled up by a factor of two, to the maximum rate supported by the Level-1 calorimeter and muon data sources, and the L1Topo FPGAs and the optical receivers. In Phase-II the distribution of the clock and timing signals is expected to change, as is the read-out. Subject to the availability of a suitable backplane interconnect, it will be possible to use the L1Topo module in conjunction with the Hub modules developed for the Level-1 calorimeter feature extractors (see Section 3.4.5) for timing distribution and read-out, by connectivity provided in zone 2 of the ATCA backplane.

5.4 Level-1 Trigger Latency

The total Level-1 latency is composed of two parts: first, the Level-1 *trigger latency*, which is defined as the time from the collision to the output of the corresponding L1A signal at the output of the CTP. Secondly, there is the Level-1 *read-out latency*, which designates the time it takes for the L1A signal to reach the sub-detector front-ends. While the Level-1 trigger latency is common to all trigger signals, determined by the slowest one, the read-out latency depends on the particular sub-detector.

For the Phase-I upgrade, extra trigger latency is needed to accommodate the many changes to the Level-1 Trigger system: eFEX and jFEX in the calorimeter trigger, muon New Small Wheel and sector logic, new MUCTPI, L1Topo and new CTP. Many of the new systems use high-speed digital optical links, which are one of the significant consumers of the latency budget.

Several ATLAS detector subsystems, including large parts of the Inner Detector and the Liquid Argon Calorimeter, will continue to use the same front-end and read-out electronics up to the end of Phase-I running. This means that the overall Level-1 trigger latency must remain within the original design values until that time. As no major changes in the Level-1 read-out latency are foreseen during the Phase-I upgrade, it is the Level-1 trigger latency that must remain in a given budget. The following will show a measurement of the Level-1 trigger latency during Run 1. It will also describe a few BC of reserve due to artificial delays in the Level-1 trigger system, depending on the trigger path. In addition, it will turn out, that there is still a global reserve of 20 BC, coming from reserves in the sub-detector read-out.

Beam pick-up Delay Contribution	Delay (BC)
Time-of-Flight	-23.4
Cables, electronics	40.3
CTPIN delay	45.0
CTP	5.0
Total	66.9

Table 13: Level-1 trigger latency during Run 1 as determined from the ATLAS beam pick-up system

The Level-1 trigger latency during Run 1 operation has been estimated using the ATLAS beam pick-up detectors (BPTX) to be 67 BC. Table 13 shows the various latency contributions of the measurement. The time-of-flight contribution is negative, because the beam pick-up signal is generated before the collision takes place. The BPTX were used because they provide a very simple system, in which all latency components can easily be measured. Using an artificial input delay in the CTPIN, they were aligned in time with all the other trigger systems. This alignment was checked constantly throughout the running period.

	L1Calo		L1Muon	
	Cluster	Jet/Energy	Barrel	Endcap
MIOCT	n/a	n/a	2 BC	6 BC
CTPIN	4 BC	6 BC	1 BC	1 BC
Reserve	4 BC	6 BC	3 BC	7 BC

Table 14: Artificial latency reserves in the Level-1 trigger system during Run 1

This latency value still includes a few artificial delays, listed in Table 14, which could be reduced to compensate for latency increases during the Phase-I upgrade.

System	Latency Reserve	Comments
Pixel	130 BC	
SCT	20 BC	
TRT	124 BC	
LAr	10+ BC	144 deep front-end memory, shared between Level-1 pipeline and derandomiser
TileCal	150 BC	
MDT	n/a	not applicable, since MDT is TDC-based
TGC	32 BC	
RPC	150 BC	

Table 15: Sub-detector read-out latency reserves during the Run 1 period

An ATLAS-wide survey, undertaken in 2010, provided the measurements of reserve read-out latency of the various subsystems, which is illustrated in Table 15. It may be seen that the SCT and LAr calorimeter define the maximum permissible latency increase. In the case of the SCT, spare latency of 20 BC is available. It has been established in tests that all of ATLAS can run with the L1A artificially delayed up to this figure. In the case of the LAr Calorimeter, a common memory is used for latency pipeline and read-out derandomising buffer. With the number of 5 read-out samples during Run 1, an additional latency beyond the current 10 tick reserve would require a reduction of the LAr read-out derandomiser FIFO with a corresponding increase of the complex deadtime, which would not allow running at 100 kHz L1A rate. By reducing the number of LAr read-out samples to 4, which is foreseen for Run 2, LAr will be able to accept a 100 kHz L1A rate and provide a latency reserve of more than 20 BC at the same time. However, when using only 4 read-out samples, a small degradation of the calorimeter resolution and a small impact on the data quality judgement is expected. These effects are being studied in detail. In summary, this means that the overall Level-1 trigger latency budget is $67 \text{ BC} + 20 \text{ BC} = 87 \text{ BC}$. It is likely that ATLAS will choose to run at this maximum latency setting already during Run 2.

ATLAS has established a detailed Level-1 latency calculation, taking proper account of the parallel calculation in different trigger stages. The calculation is based on measurements in the running trigger system, complemented by detailed estimates for new electronics. A summary of the estimated latency in different sections of the trigger is given in Table 16.

The estimated latency of the LAr analogue path, as well as the TGC and RPC path, come uncomfortably close to the Level-1 latency budget of 87 BC. With a contingency of only about 3 BC, the latency of every component needs to be scrutinised in the design and production process, in order to stay within the latency budget.

L1Calo				L1Muon			
LAr DPS	(44.2)	LAr analogue	22.5	TGC (incl. NSW)	41.0	RPC	43.0
Tile option 3	45.0						
		Pre-Processor	16.4	TGC SL	16.0	RPC SL	13.0
eFEX	(13.5)	Cluster	30.1	MUCTPI	7.9		7.9
jFEX	14.0	Jet/Energy	(21.5)				
L1Topo	11.4		7.0		11.4		11.4
CTP direct	8.0		8.0		8.0		8.0
Total	78.4		84.0		84.3		83.3

Table 16: Expected Phase-I upgrade Level-1 trigger latency contributions in units of BC for L1Calo and L1Muon trigger paths. Values in parentheses indicate partial trigger paths that are not on the critical path.

References

- [5.1] D. Berge, J. Burdalo, N. Ellis, P. Farthouat, S. Haas, J. Lundberg, S. Maettig, A. Messina, T. Pauly, D. Sherman, and R. Spiwoкс, *Hardware studies for the upgrade of the ATLAS Central Trigger Processor*, .
- [5.2] Avago, *MicroPODTM and MiniPODTM 120G Transmitters/Receivers*, (online).
http://www.avagotech.com/pages/minipod_micropod.
- [5.3] S. Ask, D. Berge, P. Borrego-Amaral, D. Caracinha, N. Ellis, et al., *The ATLAS Central Level-1 Trigger Logic and TTC System*, *JINST* **3** (2008) P08002.
- [5.4] D. Berge, N. Ellis, P. Farthouat, S. Haas, P. Klofver, A. Krasznahorkay, A. Messina, T. Pauly, G. Schuler, R. Spiwoкс, and T. Wengler, *The ATLAS Level-1 Muon to Central Trigger Processor Interface*, .
- [5.5] S. Haas, *The ATLAS Level-1 Muon to Central Trigger Processor Interface*, .
- [5.6] ATLAS Collaboration, *Letter of Intent for the Phase-I Upgrade of the ATLAS Experiment*, Tech. Rep. CERN-LHCC-2011-012. LHCC-I-020, CERN, Geneva, Nov, 2011.
<https://cds.cern.ch/record/1402470>.
- [5.7] E. Simioni, G. Anders, B. Bauss, D. Berge, V. Büscher, T. Childers, R. Degele, E. Dobson, A. Ebling, N. Ellis, P. Farthouat, C. Gabaldon, B. Gorini, S. Haas, W. Ji, M. Kaneda, S. Mättig, A. Messina, C. Meyer, S. Moritz, T. Pauly, R. Pottgen, U. Schäfer, R. Spiwoкс, S. Tapprogge, T. Wengler, and V. Wenzel, *Topological and Central Trigger Processor for 2014 LHC luminosities*, Tech. Rep. ATL-DAQ-PROC-2012-041, CERN, Geneva, Jul, 2012.
<https://cds.cern.ch/record/1459209>.
- [5.8] Xilinx, Inc., *Virtex-7 FPGA Family*, (online).
<http://www.xilinx.com/products/silicon-devices/fpga/virtex-7/>.

6 Online and High-Level Trigger Software

6.1 Introduction

This section covers the two major software parts of the TDAQ system: the “Online software” which comprises all the infrastructure needed to operate the system and the dataflow software; and the “HLT software” which includes all the HLT algorithms and the event-classification and selection framework in which they run. The latter is built on top of the ATLAS offline framework. The software deployed online contains both online and offline releases along with a small layer that connects the two. The online software provides a base on top of which the different sub-detector communities build their own software projects for DAQ and monitoring. The HLT software is also run as part of the offline simulation to provide the trigger response to simulated data, using the same algorithms and steering code as online. The Level-1 trigger simulation that also runs in this context is outside the scope of this section.

The following subsections provide a brief overview of the system used in Run 1 and plans for Run 2, then the motivation for the Phase-I software upgrades, the baseline assumptions, and a work breakdown with effort estimates.

6.2 Description of the Run 1 and Run 2 Systems

An overview of the TDAQ system used for data taking up to the end of Run 1 is given in Section 1 and shown in Figure 1.

The trigger is more fully described in Ref. [6.1]. The HLT starts from “Regions of Interest” (ROIs) passed from Level-1. Guided by these regions, the Level-2 trigger (L2) performs partial reconstruction of events using specialised software algorithms with access to all detectors at full granularity. Full scans of tracking or calorimeter data are performed on a small fraction of events as needed for specific triggers. Events accepted by L2 are then retrieved from the Read-Out System (ROS), fully assembled in the Event Builder, and passed to a third level, the “Event Filter” (EF). Here, reconstruction is performed using algorithms based more closely on the offline reconstruction software which can work on ROIs or the full event. Selected result data from L2 are transferred to the EF in order to seed the EF reconstruction. The HLT software runs on a large farm of rack mounted commodity computers in one of the surface buildings above the ATLAS detector. Usually one instance of the software is run per processor (CPU) core.

The HLT “steering” [6.2][6.3] and the Trigger Configuration provide a dynamic, flexible system that can cope with a wide range of operational scenarios. They run within the offline Athena [6.4] framework which is based on the Gaudi [6.5] framework, in order to simplify sharing of algorithms, tools and services with offline reconstruction. The HLT steering adds trigger-specific algorithm execution scheduling, monitoring and ROI-based data organisation to Athena. It classifies the event and takes the accept/reject decision at each level. The configuration allows sequences of algorithms to be built up into “signatures” that represent reconstructed objects (e.g. leptons, jets) and global quantities (e.g. missing transverse energy) and to place requirements on their properties (e.g. transverse energy, isolation, and shape). A trigger menu is used to define the combinations of these which are required to accept the event. The same code runs both online and in simulation. Together, the steering, configuration and monitoring comprise most of the “core software” of the trigger.

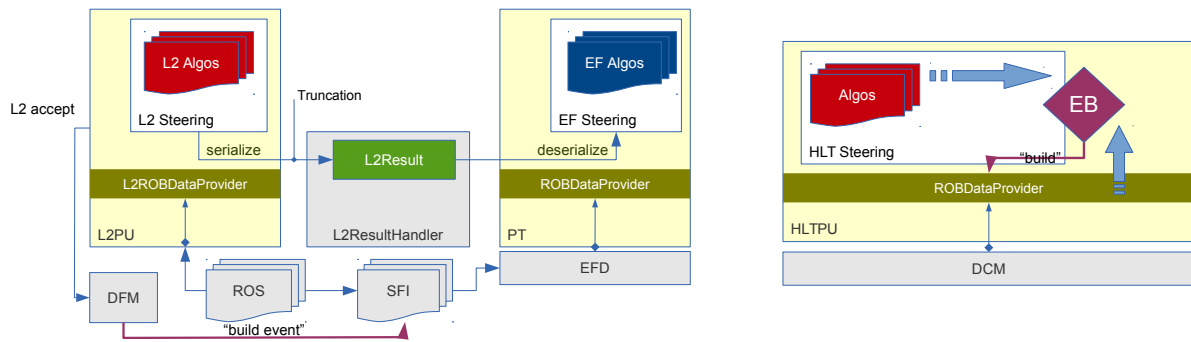


Figure 54: The separate L2 and EF software design in Run 1 (left) and the simplified design for the merged HLT in Run 2 (right). In each case the HLT algorithms are shown at the top and their interaction through interface layers with dataflow components is below. The components are explained in the referring text.

The dataflow software handles all data transfers within the system. For L2, it provides the functionality to request a subset of the full event data from the read-out system. A set of specialised event building, or Sub-Farm Input (SFI), nodes take care of the full event assembly after the L2 decision and moving the events to the Event Filter farm. The specialised data logging, or Sub-Farm Output (SFO), nodes buffer the events accepted by the Event Filter, and safely transport them to mass storage in the CERN computer centre.

The dataflow software components known as the Level-2 Processing Unit (L2PU) and Event Filter Processing Task (EFPT) provide interfaces between the offline Athena framework, under which the HLT steering and algorithms run, and the dataflow components they need to access at L2 and the EF.

A programme of work to evolve the system during LS1 is under way. The main ideas are to take advantage of rolling replacements to hardware explained in section 7 to remove the bottlenecks, simplify the software and gain flexibility (as explained below), and to prepare algorithms and signature strategies for the challenges posed by higher energy and luminosity in Run 2.

The read-out system and network are being replaced with faster equipment which will provide a needed increase in bandwidth through this critical part of the system. Although the HLT farm capacity was not a limiting factor in Run 1, the increase in pile-up and need to use more offline reconstruction code in the HLT will require significantly more computing power for Run 2. Replacement of part of the HLT farm with faster processors and the possibility to fill additional racks with extra computing resources will address this. Estimates of the required computing capacity for Run 3 are presented in Section 7.3.

For Run 2, a simplified design of the dataflow software has been implemented, and the HLT steering and algorithms are being adapted to take advantage of it. Figure 2 shows the overall TDAQ system design and Figure 54 shows how the corresponding software design has evolved. The functions of event building (SFI), L2 and EF will be merged into a single HLT node, with the event building decision (EB) taken internally. The functionality of the Dataflow Manager (DFM) and Event Filter Dataflow (EFD) components and interaction with the ROS are taken over by the Data Collection Manager (DCM).

The merger of the HLT will save processing time and network bandwidth: it removes the need to pack, transfer and unpack L2 result data from L2 to EF; data for L2-accepted

events will not have to be requested twice (for L2 processing and again for event building), and it avoids repeating some reconstruction steps that are common between the two levels, for example low-level detector data unpacking. It is also more flexible for assignment of resources. The HLT processing unit (HLTPU) component which replaces the L2PU and EFPT has been redesigned to share memory between processes by using the “Copy-on-Write” virtual memory management feature of current operating systems. The aim of this is to reduce the incremental memory cost of more processes and thus extend the potential for scaling to the anticipated many-core systems that look set to become available in the next few years.

The rest of this section describes the Phase-I upgrade to this system.

6.3 Motivation

Changes to the online and HLT software are driven by several factors: new physics requirements, detector upgrades and technology evolution. The LHC upgrade and the consequent higher energy, luminosity and increased pile-up will put new requirements on both the DAQ system in terms of rates, data bandwidth, and storage, as well as on the HLT selection to increase efficiencies and reduce the background rates in this environment.

For Phase-I, ATLAS will upgrade several detectors and add new ones. The physics motivation for these upgrades was described earlier in Section 2. This implies support for new front-end data formats, geometry and conditions, read-out software and the corresponding trigger algorithms. The Level-1 upgrades bring essential improvements in rejection power to the Level-1 trigger systems through, for example, the finer granularity of the calorimeter data and the added coverage and resolution of the New Small Wheel for the muon trigger. This means that some of the early, fast rejection currently obtained by the HLT can be moved to Level-1. To maintain sufficient rejection, the HLT will require new approaches: for example more accurate reconstruction so that tighter cuts can be applied, and more analysis-specific selections.

The ATLAS DAQ and HLT system relies on commodity computing and network technologies. The evolution of these technologies up to the end of Run 3 has to be taken into account. LS2 will see another replacement of the core networking switches with a new generation capable of higher performance. DAQ and HLT computers will be replaced multiple times during their normal upgrade cycle as part of the on-going M&O.

While vendors are reluctant to provide a roadmap for this timescale, there are several obvious trends that will change the way software is written. One is the continued increase of the number of CPU cores per motherboard. The software has been able to take advantage of these due to the natural “event parallelism” by running multiple copies of the trigger applications per machine and treating each core as an individual computer. However, this strategy will come to an end when the corresponding memory needed is no longer affordable or memory bandwidth becomes a bottleneck for large applications and data.

There are some indications that the computing industry will move away from the current CPU model to much higher numbers of small, low power cores with limited memory per core but powerful vector instructions, and a lack of cache coherency. The standard PC hardware of the future may be systems built purely from such CPUs, or a mixture of these and the current, comparatively “large” CPUs.

Another trend is the use of more specialised hardware for specific tasks, the most prominent one being the use of graphical processing units (GPUs) for general programming tasks. The Intel Xeon Phi co-processor is a less radical departure from common programming mod-

els, but still requires significant effort to restructure the existing code for optimal performance. The trend of integrating GPUs and potentially other co-processors onto the CPU die implies that in the future these systems become more of a commodity than a specialised configuration for high performance computing only. How to efficiently share these hardware extensions among multiple clients on a machine, or how to write code that will run transparently on these extensions and traditional CPUs is an active area of investigation.

The future software must be able to use any of these new technologies efficiently in order to have the large amount of computing power available for the HLT farm at a realistic cost.

Finally, apart from commodity hardware the ATLAS software also makes use of a large collection of third party software in the form of libraries and tools. While the majority of these are open source, it is difficult to justify their use if the corresponding project itself is not active and evolving over time. Furthermore, languages, programming paradigms, and libraries fall out of favour over time and see reduced support. Some of the central components of the DAQ system have been designed almost 15 years ago and use technologies one might not choose today. However, fundamental changes or replacements to these central components are extremely invasive. The end of LS1 was considered too soon to succeed with such a change.

6.4 Baseline System, Assumptions and Options

This subsection describes the assumptions made about hardware and software and the baseline software design for implementation as part of Phase-I. It will be used for planning purposes with alternatives retained as options in case they are needed.

During Run 3 it is assumed that the software will run on commodity PCs evolved from but recognisably similar to those available now. A PC will have many effective cores, many more than the 12-24 currently typical. These cores may be a mix of real and virtual cores (hyperthreading) and specialised cores like GPUs or Xeon Phis. Sufficient memory as well as I/O and memory bandwidth will be available to exploit these with ATLAS software. Systems will be chosen taking into account the cost/performance ratio for relevant benchmarks, power and infrastructure requirements rather than just the ultimate outright performance. The possibility of a more radical departure from this hardware evolution will be kept under review.

The HLT farm will be a heterogeneous system as a result of the policy of rolling replacement with a lifetime of 5 years. This means that for example in 2019 the system will be using computers purchased between 2014 and 2019 with the main pre-Run 3 purchase planned for Q2 2018. Grid systems on which trigger simulation must run will be even more varied, so that all ATLAS software will have to deal with a variety of hardware.

The purchase of additional co-processors is considered as an option. It may turn out that the balance of co-processors to CPU cores in the normal PC configuration is not optimal or that it is advantageous to augment older PCs. It would have to be shown that a higher throughput can be achieved with more co-processors, within the given effort, financial cost and infrastructure constraints.

For the HLT software framework, the baseline assumption is that the current framework will be replaced by a new framework co-developed in common with the ATLAS offline software. There is a risk that a common set of requirements for the framework or a common implementation of them cannot be identified, in which case a separate framework for the trigger will need to be developed; this risk has been mitigated through early engagement with the project. This new framework will support concurrent execution at multiple levels so it can

be used to get the best out of the prevailing technology mix in the HLT farm. This will help to maximise throughput on many-core systems without high memory costs, and broker the use of co-processor resources to accelerate critical parts of the code.

Many of the algorithmic approaches used in the trigger now will remain useful for Run 3, although in some cases significant changes to algorithms can be anticipated. For example, algorithm and data code will be modified to make use of instruction-level parallelism and data parallelism features of CPUs (for example Advanced Vector Extensions⁸). More offline reconstruction code will be adapted for the trigger to achieve the higher rejection needed while maintaining high efficiency.

