Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System - EF Tracking Amendment

The ATLAS Collaboration

Reference: v1.2

Created: March 1, 2022
Last modified: March 1, 2022
Prepared by: The ATLAS Collaboration

CERN-LHCC-2022-004 / ATLAS-TDR-029-ADD-1
01/03/2022

Abstract
This Technical Design Report Amendment describes revised plans for Event Filter Tracking in the upgrade of the ATLAS Trigger and Data Acquisition system for the High Luminosity LHC. The motivation to change the baseline for Event Filter Tracking is explained. Next, a description of the requirements for Event Filter Tracking and the definition of the proposed baseline to meet these requirements are presented. Several demonstrations using various hardware and software are reported in support of this proposal. Finally, the organization and resources needed to deliver the new baseline are set out.

The ATLAS Collaboration

Argentina
Departamento de Física (FCEN) and IFIBA, Universidad de Buenos Aires and CONICET, Buenos Aires, Argentina
Instituto de Física La Plata, Universidad Nacional de La Plata and CONICET, La Plata, Argentina

Australia
Department of Physics, University of Adelaide, Adelaide, Australia
School of Physics, University of Melbourne, Victoria, Australia
School of Physics, University of Sydney, Sydney, Australia

Austria
Institut für Astro- und Teilchenphysik, Leopold-Franzens-Universität, Innsbruck, Austria

Azerbaijan
Institute of Physics, Azerbaijan Academy of Sciences, Baku, Azerbaijan

Belarus
B.I. Stepanov Institute of Physics, National Academy of Sciences of Belarus, Minsk, Belarus
Research Institute for Nuclear Problems of Byelorussian State University, Minsk, Belarus

Brazil
Departamento de Engenharia Elétrica, Universidade Federal de Juiz de Fora (UFJF), Juiz de Fora, Brazil
Universidade Federal do Rio De Janeiro COPPE/EE/IF, Rio de Janeiro, Brazil
Instituto de Física, Universidade de São Paulo, São Paulo, Brazil
Universidade Federal de São João del Rei (UFSJ), São João del Rei, Brazil

Canada
Department of Physics, University of Alberta, Edmonton AB, Canada
Department of Physics, Carleton University, Ottawa ON, Canada
Department of Physics, McGill University, Montreal QC, Canada
Group of Particle Physics, University of Montreal, Montreal QC, Canada
Department of Physics, Simon Fraser University, Burnaby BC, Canada
Department of Physics, University of Toronto, Toronto ON, Canada
TRIUMF, Vancouver BC, Canada
Department of Physics and Astronomy, York University, Toronto ON, Canada
Department of Physics, University of British Columbia, Vancouver BC, Canada
Department of Physics and Astronomy, University of Victoria, Victoria BC, Canada

Chile
Departamento de Física, Pontificia Universidad Católica de Chile, Santiago, Chile
Universidad de la Serena, La Serena, Chile
Universidad Andres Bello, Department of Physics, Santiago, Chile

Instituto de Alta Investigación, Universidad de Tarapacá, Arica, Chile
Departamento de Física, Universidad Técnica Federico Santa María, Valparaíso, Chile

China
Institute of High Energy Physics, Chinese Academy of Sciences, Beijing, China
Physics Department, Tsinghua University, Beijing, China
Department of Physics, Nanjing University, Nanjing, China
University of Chinese Academy of Science (UCAS), Beijing, China
Department of Modern Physics and State Key Laboratory of Particle Detection and Electronics, University of Science and Technology of China, Hefei, China
Institute of Frontier and Interdisciplinary Science and Key Laboratory of Particle Physics and Particle Irradiation (MOE), Shandong University, Qingdao, China
School of Physics and Astronomy, Shanghai Jiao Tong University, Key Laboratory for Particle Astrophysics and Cosmology (MOE), SKLPPC, Shanghai, China
Tsung-Dao Lee Institute, Shanghai, China
Department of Physics, Chinese University of Hong Kong, Shatin, N.T., Hong Kong, China
Department of Physics, University of Hong Kong, Hong Kong, China
Department of Physics and Institute for Advanced Study, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China

Colombia
Facultad de Ciencias y Centro de Investigaciónes, Universidad Antonio Nariño, Bogotá, Colombia
Departamento de Física, Universidad Nacional de Colombia, Bogotá, Colombia

Czech Republic
Palacký University, Joint Laboratory of Optics, Olomouc, Czech Republic
Institute of Physics of the Czech Academy of Sciences, Prague, Czech Republic
Czech Technical University in Prague, Prague, Czech Republic
Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic

Denmark
Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark

France
LAPP, Univ. Savoie Mont Blanc, CNRS/IN2P3, Annecy , France
LPC, Université Clermont Auvergne, CNRS/IN2P3, Clermont-Ferrand, France
LPSC, Université Grenoble Alpes, CNRS/IN2P3, Grenoble INP, Grenoble, France
IJCLab, Université Paris-Saclay, CNRS/IN2P3, 91405, Orsay, France
CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France
LPNHE, Sorbonne Université, Université de Paris, CNRS/IN2P3, Paris, France
IRFU, CEA, Université Paris-Saclay, Gif-sur-Yvette, France

Georgia
E. Andronikashvili Institute of Physics, Iv. Javakhishvili Tbilisi State University, Tbilisi, Georgia

High Energy Physics Institute, Tbilisi State University, Tbilisi, Georgia

Germany
Institut für Physik, Humboldt Universität zu Berlin, Berlin, Germany
Physikalisches Institut, Universität Bonn, Bonn, Germany
Deutsches Elektronen-Synchrotron DESY, Hamburg and Zeuthen, Germany
Lehrstuhl für Experimentelle Physik IV, Technische Universität Dortmund, Dortmund, Germany
Institut für Kern- und Teilchenphysik, Technische Universität Dresden, Dresden, Germany
Physikalisches Institut, Albert-Ludwigs-Universität Freiburg, Freiburg, Germany
II. Physikalisches Institut, Georg-August-Universität Göttingen, Göttingen, Germany
II. Physikalisches Institut, Justus-Liebig-Universität Giessen, Giessen, Germany
Kirchhoff-Institut für Physik, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany
Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany
Institut für Physik, Universität Mainz, Mainz, Germany
Fakultät für Physik, Ludwig-Maximilians-Universität München, München, Germany
Max-Planck-Institut für Physik (Werner-Heisenberg-Institut), München, Germany
Department Physik, Universität Siegen, Siegen, Germany
Fakultät für Physik und Astronomie, Julius-Maximilians-Universität Würzburg, Würzburg, Germany
Fakultät für Mathematik und Naturwissenschaften, Fachgruppe Physik, Bergische Universität Wuppertal, Wuppertal, Germany

Greece
Physics Department, National and Kapodistrian University of Athens, Athens, Greece
Physics Department, National Technical University of Athens, Zografou, Greece
National Centre for Scientific Research "Demokritos", Agia Paraskevi, Greece
Department of Physics, Aristotle University of Thessaloniki, Thessaloniki, Greece

Israel
Department of Physics, Technion, Israel Institute of Technology, Haifa, Israel
Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel Aviv, Israel
Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot, Israel

Italy
Dipartimento di Fisica e Astronomia A. Righi, Università di Bologna, Bologna, Italy
INFN Sezione di Bologna, Italy
Dipartimento di Fisica, Università della Calabria, Rende, Italy
INFN Gruppo Collegato di Cosenza, Laboratori Nazionali di Frascati, Italy
INFN e Laboratori Nazionali di Frascati, Frascati, Italy
Dipartimento di Fisica, Università di Genova, Genova, Italy
INFN Sezione di Genova, Italy