The overall dataflow architecture will remain broadly the same as that used in Run 2 and described in Section 6.2. There are no plans to fundamentally change the online software architecture. As mentioned in the previous section, reviews early on will identify whether to replace centrally used components and what the candidate technologies are. Impact on clients (e.g. detector software) will be kept minimal. Parallelism in DAQ components will be improved to take advantage of many-core architectures. However, there does not appear to be an obvious use case for GPUs in the online software.

6.5 Work Breakdown and Effort Estimates

6.5.1 High-Level Trigger selection software tasks

Core software The core software of the HLT includes the framework, configuration, monitoring, and unpacking and formatting of trigger data.

The need to replace the current HLT software framework (the HLT steering and Gaudi) has been motivated in Section 6.3 and the baseline plan to work on this jointly with offline is stated in Section 6.4. There has been a good experience of this approach with Gaudi. The effort estimate is based on this and the effort needed to develop the HLT Steering for Run 1.

ATLAS is currently investigating the prototype of the Concurrent Framework Project (CF4Hep) [6.6], and a group of experts has been set up to identify the common requirements between HLT and reconstruction for algorithm scheduling. In addition to the desired common scheduling, there is scope for collaboration on data I/O while some other trigger-specific framework functionality will still be needed.

The use of instruction-level parallelism at the sub-algorithm level will be very important. This will be used for the most time-consuming parts of the software where possible, to take advantage of co-processors, GPUs and under-used CPU cores. Framework services will be developed to support this across the inhomogeneous hardware portfolio of the HLT farm. This activity ties in with both the framework and the more general task to “Evaluate and exploit new technologies”, described below.

The specific tasks of the core software work package are:

- Requirements, design and prototyping, and eventual implementation of a new framework, in collaboration with offline and other experiments;
- Interface between online software and the offline/HLT framework;
- HLT-specific features/extensions of the new framework;
- Offloading work at the sub-algorithm level to the GPU/other co-processor/idle cores;
- Migration of signatures and algorithms to the new framework;

⁸Intel’s planned AVX-512 instructions will enable the processing of 8 double precision numbers in the same instruction.

- Revision of trigger configuration to support changes to the Level-1 hardware and HLT software;
- Revision of monitoring (especially resource monitoring) to handle parallel, asynchronous component execution;
- Documentation and training specific to the HLT for the new framework;
- Tools and techniques for parallel software validation and debugging, trigger specific support and training;
- Validation and monitoring of new software functionality and performance, offline and online;
- Support for FTK in the HLT steering.

The initial setting of requirements, design and prototyping is already under way in 2013-14. The majority of the framework development will be done during 2015-16 after which the emphasis switches to the large task of porting algorithms, validation and monitoring.

Evaluate and exploit new technologies The aim of this activity is to follow the emergence of new hardware and software technologies and research their potential application in the ATLAS TDAQ system. To minimise effort, this is carried out in collaboration with ATLAS off-line computing and with other LHC experiments, taking advantage of the expertise and early technology access of CERN Openlab. Input will be provided to each HLT farm PC purchase, which will happen several times between now and Run 3. Towards the end of the upgrade programme, work culminates in implementation of any adopted hardware technologies and the best software practices in key parts of the TDAQ system and software.

The main work packages are:

- Evaluate CPU and co-processor/external accelerator developments as they become available from manufacturers. A decision on HLT farm hardware will be taken early in 2017 to allow time for purchasing and commissioning.
- Software optimisation: use profiling tools and techniques, expert code inspection and code redesign to make better use of evolving data and instruction parallelism features of CPUs, provide demonstrations and tutorials for wider developer community.
- Look at new compilers, languages and libraries to facilitate optimal use of new hardware and parallel programming techniques.
- In the last two years of the upgrade project, after HLT farm purchasing decisions have been taken, lead the implementation and testing of best practices for the chosen hardware.

A hardware evaluation project seeking to understand the potential of GPUs is already under way: studies porting L2 tracking algorithms to run on GPUs [6.7] have shown promising speed improvements. The evaluation will take into account that FTK also reduces the CPU time spent on tracking.

Trigger menus and algorithms This is a large scale task due to the number of signatures and algorithms needed. Each of more than 1200 different configurations has to be tuned for the required performance.

It is estimated that around half the algorithmic code will be updated. Factors such as the increased capabilities of Level-1, the availability of new offline reconstruction and selection techniques which must be reflected in the trigger to maintain high efficiency, and the need to

increase robustness against high pile-up, will all lead to significant work on the adoption and adaptation of offline algorithms for the HLT.

Experience has shown that the detector-specific data preparation code (unpacking and re-formatting of the raw data) is far more critical in the trigger than in offline reconstruction. This code will need revision due to detector changes, increased pile-up and changes to corresponding offline code. It requires detector knowledge combined with trigger and programming expertise to write dedicated fast code for the HLT.

Signature, algorithm and data preparation work peaks in the last 2-3 years of the project after the new framework is available, decisions on new software and hardware technologies have been made, and to benefit from the experience of Run 2 data.

FTK will provide a major change of the inner detector tracking information available to the HLT. The HLT tracking code and strategies are being revised to make the use of this in a variety of ways. Signatures will be re-configured to make the most of this. This will be done in 2013-16 since FTK will be commissioned during Run 2.

Trigger simulation Simulation of the FTK is a major challenge as explained in the FTK TDR [6.8]. With the prospect of the HLT tracking relying on the FTK for input for some part of Run 2, full simulation of the HLT by definition requires a usable FTK simulation. A serious effort is expected on this from FTK, HLT and computing experts to achieve a simulation strategy and software that fits well in the ATLAS computing model and does not have prohibitive resource requirements.

A fast simulation of the trigger based on parametrisations derived from data will become an essential tool as high pile-up makes full simulation prohibitively expensive for many analyses. However, working in or close to the turn on region of triggers will become more common as analysers naturally seek to maximise the use of the available data. This will require robust validation procedures. A flexible approach to trigger simulation, in common with the Integrated Simulation Framework [6.9] will be explored.

For precision data analysis, a realistic trigger simulation is necessary, but simulating the complex and changing trigger is very challenging. The software is constantly evolving so already in Run 2 it will be significantly different from that used for taking data during Run 1. A proof-of-concept simulation chain has been demonstrated [6.10] that uses an old online software release and conditions to “simulate” the real trigger for comparison with data of the same period. It can be run in virtual machines running old operating system versions if necessary. The solution is built on virtual machine tools developed by CERN-IT [6.11] and has a lot in common with the approach for long-term data preservation being developed by the LHC experiments. This idea needs to be developed into a production-quality product and given long-term support and maintenance.

6.5.2 Data acquisition and control online software tasks

HLT Processing Unit The HLT processing unit (HLTPU) is the application in the online software that dynamically loads and executes the HLT steering. It interacts with the online software infrastructure services such as run control, configuration, error logging and translates between the two software domains, Online and HLT.

The HLTPU will have to take full advantage of the new concurrent HLT framework. This will further reduce the memory requirements by increasing the resource sharing. It has to be studied if a multi-threaded framework will find enough parallelism in the proposed work flow

of running multiple events at the same time. If that is not the case a combination of multi-process/multi-threaded applications will be necessary, implying a possibly more complex scheduling strategy on a single node.

Online core software and infrastructure All the core parts of the online software use a common middleware layer to communicate. APIs hide the underlying implementation based on an open source implementation of CORBA. This simplifies the development and deployment of new services tremendously since it takes care of all the low-level tasks and meets all the performance requirements.

However, CORBA is an old technology from the 1990s. It could be replaced by a more modern approach such as web services which use the widely supported standard HTTP protocols, or by a CORBA-like solution but with some useful new features such as asynchronous message queues. The risks of keeping CORBA are that the technology could become obsolete on the timescale of the end of Run 3, and lack support or updates from the open source community. The impact of change can be minimised if the existing APIs are retained. An evaluation will be made during 2015 with a decision at the end of that year whether or not to replace it.

Configuration, control and monitoring The run control implementation and its associated expert system has been redone during LS1. No major architectural changes are planned for Phase-I.

The configuration for the whole online system is stored in a custom object-oriented database. It provides a persistent representation based on XML and a network-based access through servers based on the common CORBA middleware. Scalability is achieved by a master-slave model where changes are propagated from the master to the read-only slave servers. More than 16000 applications are configured during the start of a run. Clients access the data through an object-oriented interface that reflects the schema of the database. Graphical editors are available for modifying the database. A relational database (Oracle) back-end is used only for archiving purposes.

All components of the configuration database are developed and maintained in-house. Using an existing solution for the database storage and replication could free up the required manpower for other tasks. The main drawback of most commercial or open source databases is that they typically provide a relational and not an object-oriented interface. The latter are a rather negligible part of the market. Changing the current object-based API to the configuration database would have repercussions for all existing clients.

As for the CORBA infrastructure, an evaluation period is foreseen where possible replacements can be studied. However, any change in this area will take more effort than the former and the changes will be more disruptive. Therefore these evaluations will start before 2015.

The Information Service (IS) is a central part of the monitoring infrastructure of ATLAS. It provides an object-oriented data model together with a publish/lookup/subscribe model to interact with the information. All operation-related data like counters as well as the histograms of all applications are stored in IS. Additional applications on top of IS assemble and reduce the information to the essentials that are presented to a shifter. The online data quality framework is based on the information contained within this system.

The original IS design has a few drawbacks that hinder its scalability. For example, objects are always associated with a specific server process in the system. The definition and location of IS servers has to be planned manually for best performance, and the applications have to

know which server to use. Making the number and location of servers transparent to the clients would make their use of IS easier, as well as provide a way to optimise the servers in the system based on on-going experience and running conditions. All clients would be affected by such an API change, so this has not been attempted for LS1. Again the feasibility and desirability of such a change will be determined during an initial evaluation period in 2015.

Originally many graphical interfaces for the online software were developed from scratch. A variety of libraries and toolkits are used for this, such as OpenMotif, Tcl/Tk, Qt 3 and 4, Java. A disadvantage of these desktop-based GUIs is that they cannot be easily used from outside the ATLAS experiment. Recent developments have made much of the online information available via web services. This trend is expected to continue, and most custom GUIs will be replaced with HTML/JavaScript-based interfaces.

Dataflow and event format The replacement of core dataflow switches during LS2 and the availability of cheaper 10 Gb/s links throughout the system make several simplifications to the existing architecture possible. It is expected that standard Ethernet technology will remove any network limitations on the event building rate. This opens the way to further simplifications of the dataflow and the possibility of more full-scan algorithms subject to the available computing resources.

The option to implement “deferred triggers” will be considered: during the initial high instantaneous luminosity phase, events are stored in the online system to be processed at a later stage when CPU usage has decreased.

The file transfer from the ATLAS TDAQ system to the Tier-0 will undergo a complete revision since the current solution (rfcp) is being phased out. Such a change will also be an opportunity to review the existing software and the file metadata handshake mechanism which is responsible for ensuring that raw files are stored safely before they are deleted in the TDAQ system.

The raw event file format is a critical component of the DAQ system and all the following processing steps. It has to be guaranteed that ATLAS can read raw event data from the beginning of data taking until the end of ATLAS. On the other hand there are constant changes and new requirements that make full backward compatibility a struggle.

The current raw event file format is a custom binary format developed inside ATLAS. The possibility of replacing it with a more widely used standard format will be studied. Initial tests with such a format, HDF5, have shown promise considering the flexibility of a self-describing format, however, there were also some bottlenecks regarding multi-threaded write access to such files. This evaluation of such file formats will continue on a low level. Any changes in this area will have to be synchronised carefully with the rest of ATLAS.

The online software has to be available well in advance of the full HLT for installation, commissioning and cosmic data taking. All major new features will have to be available essentially at the beginning of LS2. This implies that their development will overlap with the on-going maintenance and operation of Run 2.

6.5.3 Effort estimates

Table 17 summarises the estimated effort required by the Phase-I software upgrade outlined above. Alongside this there will still be the need for significant M&O effort to operate, support and maintain the Trigger and DAQ system during Run 2.

Work package	2013	2014	2015	2016	2017	2018	Total
Trigger							
Core software	2.1	3.9	6.5	9.3	8.5	4.3	34.6
Evaluate and exploit new technologies	1.6	1.1	2.0	3.0	3.0	1.5	12.2
Menus & algorithms	3.0	3.0	7.5	9.5	13.0	12.0	48.0
Simulation	1.5	1.3	1.3	1.3	1.3	1.3	7.8
Online							
HLT Processing Unit	0.0	0.0	1.0	0.5	0.4	0.4	2.3
Core Software, Infrastructure	0.0	1.0	1.0	3.0	4.0	1.0	10.0
Configuration, Control, Monitoring	0.0	1.0	2.5	7.3	6.0	0.3	17.0
Dataflow, Event Format	0.0	0.0	1.2	2.4	3.4	1.2	8.2
Detector Software and Tools	0.0	0.0	0.5	1.0	3.0	1.5	6.0
Evaluate and exploit new technologies	0.0	0.5	1.0	1.0	1.0	0.5	4.0
Totals	8.2	9.8	24.4	38.3	43.6	24.0	150.1

Table 17: Effort estimates in FTE years for the software upgrade

6.5.4 Milestones

- Q4 2015: Decisions on need to renew major online software components and software technology choices
- Q4 2016: New HLT steering framework available for testing and algorithm migration
- Q2 2017: Decision on PC architecture for new HLT farm nodes, plus any co-processors or other new technologies to be added
- Q4 2017: Final software release ready for testing
- Q4 2018: Commissioning complete

6.6 Summary

The trigger and online software worked very well during Run 1 and is being developed to meet the new challenges of Run 2. Significant changes are necessary for Run 3, motivated by the LHC luminosity increase, detector upgrades, FTK and the Level-1 trigger improvements, and evolution of software and hardware technology.

A new software framework, jointly developed between the trigger and offline computing, will have the flexibility to maximise throughput with efficient memory use on the many-core CPUs and co-processors expected from industry trends. Dataflow software and HLT algorithms will both be adapted to this new framework. Algorithms and data will be rewritten to exploit the growing instruction- and data-parallelism features of CPUs. Continuous evaluation of new software and hardware technologies will inform the HLT farm purchasing decisions and final software optimisation.

In the face of higher pile-up and the movement of L2-like rejection to the upgraded L1, HLT algorithms will need to maintain high efficiency and rejection through optimised use of offline-like reconstruction and selection.

Fundamental DAQ components for communication (CORBA), configuration, information, graphical user interfaces and the event format will all be reviewed in the next two years and a decision be made about whether or not to replace them during Phase-I. Network technology advances are expected to yield further opportunities to simplify the dataflow design.

A work breakdown and fine-grained effort estimate gives totals over 2013-2018 of about 102 FTE years for the trigger and 48 FTE years for the online software, with a profile that peaks in 2017.

References

- [6.1] ATLAS Collaboration, *Performance of the ATLAS Trigger System in 2010*, *The European Physical Journal C* **72** (2012) no. 1, 1–61.
<http://dx.doi.org/10.1140/epjc/s10052-011-1849-1>.
- [6.2] ATLAS Collaboration, S. George, *The ATLAS high level trigger configuration and steering software: Experience with 7-TeV collisions*, PoS **ICHEP2010** (2010) 487. http://pos.sissa.it/archive/conferences/120/487/ICHEP%202010_487.pdf.
- [6.3] J. Stelzer, *The ATLAS High Level Trigger Configuration and Steering, Experience with the First 7 TeV Collisions*, Tech. Rep. ATL-DAQ-PROC-2011-013, CERN, Geneva, Mar, 2011.
<https://cds.cern.ch/record/1337002>.
- [6.4] P. Calafiura, W. Lavrijsen, C. Leggett, M. Marino, and D. Quarrie, *The athena control framework in production, new developments and lessons learned*, Proceedings of CHEP 2004, Interlaken (2005) 456–458. <http://indico.cern.ch/materialDisplay.py?contribId=108&sessionId=6&materialId=paper&confId=0>.
- [6.5] G. Barrand, I. Belyaev, P. Binko, M. Cattaneo, R. Chytracsek, et al., *GAUDI - A software architecture and framework for building HEP data processing applications*, *Comput.Phys.Commun.* **140** (2001) 45–55.
- [6.6] CERN Concurrent Framework Project (CF4Hep), (online).
<https://concurrency.web.cern.ch/GaudiHive>.
- [6.7] D. Emeliyanov and J. Howard, *GPU-Based Tracking Algorithms for the ATLAS High-Level Trigger*, *Journal of Physics: Conference Series* **396** (2012) no. 1, 012018.
<http://stacks.iop.org/1742-6596/396/i=1/a=012018>.
- [6.8] ATLAS Collaboration, *Fast TracKer (FTK) Technical Design Report*, Tech. Rep. CERN-LHCC-2013-007. ATLAS-TDR-021, CERN, Geneva, Jun, 2013.
<https://cds.cern.ch/record/1552953>.
- [6.9] J. Chapman, P. Clark, M. Duehrssen, M. Elsing, D. Froidevaux, R. Harrington, R. Jansky, R. Langenberg, R. Mandrysch, Z. Marshall, E. Ritsch, and A. Salzburger, *Concepts and Plans towards fast large scale Monte Carlo production for the ATLAS Experiment*, *Journal of Physics: Conference Series* (May, 2013) .
<https://cds.cern.ch/record/1547511>.
- [6.10] G. G. for the ATLAS Collaboration, *ATLAS Trigger Simulation with Legacy Code using Virtualization Techniques*, Tech. Rep. ATL-DAQ-SLIDE-2013-239, CERN, Geneva, May, 2013. <https://cds.cern.ch/record/1547287>.
- [6.11] Buncic, P., Aguado Sánchez, C., Blomer, J., Harutyunyan, A., and Mudrinic, M., *A practical approach to virtualization in HEP*, *Eur. Phys. J. Plus* **126** (2011) no. 1, 13.
<http://dx.doi.org/10.1140/epjp/i2011-11013-1>.

7 Data Acquisition and High-Level Trigger

7.1 Evolution during LS1

7.1.1 Read-out system

Off-detector read-out The off-detector read-out chain of the ATLAS sub-detectors, see Figure 1 in the introduction (Section 1), consists of sub-detector specific Read-Out Drivers (RODs), which output event fragments to the Read-Out System (ROS) via point-to-point optical Read-Out Links (ROLs) [7.1]. In Run 1 the ROS was implemented with 4U high rack-mountable server class PCs, each equipped with up to four custom 64-bit PCI cards, the ROBINS [7.2]. Each ROS PC is also equipped with a dual port 1 Gb/s Ethernet Network Interface Card (NIC) for connection to the TDAQ data collection network. Via these links, each ROS PC receives: requests for data fragments; delete requests; and forwards requested data. In total there are 155 PCs, most having four ROBIN cards which are each connected up to three ROLs.

Evolution to date During Run 1 the performance requirements made on the ROS have surpassed the design specifications. The degree to which the performance requirements have evolved cannot be captured in a single number. However, a metric by which the ROS performance is measured is the maximum fraction of the data received per ROL that can be read out to the HLT. During Run 1 this fraction has increased from 10% to 21%. A further example of the performance increase required of the ROS, is that required to support the introduction at Level-2 of E_T^{miss} selection based on the energy sums calculated by all the calorimeter RODs (and included in the event fragment sent to the ROS PCs). The HLT requests these sums for all 12 ROLs connected to a single ROS PC. The processing performed by a ROS PC for such a request is equivalent to an event building request, with the additional requirement of decoding the energy sums. In Run 1 the Level-2 E_T^{miss} rate has been ~ 10 kHz and the event building request rate ~ 6 kHz, representing a factor five increase in load for the ROS PC for these types of requests, referred to as "full readout requests". The increase in performance requirements were addressed in 2011 as part of the rolling replacement policy. However, due to operational issues, the rolling replacement was only performed for 78% of the installed systems.

Rolling replacement during LS1 During LS1, ATLAS will install the IBL, introduce the FTK and increase the read-out capability of the Pixels, SCT and MDTs by increasing the number of ROLs. These changes require a 25% increase in the number of deployed ROS PCs. At the same time the available rack space in USA15 will not increase proportionally thus making it a requirement for the ROS to be more compact. As already described in Section 2.4 the HLT menu for Run 2 will be significantly different compared to that of Run 1, further increasing the performance and functional requirements beyond the capabilities of the currently deployed ROS. To avoid the ROS being the limiting factor in the physics reach of ATLAS, a new ROS will be deployed in LS1 by bringing forward the next rolling replacement of the ROS.

Target performance of the new ROS The required performance of the new ROS is determined by the rate at which data is requested by the HLT and in turn depends on the trigger menu and therefore the Physics program of ATLAS in Run 2 and Run 3. The experiences of Run 1 demonstrated the need for continued flexibility and therefore headroom in the performance of the ROS to avoid limitations that could cause reduced physics output of the

experiment. Taking into account the motivations described in Section 2.4 the new main performance requirements on the ROS at a 100 kHz Level-1 accept rate are given in Table 18. Although the Level-1 trigger accept rate in Run 2 and Run 3 will be 100 kHz, for reasons of

	Run		Initial design
	2–3	1	(2003)
Full read-out rate (kHz)	17	16	3.5
ROI request rate (kHz)	41	24	21
Read-out fraction/ROL (%)	25	21 ^a	10

^aExtrapolated to 100 kHz Level-1 rate.

Table 18: The required request rate per ROS PC and the readout fraction per ROL for Run 2 and Run 3. Also shown are the values observed during Run 1 and the values specified in the original DAQ/HLT TDR [7.3].

contingency the new ROS is required to operate at a maximum Level-1 accept rate of 120 kHz for event fragment sizes up to 1.2 kB per ROL. This exceeds by 20% the maximum Level-1 rate at which the Pixel and SCT detectors are able to be read out at a μ of 80. Further, in Run 2 and Run 3 the event size (see Figure 57) at a μ of 80 is projected to be up to 2.5 MB⁹, compared to 1.5 MB in Run 1. Subsequently, the requirement on the throughput of the ROS has increased not only due to the increased request rate but also due to the increase in the average event size and the requests for calorimeter energy sums (Level-2 E_T^{miss} requests) are replaced by requests for the full data. Based on the main requirements given in Table 18, for a ROS PC receiving data from 12 ROLs, each running at 200 MB/s and at a 100 kHz Level-1 trigger rate, the target value for the read out fraction per ROL is 50%. This is to be compared to the requirement of 25%. An associated requirement on the ROBIN is that it should allow a 100% read-out, if required during Run 2 or Run 3, only by the upgrading of PCs and/or of firmware. A further requirement is that a 100% read-out fraction should be possible for a reduced number of ROLs without requiring an upgrade of the firmware and/or of the ROS PC. The rest of this section describes the design of the ROS aimed at meeting these requirements and the results of initial measurements performed on prototype hardware.

The new ROBIN The new ROBIN, referred to as RobinNP, is a PCI Express (PCIe) card supporting 12 ROLs via three QSFP transceivers. It implements an eight lane second Generation (commonly referred to as Gen2) interface, that has a maximum effective bandwidth of 3 GB/s, which in principle allows a 100% read-out of all data input on 12 ROLs running at 200 MB/s. A single Xilinx Virtex-6 series FPGA [7.4] handles data input and buffering in the on-board memory and receives and processes requests to transfer event data from buffer memory to the PC memory. The ROS PC must keep track of where data is stored (“indexing”) at a rate equal to the product of the number of ROLs being supported and the Level-1 trigger accept rate, i.e. at the maximum of 1440 kHz for a single RobinNP. Preliminary tests have shown that this performance can be handled by a single core of a current CPU.

The RobinNP functionality will be implemented using the Combined Read-Out Receiver Card (C-RORC) [7.5] developed by the ALICE collaboration. A prototype of the card is shown in Figure 55. A block diagram of the main functional elements of the RobinNP and their interactions with the host PC is shown in Figure 56. The RobinNP firmware consists of two

⁹This does not take into account possible reductions to be implemented by the detector systems.




Figure 55: Photo of a prototype of the ALICE C-RORC card. The Virtex-6 FPGA is in the middle of the board, at the left the cages for three QSFP transceivers can be seen. The connector at the top is a standard FMC connector, at the right there are two SO-DIMM memory slots, in each a 4 GB module will be installed.

functional modules, called “ROBgroups”, each of which interfaces six input data channels to one of two on-board physical memory modules and therefore can be considered as a group of six ROBs. Each ROBgroup also provides page handling and a DMA path for event fragments to be transferred to the host memory. As such, each ROBgroup is functionally independent of the other (with the exception of the PCIe interface). Splitting in this manner provides for an efficient use of FPGA resources. A commercially available “core” implements some of the PCIe protocol and provides DMA engines, thus reducing the amount of firmware to be developed.

The new ROS PCs The server class ROS PCs will be, subject to the outcome of ongoing prototyping studies, either 2U or 3U high addressing the requirement for a more compact solution. The baseline foresees a single RobinNP, connected to up to 12 ROLs, in each ROS PC, but two or even more RobinNPs per ROS PC is an option. As the existing optical fibres used for the ROLs are equipped with LC connectors and the RODs use SFP transceivers, a patch panel and fan-in cable will be used to connect four LC pairs to one QSFP transceiver. Each ROS PC will connect by two 10 Gb/s Ethernet links to the data collection network.

Prototyping studies Initial firmware development is performed on a PCIe eight-lane Gen2 Xilinx Virtex-6 FPGA¹⁰ development board [7.6]. In addition to the Virtex-6 FPGA, this board has two slots for DDR2 SO-DIMMs and the IP core for a DMA controller [7.7]. Optical I/O for testing S-LINK interface functionality is possible via an FMC mezzanine equipped with QSFP+ transceivers, an FM-S28 mezzanine [7.8], and optical patch cords.

Test measurements have been performed in which firmware implements a PCIe Gen1 interface and emulates 12 inputs being served by a single ROBgroup. Upon request ROB fragments with programmable payload size are generated and transferred under DMA control, over PCIe, to the memory of a prototype ROS PC, see Figure 56. Two test programs, running on separate PCs, are used to emulate HLT nodes, i.e. the generation of data requests and

¹⁰Type XC6VLX240T-2FFG1759C

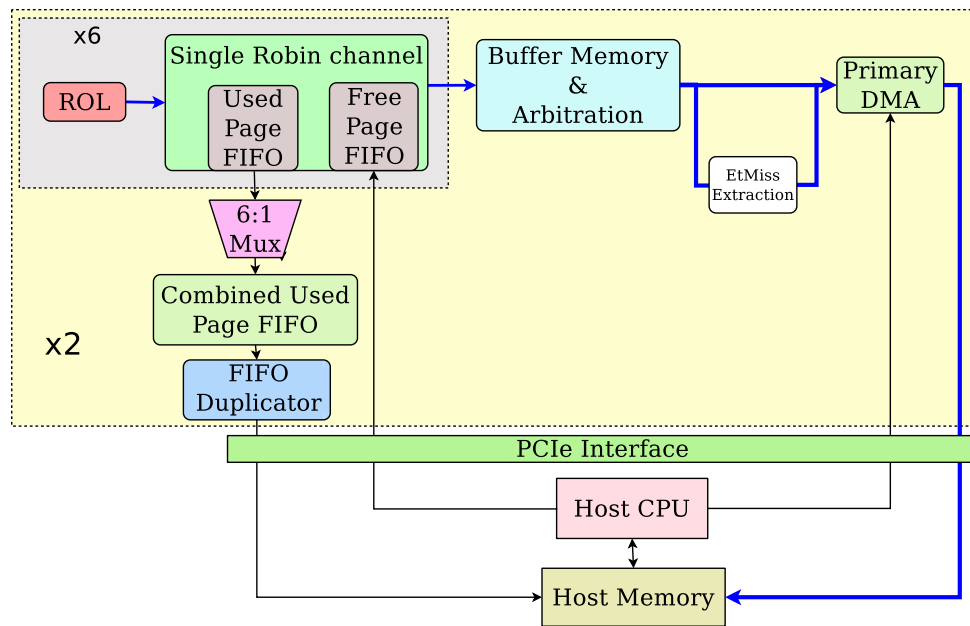


Figure 56: Block Diagram of the RobinNP firmware functionality, showing one of the two ROBgroups. Each ROBgroup interfaces to six ROLs and to a single buffer memory (in the form of a 4 GB SO-DIMM module). Also shown are the interactions with the CPU and host memory of the PC in which the RobinNP is installed.

delete commands. These PCs are each connected via a 10Gb/s Ethernet link to the prototype ROS PC. For a Level-1 accept rate of 100kHz and an event fragment size up to 0.8kB a read-out fraction of 70% has been measured. The measured value of the read-out fraction decreases to 45% for an event fragment size of 1.6kB. These results are determined principally by the performance of software executed on the prototype ROS PC, which remains to be optimised. In separate performance tests of the eight lane PCIe Gen1 interface (implemented in firmware) a bandwidth of more 1.4GB/s has been measured. This would allow read-out fractions of 100% and 73% for event fragment sizes of 1 and 1.6kB respectively. By implementing the PCIe Gen2 interface in the firmware, the bandwidth of the interface can be doubled. The measurements therefore indicate that the target read-out fraction per ROL of 50% is achievable.

Planning and major milestones An initial design review was completed in March 2013 and a combined final design and production readiness review will precede the production which is currently scheduled to start in March 2014. Pre-series modules will be produced in the fourth quarter of 2013 and will undergo exhaustive acceptance tests in both ATLAS and ALICE, and an extended period of long tests in USA15 for ATLAS, prior to the start of full production. The installation and commissioning will be completed in September 2014.

7.1.2 DAQ networks

Evolution of operational needs While the performance requirements associated to the control and monitoring functionality of the TDAQ system have not evolved since the beginning of operations, the performance requirements, i.e. bandwidth, associated with the movement of the event data has evolved during Run 1, will increase at the start of Run 2 and will fur-

ther increase as pile-up increases. Figure 57 shows the measured and extrapolated event size versus pile-up. As can be observed, the event size at a pile-up value of 80 will increase by 50% compared to Run 1, i.e. at a pile-up of 34, and as described in Section 7.1.1, the quantity of data to be read out of the ROS will in addition increase due to an increase in the rate at which the HLT will request data (see Table 18). Therefore, the bandwidth out of a ROS PC will surpass the bandwidth of two 1 Gb/s links implemented in Run 1.

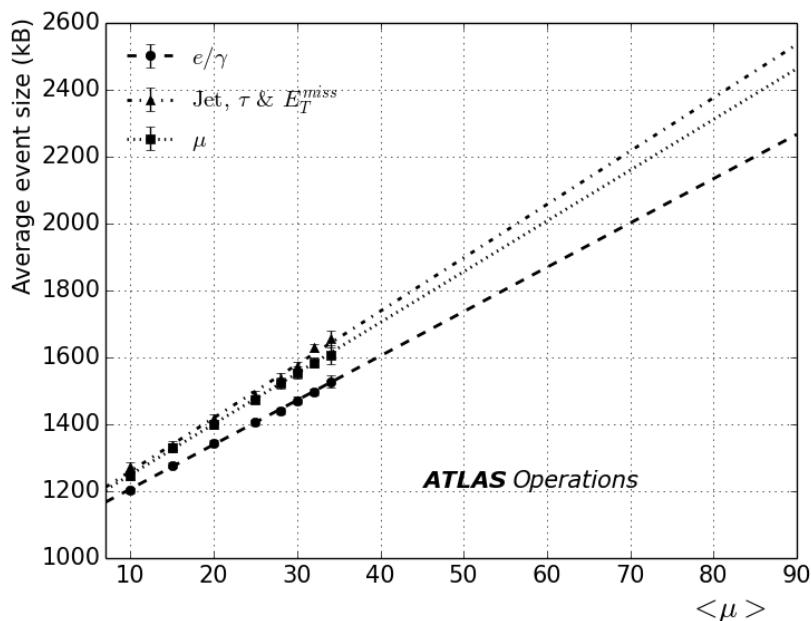


Figure 57: The measured total event size versus μ and its projection to a μ of 90 for different types of HLT selection

Additional performance requirements arise due to the intent of ATLAS to operate the HLT at an accept rate of 1 kHz, this is to be compared to the rate of 400 Hz at the end of Run 1. Combined with the increase in the average event size described earlier, the network bandwidth from the HLT, through the SFO system, to CERN's central data recording service, CASTOR, is required to increase by at least a factor of four to 2.5 GB/s. This is to be compared to 0.6 GB/s at the end of Run 1.

Finally, the HLT compute farm represents a significant computing resource within ATLAS and as such its use must be maximised. To this end a new requirement is to make available the HLT compute farm during long shutdown periods to be exploited as a private ATLAS grid-site. This translates into not only providing the necessary bandwidth for job submission, control and monitoring of the grid-site, but also to secure network access so that the traffic associated to this functionality does not affect the operations of ATLAS at point-1 and more importantly does not open up a network security hole.

Evolution of network architecture During LS1 the core routers of the DAQ networks will be refurbished, in line with the rolling replacement policy. The new core routers will reflect the trend to increased line speeds, i.e. to 10 Gb/s, in port density and will incorporate the latest enhancements to router functionality, e.g. in the domain of quality of service. Figure 58 shows a schematic diagram of the data collection network architecture after LS1. It consists of two overlapping class B sub-networks, allowing the usage of quality of service mechanisms in order to prioritise the different types of traffic as required. The core of each sub-network is

implemented by a router, labelled DC1 and DC2 in Figure 58, able to support up to 768 10 Gb/s ports with a data forwarding capacity of 6.4 Tb/s [7.9]. Each ROS PC will connect directly to the networks via two 10 Gb/s links. This bandwidth provision, 2.5 GB/s, is to be compared to the ROS target design requirement of 1.2 GB/s (and the requirement as understood today of 0.6 GB/s for a read-out fraction of 25%) thus meeting the need for increased bandwidth out of the ROS and eliminating the need for the layer of switches that concentrated the 1 Gb/s links into 10 Gb/s uplinks. Similarly, more performant SFO nodes will be directly connected to the core routers by 10 Gb/s links. Complementing these routers are concentrator switches that connect low-bandwidth devices, such as the HLT nodes and legacy SFO nodes, to the network. The connection between the edge router, used by the SFO nodes to transfer data to CASTOR, and the computing centre has been extended from 20 to 80 Gb/s.

As described in Section 6.2, the functionality of collecting ROI data and event building will be merged into a single software process replicated on each HLT processing node. With this evolution in the dataflow software and the operational experience gained to-date, the functionality of the back-end network¹¹ in this architecture is now provided by an enhanced data collection network, thus removing the need for a separate back-end network with its associated core routers, see Figures 1 and 2.

When operating as a grid-site, job submission, control and monitoring of the HLT farm is provided via the network. So as to completely isolate the network of the HLT farm each HLT rack is now connected via an additional dedicated 1 Gb/s link to the edge router. On this router strict access lists and restricted routing will be used in order to insure the security of the existing ATLAS networks and of the event output path.

Scalability In addition to the criteria of bandwidth the core routers of the TDAQ networks have been chosen taking into account their scalability with respect to the number of connections that can be supported. The known number of ROS PCs post LS1 is 190 (including the additional ROS PCs required for the FTK, IBL and expanded read-out of the Pixel, SCT and MDT detectors) and the number of HLT racks that can be deployed is 80. Taking these numbers of 10 Gb/s ports and the other nodes that will be required to connect to each core router (see Figure 58), the total number of 10 Gb/s ports per core router required post LS1 is estimated to be 250. This is to be compared to the 768 10 Gb/s connections that the core routers can be scaled to support. In particular, for the data collection network, it will be possible to accommodate up to 80 HLT racks each connected to the data collection core routers via two 10 Gb/s links, by equipping the chassis with high port density blades. The control network has a similar level of connection scalability.

7.1.3 Output to mass storage

In order to operate at an average rate of 1 kHz physics output to mass storage, at the event size expected for a value of μ equal to 80 (up to ~ 2.5 MB, see Figure 57) the system that provides the interface between the HLT and mass storage, the SFO (sub-farm output), will have to be upgraded. At peak luminosity, assuming a physics rate to disk of 2.2 kHz and an event size of 2.5 MB, this system is required to support a throughput of 5.5 GB/s. This represents a factor of five increase in the required performance compared to Run 1. At the same time, to

¹¹The back-end network is used to transport complete events from the event builder nodes to the Event Filter processing computers, and the events accepted by the Event Filter from those computers to the SFOs, for storage to disk.

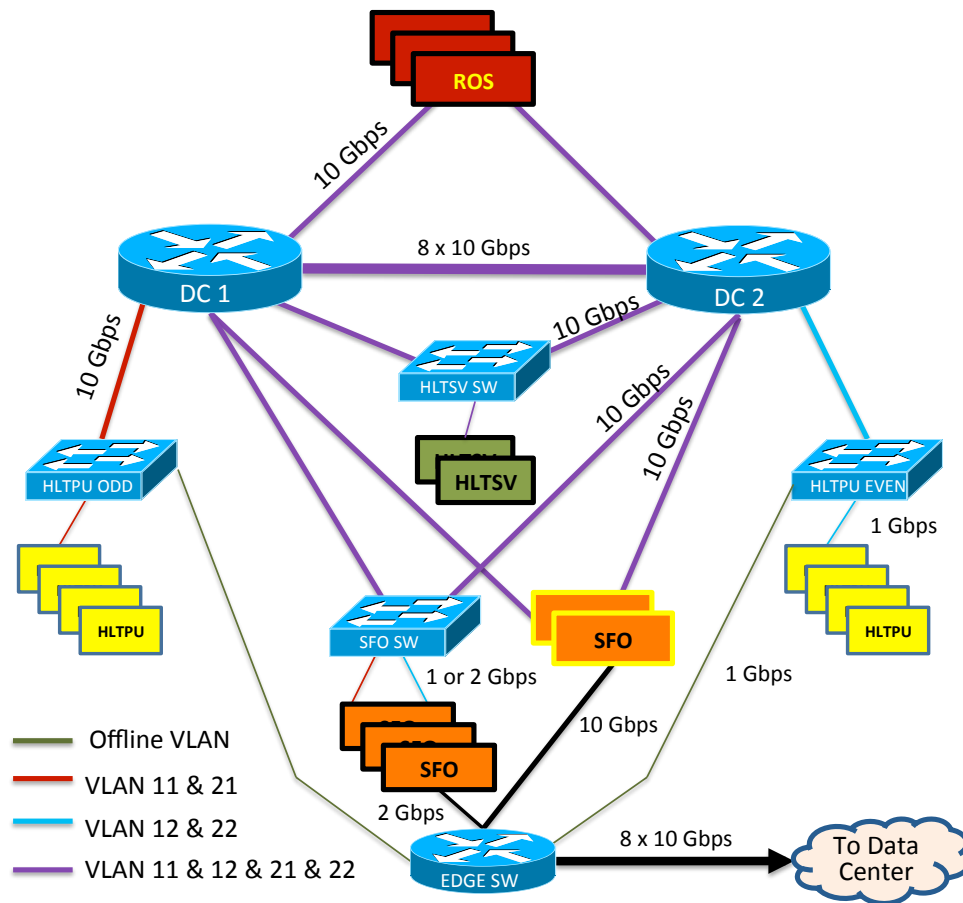


Figure 58: The data collection network after LS1. The core routers are equipped with 10 Gb/s ports only. Direct 10 Gb/s link connect each ROS PC to each core router (for redundancy). Connectivity for each HLT node is provided by a rack level switch, that is connected with a 10 Gb/s uplink to one of the routers and to each node via a 1 Gb/s link. SFOs may have a mixed configuration, with some being directly connected to the data collection network cores and some being aggregated via concentrator switches.

guarantee 24 hours of storage at Point-1 in the case of downtime of the CERN mass storage services, at least ~ 400 TB of effective storage volume is required. This includes the raw data volume for continuous operation at 1 kHz average output, the contribution of file transfer and deletion latencies, measured to be on average six hours during Run 1, and the margin needed to guarantee the throughput characteristics. The latter reflects the tendency of file-system throughput performance to decrease with the occupancy.

The functional operation of the SFO system is not expected to change. Streaming, local storage and off-line transfer functionalities will be provided in a context where global or shared file access across the HLT compute farm is not required. Thus each compute element comprising the SFO system will operate independently and concurrently. Based on this operational model, a direct-attached storage solution, as implemented to-date, continues to be the simplest, scalable and most cost effective solution to implementing the SFO system.

On the time scale of Run 3, the underlying storage technology is not expected to change significantly. Magnetic hard drives will grow in terms of information density but not in throughput performance. Solid-state drives offer increased throughput, but it is expected that they will continue to have a significantly higher cost per unit of stored information. Therefore

we can estimate that ~ 350 hard disk drives will be required to provide the expected input and output throughput. Considering the typical 10%-25% throughput and capacity penalty of the RAID technologies used to ensure redundancy and safety of the temporary stored data, the total number of hard drives to be deployed will be larger than 400. The smallest commercially viable disk drive size compatible with the total volume requirement is 2 TB.

Based on product trends it is estimated that a baseline system composed of 12 storage servers each with 72 TB of raw storage capacity and multiple 10 Gb/s Ethernet links is required to obtain the required performance. This estimation assumes that the required throughput for calibration data is 1% of the average physics output (see Section 2.4.3).