INFN Gruppo Collegato di Udine, Sezione di Trieste, Udine, Italy
ICTP, Trieste, Italy
Dipartimento Politecnico di Ingegneria e Architettura, Università di Udine, Udine, Italy
INFN Sezione di Lecce, Italy
Dipartimento di Matematica e Fisica, Università del Salento, Lecce, Italy
INFN Sezione di Milano, Italy
Dipartimento di Fisica, Università di Milano, Milano, Italy
INFN Sezione di Napoli, Italy
Dipartimento di Fisica, Università di Napoli, Napoli, Italy
INFN Sezione di Pavia, Italy
Dipartimento di Fisica, Università di Pavia, Pavia, Italy
INFN Sezione di Pisa, Italy
Dipartimento di Fisica E. Fermi, Università di Pisa, Pisa, Italy
INFN Sezione di Roma, Italy
Dipartimento di Fisica, Sapienza Università di Roma, Roma, Italy
INFN Sezione di Roma Tor Vergata, Italy
Dipartimento di Fisica, Università di Roma Tor Vergata, Roma, Italy
INFN Sezione di Roma Tre, Italy
Dipartimento di Matematica e Fisica, Università Roma Tre, Roma, Italy
INFN-TIFPA, Italy
Università degli Studi di Trento, Trento, Italy

Japan
KEK, High Energy Accelerator Research Organization, Tsukuba, Japan
Graduate School of Science, Kobe University, Kobe, Japan
Faculty of Science, Kyoto University, Kyoto, Japan
Kyoto University of Education, Kyoto, Japan
Research Center for Advanced Particle Physics and Department of Physics, Kyushu University, Fukuoka , Japan
Graduate School of Science and Kobayashi-Maskawa Institute, Nagoya University, Nagoya, Japan
Ochanomizu University, Otsuka, Bunkyo-ku, Tokyo, Japan
Graduate School of Science, Osaka University, Osaka, Japan
Department of Physics, Shinshu University, Nagano, Japan
International Center for Elementary Particle Physics and Department of Physics, University of Tokyo, Tokyo, Japan
Department of Physics, Tokyo Institute of Technology, Tokyo, Japan
Division of Physics and Tomonaga Center for the History of the Universe, Faculty of Pure and Applied Sciences, University of Tsukuba, Tsukuba, Japan
Waseda University, Tokyo, Japan

Morocco
Faculté des Sciences Ain Chock, Réseau Universitaire de Physique des Hautes Energies - Université Hassan II, Casablanca, Morocco

Faculté des Sciences, Université Ibn-Tofail, Kénitra, Morocco
Faculté des Sciences Semlalia, Université Cadi Ayyad, LPHEA-Marrakech, Morocco
LPMR, Faculté des Sciences, Université Mohamed Premier, Oujda, Morocco
Faculté des sciences, Université Mohammed V, Rabat, Morocco
Mohammed VI Polytechnic University, Ben Guerir, Morocco

Netherlands
Institute for Mathematics, Astrophysics and Particle Physics, Radboud University/Nikhef, Nijmegen, Netherlands
Nikhef National Institute for Subatomic Physics and University of Amsterdam, Amsterdam, Netherlands

Norway
Department for Physics and Technology, University of Bergen, Bergen, Norway
Department of Physics, University of Oslo, Oslo, Norway

Poland
AGH University of Science and Technology, Faculty of Physics and Applied Computer Science, Krakow, Poland
Marian Smoluchowski Institute of Physics, Jagiellonian University, Krakow, Poland
Institute of Nuclear Physics Polish Academy of Sciences, Krakow, Poland

Portugal
Laboratório de Instrumentação e Física Experimental de Partículas - LIP, Lisboa, Portugal
Departamento de Física, Faculdade de Ciências, Universidade de Lisboa, Lisboa, Portugal
Departamento de Física, Universidade de Coimbra, Coimbra, Portugal
Centro de Física Nuclear da Universidade de Lisboa, Lisboa, Portugal
Departamento de Física, Universidade do Minho, Braga, Portugal
Departamento de Física Teórica y del Cosmos, Universidad de Granada, Granada (Spain), Spain
Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal

Romania
Transilvania University of Brasov, Brasov, Romania
Horia Hulubei National Institute of Physics and Nuclear Engineering, Bucharest, Romania
Department of Physics, Alexandru Ioan Cuza University of Iasi, Iasi, Romania
National Institute for Research and Development of Isotopic and Molecular Technologies, Physics Department, Cluj-Napoca, Romania
University Politehnica Bucharest, Bucharest, Romania
West University in Timisoara, Timisoara, Romania

Russia
Joint Institute for Nuclear Research, Dubna, Russia
P.N. Lebedev Physical Institute of the Russian Academy of Sciences, Moscow, Russia
National Research Nuclear University MEPhI, Moscow, Russia
D.V. Skobeltsyn Institute of Nuclear Physics, M.V. Lomonosov Moscow State University, Moscow, Russia

Budker Institute of Nuclear Physics and NSU, SB RAS, Novosibirsk, Russia
Novosibirsk State University Novosibirsk, Russia
Institute for High Energy Physics of the National Research Centre Kurchatov Institute, Protvino, Russia
Institute for Theoretical and Experimental Physics named by A.I. Alikhanov of National Research Centre "Kurchatov Institute", Moscow, Russia
Konstantinov Nuclear Physics Institute of National Research Centre "Kurchatov Institute", PNPI, St. Petersburg, Russia
Tomsk State University, Tomsk, Russia

Serbia
Institute of Physics, University of Belgrade, Belgrade, Serbia

Slovak Republic
Faculty of Mathematics, Physics and Informatics, Comenius University, Bratislava, Slovak Republic
Department of Subnuclear Physics, Institute of Experimental Physics of the Slovak Academy of Sciences, Kosice, Slovak Republic

Slovenia
Department of Experimental Particle Physics, Jožef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana, Slovenia

South Africa
Department of Physics, University of Cape Town, Cape Town, South Africa
iThemba Labs, Western Cape, South Africa
Department of Mechanical Engineering Science, University of Johannesburg, Johannesburg, South Africa
National Institute of Physics, University of the Philippines Diliman (Philippines), Philippines
University of South Africa, Department of Physics, Pretoria, South Africa
School of Physics, University of the Witwatersrand, Johannesburg, South Africa

Spain
Institut de Física d'Altes Energies (IFAE), Barcelona Institute of Science and Technology, Barcelona, Spain
Departamento de Física Teorica C-15 and CIAFF, Universidad Autónoma de Madrid, Madrid, Spain
Departamento de Física Teórica y del Cosmos, Universidad de Granada, Granada (Spain), Spain
Instituto de Física Corpuscular (IFIC), Centro Mixto Universidad de Valencia - CSIC, Valencia, Spain

Sweden
Department of Physics, Stockholm University, Sweden
Oskar Klein Centre, Stockholm, Sweden

Fysiska institutionen, Lunds universitet, Lund, Sweden
Department of Physics, Royal Institute of Technology, Stockholm, Sweden
Department of Physics and Astronomy, University of Uppsala, Uppsala, Sweden

Switzerland
Albert Einstein Center for Fundamental Physics and Laboratory for High Energy Physics, University of Bern, Bern, Switzerland
CERN, Geneva, Switzerland
Département de Physique Nucléaire et Corpusculaire, Université de Genève, Genève, Switzerland

Taiwan
Department of Physics, National Tsing Hua University, Hsinchu, Taiwan
Institute of Physics, Academia Sinica, Taipei, Taiwan

Turkey
Department of Physics, Ankara University, Ankara, Turkey
Istanbul Aydin University, Application and Research Center for Advanced Studies, Istanbul, Turkey
Division of Physics, TOBB University of Economics and Technology, Ankara, Turkey
Bahcesehir University, Faculty of Engineering and Natural Sciences, Istanbul, Turkey
Istanbul Bilgi University, Faculty of Engineering and Natural Sciences, Istanbul, Turkey
Department of Physics, Bogazici University, Istanbul, Turkey
Department of Physics Engineering, Gaziantep University, Gaziantep, Turkey

United Arab Emirates
New York University Abu Dhabi, Abu Dhabi, United Arab Emirates
United Arab Emirates University, Al Ain, United Arab Emirates
University of Sharjah, Sharjah, United Arab Emirates

United Kingdom
School of Physics and Astronomy, University of Birmingham, Birmingham, United Kingdom
Cavendish Laboratory, University of Cambridge, Cambridge, United Kingdom
SUPA - School of Physics and Astronomy, University of Edinburgh, Edinburgh, United Kingdom
SUPA - School of Physics and Astronomy, University of Glasgow, Glasgow, United Kingdom
Physics Department, Lancaster University, Lancaster, United Kingdom
Oliver Lodge Laboratory, University of Liverpool, Liverpool, United Kingdom
School of Physics and Astronomy, Queen Mary University of London, London, United Kingdom
Department of Physics, Royal Holloway University of London, Egham, United Kingdom
Department of Physics and Astronomy, University College London, London, United Kingdom
School of Physics and Astronomy, University of Manchester, Manchester, United Kingdom

Department of Physics, Oxford University, Oxford, United Kingdom
Particle Physics Department, Rutherford Appleton Laboratory, Didcot, United Kingdom
Department of Physics and Astronomy, University of Sheffield, Sheffield, United Kingdom
Department of Physics and Astronomy, University of Sussex, Brighton, United Kingdom
Department of Physics, University of Warwick, Coventry, United Kingdom

United States of America
High Energy Physics Division, Argonne National Laboratory, Argonne IL, United States of America
Department of Physics, University of Arizona, Tucson AZ, United States of America
Department of Physics, University of Texas at Arlington, Arlington TX, United States of America
Department of Physics, University of Texas at Austin, Austin TX, United States of America
Physics Division, Lawrence Berkeley National Laboratory and University of California, Berkeley CA, United States of America
Department of Physics, Boston University, Boston MA, United States of America
Department of Physics, Brandeis University, Waltham MA, United States of America
Physics Department, Brookhaven National Laboratory, Upton NY, United States of America
California State University, CA, United States of America
Enrico Fermi Institute, University of Chicago, Chicago IL, United States of America
Nevis Laboratory, Columbia University, Irvington NY, United States of America
Physics Department, Southern Methodist University, Dallas TX, United States of America
Physics Department, University of Texas at Dallas, Richardson TX, United States of America
Department of Physics, Duke University, Durham NC, United States of America
Laboratory for Particle Physics and Cosmology, Harvard University, Cambridge MA, United States of America
Department of Physics, Indiana University, Bloomington IN, United States of America
University of Iowa, Iowa City IA, United States of America
Department of Physics and Astronomy, Iowa State University, Ames IA, United States of America
Louisiana Tech University, Ruston LA, United States of America
Department of Physics, University of Massachusetts, Amherst MA, United States of America
Department of Physics, University of Michigan, Ann Arbor MI, United States of America
Department of Physics and Astronomy, Michigan State University, East Lansing MI, United States of America
Department of Physics and Astronomy, University of New Mexico, Albuquerque NM, United States of America
Department of Physics, Northern Illinois University, DeKalb IL, United States of America
Department of Physics, New York University, New York NY, United States of America
Ohio State University, Columbus OH, United States of America

Homer L. Dodge Department of Physics and Astronomy, University of Oklahoma, Norman OK, United States of America
Department of Physics, Oklahoma State University, Stillwater OK, United States of America
Institute for Fundamental Science, University of Oregon, Eugene, OR, United States of America
Department of Physics, University of Pennsylvania, Philadelphia PA, United States of America
Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh PA, United States of America
Santa Cruz Institute for Particle Physics, University of California Santa Cruz, Santa Cruz CA, United States of America
Department of Physics, University of Washington, Seattle WA, United States of America
SLAC National Accelerator Laboratory, Stanford CA, United States of America
Departments of Physics and Astronomy, Stony Brook University, Stony Brook NY, United States of America
Department of Physics and Astronomy, Tufts University, Medford MA, United States of America
Department of Physics and Astronomy, University of California Irvine, Irvine CA, United States of America
Department of Physics, University of Illinois, Urbana IL, United States of America
Department of Physics, University of Wisconsin, Madison WI, United States of America
Department of Physics, Yale University, New Haven CT, United States of America

Contents

Executive Summary

The Phase-II ATLAS Trigger and Data Acquisition System (TDAQ) Technical Design Report (TDR) [1] documented the strategy and the design of the ATLAS TDAQ system for the High Luminosity upgrade of the Large Hadron Collider (HL-LHC). This document amends the TDR and describes the proposed new baseline for the Event Filter (EF) Tracking project.

Recent advances in software tracking and heterogeneous computing systems (which integrate multiple types of computational units), design choices for the Inner Tracker (ITk), and the decision to drop the option of a low-latency L1 hardware track trigger fundamentally changed the assumptions behind the technology choice for EF Tracking in the TDR. Thus, a decision process was set up to consider software-only, heterogeneous commodity, and custom systems according to: technical feasibility, estimated tracking performance, operational procedures, opportunities for improvement, risks, and resource requirements. An independent ATLAS review committee recommended that TDAQ commit to a commercial solution for EF tracking and pursue accelerators to potentially mitigate risks related to power and cost.

The physics goals and consequent EF Tracking use-cases remain as documented in the TDR. High pile-up conditions at the HL-LHC present a significant challenge. Track and vertex reconstruction are key tools to address this challenge. A combination of regional tracking at 1 MHz and full-scan tracking at 150 kHz will be used to achieve the required rejection. The full-scan tracking rate has been increased with respect to the original TDR to accommodate newly identified potential requirements for missing transverse momentum, particle flow reconstruction, large-radius tracking (LRT) for long-lived particles (LLPs), B-physics triggers and Trigger-object Level Analyses (TLAs).

The proposed EF design is a flexible, heterogeneous commercial system consisting of CPU cores and possibly accelerators, to perform the EF reconstruction which includes the compute-intensive ITk track reconstruction. Demonstrators of tracking on several types of commodity hardware (CPU, without and with FPGA and GPU accelerators) provide confidence that the final system will meet all necessary specifications.

A management plan and organization breakdown structure is in place to deliver the fully commissioned EF Tracking system by the end of Long Shutdown 3. The CORE cost of the full EF farm including tracking needs is estimated to be 8.8 MCHF. After establishing the requirements and specifications, two demonstrator phases will aim to develop and optimize particular aspects of EF Tracking on a range of hardware. A technology choice will

be made at the time of the Final Design Review (FDR) in 2025, followed by the final system preparation for the overall EF Production Readiness Review (PRR). The technology decision will be based on the trigger performance, the required power, rack space, cost, flexibility, scalability, and operations considerations. The required effort for the project is estimated to be 18 FTE/year (including 8 FTE/year of software/firmware engineering during the initial demonstrator phases), for a total of 90 FTE-years. An initial analysis of risks is also described.

1 Introduction

1.1 Context of this Document

The TDR for the Phase-II upgrade of the TDAQ [1] was completed in 2017 and approved by the LHC Experiments Committee (LHCC) in 2018. It includes the baseline for EF Tracking as a Hardware-based Tracking for the Trigger (HTT) based on associative memory ASICs and FPGAs on custom ATCA boards, with the possibility of an evolution into a Level-1 Trigger. Since then, this baseline has been pursued according to the plans developed by the ATLAS Collaboration and tracked by the Collaboration and by the Phase-II Upgrade Cost Group (P2UG). However, a number of developments, which will be explained in the following subsections, have led the ATLAS Collaboration to review and ultimately to decide to change the baseline for EF Tracking. This document supersedes the original Phase-II TDAQ TDR on the matter of EF Tracking. It does not seek to provide an update on the rest of the system described in the TDR, even though that may have evolved through the normal process of specifications and design reviews that have followed.

The specific goals of this document are as follows:

- Describe the new baseline design for EF Tracking and the associated work plan.
- Describe the organization and resources needed to deliver the project.
- Give clarity to TDAQ institutes about the new project and how they can contribute.[1]
- Provide information needed to seek approval for this new baseline from the ATLAS Collaboration, and CERN through the LHCC.
- Provide documentation to support approaches to funding agencies.

1.2 Recent Advances Related to Track Reconstruction in Software

Over the past years ATLAS has carried out an intensive program of R&D to finalize the ITk layout and to optimize the Phase-II track reconstruction software. Since the Phase-II TDAQ TDR [1], the design of the ITk has been refined considerably. In particular, the five-layer pixel system underwent major revision, not only to simplify the detector layout and construction, but in particular to optimize the ring design for tracking performance and to

[1] While true at this draft stage, institute interest will be reported in section 3.4 in a future version so this point will be removed from the document before it is submitted to the LHCC.

minimize the CPU required for reconstruction. The innovative ring design of the inclined barrel sections, the pixel end-caps and inner ring sections allows for a strategic placement of track measurements along the trajectories, limiting the distance between adjacent hits and keeping the required hit coverage for efficient and high purity tracking at all η. Figure 1.1 shows the ITk layout that is the result of this development process.