7.2 Upgrades in LS2

7.2.1 Detector read-out

Requirements The current implementation of the read-out architecture does not have intrinsic limitations, apart from the throughput of the ROLs, preventing it meeting the increased read-out performance requirements after the Phase-I upgrades. Prior to LS2, the rolling replacement of the ROS in LS1 (see Section 7.1.1) in combination with the installation of additional ROLs will suffice for the read-out of the detectors. In LS2, the muon New Small Wheels (NSWs) [7.10] will be installed and the Liquid Argon detector's Front-End (FE) electronics upgraded so as to provide finer granularity data to the upgraded Level-1 calorimeter trigger (see Section 3 and Ref. [7.11]). These upgrades will use GBT links [7.12]. This link technology provides for the multiplexing of different types of traffic by the means of so-called e-links, which function as independent data links sharing a single physical connection. These links will be used for: the connection to the TTC system; communication with the DCS and TDAQ systems for control, configuration and monitoring; and the transport of event data to the DAQ system.

The read-out of the systems mentioned will be in operation through to the end of Run 4 and is therefore required to be forward compatible with the proposed Phase-II read-out architecture, mitigating the need for further upgrade prior to Run 4. The final design and baseline implementation of the Phase-II read-out will only be specified in the TDAQ Phase-II Technical Design Report, currently expected in 2016. The emphasis will be on the continued use of commodity computing hardware and software, thus allowing to leverage technology trends in this domain and remain open to changes in the requirements. Notably, the Level-1 accept rate for operations beyond LS3 is currently only specified to be at least 200 kHz due to the open questions regarding the detector capabilities. The uncertainties in the physics scenarios are expected to always be present.

Proposed upgrade The proposed read-out architecture is shown in Figure 59 and in more detail for its Phase-I implementation in Figure 60. Compared to today's read-out architecture (also shown in Figure 59), the main architectural change is the introduction of commodity switching as early as possible in the detector read-out. This replaces the static point-to-point connections between the detector FE electronic systems and the ROS in today's read-out implementation.

The new read-out architecture offers several additional advantages:

- the capability to reconfigure the FE to ROS topology in relation to the bandwidth and computing power required by the actual detector occupancy, as well as with respect to

the performance demanded of specific ROS PCs by the HLT selection algorithms i.e. the rate at which event data is requested;

- processing related to detector specific event data handling and associated error detection and handling could be migrated to COTS processors;
- processing related to control, configuration, DCS and monitoring is migrated away from custom electronics into COTS processors;
- the capability of embedding the TTC protocol and its possible evolution behind a common and adaptable interface, which would avoid the need of the upgrade of the detector FE read out otherwise;
- increased level of commonality between detector chains, which simplifies commissioning and long-term operations and maintenance;
- commodity switching allows straightforward communication of DCS and TDAQ control and monitoring systems with the interfaces to the GBT links.

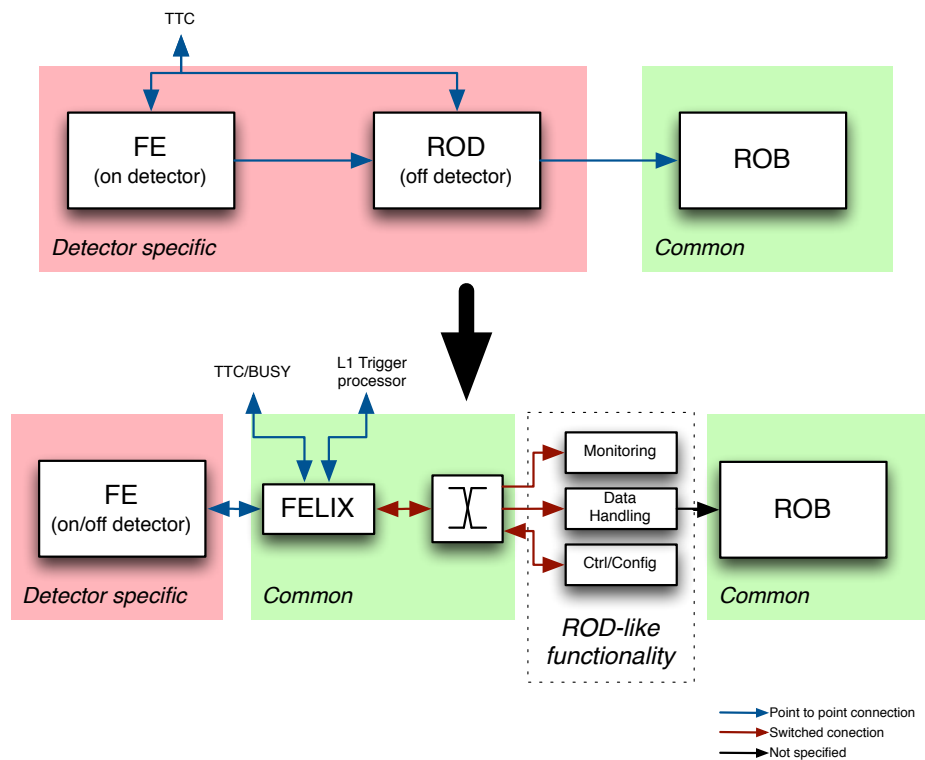


Figure 59: Evolution of the global read-out architecture from the present situation to that foreseen to be implemented for Phase-II

As shown in Figure 59, the introduction of commodity switching as close as possible to the FE link interfacing is achieved by the functional block labelled Front-End Link Interface eXchange (FELIX). The latter receives the detector data via FE links and multiplexes this data on to a network built with commodity switching technology. In the data path towards the FE electronics it provides for the de-multiplexing of: communication between the DCS system and the on-detector control and monitoring circuitry; configuration of the FE electronics, e.g. the setting of thresholds and the enabling or disabling of individual channels. A baseline implementation of a FELIX node will be based on:

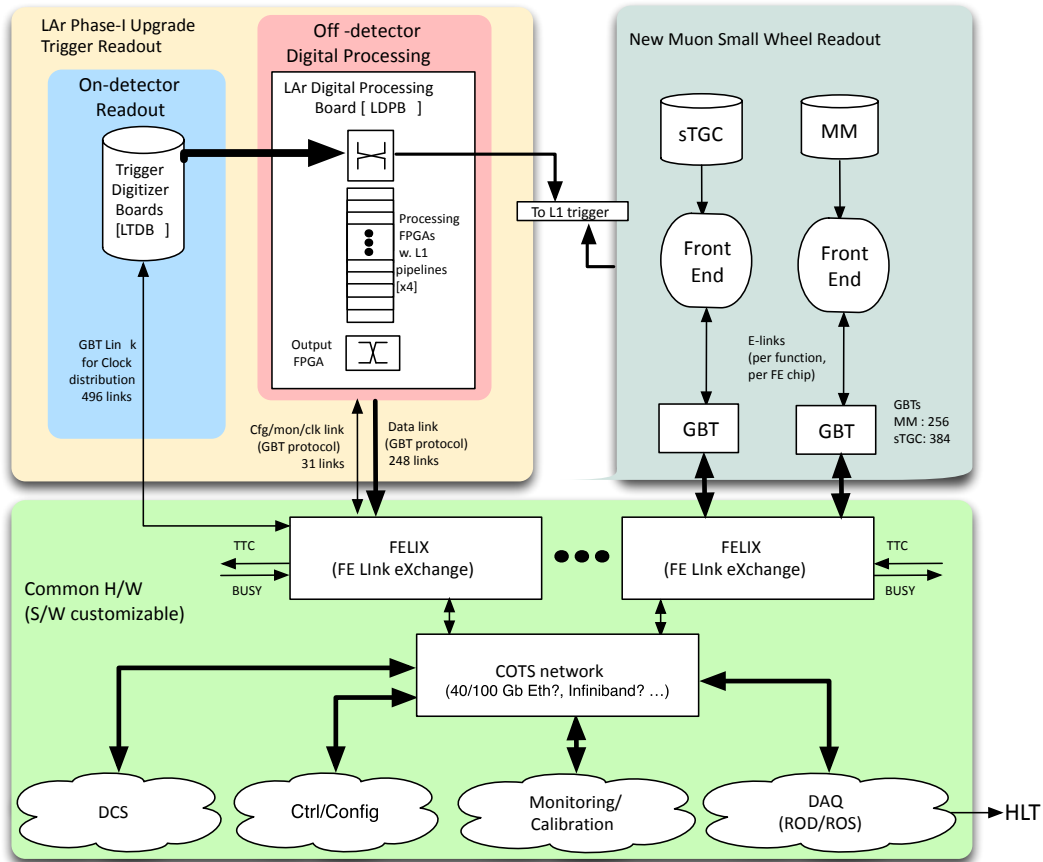


Figure 60: Deployment of FELIX for Phase-I

- A rack mountable server class PC with at least three PCIe Gen3 slots
- One or two PCIe cards each hosting an FPGA that has high-density optical interfaces for interfacing to at least 24 GBT links, sufficient resources for the handling of the data flowing through these links and the necessary PCIe bandwidth
- A PCIe dual NIC implementing a commodity network technology, e.g. 10 or 40 Gb Ethernet or 56 Gb Infiniband.

Suitable PCIe cards hosting a FPGA are commercially available and provide an FPGA Mezzanine Card (FMC) connector. For initial studies interfacing to the TTC and BUSY output could be implemented by an FMC, such as that developed in the context of the CERN GLIB project [7.13]. In the final implementation it is foreseen that one GBT link per PCIe card will be used for providing a clock, TTC signals and to set or remove the BUSY. Each link allocated to the provision of this functionality will be connected to the same type of PCIe card housed in separate and suitable rack mountable housing with an FMC interfacing to the TTC system and providing a NIM output for the BUSY signal. Via each of its 24 GBT links each of these systems will provide a stable clock and TTC signals to up to 24 FELIX nodes.

Table 19 shows the number of GBT links to be deployed for the NSWs and LAr Phase-I upgrades and the subsequent number of FELIX systems. The deployment and operation of the architecture described will be possible in Phase-I as the interface between detector specific electronics and the TDAQ system is encapsulated in the ROS. ROS units connecting to the new

		GBT	FELIX
NSW	sTGC detector	384	18
	sTGC pad trigger	32	2
	MM detector	256	12
LAr	LTDB	496	24
	LDPB	156	8

Table 19: Size of FELIX system for 23 GBT link interfaces per FELIX

read out switched network but supporting the Phase-I HLT request protocol can be provided without particular complexity.

Planning and major milestones The main milestones are:

- 3Q2013 – 4Q2014 design specification
- 2014: prototyping studies
- 4Q2014: Initial design review
- 4Q2015: FDR/PRR and start of production

For testing chambers after production and after installation on the NSWs, prototype FELIX cards will be used. These may be based on commercially available FPGA development kits and/or the C-RORC/RobinNP described in Section 7.1.1).

7.3 Evolution of HLT Compute Power

The currently deployed DAQ/HLT architecture has been designed to handle a maximum Level-1 trigger rate of 100 kHz at a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, corresponding to a pile-up (μ) of approximately 27. In Run 1, the LHC reached a peak luminosity of $0.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, a μ value of 34, and ATLAS has operated at a Level-1 rate of up to ~ 70 kHz. During a typical fill of the LHC machine, the value of μ has varied over the range of 34–12. The measured processing time over this range is shown in Figure 61. It can be seen that the maximum, average event processing times are 75 ms and 900 ms for the Level-2 and Event Filter respectively. The percentage of the average processing time, measured in 2012, in terms of processing the tracking, calorimetry and muon information is shown in Table 20 for both the Level-2 and Event Filter. In 2012, approximately 70% of the average event processing time

	Tracking	Calorimetry	Muon	Other
Level-2	73	17	7	3
Event Filter	52	24	16	8

Table 20: Percentage of the maximum–average event processing time in terms of the type of processing at a $\langle \mu \rangle$ of 35

at the HLT can be attributed to selection algorithms primarily based on track reconstruction. For Run 2 and Run 3 the FTK will perform the primary track reconstruction for the HLT. Based on the motivations presented in Section 2.4.3 it is estimated that the FTK will reduce the contribution of track reconstruction to the average event processing time by 50%. As documented in Ref. [7.14], the FTK will be fully operational for data taking in 2016.




Figure 61: Level-2 (a) and Event Filter (b) processing time measured in 2012 as function of pile-up using an Intel X5650 processor with the results of a quadratic and linear fit to the Level-2 and Event Filter respectively. The observed step in Event Filter processing time occurs when more time consuming selection algorithms are used when $\langle\mu\rangle$ falls below a defined value.

Compared to the HLT selection algorithms used in Run 1, as described in Section 6, to achieve a rate reduction factor of 200 at a μ of 80, in addition to tighter selection criteria, additional selection algorithms as well as more sophisticated offline selection algorithms will be deployed in the HLT. For example, to increase the calorimeter based rejection in multi-jet and E_T^{miss} events a topological, variable-size clustering algorithm will be used, which due to processing time constraints is currently only deployed in the Event Filter. Results of benchmarking indicate that the processing time of the topo-cluster algorithm is approximately 200 ms per event. Taking into account improvements in the performance of this algorithm, and that it is currently estimated to only be executed for 23% of the Level-1 accepted events (see Section 2.4.3), this additional algorithm will contribute to an increase in the average event processing time of 10%, subject to evolution in the trigger menus during operations.

Table 20 shows that the third largest contribution to the average event processing time was the processing of data from the muon detectors. Initial results from ongoing studies and optimisation of the muon selection algorithms indicate that the merging of the Level-2 selection and Event Filter selection into a single HLT selection will reduce the average event processing time by 30%.

In addition to the optimisation of the selection algorithms and the deployment of the FTK, evolution in computing technology is expected to reduce the average event processing time. The per-core compute performance has increased by approximately two HEP-SPEC06 units [7.15] per year. A similar per core performance improvement is also observed using trigger selection algorithms executed on simulated and real data.

Taking into account the above considerations a projection of the average HLT event processing time versus μ , at a Level-1 accept rate of 100 kHz, is shown in Figure 62. Based on this projection the required number of computer cores, of the type expected to be available in 2018, at a $\langle\mu\rangle$ value of 80 is 30,000. This estimate includes the effect of the expected code speed-up, an additional contribution of 10% to the average processing time to account for data I/O and the computing required for online software infrastructure and HLT performance monitoring. The resource estimates given above are subject to large uncertainties:

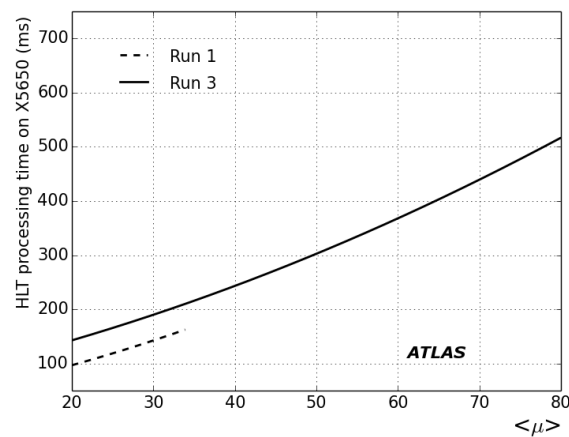


Figure 62: The projected average event processing time in Run 2 as a function of $\langle\mu\rangle$

- The detailed Trigger menu for Run 3, i.e. beyond 2018, is expected to evolve based on the data collected by ATLAS in the years 2015–2017.
- The models for the processing time extrapolation are based on the features of currently known operational data.
- In extrapolating to higher values of μ it has been assumed that the ratio of selection based on tracking, calorimeter, muon and other remains constant for different values of μ .
- It is assumed that the FTK performs “full” track reconstruction at the Level-1 rate of 100 kHz.

References

- [7.1] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, [JINST 3 \(2008\) S08003](#).
- [7.2] R. Cranfield, G. Crone, D. Francis, B. Gorini, B. Green, et al., *The ATLAS ROBIN*, [JINST 3 \(2008\) T01002](#).
- [7.3] ATLAS Collaboration, *ATLAS High-Level Trigger, Data Acquisition and Controls: Technical Design Report*, Tech. Rep. CERN-LHCC-2003-022, ATLAS-TDR-016, CERN, Geneva, Jun, 2003. <https://cds.cern.ch/record/616089>.
- [7.4] Xilinx, Inc., *Virtex-6 FPGA Family*, (online). <http://www.xilinx.com/products/silicon-devices/fpga/virtex-6/>.
- [7.5] H. Engel and U. Keschull, *ALICE CRORC as CBM FLES Interface Board Prototype*, (online). <http://www-alt.gsi.de/informationen/wti/library/scientificreport2011/PAPERS/PHN-NQM-CBM-39.pdf>.
- [7.6] PLDA, Inc., *Standard length Board XpressV6*, (online). <http://www.plda.com/products/boards/standard/xilinx/xpressv6>.
- [7.7] PLDA, Inc., *EZDMA for Xilinx*, (online). <http://new.plda.com/products/fpga-ip/xilinx/fpga-ip-pcie/ezdma-xilinx>.

- [7.8] *Faster Technologies, Inc., FM-S28 Dual QSFP/QSFP+ transceiver FMC*, (online). <http://new.plda.com/products/fpga-ip/xilinx/fpga-ip-pcie/ezdma-xilinx>.
- [7.9] *Brocade MLX Series Routers*, (online). http://www.brocade.com/downloads/documents/data_sheets/product_data_sheets/MLX_Series_DS.pdf.
- [7.10] ATLAS Collaboration, *New Small Wheel Technical Design Report*, Tech. Rep. CERN-LHCC-2013-006. ATLAS-TDR-020, CERN, Geneva, Jun, 2013. <https://cds.cern.ch/record/1552862>.
- [7.11] ATLAS Liquid Argon Calorimeter Group, *ATLAS Liquid Argon Calorimeter Technical Design Report: Trigger Electronics for the LHC Phase-I Upgrade*, Tech. Rep. CERN-LHCC-2013-???. ATLAS-TDR-??, CERN, Geneva, (in preparation). <https://twiki.cern.ch/twiki/bin/view/LAr/LArPhaseITDR>.
- [7.12] P. Moreira, R. Ballabriga, S. Baron, S. Bonacini, O. Cobanoglu, F. Faccio, T. Fedorov, R. Francisco, P. Gui, P. Hartin, K. Kloukinas, X. Llopart, A. Marchioro, C. Paillard, N. Pinilla, K. Wyllie, and B. Yu, *The GBT Project*, . <https://cds.cern.ch/record/1235836>.
- [7.13] *CERN GLIB-project public page*, (online). <https://espace.cern.ch/project-GBLIB/public/default.aspx>.
- [7.14] ATLAS Collaboration, *Fast TracKer (FTK) Technical Design Report*, Tech. Rep. CERN-LHCC-2013-007. ATLAS-TDR-021, CERN, Geneva, Jun, 2013. <https://cds.cern.ch/record/1552953>.
- [7.15] *Homepage HEP-SPEC06 Benchmark*, (online). <http://w3.hepix.org/benchmarks/doku.php/>.

8 Resources, Organisation, and Workplan

8.1 Participating Institute Responsibilities

Tables 21 and 22 list the institutes participating in the TDAQ Phase-I upgrade and their areas of principal interest. For the hardware construction projects – the Level-1 Central Trigger, Level-1 Calorimeter Trigger, Level-1 Muon Trigger, and the DAQ/HLT Dataflow – a cross indicates that the institute plans to develop and provide hardware, firmware, or software closely related to the hardware. For the HLT CPU column, a cross indicates that the institute is contributing financially to the purchase of HLT CPUs. The remaining software upgrades within TDAQ are indicated in the Online & Trig SW column, where a cross indicates a contribution to development of the online or HLT software. Table 23 gives a more detailed breakdown of these (non hardware-related) software contributions planned by TDAQ institutes.