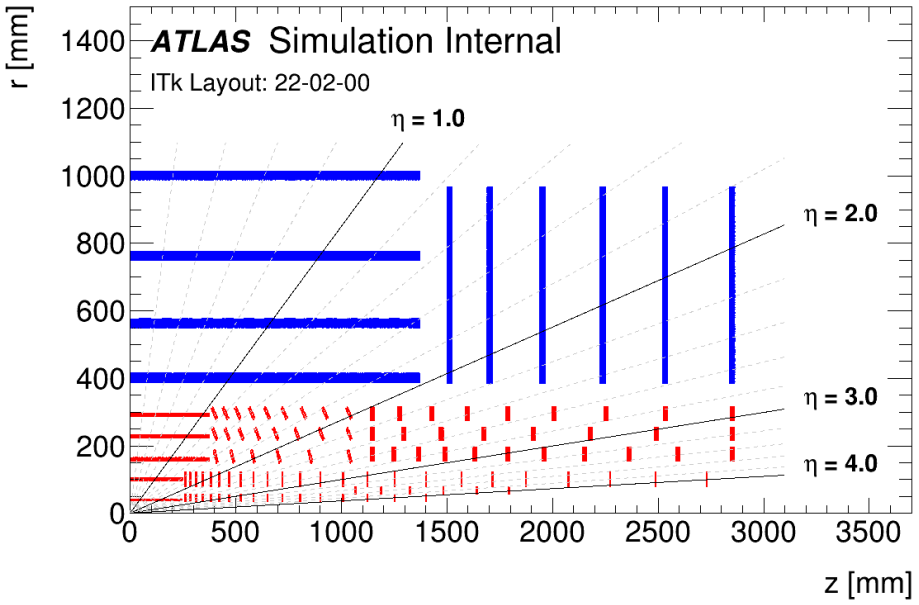

Figure 1.1: *A schematic view of the "ITk Layout" used for the studies presented in this document, with a five-layer Pixel detector (red) surrounded by the Strip detector (blue). Only the positions of the active sensors are shown. The final "ITk Layout" can be found in Reference [2].*

At the same time, a novel approach for track reconstruction for high pile-up events in the ITk was developed that fully exploits the precision and redundancy of the five-layer Pixel System [3]. The new approach starts with a search for track candidate seeds from hit triplets in the Pixel layers and makes use of the additional Pixel layers to very efficiently reject fake combinations, before even starting the CPU intensive combinatorial Kalman filter based track finder for each seed combination. The Pixel System hermetically covers the full η acceptance of the ITk and provides by itself enough redundancy for a very high track-finding efficiency.

At the time of the Phase-II TDAQ TDR, the track reconstruction using software took approximately 300 HS06 $\times$ sec [1] for the so-called "Inclined Duals" layout of the ITk. With the latest fast ITk reconstruction software prototype, and with the updated ITk layout, the time for reconstructing an event went down by more than a factor of 8 without major compromises in tracking performance, as will be shown in this document. Figure 1.2 shows a comparison of the CPU needed for track reconstruction as a function of average pile-up

using the tracking software at the time of the Phase-II TDAQ TDR and the latest software prototype.

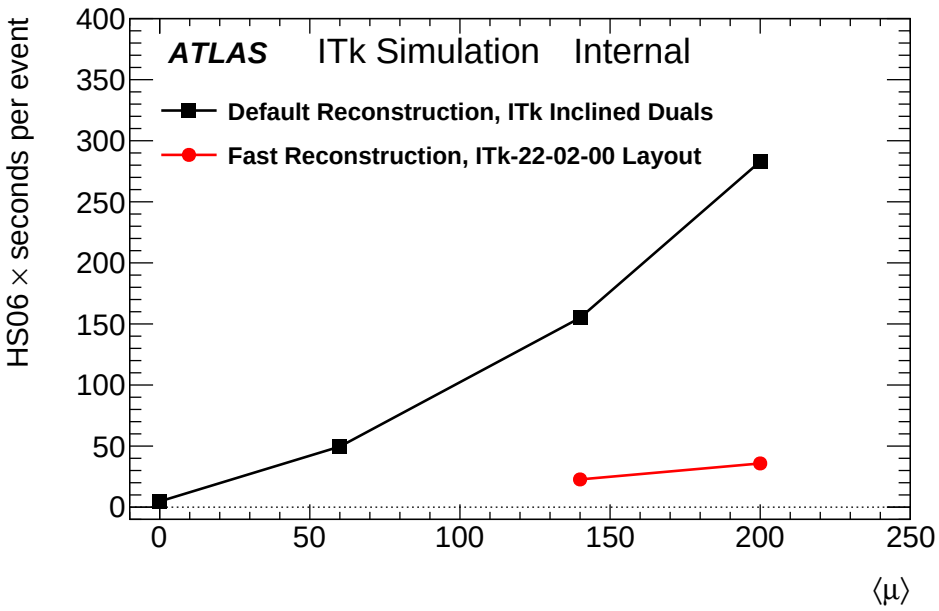

Figure 1.2: *The total CPU time required in* HS06 × sec *to reconstruct a* $t\bar{t}$ *event in the ITk, as a function of the average event pile-up. Shown are the results obtained for the "Inclined Duals" detector layout and using the tracking software at the time of the Phase-II TDAQ TDR and the latest results for the fast ITk track reconstruction using the updated ITk detector design.*

In addition to improvements in the software tracking, delays to the LHC schedule and recent developments on the CPU market have resulted in a significant reduction in predicted cost per HEP-SPEC06 (HS06) [4]: it is now 1.3 CHF/HS06 in 2027, compared to the Phase-II TDAQ TDR prediction of 2 CHF/HS06 for 2026. Figure 1.3 shows the past and future projected server price evolution as reported by CERN IT [5].

1.3 Recent Advances Related to Commercial Compute Accelerators

Since the Phase-II TDAQ TDR [1], industry trends for data centers are shifting toward a model where a CPU is no longer the sole and primary unit of compute for many workloads such as data streaming, analytics, and artificial intelligence applications. These *heterogeneous* systems integrate multiple types of computational units such as multi-core CPUs, GPUs, DSPs, FPGAs and ASICs to perform the required computations more quickly (lower latency) and/or achieve higher performance with lower power consumption to satisfy electrical power and cooling constraints, as well as rack-space limitations. The concept of a

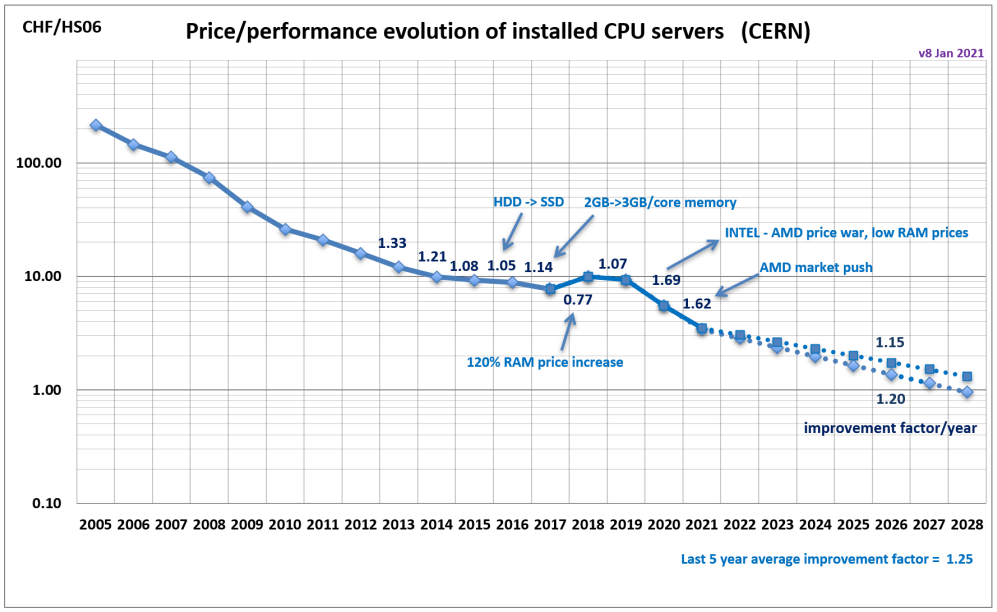

Figure 1.3: *Price/performance evolution of installed CPU servers by CERN IT [5]. The data point labels show the price ratio of the previous to the current year. Two scenarios are shown with an assumed 15% and 20% improvement per year, respectively.*

heterogeneous commodity system lends itself well to consider promising future technologies.