Institute	Country	Central Trigger		Level-1 Calo Trigger							Level-1 Muon Trigger				DAQ/HLT				
		CTP	Topology Proc MUCTPI	eFEX	gFEX	jFEX	Hub,Rod,Opt	nMCM	CMX	TileIn	Muon Endcap	Tile D-Layer	NSW Trig Proc	Bar SL/MUCTPI	Dataflow	Online & Trig SW	HLT CPU		
La Plata	Argentina																x		
Rio de Janeiro UF	Brazil										x								
UAN Bogota	Colombia																x		
Copenhagen NBI	Denmark	x																	
Saclay	France											x							
Berlin HU	Germany	x			x												x		
Heidelberg KIP						x		x											
Mainz		x			x				x										
Weizmann	Israel											x		x					
Genova	Italy																x		
Lecce																	x		
Napoli													x						
Pavia																		x	
Roma I														x				x	
Roma II														x					
Hiroshima IT	Japan										x	x			x	x			
KEK											x	x					x		
Kobe												x	x					x	
Kyoto												x	x						
Nagasaki																		x	
Nagoya												x	x						
Osaka												x	x						
Shinshu												x	x						
Tokyo ICEPP												x	x	x					x
Tokyo MU												x	x						
Tokyo Tech																			x

Table 21: List of participating institutes and areas of interest, countries A-J

Institute	Country	Central Trigger			Level-1 Calo Trigger							Level-1 Muon Trigger				DAQ/HLT		
		CTP	Topology Proc	MUCTPI	eFEX	gFEX	jFEX	Hub,Rod,Opt	nMCM	CMX	TileIn	Muon Endcap	Tile D-Layer	NSW Trig Proc	Bar SL/MUCTPI	Dataflow	Online & Trig SW	HLT CPU
NIKHEF	Netherlands	x													x	x		
Crakow INP Jagiellonian	Poland		x													x		
Portugal LIP	Portugal															x		
Bucharest	Romania												x			x		
Moscow SU	Russia															x	x	
Barcelona	Spain															x		
Stockholm	Sweden		x			x					x							
CERN	Switzerland	x		x											x	x	x	
Birmingham	United Kingdom				x		x											
Cambridge					x		x											
Edinburgh																x		
London QMUL					x		x										x	
London RHBNC																	x	
London UC																	x	
Manchester																	x	
RAL						x		x									x	
Sussex																	x	
Warwick																	x	
Argonne	USA						x										x	
Arizona													x					
Brookhaven BNL						x		x					x		x			
Boston																	x	
Harvard													x					
Michigan SU								x		x								
Oregon								x									x	
Pennsylvania						x		x									x	
SLAC																	x	
Stony Brook								x			x							
UC Irvine														x		x	x	
Wisconsin																	x	

Table 22: List of participating institutes and areas of interest, countries N-U

Institute	Country	Trigger				DAQ/HLT				
		Core software	New technologies	Menus & algorithms	Simulation	HLT Processing	Core Software	Configuration & Control	Dataflow	New technologies
Berlin HU	Germany	x						x		
Genova	Italy					x				
Lecce						x				
Napoli								x		
Pavia										x
Roma I						x		x	x	
Hiroshima IT	Japan							x	x	x
KEK		x		x	x					
Kobe		x		x	x	x		x		
Nagasaki					x					
Tokyo ICEPP										x
Tokyo Tech		x		x	x					
NIKHEF	Netherlands			x						
Portugal LIP	Portugal		x	x	x					
Bucharest	Romania							x		
Moscow SU	Russia									
Barcelona	Spain			x						
CERN	Switzerland	x	x	x	x	x	x	x	x	x
Edinburgh	United Kingdom	x	x	x	x					
London RHBNC			x		x					x
London UC			x		x	x				x
Manchester					x	x				
RAL			x	x	x	x				
Sussex			x		x	x				
Warwick			x		x	x				
Argonne	USA							x	x	
Boston					x	x				
Oregon							x	x		
Pennsylvania			x		x					
SLAC				x	x		x		x	
UC Irvine					x	x	x	x	x	x
Wisconsin			x	x			x			

Table 23: List of participating institutes and areas of TDAQ Software Contribution

8.2 Management Organisation

The TDAQ management structure is shown in Figure 63. TDAQ is led by a TDAQ Management Team (TDMT) of two members with complementary roles, each member being elected by the Trigger and DAQ Institute Board (TDIB) according to ATLAS rules and serving a two year term. The terms of office are arranged in antiphase, the TDMT member in their second year of office acting as the TDAQ Project Leader. The TDAQ Project Leader is a member of the Executive Board and directly responsible to the ATLAS Executive Board and Upgrade Steering Committee for all aspects of the TDAQ upgrade.

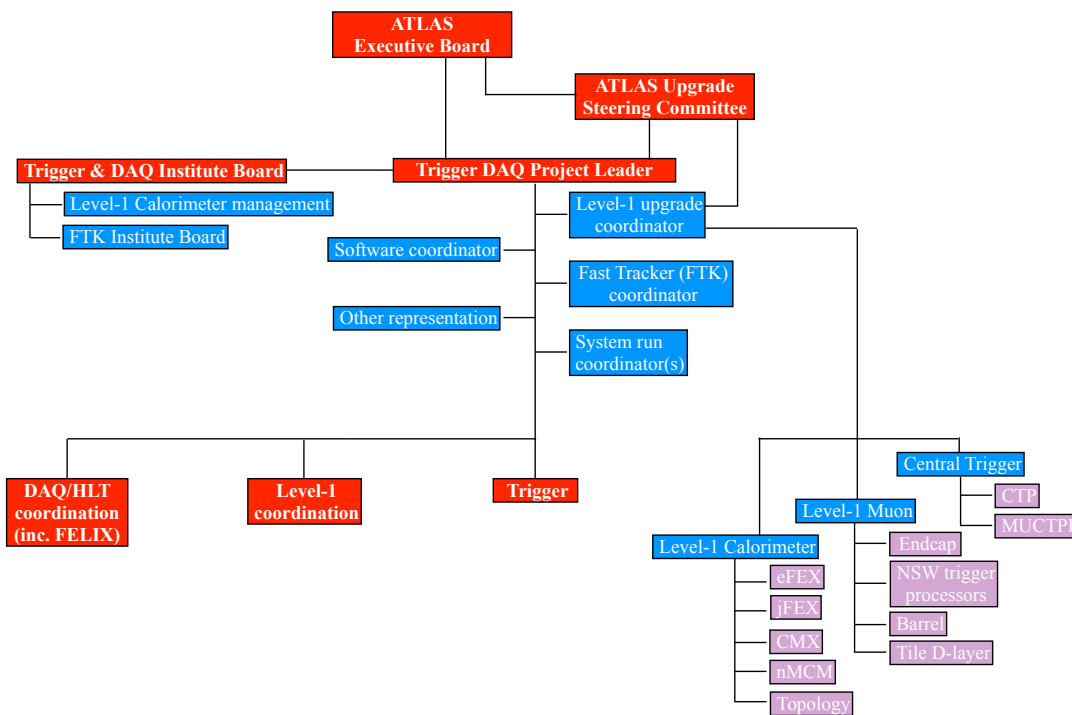


Figure 63: Simplified representation of the current TDAQ Management Structure, highlighting the upgrade activities

The major construction work in the TDAQ upgrade is related to the Level-1 Trigger, and falls under the responsibility of the Level-1 Upgrade Coordinator. This person reports to the ATLAS Upgrade Steering Committee together with the TDAQ Project Leader, and also reports to the TDMT and the TDAQ Steering Group.

Overall coordination of the work for the TDAQ upgrade is provided by the TDAQ Steering Group, the structure of which may evolve after the LS1 work is completed. The members, who are coordinators of the major TDAQ activities, are nominated by the TDMT and approved by the TDAQ Institute Board. The current members include system run coordinators, L1Calo, L1Central, L1Muon endcap and barrel, DAQ/HLT, Trigger, TDAQ Software Coordination, L1Track, FTK, Physics (representative in Upgrade Physics Steering Group), Upgrade Simulation and Upgrade Detector Read-out. Ex-officio members of the TDAQ steering group are ATLAS management, Upgrade, Trigger activity, Run and the Computing coordinators.

Institutes participating in the upgrade by providing financial resources or staff attend the TDIB on an equal basis with institutes participating in operation of the running TDAQ system. The TDIB takes decisions on major technical issues and on sharing of resources

and responsibilities, following recommendations from the TDAQ Steering Group and TDAQ Management.

8.3 Construction and Software Effort

The Work Breakdown Structure for the Phase-I TDAQ upgrade is shown in Table 24, which also gives the construction and software effort required. This is in units of Full-Time-Equivalent years (i.e. one person working full-time for the whole 6-year period would contribute 6 FTE-years), including qualified physicist, engineer and technical effort but not including effort contributions from students.

The total effort of 299.7 FTE-years is contributed by the 51 participating institutes, representing an average of 5.8 FTE-years per institute from 2013 to 2018. 129.1 FTE-years (approximately half the total) is used by the 31 institutes involved in Level-1 trigger activities, averaging 4.2 FTE-years per institute; 102.6 is used in 22 institutes for Trigger software, averaging 4.7 FTE-years per institute; 47.5 are provided by 17 institutes participating in DAQ/HLT software, averaging 2.8 FTE-years per institute; and the remaining 20.5 is provided by 5 institutes participating in Dataflow, averaging 4.1 FTE-years per institute. In each case, a detailed breakdown of the staff effort is included in the corresponding chapter.

8.4 Schedule and Milestones

This section lists in tables 25, 26, 27, and 28 the major milestones for the principal components of the TDAQ upgrade. Detailed schedules are given in the respective chapters.

8.5 Cost and Resources

The estimated CORE cost of the upgrade is shown in total and year-by-year in Table 29. These costs are shown converted to Swiss Francs at average interbank exchange rates for the whole of July 2013. Details of financial contributions by funding agency will be provided separately in the TDAQ Phase-I Upgrade Memorandum of Understanding.

WBS	System	Effort(FTE Years)
1	Central Trigger	
1.1	Central Trigger Processor	6.0
1.2	Level-1 Topology Processor	13.0
1.3	MUCTPI	9.0
		28.0
2	Level-1 Calorimeter Trigger	
2.1	eFEX system	27.7
2.2	jFEX system	31.0
2.3	optional gFEX	see note
2.4	Hub, ROD & Optical plant	4.4
2.5	Tile Input	3.0
2.6	PreProcessor (nMCM)	11.0
2.7	Extended Merger Module (CMX)	3.0
		80.1
3	Level-1 Muon Trigger	
3.1	Endcap Sector Logic	12.5
3.2	Tile D-layer muon interface	8.0
3.3	NSW Trigger Processor	tbs
3.4	Barrel MUCTPI interface	0.5
		21.0
4	DAQ/HLT	
4.1	Dataflow	20.5
4.2	Online & HLT Software	
4.2.1	Trigger	
4.2.1.1	Core software	34.6
4.2.1.2	New technologies	12.2
4.2.1.3	Menus & algorithms	48.0
4.2.1.4	Simulation	7.8
		102.6
4.2.2	DAQ/HLT	
4.2.2.1	HLT processing	2.3
4.2.2.2	Core software	10.0
4.2.2.3	Configuration & control	17.0
4.2.2.4	Dataflow	8.2
4.2.2.5	Detector software & tools	6.0
4.2.2.6	New technologies	4.0
		47.5
	Total	299.7

Note: Effort for the Optional gFEX depends on design details.

Table 24: Work Breakdown Structure for the TDAQ Upgrade, showing effort required for each item. The table includes both hardware and software effort. See text for further details.

WBS	System	Milestone	Date
1.1	Central Trigger Processor	IDR	Q4 2012
		Commissioning complete	Q4 2014
1.2	Level-1 Topology Processor	PDR	Q2 2012
		PRR	Q3 2013
		Commissioning complete	Q4 2014
1.3	MUCTPI	PDR	Q1 2016
		FDR	Q1 2017
		Commissioning complete	Q2 2018

Table 25: Milestones for the Level-1 Central Trigger

WBS	System	Milestone	Date
2.1	eFEX system	PDR	Q3 2013
		PRR	Q2 2016
		Commissioning complete	Q4 2018
2.2	jFEX system	PDR	Q3 2013
		PRR	Q4 2016
		Commissioning complete	Q4 2018
2.4	Hub, ROD & Optical plant	Hub FDR	Q1 2017
		ROD FDR	Q1 2016
		EM Fibre inputs complete	Q4 2018
		Had Fibre Inputs complete	Q4 2016
2.5	Tile Input	IDR	Q3 2013
		PDR	Q4 2014
		PRR	Q4 2015
		Commissioning complete	Q4 2018
2.6	PreProcessor (nMCM)	PRR	Q4 2012
		Commissioning complete	Q2 2014
2.7	CMX Module	PDR	Q2 2011
		PRR	Q1 2014
		Commissioning complete	Q4 2014
2.3	optional gFEX	Architecture decision	Q4 2013

Table 26: Milestones for the Level-1 Calorimeter Trigger

WBS	System	Milestone	Date
3.1	Endcap Sector Logic	PDR	Q4 2015
		PRR	Q4 2016
		Commissioning complete	Q2 2018
3.2	Tile D-layer Muon Interface	PDR	Q3 2013
		PRR	Q2 2014
		Commissioning complete	Q4 2014
3.3	NSW Trigger Processor	PRR	Q4 2015
		Production complete	Q3 2017
		Commissioning complete	Q2 2018
3.4	Barrel MUCTPI Interface	PDR	Q4 2014
		PRR	Q4 2016
		Commissioning complete	Q2 2018

Table 27: Milestones for the Level-1 Muon Trigger

WBS	System	Milestone	Date
4.1	Dataflow	RobinNP production start	Q1 2014
		RobinNP commissioning complete	Q4 2014
		FELIX production start	Q4 2015
4.2	Online & HLT Software	Select components to renew	Q4 2015
		New HLT steering framework available	Q4 2016
		Fix PC architecture for HLT farm	Q1 2017
		Final software release	Q4 2017
		Commissioning complete	Q4 2018
4.3	HLT compute power	Selection by tender (see note)	Q3 2018
		Installation complete	Q4 2018

Note: The HLT CPU schedule assumes all processors are needed in early 2019.

Table 28: Milestones for the High-Level Trigger and DAQ

WBS	System	CORE (kCHF)	Possible addition	Common items	CORE Cost per year (kCHF)						
					2013	2014	2015	2016	2017	2018	
1	Central Trigger										
1.1	Central Trigger Processor	90.0			90.0						
1.2	Topology	503.1			63.6	219.4				220.1	
1.3	MUCTPI	342.0								342.0	
2	Level-1 Calo										
2.1	eFEX system	1,370.0			147.3	147.3	294.6		339.0	294.6	147.3
2.2	jFEX system	786.5							371.1	415.5	
2.3	optional gFEX system		389.3								
2.4	Hub, ROD & optical plant	700.5	53.7				18.0		270.5	412.0	
2.5	Tile Input	139.5			3.1	3.1	6.1		118.1	6.1	3.1
2.6	PreProcessor (nMCM)	500.0			400.0	100.0					
2.7	Extended Merger Module (CMX)	153.6			47.8	105.8					
3	Level-1 Muon										
3.1	Endcap Sector Logic	500.0					30.0		325.0	135.0	10.0
3.2	Tile D-layer muon interface	264.8			264.8						
3.3	NSW Trigger Processor	679.7							235.6	444.1	
3.4	Barrel MUCTPI interface	72.8								72.8	
4	DAQ/HLT										
4.1	Dataflow (see note 1)	1,219.4							853.6	365.8	
4.3	HLT compute power			3,208.3							
	Total	7,321.9	443.0	3,208.3	751.8	840.4	348.7		2,512.9	2,708.0	160.4

- Notes: 1. The Dataflow cost is subject to change with optimisation between the different ATLAS upgrade projects.
2. Common item funds will be spent in 2018.

Table 29: CORE Cost of the TDAQ Upgrade

A Glossary

ATCA Advanced Telecommunications Computing Architecture (industry standard)	HLT High-Level Trigger
BC Bunch Crossing	HLTPU HLT Processing Unit
BCID Bunch Crossing Identification	IB (Sector Logic to MUCTPI) Interface Board (L1Muon)
BCM Beam Condition Monitor	IBL Insertable b-Layer (Pixel detector upgrade during LS1)
BCMuX Beam Crossing Multiplexing	IDR Initial Design Review
BER Bit Error Rate	IPbus protocol implementing register level access over Ethernet to modules
BSM (Physics) Beyond the Standard Model	IPMB Intelligent Platform Management Bus
BW Big Wheel (muon endcap detector)	IPMC Intelligent Platform Management Controller
CMM Common Merger Module (L1Calo)	IPMI Intelligent Platform Management Interface
CMX Common Merger Module extended (L1Calo)	IS Information Service (TDAQ communication infrastructure based on CORBA)
CORBA Common Object Request Broker Architecture (middleware software standard)	JEM Jet/Energy processor Module (L1Calo)
CP Cluster Processor (L1Calo)	JEP Jet/Energy Processor (L1Calo)
CPM Cluster Processor Module (L1Calo)	jFEX Jet Feature Extractor (L1Calo)
CRC Cyclic Redundancy Check	L0A Level-0 Accept (at Phase-II only)
CTP Central Trigger Processor	L1A Level-1 Accept (Phase-I or Phase-II)
CTPCORE CTP Core Module	L1Calo Level-1 Calorimeter Trigger
CTPIN CTP Input Module	L1Muon Level-1 Muon Trigger
CTPOUT CTP Output Module	L1Topo Level-1 Topological processor
DAQ Data Acquisition	L2 Level-2 (part of the Run 1 HLT)
DCS Detector Control System	LAr Liquid Argon (calorimeter)
DDR Double-Data Rate	LS1 Long Shutdown 1 of the LHC (currently in 2013/14)
DPS Digital Processing System (L1Calo, prepares digital eFEX and jFEX inputs)	LS2 Long Shutdown 2 of the LHC (anticipated for 2018)
EF Event Filter (part of the Run 1 HLT)	LS3 Long Shutdown 3 of the LHC (anticipated for 2022)
eFEX Electromagnetic Feature Extractor (L1Calo)	LUT Look-Up Table
EM Electromagnetic	LVDS Low-Voltage Differential Signalling (industry standard)
FCAL Forward Calorimeter	MCM Multi-Chip Module (L1Calo)
FDR Final Design Review	MDT Monitored Drift Tubes (muon barrel detector)
FE (Detector) Front-End	MicroPODTM/MiniPODTM 12-channel 10 Gb/s high-density optical transceivers (proprietary)
FELIX Front-End Link eXchange (interfaces custom front-end GBT links to commercial networks and the TTC system)	MUCTPI (Level-1) Muon-to-CTP Interface
FEX Feature Extraction	nMCM New Multi-Chip Module (L1Calo)
FMC FPGA Mezzanine Card	NSW New Small Wheel (Phase-I muon detector upgrade)
FPGA Field Programmable Gate Array	PCB Printed Circuit Board
FTE Full-Time Equivalent (effort unit)	PCI Peripheral Component Interconnect (industry standard computer bus)
FTK Fast Tracker	PDR Preliminary Design Review
GBT Gigabit Transceiver (Multi-Gb/s data transmission link for DAQ, TTC and DCS data)	PIT Pattern-In-Time bus (CTP)
gFEX Global Feature Extractor module (L1Calo)	
GPU Graphics Processing Unit	

PPM Pre-Processor Module (L1Calo)	SD Secure Digital (memory card format)
PRR Production Readiness Review	SFO Sub-Farm Output (HLT data logger)
QSFP Quad (4-channel) SFP (industry standard)	SFP Small Form-Factor Pluggable Transceiver (industry standard)
ROB Read-Out Buffer	SL Sector Logic (L1Muon)
ROBIN Read-Out Buffer Input card (part of the ROS)	SM Standard Model
RobinNP ROBIN "No-Processor" (second generation ROBIN card)	SNR Signal-to-Noise Ratio
ROD Read-Out Driver	SuperCell LAr calorimeter region formed by combining E_T from a number of adjacent cells in $\eta - \phi$
ROI Region Of Interest	TDAQ Trigger and Data Acquisition
ROIB Region Of Interest Builder	TDR Technical Design Report
ROL Read-Out Link	TGC Thin Gap Chambers (muon endcap detector)
ROS Read-Out System	TMDB Tile-Muon Digitiser Boards
RPC Resistive Plate Chambers (muon barrel detector)	TOB Trigger Object
RTDP Real-Time Data Path	TTC Trigger, Timing & Control
RTM Rear Transition Module (ATCA)	USA15 Underground electronics cavern near the ATLAS detector
Run 1 Data-taking period ending in Feb 2013	VBF Vector-Boson Fusion (Higgs production process)
Run 2 Data-taking period starting after LS1 (anticipated for 2015)	WBS Work Breakdown Structure
Run 3 Data-taking period starting after LS2 (anticipated for 2019)	

B List of Authors

Argentina

Buenos Aires

M.L. Gonzalez Silva, G. Otero y Garzon, R. Piegai, H. Reisin, G. Romeo, S. Sacerdoti

La Plata

M.J. Alconada Verzini, F. Alonso, X.S. Anduaga, M.T. Dova, F. Monticelli, M.F. Tripiana

Armenia

Yerevan

H. Hakobyan, G. Vardanyan

Australia

Adelaide

P. Jackson, N. Soni, M.J. White

Melbourne

E.L. Barberio, A.J. Brennan, S. Diglio, K. Hamano, D. Jennens, T. Kubota, A. Limosani, G. Nunes Hanninger, F. Nuti, Q.T. Shao, K.G. Tan, G.N. Taylor, W.M. Thong, M. Volpi

Sydney

A. Bangert, C.W. Black, C. Cuthbert, G.-Y. Jeng, N.D. Patel, A.F. Saavedra, M. Scarcella, K.E. Varvell, I.J. Watson, B. Yabsley

Austria

Innsbruck

S. Franz, P. Jussel, E. Kneringer, W. Lukas, K. Nagai, E. Ritsch, A. Usanova

Azerbaijan

Baku

O. Abdinov, F. Ahmadov, N. Huseynov, F. Khalil-zada

Brazil

Rio de Janeiro UF

Y. Amaral Coutinho, V. Araujo Ferraz, L.P. Caloba, C.T. Ciodaro Xavier, R. Coura Torres, R. Goncalves Gama, C. Maidantchik, F. Marroquim, A.A. Nepomuceno, J.M. Seixas

UFJF

A.S. Cerqueira, L. Manhaes de Andrade Filho, V.B. Schettino, M.V. Silva Oliveira, J. Vieira De Souza

UFSJ

M.A.B. do Vale

USP

M. Donadelli, M.A.L. Leite

Canada

Alberta

A.I. Butt, K. Chan, P. Czodrowski, D.M. Gingrich, R.W. Moore, J.L. Pinfold, A. Saddique, A. Sbrizzi, HS. Subramania, F. Vives Vaque

Carleton

A. Bellerive, G. Cree, D. Di Valentino, T. Koffas, J. Lacey, J.F. Marchand, T.G. McCarthy, F.G. Oakham, F. Tarrade, R. Ueno, M.G. Vincter, K. Whalen

McGill

C. Belanger-Champagne, B. Chapleau, S. Cheatham, F. Corriveau, R.A. Keyes, R. Mantifel,
S.H. Robertson, M.C. Stockton, M. Stoebe, B. Vachon, K. Wang, A. Warburton

Montreal

J-F. Arguin, N. Asbah, G. Azuelos, J. Bouchami, F. Dallaire, M. Davies, L. Gauthier, M. Giunta,
C. Leroy, R. Rezvani, P. Soueid

Simon Fraser Burnaby

E. Dawe, J. Godfrey, J. Kvita, D.C. O'Neil, M. Petteni, B. Stelzer, A.J. Tanasijczuk, H. Torres,
M. Trottier-McDonald, J. Van Nieuwkoop, M.C. Vetterli

TRIUMF

G. Azuelos, A. Canepa, S.V. Chekulaev, D. Fortin, D.M. Gingrich, A. Koutsman, F.G. Oakham,
E. Perez Codina, P. Savard, D. Schouten, R. Seuster, O. Stelzer-Chilton, R. Tafirout, I.M. Trigger,
M.C. Vetterli

Toronto

O.S. AbouZeid, B. Brelier, B. Fatholahzadeh, N. Ilic, J. Keung, P. Krieger, G. Mc Goldrick, R.S. Orr,
R. Polifka, M.S. Rudolph, P. Savard, S. Schramm, P. Sinervo, T. Spreitzer, J. Taenzer, R.J. Teuscher,
P.D. Thompson, W. Trischuk, N. Venturi

Vancouver UBC

W. Fedorko, C. Gay, Z. Gecse, S.B. King, A. Lister, C.W. Loh, S. Swedish, S. Viel

Victoria

J. Albert, V. Bansal, F. Berghaus, F.U. Bernlochner, C. David, M. Fincke-Keeler, R. Keeler,
R. Kowalewski, M. Lefebvre, C.P. Marino, R.A. McPherson, E.A. Ouellette, J. Pearce, R. Sobie

York

J.A. Benitez Garcia, A.C. Florez Bustos, G. Palacino, W. Taylor

Chile

Santiago

E. Carquin, G. Cottin, M.A. Diaz, B. Panes, F. Quinonez, D. Romero Maltrana

Valparaiso

W.K. Brooks, S. Kuleshov, R. Pezoa, F. Prokoshin, R. White

China

Beijing

Y. Bai, Y. Fang, S. Jin, F. Lu, Q. Ouyang, L.Y. Shan, J. Wang, D. Xu, L. Yao, X. Zhuang, H. Zhu

Hefei

J. Gao, L. Guan, L. Han, Y. Jiang, B. Li, J.B. Liu, K. Liu, M. Liu, Y. Liu, H. Peng, L. Xu, Z. Zhao, Y. Zhu

Nanjing

S. Chen

Shandong

L. Chen, C. Feng, P. Ge, L.L. Ma, X. Zhang, C.G. Zhu

Shanghai

H. Yang

Colombia

UAN Bogota

M. Losada, L. Mendoza Navas, G. Navarro

Czech Republic

Olomouc

P. Hamal, M. Hrabovsky, L. Nozka

Prague AS

J. Chudoba, J. Hejbal, T. Jakoubek, O. Kepka, A. Kupco, V. Kus, M. Lokajicek, R. Lysak, M. Marcisovsky, M. Mikestikova, M. Myska, S. Nemecek, D. Roda Dos Santos, P. Ruzicka, P. Sicho, P. Staroba, M. Svatos, M. Tasevsky, V. Vrba

Prague CTU

K. Augsten, P. Gallus, J. Gunther, J. Jakubek, Z. Kohout, V. Kral, S. Pospisil, F. Seifert, V. Simak, T. Slavicek, K. Smolek, M. Solar, J. Solc, B. Sopko, V. Sopko, M. Suk, D. Turecek, V. Vacek, M. Vlasak, P. Vokac, Z. Vykydal, M. Zeman

Prague CU

P. Balek, P. Berta, K. Cerny, I. Chalupkova, T. Davidek, J. Dolejsi, Z. Dolezal, E. Fullana Torregrosa, P. Kodys, R. Leitner, J. Novakova, V. Pleskot, M. Rybar, D. Scheirich, M. Spousta, T. Sykora, P. Tas, S. Todorova-Nova, S. Valkar, V. Vorobel

Denmark

Copenhagen NBI

A. Alonso, H. Bertelsen, M. Dam, M. Dano Hoffmann, G. Galster, K. Gregersen, J.B. Hansen, J.D. Hansen, P.H. Hansen, S. Heisterkamp, S. Jakobsen, M.D. Joergensen, A.E. Loevschall-Jensen, S. Mehlhase, J. Monk, T.C. Petersen, A. Pingel, M. Simonyan, L.A. Thomsen, C. Wiglesworth, S. Xella

France

Annecy LAPP

N. Berger, M. Delmastro, L. Di Ciaccio, T.K.O. Doan, S. Elles, C. Goy, T. Hryn'ova, S. Jézéquel, H. Keoshkerian, I. Koletsou, R. Lafaye, J. Levêque, V.P. Lombardo, N. Massol, H. Przysieszniak, E. Sauvan, M. Schwoerer, O. Simard, T. Todorov, I. Wingerter-Seez

Clermont - Ferrand

D. Boumediene, E. Busato, D. Calvet, S. Calvet, J. Donini, E. Dubreuil, N. Ghodbane, Ph. Gris, C. Guicheney, H. Liao, D. Pallin, D. Paredes Hernandez, F. Podlyski, C. Santoni, T. Thevenaux-Pelzer, L. Valery, F. Vazeille

FR-CCIN2P3 Tier-1

G. Rahal

Grenoble LPSC

S. Albrand, J. Brown, Q. Buat, B. Clement, J. Collot, S. Crépe-Renaudin, B. Dechenaux, T. Delemontex, P.A. Delsart, C. Gabaldon, M.H. Genest, J-Y. Hostachy, E. Laisne, B.T. Le, F. Ledroit-Guillon, A. Lleres, A. Lucotte, F. Malek, C. Monini, J. Stark, X. Sun, B. Trocmé

LPNHE-Paris

T. Beau, M. Bomben, G. Calderini, F. Crescioli, O. Davignon, S. De Cecco, A. Demilly, F. Derue, M.W. Krasny, D. Lacour, B. Laforge, S. Laplace, O. Le Dortz, G. Lefebvre, K. Liu, B. Malaescu, G. Marchiori, I. Nikolic-Audit, J. Ocariz, C. Rangel-Smith, M. Ridel, L. Roos, S. Trincaz-Duvoid, F. Vannucci

Marseille CPPM

L. Alio, M. Barbero, C. Bertella, N. Bousson, L. Chen, J.C. Clemens, Y. Coadou, F. Djama, L. Feligioni, J. Gao, D. Hoffmann, F. Hubaut, E.B.F.G. Knoop, E. Le Guirriec, B. Li, J. Maurer, C. Meessen, E. Monnier, S.G. Muanza, Y. Nagai, P. Pralavorio, A. Rozanov, T. Serre, M. Talby, N. Tannoury, E. Tiouchichine, S. Tisserant, J. Toth, F. Touchard, M. Ughetto, L. Vacavant

Orsay LAL

S. Abdel Khalek, A. Bassalat, S. Binet, C. Bourdarios, D. Charfeddine, J.B. De Vivie De Regie, L. Dufлот, M. Escalier, L. Fayard, D. Fournier, J.-F. Grivaz, T. Guillemin, S. Henrot-Versille, J. Hrivnac,

L. Iconomidou-Fayard, M. Kado, N. Lorenzo Martinez, A. Lounis, N. Makovec, L. Poggioli, P. Puzo, A. Renaud, D. Rousseau, G. Rybkin, J.B. Sauvan, A.C. Schaffer, E. Scifo, L. Serin, S. Simion, R. Tanaka, H.L. Tran, D. Zerwas, Z. Zhang

Saclay CEA

H. Abreu, H. Bachacou, F. Balli, F. Bauer, N. Besson, J.-B. Blanchard, N.M. Bolnet, M. Boonekamp, L. Chevalier, F. Deliot, J. Ernwein, A.I. Etievre, A. Formica, P.F. Giraud, H.M.X. Grabas, C. Guyot, S. Hassani, W. Kozanecki, E. Lançon, J.F. Laporte, C. Maiani, P. Mal, J.A. Manjarres Ramos, B. Mansoulie, H. Martinez, N. Meric, J.-P. Meyer, L. Mijović, V. Nguyen Thi Hong, R. Nicolaidou, A. Ouraou, E. Protopapadaki, C.R. Royon, L. Schoeffel, Ph. Schune, Ph. Schwemling, J. Schwindling, D. Tsionou, N. Vranjes, M. Xiao

Georgia

Tbilisi IP

E.G. Tskhadadze

Tbilisi SU

T. Djobava, J. Khubua, M. Mosidze

Germany

Berlin HU

E. Bergeaas Kuutmann, F.M. Giorgi, S. Grancagnolo, G.H. Herbert, R. Herrberg-Schubert, I. Hristova, O. Kind, H. Kolanoski, H. Lacker, M. Leyton, T. Lohse, A. Nikiforov, L. Rehnisch, P. Rieck, H. Schulz, D. Wendland, M. zur Nedden

Bonn

T. Abajyan, O. Arslan, M. Backhaus, P. Bechtle, I. Brock, M. Cristinziani, W. Davey, K. Desch, J. Dingfelder, W. Ehrenfeld, G. Gaycken, Ch. Geich-Gimbel, J. Glatzer, L. Gonella, P. Haefner, S. Hageboeck, M. Havranek, D. Hellmich, S. Hillert, F. Huegging, J. Janssen, G. Khorauli, P. Koevesarki, V.V. Kostyukhin, J.K. Kraus, J. Kroseberg, H. Krüger, C. Lapoire, M. Lehmacher, A.M. Leyko, J. Liebal, C. Limbach, T. Loddenkoetter, S. Mergelmeyer, K. Mueller, G. Nanava, T. Nattermann, A.-E. Nuncio-Quiroz, D. Pohl, S. Psoroulas, B. Sarrazin, S. Schaepe, M.J. Schultens, T. Schwindt, F. Scutti, J.A. Stillings, J. Therhaag, J.-W. Tsung, K. Uchida, M. Uhlenbrock, P. Urquijo, A. Vogel, E. von Toerne, P. Wagner, T. Wang, N. Wermes, P. Wienemann, L.A.M. Wiik-Fuchs, K.H. Yau Wong, R. Zimmermann, S. Zimmermann

DESY

S. Argyropoulos, I. Bloch, S. Borroni, J.A. Dassoulas, J. Dietrich, V. Ferrara, M. Filipuzzi, C. Friedrich, A. Glazov, L.S. Gomez Fajardo, J. Goncalves Pinto Firmino Da Costa, K.-J. Grah, I.M. Gregor, A. Grohsjean, M. Haleem, C. Hengler, K.H. Hiller, A. Huettmann, M. Jimenez Belenguer, J. Katzy, T. Kuhl, C. Lange, M. Lisovyi, E. Lobodzinska, D. Ludwig, M. Medinnis, S. Mättig, K. Mönig, T. Naumann, R. Peschke, R.F.Y. Peters, E. Petit, S.M. Piec, V. Radescu, I. Rubinskiy, G. Sedov, S. Shushkevich, D. South, M. Stanescu-Bellu, M.M. Stanitzki, P. Starovoitov, N.A. Styles, K. Tackmann, P. Vankov, J. Wang, C. Wasicki, M.A. Wildt, E. Yatsenko, E. Yildirim

Dortmund

M. Bunse, I. Burmeister, H. Esch, C. Gössling, J. Jentsch, C.A. Jung, R. Klingenberg, T. Wittig

Dresden

P. Anger, F. Friedrich, J.P. Grohs, C. Gumpert, M. Kobel, K. Leonhardt, W.F. Mader, M. Morgenstern, X. Prudent, C. Rudolph, U. Schnoor, F. Socher, P. Steinbach, A. Straessner, A. Vest, S. Wahrenund

Freiburg

G. Aad, S. Amoroso, T. Barber, C. Betancourt, M. Boehler, R. Bruneliere, F. Buehrer, V. Consorti, A. Di Simone, M. Fehling-Kaschek, M. Flechl, C. Giuliani, G. Herten, K. Jakobs, P. Jenni, A.K. Kopp, S. Kuehn, K. Köneke, S. Lai, U. Landgraf, K. Lohwasser, R. Madar, K. Mahboubi, U. Parzefall, M. Rammensee, T.C. Rave, Z. Rurikova, N. Ruthmann, C. Schillo, E. Schmidt, M. Schumacher,

F. Siegert, K. Stoerig, J.E. Sundermann, K.K. Temming, S. Thoma, V. Tsiskaridze, F.C. Ungaro, M. Venturi, H. von Radziewski, T. Vu Anh, C. Weiser, M. Werner, S. Winkelmann, S. Zimmermann

Giessen

M. Düren, K. Kreutzfeldt, H. Stenzel

Goettingen

K. Bierwagen, U. Blumenschein, D. Evangelakou, M. George, L. Graber, J. Grosse-Knetter, M. Hamer, C. Hensel, G. Kawamura, M. Keil, N. Krieger, K. Kroeninger, B. Lemmer, E. Magradze, G. Mchedlidze, J. Meyer, J. Morel, M. Moreno Llácer, O. Nackenhorst, J. Nadal, R.F.Y. Peters, A. Quadt, A.L.S. Schorlemmer, L. Serkin, E. Shabalina, T. Vazquez Schroeder, J. Weingarten, Z. Zinonos

Heidelberg KIP

R. Achenbach, G. Anders, V. Andrei, A. Baas, O. Brandt, Y. Davygora, T.A. Dietzsch, M. Dunford, P. Hanke, J.I. Hofmann, E.-E. Kluge, H. Laier, V.S. Lang, K. Meier, F. Mueller, S. Poddar, V. Scharf, K. Schmitt, H.-C. Schultz-Coulon, R. Stamen, M. Wessels

Heidelberg PI

C.F. Anders, M. Giulini, G. Kasieczka, T. Klimkovich, R. Narayan, S. Schaetzel, S. Schmitt, A. Schoening

Heidelberg ZITI

T. Colombo, A. Kugel, N. Schroer

Mainz

O. Arnaez, B. Bauss, W. Blum, V. Büscher, R. Caputo, F. Ellinghaus, O.C. Endner, E. Ertel, F. Fiedler, C. Goeringer, T. Heck, M. Hohlfeld, P.J. Hsu, T.A. Hülsing, K.B. Jakobi, W. Ji, C. Kahra, A. Kaluza, M. Karnevskiy, P.K. Kiese, K. Kleinknecht, S. Koenig, L. Köpke, M. Lungwitz, S. Maldaner, L. Masetti, J. Mattmann, C. Meyer, D. Moreno, S. Moritz, T. Mueller, A. Neusiedl, R. Poettgen, S. Rave, A. Reiss, H.G. Sander, J. Schaeffer, C. Schmitt, M. Schott, C. Schroeder, N. Schuh, U. Schäfer, E. Simioni, S. Tapprogge, P. Urrejola, V. Wenzel, S.J. Wollstadt, C. Zimmermann

Munich LMU

S. Adomeit, S. Becker, O. Biebel, J. Bortfeldt, P. Calfayan, B.K.B. Chow, J. de Graat, G. Duckeck, J. Ebke, J. Elmsheuser, C. Heller, R. Hertenberger, F. Legger, J. Lorenz, A. Mann, C. Meineck, J. Mitrevski, T. Nunnemann, L.B. Oakes, F. Rauscher, P. Reznicek, A. Ruschke, M.P. Sanders, D. Schaile, J. Schieck, C. Schmitt, D. Vladoiu, R. Walker, J.Z. Will, J. Wittkowski, A. Zibell

Munich MPI

T. Barillari, S. Bethke, B. Bittner, J. Bronner, G. Compostella, G. Cortiana, M.J. Flowerdew, P. Giovannini, M. Goblirsch-Kolb, T. Ince, A.E. Kiryunin, S. Kluth, O. Kortner, S. Kortner, H. Kroha, A. Macchiolo, A. Manfredini, S. Menke, H.G. Moser, M. Nagel, R. Nisius, H. Oberlack, C. Pahl, R. Richter, D. Salihagic, R. Sandstroem, P. Schacht, Ph. Schwegler, F. Sforza, S. Stern, S. Stonjek, S. Terzo, H. von der Schmitt, P. Weigell, A. Wildauer, D. Zanzi

Siegen

N.B. Atlay, P. Buchholz, H. Czirr, I. Fleck, B. Gaur, K. Grybel, I. Ibragimov, K. Ikematsu, M. Rammes, O. Rosenthal, V. Sipica, W. Walkowiak, M. Ziolkowski

Wuerzburg

P. Fleischmann, A. Redelbach, M. Schreyer, G. Siragusa, R. Ströhmer, J.Y.C. Tam, T. Trefzger, S.W. Weber

Wuppertal

M. Barisonzi, K. Becker, T.A. Beermann, J. Boek, T.T. Boek, T. Cornelissen, D. Duda, G. Ernis, J. Fischer, S. Fleischmann, T. Flick, K. Hamacher, T. Harenberg, T. Heim, D. Hirschbuehl, S. Kersten, A. Khoroshilov, S. Kohlmann, P. Mättig, M. Neumann, S. Patariaia, M. Sandhoff, G. Sartisohn, W. Wagner, C. Zeitnitz

Greece

Athens

S. Angelidakis, A. Antonaki, S. Chouridou, D. Fassouliotis, N. Giokaris, P. Ioannou, K. Iordanidou, C. Kourkoumelis, A. Manousakis-Katsikakis, N. Tsirintanis

Athens NTU

T. Alexopoulos, M. Byszewski, M. Dris, E.N. Gazis, G. Iakovidis, K. Karakostas, N. Karastathis, S. Leontsinis, S. Maltezos, K. Ntekas, E. Panagiotopoulou, Th.D. Papadopoulou, G. Tsipolitis, S. Vlachos

Thessaloniki

K. Bachas, C. Gentsos, S. Gkaitatzis, I. Gkialas, D. Iliadis, K. Kordas, V. Kouskoura, S. Nikolaidis, I. Nomidis, K. Papageorgiou, C. Petridou, D. Sampsonidis, O. Sidiropoulou, K.L. Sotiropoulou

Israel

Technion Haifa

A. Di Mattia, R. Kopeliansky, N. Lupu, E. Musto, Y. Rozen, S. Tarem

Tel-Aviv

H. Abramowicz, G. Alexander, N. Amram, G. Bella, O. Benary, Y. Benhammou, E. Etzion, A. Gershon, O. Gueta, N. Guttman, Y. Munwes, Y. Oren, I. Sadeh, Y. Silver, A. Soffer, N. Taiblum