While FPGAs have a 30-year history throughout the electronics industry, their use as server accelerators in data centers is a relatively new but growing trend. For server use, FPGAs are packaged as accelerator cards that plug into a Peripheral Component Interconnect Express (PCIe) slot in the server's motherboard; these commercial accelerators are widely available today. The computing performance and capabilities of GPUs have drastically improved in recent years. Modern compilers enable writing feature-rich code that can be executed by GPUs and extensions of the C++ language are being developed to allow implementing "offloaded" calculations that may be executed by GPUs using a standard C++ formalism. Future technologies that may be extremely compelling include unified memory architectures where hybrid computing units with multiple types of CPU/GPU/FPGA cores are co-located on the same chip and make use of a shared high-speed memory.

Developments in programming interfaces for a wide selection of computing hardware including CPUs, GPUs and FPGAs also make heterogeneous commodity systems attractive. Projects such as oneAPI/SYCL [6] are designed to make it possible to run a single codebase on as many different types of hardware as possible, and to lower the barrier of entry for writing code for non-CPU hardware.

1.4 Event Filter Tracking Technology Decision

In 2020 a series of steps and decision points were set out by the ATLAS P2UG Committee and the TDAQ and ITk projects to review the proposed "evolution" of the Phase-II TDAQ system described in the Phase-II TDAQ TDR [1]. As a result, a decision was taken not to retain the "evolution" option. In summary, the evolution scenario foresaw a significant increase of the Level-0 (L0) trigger rate (up to 4 MHz) and the concurrent introduction of a second stage of fast selection (L1), mainly based on tracks reconstructed from selected data from the ITk, to reduce the readout rate to an acceptable level (< 1 MHz). The combined latency of L0 and L1 trigger stages was constrained by the size of the detector front-end (FE) buffers, as most detectors would only be read out in conjunction with a L1 trigger signal. The decision not to pursue the "evolution" scenario was based primarily on the risks in the development of the ITk Pixel FE ASIC, the reduced physics motivations due to the stronger than anticipated limitations of throughput for the ITk Pixel Detector readout, risks related to L1 latency and the technical challenges in commissioning an L0/L1 TDAQ system alongside the L0 TDAQ system.

The major implication for TDAQ was to eliminate the need to design a custom-based track trigger for L1. In a statement to the collaboration, the Extended TDAQ Steering Group (eTDSG) concluded: "The decision of ATLAS to exclude the possible evolution to a L0/L1 trigger system for the Phase-II upgrade, has removed the low latency requirement that drove the TDR custom-based track trigger design." Consequently, the "TDAQ hooks" described in the Phase-II TDAQ TDR were deemed no longer necessary.

It also became apparent that some of the assumptions behind the original Phase-II TDAQ TDR were in need of revision, due to recent advances in track reconstruction software (Section 1.2) and commercial accelerator technology (Section 1.3). The L1 decision and revised assumptions prompted TDAQ to review the Event Filter tracking technology, and in particular to re-evaluate which of the custom- and commodity-based approaches would be more appropriate.

Three task forces were formed to evaluate the technical feasibility, estimated tracking performance, operational procedures, opportunities for improvement, risks, and resource requirements for software-only, heterogeneous commodity, and custom EF tracking solutions. On the specific criterion of cost, the results from the task forces showed that the estimated system cost from the FPGA-based heterogeneous commodity prototype is lower than that of the software-only equivalent, although with potentially large uncertainties.

An independent ATLAS committee reviewed the reports from these three task forces and recommended that "ATLAS commit to a commercial solution for EF tracking at HL-LHC," including further optimizations targeting the use of accelerators to potentially mitigate risks related to power and cost: "TDAQ should continue investigating using hardware accelerators to optimize the EF farm. The Heterogeneous commodity task force has largely demonstrated proof-of-concept, and a heterogeneous solution (including FPGAs and/or

GPUs) could lead to substantial power and cost savings." The eTDSG endorsed this recommendation, following which the TDAQ Institutional Board (TDIB) approved the drafting of this TDR amendment. Following approval by the TDIB and the ATLAS Executive Board, it was submitted to the LHCC for review.

The new Phase-II TDAQ architecture baseline is shown in Fig. 1.4.

1.5 Outline of this Document

This document is organized into three main parts. Chapter 1 contains an introduction. Chapter 2 presents an overview of track reconstruction in the EF; the physics motivation, trigger menu, and tracking requirements; the definition of the new system; a description of the demonstrators that have been developed since the publication of the Phase-II TDAQ TDR [1]; and the program for the optimization of the system. Chapter 3 describes the project management and organization, planning and scheduling, required effort and financial resources, and risk analysis and mitigation strategies. The glossary in Appendix A defines many of the common acronyms and terms used in this document.

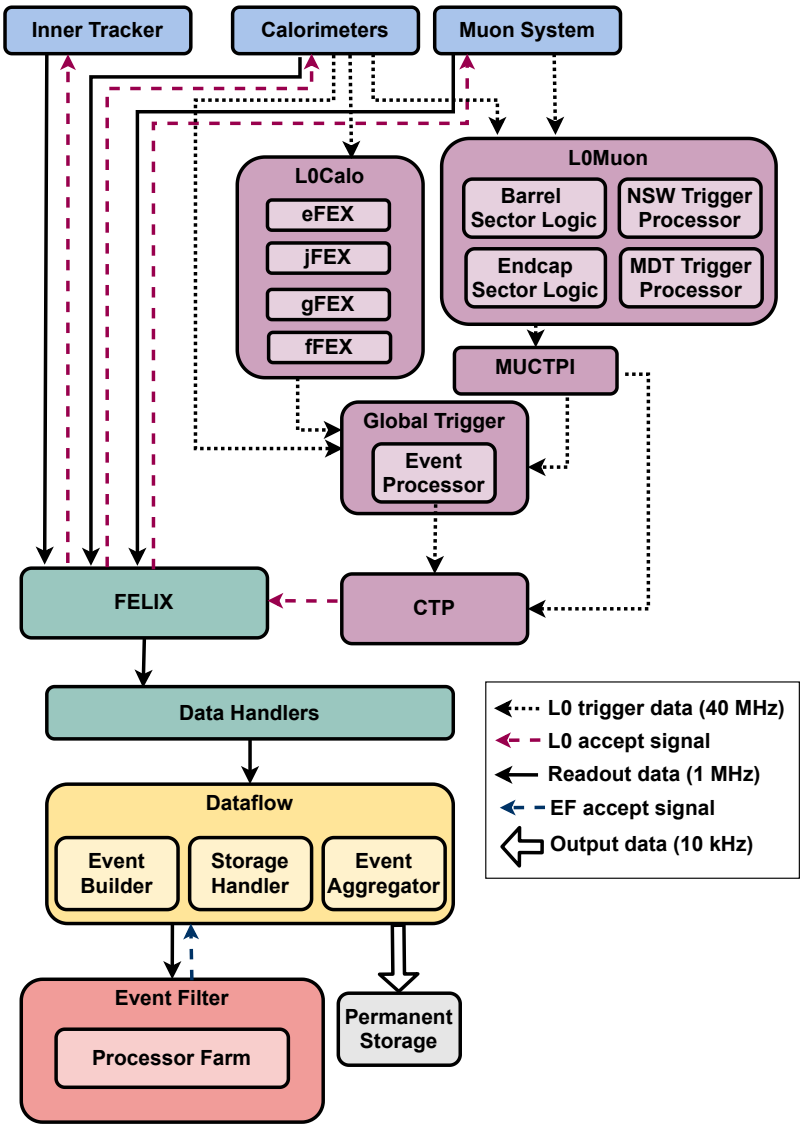

Figure 1.4: *The TDAQ Phase-II architecture with the EF updated to reflect the baseline change to use only commercial processors. The black dotted arrows indicate the Level-0 dataflow from the detector systems to the Level-0 trigger system at 40 MHz, which must identify physics objects and calculate event-level physics quantities within 10 µs. The result of the Level-0 trigger decision (L0A) is transmitted to the detectors as indicated by the red dashed arrows. The resulting trigger data and detector data are transmitted through the Data Acquisition System (DAQ) system at 1 MHz, as shown by the black solid arrows. Direct connections between each Level-0 trigger component and the Readout system are suppressed for simplicity. The EF system is composed of a heterogeneous processor farm that must reduce the event rate to 10 kHz. Events that are selected by the EF trigger decision are transferred for permanent storage.*