Weizmann Rehovot

R. Alon, L. Barak, S. Bressler, Z.H. Citron, E. Duchovni, D.M. Front, O. Gabizon, E. Gross, J. Groth-Jensen, D. Lellouch, L.J. Levinson, G. Mikenberg, A. Milov, D. Milstein, I. Roth, J. Schaarschmidt, O. Silbert, V. Smakhtin, O. Vitells

Italy

Bologna

L. Bellagamba, M. Bindi, D. Boscherini, A. Bruni, G. Bruni, M. Bruschi, D. Caforio, M. Corradi, S. De Castro, R. Di Sipio, L. Fabbri, M. Franchini, A. Gabrielli, B. Giacobbe, P. Grafström, M.K. Jha, I. Massa, L. Massa, A. Mengarelli, S. Monzani, M. Negrini, M. Piccinini, A. Polini, L. Rinaldi, M. Romano, C. Sbarra, N. Semprini-Cesari, R. Spighi, S.A. Tuppusti, S. Valentinetti, M. Villa, A. Zoccoli

Cosenza

M. Capua, G. Crosetti, L. La Rotonda, V. Lavorini, A. Mastroberardino, A. Policicchio, D. Salvatore, M. Schioppa, E. Tassi

Frascati

A. Annovi, M. Antonelli, M. Beretta, H. Bilokon, V. Chiarella, M. Curatolo, R. Di Nardo, B. Esposito, C. Gatti, P. Laurelli, G. Maccarrone, A. Sansoni, M. Testa, E. Vilucchi, G. Volpi

Genova

D. Barberis, G. Darbo, A. Favareto, A. Ferretto Parodi, G. Gagliardi, C. Gemme, E. Guido, P. Morettini, B. Osculati, F. Parodi, S. Passaggio, L.P. Rossi, C. Schiavi

Lecce

G. Chiadini, E. Gorini, N. Orlando, M. Primavera, S. Spagnolo, A. Ventura

Milano

G. Alimonti, A. Andreazza, M.I. Besana, L. Carminati, D. Cavalli, M. Citterio, S.M. Consonni, G. Costa, M. Fanti, D. Giugni, T. Lari, V. Liberali, L. Mandelli, F. Meloni, C. Meroni, L. Perini, C. Pizio, F. Ragusa, S. Resconi, R. Simoniello, A. Stabile, G.F. Tartarelli, C. Troncon, R. Turra

Napoli

A. Aloisio, M.G. Alviggi, V. Canale, G. Carlino, G. Chiefari, F. Conventi, R. de Asmundis, M. Della Pietra, C. Di Donato, A. Doria, R. Giordano, P. Iengo, V. Izzo, L. Merola, S. Patricelli, S. Perrella, E. Rossi, A. Sanchez, G. Sekhniaidze, G. Zurzolo

Pavia

P. Dondero, R. Ferrari, M. Fraternali, G. Gaudio, A. Lanza, M. Livan, A. Negri, G. Polesello, D.M. Rebuffi, A. Rimoldi, V. Vercesi

Pisa

F. Bertolucci, V. Cavasinni, S. Citraro, M. Dell'Orso, T. Del Prete, S. Donati, P. Giannetti, P. Luciano, C. Luongo, M. Piendibene, C. Roda, F. Scuri

Roma I

F. Anulli, P. Bagiacchi, P. Bagnaia, C. Bini, G. Ciapetti, A. D'Orazio, D. De Pedis, A. De Salvo, A. Di Domenico, C. Dionisi, S. Falciano, A. Gabrielli, S. Gentile, S. Giagu, V. Ippolito, M. Kuna, F. Lacava, C. Luci, L. Luminari, A. Messina, A. Nisati, E. Pasqualucci, E. Petrolò, L. Pontecorvo, M. Rescigno, S. Rosati, F. Safai Tehrani, A. Sidoti, E. Solfaroli Camillocci, M. Vanadia, R. Vari, S. Veneziano, L. Zanello

Roma II

G. Aielli, R. Cardarelli, G. Cattani, A. Di Ciaccio, G.C. Grossi, B. Liberti, F. Marchese, L. Mazzaferro, A. Salamon, R. Santonico

Roma Tre

A. Baroncelli, M. Biglietti, V. Bortolotto, F. Ceradini, B. Di Micco, A. Farilla, E. Graziani, M. Iodice, D. Orestano, F. Petrucci, C. Stanescu, M. Trovatelli

Trieste ICTP

B.S. Acharya, W.B. Quayle, C. Sandoval

Udine

B.S. Acharya, M. Alhroob, S.F. Brazzale, M. Cobal, U. De Sanctis, M.P. Giordani, M. Pinamonti, W.B. Quayle, C. Sandoval, K. Shaw, R. Soualah

Japan

Hiroshima IT

Y. Nagasaka

KEK

M. Aoki, Y. Arai, Y. Ikegami, M. Ikeno, H. Iwasaki, J. Kanzaki, T. Kohriki, T. Kondo, T. Kono, Y. Makida, S. Mitsui, K. Nagano, K. Nakamura, M. Nozaki, S. Odaka, O. Sasaki, Y. Suzuki, Y. Takubo, S. Tanaka, S. Terada, K. Tokushuku, S. Tsuno, Y. Unno, M. Yamada, A. Yamamoto, Y. Yasu

Kobe

Y. Inamaru, T. Kishimoto, T. Kitamura, H. Kurashige, R. Kurumida, T. Matsushita, A. Ochi, S. Shimizu, H. Takeda, K. Tani, I. Watanabe, Y. Yamazaki, L. Yuan

Kyoto

M. Ishino, T. Kunigo, T. Sumida, T. Tashiro

Kyoto UE

R. Takashima

Kyushu

K. Kawagoe, S. Oda, H. Otono, J. Tojo

Nagasaki

T. Fusayasu, M. Shimojima

Nagoya

S. Hasegawa, L. Morvaj, T. Ohshima, Y. Takahashi, M. Tomoto, J. Wakabayashi, K. Yamauchi

Okayama

I. Nakano

Osaka

M. Endo, K. Hanagaki, M. Hirose, J.S.H. Lee, M. Nomachi, W. Okamura, Y. Sugaya

Shinshu

Y. Hasegawa, H. Ohshita, T. Takeshita

Tokyo ICEPP

G. Akimoto, S. Asai, Y. Azuma, T. Dohmae, Y. Enari, K. Hanawa, N. Kanaya, Y. Kataoka, T. Kawamoto, S. Kazama, K. Kessoku, T. Kobayashi, Y. Komori, T. Mashimo, T. Masubuchi, H. Matsunaga, T. Nakamura, Y. Ninomiya, T. Okuyama, H. Sakamoto, Y. Sasaki, J. Tanaka, K. Terashi, I. Ueda, H. Yamaguchi, Y. Yamaguchi, S. Yamamoto, T. Yamamura, T. Yamanaka, K. Yoshihara

Tokyo MU

U. Bratzler, C. Fukunaga

Tokyo Tech

K. Higuchi, M. Ishitsuka, O. Jinnouchi, T. Kanno, D. Kobayashi, M. Kuze, R. Nagai, T. Nobe

Tsukuba

K. Hara, T. Hayashi, S.H. Kim, K. Kiuchi, F. Ukegawa

Waseda

T. Iizawa, N. Kimura, T. Mitani, Y. Sakurai, K. Yorita

Morocco

CNESTEN Rabat

Casablanca

D. Benchekroun, A. Chafaq, M. Gouighri, A. Hoummada, S. Lablak

Marrakesh

M. El Kacimi, D. Goujdami

Oujda

S. Boutouil, J.E. Derkaoui, M. Ouchrif, Y. Tayalati

Rabat

R. Cherkaoui El Moursli

Netherlands

Nijmegen

G.J. Besjes, S. Caron, V. Dao, N. De Groot, F. Filthaut, C. Galea, P.F. Klok, A.C. König, A. Salvucci

Nikhef

R. Aben, I. Angelozzi, L.J. Beemster, S. Bentvelsen, D. Berge, E. Berglund, G.J. Bobbink, A. Borga, K. Bos, H. Boterenbrood, A. Castelli, A.P. Colijn, I. Deigaard, P. de Jong, C. Deluca, L. De Nooij, P.O. Deviveiros, S. Dhaliwal, P. Ferrari, S. Gadatsch, D.A.A. Geerts, F. Hartjes, N.P. Hessey, N. Hod, O. Igonkina, P.P.M. Jansweijer, P. Kluit, E. Koffeman, H. Lee, T. Lenz, F. Linde, J. Mahlstedt, J. Mechnich, I. Mussche, K.P. Oussoren, P. Pani, D. Salek, N. Valencic, P.C. Van Der Deijl, R. van der Geer, H. van der Graaf, R. Van Der Leeuw, I. van Vulpen, W. Verkerke, J.C. Vermeulen, M. Vranjes Milosavljevic, M. Vreeswijk, H. Weits

Norway

Bergen

T. Buanes, O. Dale, G. Eigen, A. Kastanas, W. Liebig, A. Lipniacka, P.L. Rosendahl, H. Sandaker, T.B. Sjursen, B. Stugu, M. Ugland

Oslo

L. Bugge, M.K. Bugge, D. Cameron, B.K. Gjølsten, E. Gramstad, F. Ould-Saada, K. Pajchel, M. Pedersen, A.L. Read, O. Røhne, L. Smestad, S. Stapnes

Poland

Cracow AGH-UST

L. Adamczyk, T. Bold, W. Dabrowski, M. Dwuznik, I. Grabowska-Bold, D. Kisielewska, S. Koperny, T.Z. Kowalski, B. Mindur, M. Przybycien, A. Zemla

Cracow IFJ PAN

E. Banas, P.A. Bruckman de Renstrom, D. Derendarz, E. Gornicki, W. Iwanski, A. Kaczmarska, K. Korcyl, Pa. Malecki, A. Olszewski, J. Olszowska, E. Stanecka, R. Staszewski, M. Trzebinski, A. Trzupek, M.W. Wolter, B.K. Wosiek, K.W. Wozniak, B. Zabinski

Jagiellonian

Portugal

LIP

J.A. Aguilar-Saavedra, S.P. Amor Dos Santos, N. Anjos, J. Augusto, N.F. Castro, P. Conde Muiño, M.J. Da Cunha Sargedas De Sousa, M.C.N. Fiolhais, B. Galhardo, A. Gomes, P.M. Jorge, L. Lopes, J. Machado Miguens, A. Maio, J. Maneira, A. Onofre, A. Palma, R. Pedro, B. Pinto, H. Santos, J.G. Saraiva, J. Silva, A. Tavares Delgado, F. Veloso, H. Wolters

Republic of Belarus

Minsk AC

S. Harkusha, Y. Kulchitsky, Y.A. Kurochkin, P.V. Tsiareshka

Minsk NC

S. Yanush

Romania

Bucharest IFIN-HH

C. Alexa, E. Badescu, S.I. Buda, I. Caprini, M. Caprini, A. Chitan, M. Ciubancan, S. Constantinescu, C.-M. Cuciuc, P. Dita, S. Dita, O.A. Ducu, A. Jinaru, A. Olariu, D. Pantea, M. Rotaru, G. Stoicea, A. Tudorache, V. Tudorache

ITIM

G.A. Popeneciu

Politehnica Bucharest

G.L. Darlea

Timisoara WU

Russia

JINR Dubna

F. Ahmadov, I.N. Aleksandrov, E. Alexandrov, V.A. Bednyakov, I.R. Boyko, I.A. Budagov, G.A. Chelkov, A. Cheplakov, M.V. Chizhov, D.V. Dedovich, M. Demichev, G.L. Glonti, M.I. Gostkin, N. Huseynov, S.N. Karpov, M.Y. Kazarinov, D. Kharchenko, E. Khramov, V.M. Kotov, U. Kruchonak, Z.V. Krumshcheyn, V. Kukhtin, E. Ladygin, I.A. Minashvili, M. Mineev, V.D. Peshekhonov, E. Plotnikova, I.N. Potrap, V. Pozdnyakov, N.A. Rusakovich, R. Sadykov, A. Saponov, M. Shiyakova, V.B. Vinogradov, A. Zhemchugov, N.I. Zimine

Moscow FIAN

A.V. Akimov, I.L. Gavrilenko, A.A. Komar, R. Mashinistov, P.Yu. Nechaeva, A. Shmeleva, A.A. Snesarev, V.V. Sulin, V.O. Tikhomirov

Moscow ITEP

A. Artamonov, P.A. Gorbounov, V. Khovanskiy, P.B. Shatalov, I.I. Tsukerman

Moscow MEPhI

A. Antonov, K. Belotskiy, O. Bulekov, V.A. Kantserov, A. Khodinov, D. Krasnopevtsev, A. Romaniouk, E. Shulga, S.Yu. Smirnov, Y. Smirnov, E.Yu. Soldatov, V.O. Tikhomirov, S. Timoshenko

Moscow SU

A.S. Boldyrev, L.K. Gladilin, Y.V. Grishkevich, V.A. Kramarenko, V.I. Rud, S.Yu. Sivoklov,
L.N. Smirnova, S. Turchikhin

Novosibirsk BINP

A.V. Anisenkov, O.L. Beloborodova, V.S. Bobrovnikov, A.G. Bogdanchikov, V.F. Kazanin, A.A. Korol,
V.M. Malyshev, A.L. Maslennikov, D.A. Maximov, S.V. Peleganchuk, K.Yu. Skovpen, A.M. Soukharev,
A.A. Talyshev, Yu.A. Tikhonov

Petersburg NPI

O.L. Fedin, V. Gratchev, O.G. Grebenyuk, A. Kazarov, V.P. Maleev, Y.F. Ryabov, V.A. Schegelsky,
E. Sedykh, D.M. Seliverstov, V. Solovyev

Protvino IHEP

A. Borisov, S.P. Denisov, R.M. Fakhruddinov, A.B. Fenyuk, D. Golubkov, A.V. Ivashin, A.N. Karyukhin,
V.A. Korotkov, A.S. Kozhin, A.A. Minaenko, A.G. Myagkov, V. Nikolaenko, A.A. Solodkov,
O.V. Solovyanov, E.A. Starchenko, A.M. Zaitsev, O. Zenin

Serbia

Belgrade IP

A. Dimitrievska, J. Krstic, D.S. Popovic, Dj. Sijacki, Lj. Simic

Belgrade Vinca

T. Agatonovic-Jovin, I. Bozovic-Jelisavcic, P. Cirkovic, J. Mamuzic

Slovak Republic

Bratislava

R. Astalos, P. Bartos, L. Batkova, T. Blazek, P. Federic, I. Sykora, S. Tokár, T. Ženiš

Kosice

J. Antos, D. Bruncko, E. Kladiva, P. Strizenec

Slovenia

Ljubljana

V. Cindro, M. Deliyergiyev, A. Filipčič, A. Gorišek, B.P. Kerševan, G. Kramberger, I. Mandić, B. Maček,
M. Mikuš, A. Tykhonov

South Africa

Cape Town

A. Hamilton

Johannesburg

M. Arousseau, S. Ballestrero, E. Castaneda-Miranda, S. Yacoob

Witwatersrand

K. Bristow, G.D. Carrillo-Montoya, X. Chen, Y. Huang, K.J.C. Leney, B.R. Mellado Garcia, X. Ruan,
O.E. Vickey Boeriu, T. Vickey

Spain

Barcelona

M. Bosman, R. Caminal Armadans, M.P. Casado, M. Cavalli-Sforza, M.C. Conidi, A. Cortes-Gonzalez,
T. Farooque, P. Francavilla, V. Giangioffe, G. Gonzalez Parra, S. Grinstein, A. Juste Rozas, I. Korolkov,
E. Le Menedeu, M. Martinez, L.M. Mir, J. Montejo Berlingen, A. Pacheco Pages, C. Padilla Aranda,
X. Portell Bueso, I. Riu, F. Rubbo, V. Sorin, A. Succurro, S. Tsiskaridze

Granada

J.A. Aguilar-Saavedra

Madrid UA

V. Arnal, F. Barreiro, J. Cantero, H. De la Torre, J. Del Peso, C. Glasman, J. Llorente Merino, J. Terron

Valencia

S. Cabrera Urbán, V. Castillo Gimenez, M.J. Costa, F. Fassi, A. Ferrer, L. Fiorini, J. Fuster, C. García, J.E. García Navarro, S. González de la Hoz, Y. Hernández Jiménez, E. Higón-Rodríguez, A. Irlés Quiles, M. Kaci, M. King, C. Lacasta, V.R. Lacuesta, L. March, S. Marti-Garcia, V.A. Mitsou, M. Miñano Moya, R. Moles-Valls, E. Oliver Garcia, S. Pedraza Lopez, M.T. Pérez García-Estañ, E. Romero Adam, E. Ros, J. Salt, V. Sanchez Martinez, U. Soldevila, J. Sánchez, E. Torró Pastor, A. Valero, E. Valladolid Gallego, J.A. Valls Ferrer, M. Villaplana Perez, M. Vos

Sweden

Lund

S.S. Bocchetta, L. Bryngemark, A. Floderus, A.D. Hawkins, V. Hedberg, G. Jarlskog, E. Lytken, B. Meirose, J.U. Mjörnmark, O. Smirnova, O. Viazlo, M. Wielers, T.P.A. Åkesson

Stockholm

Y. Abulaiti, K. Bendtz, O. Bessidskaia, C. Bohm, C. Clement, K. Gellerstedt, S. Hellman, K.E. Johansson, K. Jon-And, H. Khandanyan, H. Kim, P. Klimek, J. Lundberg, O. Lundberg, D.A. Milstead, T. Moa, S. Molander, A. Petridis, P. Plucinski, V. Rossetti, S.B. Silverstein, J. Sjölin, S. Strandberg, M. Tylmad, B. Åsman

Stockholm KTH

J. Jovicevic, E.S. Kuwertz, B. Lund-Jensen, A.K. Morley, J. Strandberg

Uppsala

R. Brenner, C.P. Buszello, E. Coniavitis, T. Ekelof, M. Ellert, A. Ferrari, C. Isaksson, A. Madsen, D. Pelikan

Switzerland

Bern

M. Agustoni, L.S. Ancu, H.P. Beck, A. Cervelli, A. Ereditato, V. Gallo, S. Haug, T. Kruker, L.F. Marti, B. Schneider, F.G. Sciacca, S.A. Stucci, M.S. Weber

CERN

R. Abreu, M. Aleksa, C. Anastopoulos, N. Andari, G. Avolio, M.A. Baak, M. Backes, M. Battistin, O. Beltramello, M. Bianco, J. Boyd, S. Campana, M.D.M. Capeans Garrido, T. Carli, A. Catinaccio, A. Cattai, M. Cerv, C.A. Chavez Barajas, J.T. Childers, A. Dell'Acqua, A. Di Girolamo, B. Di Girolamo, F. Dittus, D. Dobos, J. Dopke, A. Dudarev, M. Dührssen, N. Ellis, M. Elsing, G. Facini, P. Farthouat, P. Fassnacht, S. Feigl, S. Fernandez Perez, S. Franchino, D. Francis, D. Froidevaux, V. Garonne, M. Ghibaudi, F. Gianotti, D. Gillberg, J. Godlewski, L. Goossens, B. Gorini, H.M. Gray, S. Haas, M. Hauschild, R.J. Hawkings, M. Heller, C. Helsens, A.M. Henriques Correia, L. Hervas, A. Hoecker, Z. Hubacek, M. Huhtinen, M.R. Jaekel, H. Jansen, P. Jenni, M. Joos, R.M. Jungst, M. Kaneda, T. Klioutchnikova, A. Krasznahorkay, K. Lantzsch, M. Lassnig, G. Lehmann Miotto, B. Lenzi, D. Macina, S. Malyukov, B. Mandelli, L. Mapelli, B. Martin, A. Messina, J. Meyer, G. Mornacchi, A.M. Nairz, Y. Nakahama, G. Negri, M. Nessi, M. Nordberg, C.C. Ohm, S. Palestini, T. Pauly, H. Pernegger, B.A. Petersen, K. Peters, K. Pommès, M.E. Pozo Astigarraga, S. Prasad, M. Raymond, C. Rembser, L. Rodrigues, S. Roe, A. Salzburger, D.O. Savu, T. Scanlon, S. Schlenker, K. Schmieden, C. Serfon, A. Sfyrla, A.D. Sicoe, C.A. Solans, G. Spigo, R. Spiwoks, G.A. Stewart, F.A. Teischinger, H. Ten Kate, L. Tremblet, A. Tricoli, C. Tsarouchas, G. Unal, W. Vandelli, D. van der Ster, N. van Eldik, M.C. van Woerden, R. Vigne, R. Voss, R. Vuillermet, P.S. Wells, T. Wengler, S. Wenig, P. Werner, H.G. Wilkens, J. Wotschack, C.J.S. Young, L. Zwalinski