2 System Considerations

2.1 Track Reconstruction in the Event Filter

The EF system is a multi-threaded asynchronous processing farm of commodity servers running a subset of offline-like reconstruction algorithms and menu-driven event selection. As shown in Fig. 1.4, EF processes receive event data from the Dataflow system and provide the accept/reject decision and selected reconstruction data for offline storage. As described in the Phase-II TDAQ TDR [1], high luminosity (and consequently high pile-up) conditions give rise to specific challenges for the object and event reconstruction algorithms planned for the EF, in particular for track reconstruction.

The EF track reconstruction chain is comprised of the preparation of the input ITk data followed by track finding (which is usually broken up into a series of algorithmic steps as described below). Track candidates and their associated properties are provided for use by subsequent trigger algorithms.

The input to the track reconstruction is the ITk raw data: the size of each event at $\langle\mu\rangle = 200$ is assumed to be 1.6 MB (Pixel) + 0.5 MB (Strip) = 2.1 MB (total). An initial data preparation stage (**ITk Data Preparation**) consists of decoding of the ITk raw data byte stream and clustering the hits in each ITk sensor. There is the option to prefilter hits and/or to create space points using the ITk Strip clusters from each side of a stave or petal, exploring the stereo angle between both sides and reducing the number of measurements to be handled in later steps in the reconstruction chain.

Next, a seed formation and pattern recognition step (**Track Seeding & Pattern Recognition**) takes a subset of hit clusters or space points to identify hit combinations that would correspond to likely track candidates. One option is to utilize Pixel hit combinations to determine the initial track seeds; this is made possible by the five-layer ITk Pixel Detector which covers the full range of $-4 < \eta < 4$. A second option is to seed tracks from the outer layers of the ITk, typically 7 or 8 Strip layers plus the outermost Pixel layer at central η.

The final stage (**Track Extension, Fitting & Ambiguity Resolution**) includes the combinatorial Kalman Filter to extend track seeds into complete track candidates (into additional layers if needed), algorithms for duplicate removal, fake rejection and resolving ambiguous track candidates, and the final high precision track fit to determine the track parameters.

The ITk geometry and magnetic field map of the detector is used at different levels of precision across the algorithms in the tracking chain, depending on the needs. The ultimate

offline ITk material maps and field description will be available to the trigger reconstruction by means of the A Common Tracking Software (Acts) tracking software framework [7].

The resulting track candidates are used by EF trigger algorithms such as primary vertex determination, *b*-tagging, charged lepton identification and particle flow. The proposed use of tracking for specific trigger signatures is described in Section 2.2.

2.2 Physics Motivation, Trigger Menu, and EF Tracking Requirements

The physics goals of the High Luminosity LHC (HL-LHC) remain as set out in the Phase-II TDAQ TDR [1]. They include a detailed exploration of the mechanism of electroweak symmetry breaking through the properties of the Higgs boson, searches for new physics through the study of rare Standard Model processes, searches for new heavy states and LLPs, searches for weakly-coupled particles, and measurements of the properties of any newly discovered particles.

The high pile-up environment at the HL-LHC will make it challenging to achieve these physics goals. Exceptional performance of the ATLAS trigger system, including high-quality tracks at the Event Filter level, is required. Figure 2.1 shows the expected luminosity profile for HL-LHC. It is expected that Run-4 will start at the end of 2027 and will reach a maximum peak luminosity of 5×10^{34} cm^{-2}s^{-1}, corresponding to an average of 140 simultaneous p-p interactions per bunch crossing (pile-up). Run-5 starts in 2032 and will reach 7.5×10^{34} cm^{-2}s^{-1} or 200 pile-up.

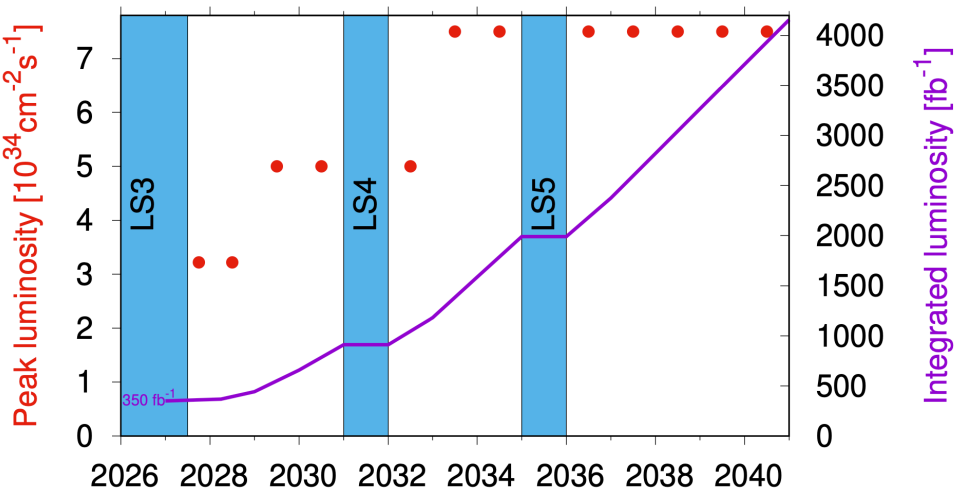

Figure 2.1: *Projected HL-LHC luminosity profile. The highest pile-up of 200 will be reached only in the second year of Run-5 (2033).*

The high pile-up conditions planned for the HL-LHC give rise to specific challenges for the trigger algorithms. The calorimeter energy resolution and calorimeter isolation effectiveness are reduced due to pile-up, while the rate of interesting events increases. Multi-object triggers also have higher rates due to coincidence candidates where the objects come from different collisions. This leads to a need for a high L0 trigger rate. Tracking information from the ITk in the trigger selection is a key ingredient to control this rate in the ATLAS Trigger strategy for HL-LHC. Tracking can be used to reduce the rate of coincidences by requiring tracks associated with objects to come from the same vertex, requiring candidates have an appropriate number of high-p_{T} tracks, and using particle flow and similar techniques to improve resolution. The tracking itself however becomes more challenging because the higher hit density of tracking detectors makes pattern recognition more computationally intensive and more susceptible to fakes.

2.2.1 Tracking in the Phase-II TDAQ TDR Trigger Menu and Beyond

The Phase-II TDAQ TDR presented a representative trigger menu which demonstrated that, with the Phase-II upgrade, ATLAS could achieve trigger selections as good as, or even better than, those available in Run-1 for most triggers, with only moderate losses for the most challenging triggers. The trigger menu from the Phase-II TDAQ TDR is shown in Table 2.1. Tracking in the EF is critical to achieving the rejection shown in the last two columns of the table.

The EF will receive input from Level-0 at 1 MHz, and will perform **regional tracking** on the majority of those events[1], in Regions of Interest (RoIs) which are defined based on objects identified by the Level-0 trigger. After trigger selections using the regional tracking, **full-scan tracking** will run at a reduced rate, over the entire ITk detector. In both cases, tracking extends to the full η range of the ITk detector.