Geneva

G. Alexandre, G. Barone, P.J. Bell, W.H. Bell, E. Benhar Noccioli, J. Bilbao De Mendizabal, F. Bucci, R. Camacho Toro, A. Clark, D. della Volpe, C. Doglioni, D. Ferrere, S. Gadomski, S. Gonzalez-Sevilla, M.P. Goulette, J. Gramling, F. Guescini, G. Iacobucci, A. Katre, A. La Rosa, B. Martin dit Latour,

P. Mermod, A. Miucci, C. Mora Herrera, D. Muenstermann, S. Nektarijevic, M. Nessi, K. Nikolics, A. Picazio, M. Pohl, G. Pásztor, K. Rosbach, S. Vallecorsa, X. Wu

Taiwan

Taipei AS

J. Abdallah, S. Hou, D.O. Jamin, C.A. Lee, S.C. Lee, B. Li, S.C. Lin, B. Liu, D. Liu, F. Lo Sterzo, R. Mazini, D.A. Soh, P.K. Teng, C. Wang, S.M. Wang, Z. Weng, L. Zhang

Turkey

Ankara

O. Cakir, A.K. Ciftci, R. Ciftci, H. Duran Yildiz, S. Kuday

Bogazici Istanbul

M. Arik, S. Istin, V.E. Ozcan

Dogus

S.A. Cetin

Gazi

M. Yilmaz

Gaziantep

A. Beddall, A.J. Beddall, A. Bingul

TOBB

S. Sultansoy

Turkey Atomic EA

United Kingdom

Birmingham

B.M.M. Allbrooke, L. Aperio Bella, H.S. Bansil, J. Bracinik, D.G. Charlton, A.S. Chisholm, A.C. Daniells, P.J.W. Faulkner, C.M. Hawkes, S.J. Head, S.J. Hillier, T. Mclaughlan, R.D. Mudd, J.A. Murillo Quijada, P.R. Newman, K. Nikolopoulos, J.D. Palmer, M. Slater, R.J. Staley, J.P. Thomas, P.D. Thompson, P.M. Watkins, A.T. Watson, M.F. Watson, J.A. Wilson

Cambridge

S. Ask, N. Barlow, J.R. Batley, F.M. Brochu, W. Buttinger, J.R. Carter, J.D. Chapman, S.T. French, J.A. Frost, T.P.S. Gillam, J.C. Hill, S. Kaneti, T.J. Khoo, C.G. Lester, V. Moeller, T. Mueller, M.A. Parker, D. Robinson, T. Sandoval, M. Thomson, C.P. Ward, S. Williams

Edinburgh

W. Bhimji, T.M. Bristow, P.J. Clark, C. Debenedetti, N.C. Edwards, F.M. Garay Walls, R.D. Harrington, A. Korn, C. Leonidopoulos, V.J. Martin, B.J. O'Brien, S.A. Olivares Pino, M. Proissl, A. Schaelicke, K.E. Selbach, B.H. Smart, A. Washbrook, B.M. Wynne

Glasgow

S.E. Allwood-Spiers, R.L. Bates, D. Britton, A.G. Buckley, P. Bussey, C.M. Buttar, A. Buzatu, C. Collins-Tooth, S. D'Auria, T. Doherty, A.T. Doyle, S. Ferrag, J. Ferrando, D.E. Ferreira de Lima, A. Gemmell, U. Gul, N.G. Gutierrez Ortiz, D. Kar, A. Knue, A. Moraes, V. O'Shea, C. Oropeza Barrera, D. Quilty, T. Ravenscroft, A. Robson, R.D. St. Denis, G. Steele, A.S. Thompson, K. Wraight, M. Wright

Lancaster

L.J. Allison, A.E. Barton, G. Borissov, E.V. Bouhova-Thacker, J.R. Catmore, A. Chilingarov, W.J. Dearnaley, H. Fox, K. Grimm, R.C.W. Henderson, G. Hughes, R.W.L. Jones, V. Kartvelishvili, R.E. Long, P.A. Love, H.J. Maddocks, M. Smizanska, J. Walder

Liverpool

P.P. Allport, A.C. Bundock, S. Burdin, M. D'Onofrio, P. Dervan, C.B. Gwilliam, H.S. Hayward, J.N. Jackson, M. Jackson, T.J. Jones, B.T. King, M. Klein, U. Klein, J. Kretzschmar, P. Laycock, A. Lehan,

S. Mahmoud, S.J. Maxfield, A. Mehta, S. Migas, J. Price, Y.J. Schnellbach, G. Sellers, J.H. Vossebeld, P. Waller

London QMUL

M. Bona, L. Cerrito, K. Ellis, G. Fletcher, J.R. Goddard, R. Hickling, M.P.J. Landon, S.L. Lloyd, T. Macey, J.D. Morris, E. Piccaro, E. Rizvi, G. Salamanna, G. Snidero, M. Teixeira Dias Castanheira

London RHBNC

M.A. Alam, T. Berry, V. Boisvert, T. Brooks, R. Cantrill, I.A. Connelly, N.J. Cooper-Smith, G. Cowan, L. Duguid, C.A. Edwards, S. George, S.M. Gibson, R. Gonçalo, B. Green, J.J. Kempster, J.G. Panduro Vazquez, Fr. Pastore, M. Rose, G. Savage, F. Spanò, P. Teixeira-Dias, J. Thomas-wilsker

London UC

S. Baker, P. Bernat, S.P. Bieniek, J.M. Butterworth, M. Campanelli, D. Casadei, R.T. Chislett, I.A. Christidi, B.D. Cooper, G.J. Crone, A.R. Davison, E. Dobson, C. Gutschow, G.G. Hesketh, E. Jansen, N. Konstantinidis, L. Lambourne, A.C. Martyniuk, J.A. Mcfayden, M. Nash, E. Nurse, M.I. Ochoa, A.D. Pilkington, R. Prabhu, P. Sherwood, B. Simmons, C. Taylor, D.R. Wardrope, B.M. Waugh, P.A. Wijeratne

Manchester

J. Almond, M. Borri, G. Brown, V. Chavda, B.E. Cox, C. Da Via, A. Forti, J. Howarth, J.M. Iturbe Ponce, K.D. Joshi, J.A. Klinger, F.K. Loebinger, J. Masik, T.J. Neep, A. Oh, M. Owen, J.R. Pater, D. Price, J.E.M. Robinson, L. Tomlinson, S. Watts, S. Webb, M.J. Woudstra, T.R. Wyatt, U.K. Yang

Oxford

R. Apolle, A.J. Barr, K. Behr, C.R. Boddy, R.M. Buckingham, A.M. Cooper-Sarkar, M. Crispin Ortuzar, A. Dafinca, E. Davies, E.J. Gallas, S. Gupta, C. Gwenlan, D. Hall, C.P. Hays, J. Henderson, J. Howard, T.B. Huffman, C. Issever, R.S.B. King, L.A. Kogan, A. Larner, A. Lewis, Z. Liang, S.S.A. Livermore, C. Mattravers, R.B. Nickerson, K. Pachal, A. Pinder, A. Robichaud-Veronneau, N.C. Ryder, C. Sawyer, D. Short, J.C.-L. Tseng, T. Vickey, G.H.A. Viehhauser, A.R. Weidberg, J. Zhong

RAL

T. Adye, R. Apolle, J.T. Baines, B.M. Barnett, I.P. Brawn, S. Burke, E. Davies, A. Dewhurst, D. Emeliyanov, B.J. Gallop, C.N.P. Gee, S.J. Haywood, J. Kirk, S. Martin-Haugh, C. Mattravers, S.J. McMahon, R.P. Middleton, W.J. Murray, M. Nash, P.W. Phillips, W. Qian, D.P.C. Sankey, W.G. Scott, M.J. Siyad, S. Taghavirad, F.J. Wickens, M. Wielers

Sheffield

D. Costanzo, T. Cuhadar Donszelmann, I. Dawson, G.T. Fletcher, M.C. Hodgkinson, P. Hodgson, P. Johansson, E.V. Korolkova, B. Lopez Paredes, P.S. Miyagawa, S. Owen, E. Paganis, K. Suruliz, D.R. Tovey, A. Tua

Sussex

V. Bartsch, A. Cerri, A. De Santo, Z.J. Grout, C.J. Potter, A. Rose, F. Salvatore, I. Santoyo Castillo, C.Y. Shehu, M.R. Sutton, I. Vivarelli

Warwick

S.M. Farrington, P.F. Harrison, M. Janus, C. Jeske, G. Jones, T.A. Martin, W.J. Murray, E. Pianori

United States of America

Albany

W. Edson, J. Ernst, S. Guindon, V. Jain

Argonne

J.T. Anderson, L. Asquith, B. Auerbach, R.E. Blair, S. Chekanov, G. Drake, E.J. Feng, A.T. Goshaw, T. LeCompte, J. Love, D. Malon, D.H. Nguyen, L. Nodulman, A. Paramonov, L.E. Price, J. Proudfoot, R.W. Stanek, P. van Gemmeren, A. Vaniachine, R. Yoshida, J. Zhang

Arizona

E. Cheu, K.A. Johns, V. Kaushik, C.L. Lampen, W. Lampl, X. Lei, R. Leone, P. Loch, R. Nayyar,
F. O'grady, J.P. Rutherford, F. Rühr, M.A. Shupe, E.W. Varnes, J. Veatch

Arlington UT

A. Brandt, D. Côté, S. Darmora, K. De, A. Farbin, J. Griffiths, H.K. Hadavand, L. Heelan, M. Maeno,
P. Nilsson, N. Ozturk, R. Pravahan, E. Sarkisyan-Grinbaum, M. Sosebee, A.R. Stradling, G. Usai,
A. Vartapetian, A. White, J. Yu

Berkeley LBNL

A.M. Bach, R.M. Barnett, J. Beringer, J. Biesiada, G. Brandt, J. Brosamer, P. Calafiura, L.M. Caminada,
F. Cerutti, A. Ciocio, R.N. Clarke, M. Cooke, K. Copic, S. Dube, K. Einsweiler, M. Garcia-Sciveres,
C. Haber, M. Hance, B. Heinemann, I. Hinchliffe, T.R. Holmes, M. Hurwitz, L. Jeanty, W. Lavrijsen,
C. Leggett, P. Loscutoff, Z. Marshall, A. Ovcharova, S. Pagan Griso, K. Potamianos, A. Pranko,
D.R. Quarrie, M. Shapiro, L.A. Skinnari, A. Sood, M.J. Tibbetts, V. Tsulaia, D. Varouchas, J. Virzi,
H. Wang, D.R. Yu

Boston

S.P. Ahlen, C. Bernard, K.M. Black, J.M. Butler, L. Dell'Asta, L. Helary, M. Kruskal, B.A. Long,
J.T. Shank, Z. Yan, S. Youssef

Brandeis

S. Aefsky, C. Amelung, G. Amundsen, G. Artoni, J.R. Bensinger, L. Bianchini, C. Blocker, L. Coffey,
R.K. Daya-Ishmukhametova, E.A. Fitzgerald, S. Gozpinar, D. Pomeroy, G. Sciolla, A. Venturini,
S. Zambito, K. Zengel

Brookhaven BNL

D.L. Adams, K. Assamagan, M. Begel, H. Chen, V. Chernyatin, R. Debbé, M. Ernst, T. Gadfort,
H.A. Gordon, X. Hu, A. Klimentov, A. Kravchenko, F. Lanni, D. Lynn, T. Maeno, H. Ma, J. Metcalfe,
E. Mountricha, P. Nevski, H. Okawa, D. Oliveira Damazio, F. Paige, S. Panitkin, M.-A. Pleier,
V. Polychronakos, S. Protopopescu, M. Purohit, S. Rajagopalan, G. Redlinger, J. Schovancova,
S. Snyder, P. Steinberg, H. Takai, M.C. Tamsett, N. Triplett, A. Undrus, T. Wenaus, S. Ye, A. Zaytsev

Chicago

J. Alison, K.J. Anderson, M. Bogdan, A. Boveia, F. Canelli, Y. Cheng, M. Fiascaris, R.W. Gardner,
Y. Ilchenko, A. Kapliy, H.L. Li, S. Meehan, C. Melachrinou, F.S. Merritt, C. Meyer, D.W. Miller,
Y. Okumura, P.U.E. Onyisi, M.J. Oreglia, B. Penning, J.E. Pilcher, M.J. Shochet, L. Tompkins, I. Vukotic,
J.S. Webster

Columbia

A. Altheimer, T. Andeen, A. Angerami, T. Bain, G. Brooijmans, Y. Chen, B. Cole, J. Guo, D. Hu,
E.W. Hughes, S. Mohapatra, N. Nikiforou, J.A. Parsons, D.V. Perepelitsa, M.I. Scherzer,
E.N. Thompson, F. Tian, P.M. Tuts, D. Urbaniec, E. Wulf, L. Zhou

Dallas SMU

T. Cao, A. Firan, J. Hoffman, S. Kama, R. Kehoe, A.S. Randle-Conde, S.J. Sekula, R. Stroynowski,
H. Wang, J. Ye

Dallas UT

J.M. Izen, X. Lou, H. Namasivayam, K. Reeves

Duke

A.T.H. Arce, D.P. Benjamin, A. Bocci, B. Cerio, K.D. Finelli, E. Kajomovitz, A. Kotwal, M.C. Kruse,
S. Li, M. Liu, S.H. Oh, C.S. Pollard, C. Wang

Hampton

V.I. Vassilakopoulos

Harvard

J. Barreiro Guimarães da Costa, A. Belloni, B. Butler, P. Catastini, G. Conti, M. Franklin, J. Huth, D. Lopez Mateos, K.M. Mercurio, C. Mills, M. Morii, H.P. Skottowe, W.R. Spearman, A.L. Yen, G. Zevi della Porta

Indiana

S. Brunet, H. Evans, P. Gagnon, F. Luehring, H. Ogren, J. Penwell, J. Poveda, D. Whittington, D. Zieminska

Iowa

D. Cinca, R.P. Gandrajula, U. Mallik, R. Mandrysch, N. Morange, Y. Pylypchenko, R. Zaidan

Iowa State

C. Chen, J. Cochran, F. De Lorenzi, F. Dudziak, N. Krumnack, S. Prell, A. Ruiz-Martinez, S. Shrestha, K. Yamamoto

Louisiana Tech

C. Bernius, Z.D. Greenwood, D.K. Jana, L. Sawyer, A. Sircar, R. Subramaniam, M.C. Tamsett

MIT

F.E. Taylor

Massachusetts

M. Bellomo, B. Brau, G. Colon, C. Dallapiccola, A. Meade, E.J.W. Moyses, P. Pais, E. Pueschel, T. Varol, D. Ventura, S. Willcoq

Michigan

A.J. Armbruster, M.A. Chelstowska, H.C. Cheng, T. Dai, E.B. Diehl, J. Dubbert, H. Feng, C. Ferretti, S. Goldfarb, D. Harper, D. Levin, L. Liu, X. Li, J.D. Long, A. McCarn, S.P. Mc Kee, H.A. Neal, N. Panikashvili, J. Qian, J. Searcy, A. Wilson, Y. Wu, L. Xu, J.M. Yu, D. Zhang, B. Zhou, J. Zhu

Michigan SU

B. Alvarez Gonzalez, G. Arabidze, R. Brock, S. Caughron, D. Edmunds, I. Ermoline, P. Ge, G. Halladjian, R. Hauser, D. Hayden, J. Huston, J. Koll, P. Laurens, J.T. Linnemann, B. Martin, B.G. Pope, B.D. Schoenrock, R. Schwienhorst, D. Shooltz, H.J. Stelzer, D. Ta, K. Tollefson, P. True, H. Zhang

NYU New York

J.B. Beacham, B. Budick, K. Cranmer, A. Haas, L. Heinrich, L. Hooft van Huysduynen, B. Kaplan, K. Karthik, R. Konoplich, S. Kreiss, G.H. Lewis, A.I. Mincer, P. Nemethy, R.M. Neves, K. Prokofiev

New Mexico

I. Gorelov, M.R. Hoferkamp, S.C. Seidel, K. Toms, R. Wang

Northern Illinois

B. Burghgrave, R. Calkins, D. Chakraborty, S. Cole, J.G. Rocha de Lima, C. Suhr, A. Yurkewicz, V. Zutshi

Ohio SU

M.J. Fisher, K.K. Gan, R. Ishmukhametov, H. Kagan, R.D. Kass, H. Merritt, J. Moss, A. Nagarkar, D.T. Pignotti, Y. Yang

Oklahoma

B. Abbott, P. Gutierrez, A. Marzin, R. Meera-Lebbai, S. Norberg, M. Saleem, H. Severini, P. Skubic, J. Snow, M. Strauss

Oklahoma SU

B. Abi, A. Khanov, F. Rizatdinova, D. Sidorov, J. Yu

Oregon

J.E. Brau, E. Brost, S. Majewski, C.T. Potter, E. Ptacek, P. Radloff, A. Reinsch, M. Shamim, N.B. Sinev, D.M. Strom, E. Torrence, F. Winklmeier

Pennsylvania

K. Brendlinger, J. Degenhardt, S. Fratina, S. Heim, E. Hines, T.M. Hong, B. Jackson, J. Kroll, J. Kunkle, C.M. Lester, E. Lipeles, D. Olivito, R. Ospanov, J. Saxon, D. Schaefer, J. Stahlman, E. Thomson, A.N. Tuna, H.H. Williams

Pittsburgh

R.M. Bianchi, J. Boudreau, C. Escobar, J. Mueller, K. Sapp, J. Su, R. Yoosoofmiya

SLAC

I. Aracena, J. Backus Mayes, T. Barklow, R. Bartoldus, H.S. Bawa, J.E. Black, J.G. Cogan, T. Eifert, B.G. Fulsom, Y.S. Gao, N. Garelli, P. Grenier, P. Hansson, M. Kagan, M. Kocian, T. Koi, A.J. Lowe, C. Malone, R. Mount, T.K. Nelson, G. Piacquadio, A. Salnikov, A. Schwartzman, D. Silverstein, E. Strauss, D. Su, M. Swiatlowski, M. Wittgen, C. Young

Santa Cruz UC

A.A. Grillo, A. Kuhl, A.M. Litke, W.S. Lockman, P.M. Manning, J. Nielsen, R. Reece, H.F.W. Sadrozinski, B.A. Schumm, A. Seiden

Seattle Washington

M. Beckingham, D. Blackburn, A. Coccaro, A.G. Goussiou, O.M. Harris, S.-C. Hsu, J.S. Keller, H.J. Lubatti, M. Marx, N. Rompotis, R. Rosten, J. Rothberg, P.H. Sales De Bruin, M.S. Twomey, M. Verducci, G. Watts

Stony Brook

A. Ahmad, C.P. Bee, A. Campoverde, K. Chen, V. Grassi, J.A. Gray, J. Hobbs, J. Jia, H. Li, B.E. Lindquist, P. Mastrandrea, R.L. McCarthy, D. Puldon, S.K. Radhakrishnan, M. Rijssenbeek, R.D. Schamberger, J. Stupak, D. Tsybychev, A. Zaman

Tufts

P.H. Beauchemin, S. Hamilton, E. Meoni, A. Napier, K. Sliwa, J. Wetter

UC Irvine

S.M. Batraneanu, A. Corso-Radu, S. Farrell, D. Gerbaudo, S. Kolos, A.J. Lankford, L. Magnoni, A.S. Mete, R. Murillo Garcia, A. Nelson, K. Rao, M. Relich, D.A. Scannicchio, M. Schernau, C.O. Shimmin, I. Soloviev, A. Taffard, B. Toggerson, G. Unel, D. Whiteson, N. Zhou

UI Urbana

M. Atkinson, A. Basye, N. Benekos, V. Cavaliere, P. Chang, J. Coggeshall, D. Errede, S. Errede, K. Lie, T.M. Liss, M.S. Neubauer, I. Vichou

Wisconsin

Sw. Banerjee, L.R. Flores Castillo, A.S. Hard, H. Ji, X. Ju, L. Kashif, A. Kruse, Y. Ming, Y.B. Pan, W. Wiedenmann, S.L. Wu, H. Yang, G. Zobernig

Yale

J. Adelman, O.K. Baker, S. Bedikian, C. Cuenca Almenar, J. Cummings, Z. Czyczula, S. Demers, J. Erdmann, F. Garberson, T. Golling, D. Guest, A. Henrichs, E. Ideal, T. Lagouri, L. Lee, A.G. Leister, A. Loginov, P. Tipton, R. Wall, B. Walsh, X. Wang