Regional Tracking

Regional tracking can be used to verify the presence of high-p_{T} tracks in single high-p_{T} lepton triggers and associating objects to a common vertex, while significantly reducing backgrounds containing objects from multiple vertices. For this purpose, the regional tracking reconstructs tracks down to 2 GeV, which was demonstrated in the TDR to be sufficient to associate jets with $p_{\mathrm{T}} \gtrsim 40$ GeV to vertices.

In order to achieve the rates presented in the Phase-II TDAQ TDR trigger menu, regional tracking is needed at approximately 1.5% $\eta \times \phi$ coverage [2], but since this is based on a

[1] Specifically, on all the triggers which show a rate reduction from L0 to "After regional tracking" in the table.

[2] The TDR had a number of 2.3% estimated for the core $|\eta| < 2.5$ region, but making an appropriate average over the full coverage using a Monte-Carlo (MC) simulation gives 1.5%. This same MC simulation was used for the regional readout task force.

Table 2.1: *Representative trigger menu for 1 MHz Level-0 rate. The offline p_{T} thresholds indicate the momentum above which a typical analysis would use the data.*

	Offline p_{T} Threshold [GeV]		Planned HL-LHC		After	Event
Trigger Selection	Run 1	Run 2 (2017)	Offline p_{T} Threshold [GeV]	L0 Rate [kHz]	regional tracking [kHz]	Filter Rate [kHz]
isolated single e	25	27	22	200	40	1.5
isolated single μ	25	27	20	45	45	1.5
single γ	120	145	120	5	5	0.3
forward e			35	40	8	0.2
di-γ	25	25	25,25		20	0.2
di-e	15	18	10,10	60	10	0.2
di-μ	15	15	10,10	10	2	0.2
$e-\mu$	17,6	8,25 / 18,15	10,10	45	10	0.2
single τ	100	170	150	3	3	0.35
di-τ	40,30	40,30	40,30	200	40	$0.5^{\dagger\dagger\dagger}$
single b-jet	200	235	180	25	25	$0.35^{\dagger\dagger\dagger}$
single jet	370	460	400			0.25
large-R jet	470	500	300	40	40	0.5
four-jet (w/ b-tags)		$45^{\dagger}$(1-tag)	65(2-tags)	100	20	0.1
four-jet	85	125	100			0.2
H_{T}	700	700	375	50	10	$0.2^{\dagger\dagger\dagger}$
$E_{\mathrm{T}}^{\mathrm{miss}}$	150	200	210	60	5	0.4
VBF inclusive			2x75 ($\Delta\eta > 2.5$ & $\Delta\phi < 2.5$)	33	5	$0.5^{\dagger\dagger\dagger}$
B-physics††				50	10	0.5
Support Triggers				100	40	2
Total rate				1066	338	10.15

† In Run 2, the 4-jet b-tag trigger operates below the efficiency plateau of the Level-1 trigger.
†† This is a place-holder for selections to be defined.
††† Assumes additional analysis-specific requirements at the Event Filter level

representative menu, a suitable margin needs to be included. The TDR requirement was specified using a data volume estimate specific to the architecture in the TDR which corresponds to 5% $\eta \times \phi$ coverage. The 5% $\eta \times \phi$ coverage requirement will be retained. This system requirement provides a margin of around a factor of 3. It should be noted that differences in configuration and performance of triggers can easily change the regional tracking needs by around a factor of 2, so this is conservative, but not excessive. For example, the TDR Level-0 Trigger (L0) menu dedicates 200 kHz each to single electron and di-τ triggers, however improved performance of those triggers' identification could reduce their rate, allowing higher rates to be dedicated to multijet or $E_{\mathrm{T}}^{\mathrm{miss}}$triggers. Similarly, changes in the trigger prioritization could also lead to more hadronic triggers. An addition of 200 kHz of hadronic triggers could lead to an additional $\approx$ 1.5% in regional tracking (doubling the TDR menu estimate). Changing the RoI size could again change the regional tracking by order 25%, and using asymmetric trigger thresholds could require tracking in lower-p_{T} jets which have higher multiplicities (see Figure 2.2).

The regional tracking estimate in the Phase-II TDAQ TDR trigger menu is calculated assuming $\eta \times \phi$ RoIs:

- e, μ, and τ: 0.2×0.2
- small-R jets: 0.8×0.8
- large-R jets: 2.0×2.0

Full-scan Tracking

Use cases for full-scan tracking include b-tagging and track-based calibration on all jets with $p_{\mathrm{T}} > 20$ GeV, calculation of variables such as the $E_{\mathrm{T}}^{\mathrm{miss}}$ soft-term, pile-up correction/mitigation, and searching for additional soft jets. Figure 2.2 from the Phase-II TDAQ TDR [1] shows that for both 4-jet and $E_{\mathrm{T}}^{\mathrm{miss}}$ events, there are on average 13 jets with $p_{\mathrm{T}} > 20$ GeV. This large number combined with the effect of the beam spot spread in z and the bending of tracks in ϕ means that essentially the entire detector is needed for b-tagging and track-based calibration of these jets. It is therefore assumed that single jet (small-R), multi-jet, di-τ, $E_{\mathrm{T}}^{\mathrm{miss}}$, and H_{T} triggers will all use full-scan tracking.

In the Phase-II TDAQ TDR trigger menu, the full-scan tracking need was estimated to be 91 kHz, with the system requirement set at 100 kHz. This was based on estimates of the tracking needs for each of the items in the representative menu. The TDR margin on this was lower than for regional tracking, but the previous baseline hardware was meant to be designed with sufficient margin to ensure that the requirements are met.

This requirement has been revisited and raised to 150 kHz full-scan tracking to provide additional margin. One scenario where additional full-scan tracking could be critical is in the $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger, where the number of objects used in the Global Trigger to construct the L0 $E_{\mathrm{T}}^{\mathrm{miss}}$ was assumed in the Phase-II TDAQ TDR to be based on a 45 GeV jet threshold. If a lower threshold is used, or if unclustered soft-terms are used it may be preferable to

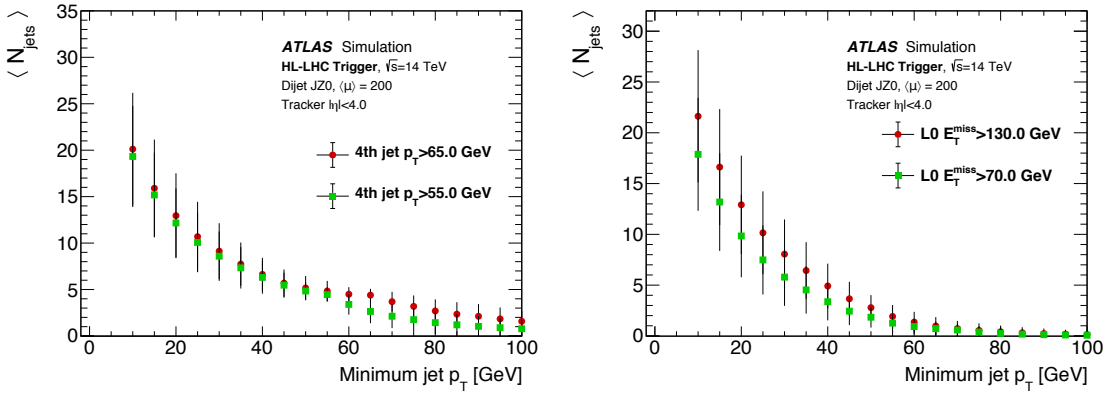

Figure 2.2: *Left: The mean number of jets in events as a function of the minimum jet p_{T} for events that are selected at Level-0 to have at least 4 jets with p_{T} above 50 GeV for two threshold scenarios. The error bars on the plot indicate the RMS of the jet distributions. The x-axis jet energy is the calibrated offline energy. Right: The mean number of jets for events selected by a Level-0 $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger for two threshold scenarios. The errors bars on the plot indicate the RMS of the jet distributions. The x-axis jet energy is the calibrated offline energy.*

do full-scan tracking on all the L0 $E_{\mathrm{T}}^{\mathrm{miss}}$ triggers, without an initial rejection from regional tracking. This would raise the full-scan tracking requirement by about 50 kHz.

Since the development of the representative trigger menu in the Phase-II TDAQ TDR, several other use-cases for full-scan tracking have become apparent. These cases would benefit from any available full-scan tracking rate. The use of LRT has grown in recent years, motivated by searches for LLPs. Some triggers (e.g. $E_{\mathrm{T}}^{\mathrm{miss}}$) might warrant full-scan LRT processing and/or displaced vertex finding. The Run-3 B-physics trigger menu includes some triggers that are seeded from all L1 triggers for $B \to K^{\star}ee$. The use of such trigger strategies will be EF processing limited and could lead to an increase in either regional or full-scan tracking. Trigger-level analyses (TLA) need offline-quality tracking to make their decisions, since no further event and object reconstruction will be performed offline. Depending on the signature, this could require either regional or full-scan tracking. Particle flow jet reconstruction requires high quality tracks from all vertices to achieve pile-up subtraction in the calorimeter.

2.2.2 Event Filter Tracking Requirements

To achieve the goal of maintaining a menu with thresholds similar to Run-1, the EF reconstruction algorithms and resulting trigger decisions must be robust against pile-up, so that rates scale approximately linearly with luminosity. Some of our most powerful tools to mitigate pileup include the use of tracking to isolate high-transverse-momentum recon-

structed objects or to associate them with a common vertex. The EF must be capable of reconstructing vertices and tracks as needed for pile-up suppression.

Overall tracking efficiency shall be close to the offline efficiency [2]; for example, central muons with $p_T > 10$ GeV shall be 98% efficient relative to offline tracking. **Regional tracking** is required at the full Level-0 rate over 5% $\eta \times \phi$ coverage, for initial rejection in single high-p_{T} lepton and multi-object triggers. Regional tracking will be performed for all charged tracks with $p_{\mathrm{T}} > 2$ GeV and $|\eta| < 4.0$. **Full-scan tracking** is required for a rate of 150 kHz. Full-scan tracking for soft jets, calculation of the $E_{\mathrm{T}}^{\mathrm{miss}}$soft term and pile-up correction and mitigation requires reconstruction of all charged tracks with $p_{\mathrm{T}} > 1$ GeV and $|\eta| < 4.0$. Full-scan tracking for b-jet identification additionally requires track quality and minimum p_{T} similar to offline track reconstruction [8]. For example, track parameter resolutions for full-scan tracking shall be within a factor of two of the offline performance. **Trigger-signature-specific tracking needs**, including likely signature-specific tunings, will be evaluated.

Requirements on efficiency, resolution and fake rates are inspired by the Run-1/2 experience and previously developed track trigger requirements within the TDAQ community to evaluate the demonstrators described in this document. The full requirements for EF Tracking will be set out in the system specifications document and reviewed in the Specification Validation Review, as described in Section 3.3.

2.2.3 Large Radius Tracking

The HL-LHC dataset will provide an important opportunity for new physics searches in final states with LLPs. These final states appear in a wide range of beyond-Standard-Model theories, arising in decays with small couplings, limited phase space, or a heavy off-shell intermediary particle. Searches for LLPs are a key target for the future and will benefit greatly from increasing luminosity thanks to their low backgrounds, but they are often not covered by typical trigger selections. To enable new triggers for some of these searches, the EF should provide LRT, which focuses on tracks with high impact parameters like those resulting from the decays of long-lived particles. This is discussed further in Sections 2.4.1 and 2.4.2.

2.2.4 Simulation of the EF Trigger Selection

The need to study the trigger reconstruction and selection on Monte Carlo (MC) simulation is an important constraint to the EF tracking. A CPU-only track reconstruction would allow to apply the same online EF trigger selection in the offline MC simulation production. FPGA or GPU based tracking implementations would require developing and maintaining a dedicated EF track reconstruction emulation for the offline MC simulation, unless such

accelerator-based systems would also become available for the full offline simulation production. An implementation of the EF track reconstruction using C++ extensions supporting heterogeneous processing technologies would ensure better portability of accelerator code to other types of processors. Such heterogeneous processing technologies are under active development in the ATLAS Computing and Software project [9].

2.3 Definition of New System

2.3.1 Flexible, Heterogeneous Commercial System

The proposed EF design is a flexible, heterogeneous commercial system consisting of CPU cores and possibly accelerators, to perform the EF reconstruction including the compute-intensive ITk track reconstruction. A program will be undertaken to optimize the specific components that perform the different tasks in the ITk track reconstruction chain. This chain involves all reconstruction steps of **ITk Data Preparation**, **Track Seeding & Pattern Recognition**, and **Track Extension, Fitting & Ambiguity Resolution** described in Section 2.1. The different steps of the ITk track reconstruction chain could be executed either on CPUs or using accelerator hardware (GPU, FPGA, etc.), which could either be installed within the EF nodes or in separate servers. The three demonstrators described in Section 2.4 span the range from a full software implementation on CPUs to a configuration which relies heavily on FPGA accelerators and to a demonstrator based on GPU accelerators. Such a heterogeneous approach gives the flexibility to select the best commercial hardware together with optimal algorithmic approaches. This approach allows to also benefit from emerging commercial technologies as described in Section 1.3.

The optimizations of the overall EF system will be driven by the trigger performance, the required cooling power, rack space, cost, flexibility, scalability, and by operations considerations. The power limit at the ATLAS experiment site on the LHC ring (Point-1) is currently planned to be 2.5 MW. The EF farm can be housed in the existing SDX1 data center and, if needed, in additional compute containers to be installed at Point-1. The use of accelerator cards could significantly lower the power and space requirements of the EF farm and provide flexibility in case additional compute capacity is needed to cover new physics use-cases.

Compute requirements in this document are reported in units of HEP-SPEC06 (HS06), which is the default benchmark for high energy physics applications [10]. The system costing presented in Section 3.2 is based on the CPU-only fast ITk track reconstruction prototype shown in Section 2.4.1 and includes the non-tracking CPU estimates described in the Phase-II TDAQ TDR [1]. Based on conservative [3] CPU power estimates of 0.25 (0.20)

[3] The 2021 TDAQ server purchase (AMD EPYC 7302) achieves 0.3 W/HS06 and servers with 0.2 W/HS06 are already on the market.

W/HS06 for Run-4 (Run-5), such a system would require 1.9 MW (2.3 MW) of power and cooling in Run-4 (Run-5).

2.3.2 Opportunities Beyond Tracking

Additional Uses in Trigger

While tracking is the most computing intensive task in the EF, the use of accelerators could also prove beneficial for other domains, i.e. calorimeter and muon reconstruction in the EF. Parts of these have been studied in earlier GPU prototypes for the use in the Run-3 EF farm [11]. At the time they were not considered cost-effective compared to CPUs but work has already restarted to re-evaluate these and other algorithms on current accelerator hardware. Studies (including necessary demonstrators) will be coordinated in order to converge on the optimal system for the EF overall.

Uses of System Outside of Data Taking

A large commodity EF farm will provide significant additional resources to the offline simulation production (the resource known as Sim@P1 [12]). Figure 2.3 shows the typical EF farm usage during Run-2 data-taking (left) and the current long shutdown (right). Assuming a 3:1 ratio of data-taking versus shutdown years during Phase-II, it is expected that the EF farm will provide on average 42% of its compute capacity to offline production. Thanks to operational improvements in switching the farm from data-taking to offline production, the compute available to Sim@P1 is likely to increase even further already during Run-3. Moreover, with the use of software containers and container orchestration for the EF applications in Run-4, it should be possible to run MC production even during data-taking when processing resources are not needed by the EF.

Today, the usage of the EF farm for offline Geant4 simulation production would be restricted to the CPU servers, because Geant4 does not support FPGA or GPU accelerators. Applying the accelerator based trigger reconstruction to simulated events could make use of the co-processors in the EF farm. Ongoing R&D by the offline software project to improve the support for heterogeneous computing for other offline applications, like QCD generators, fast simulation, or even Geant4, may lead to further opportunities to explore accelerators if installed in the EF farm.

Exploration of New Technologies for Broader ATLAS Use

We expect the current trend towards power-efficient, heterogeneous computing architectures to continue over the next decade. ATLAS is pursuing opportunities to run its software applications on systems including massively-parallel, domain-specific accelerators

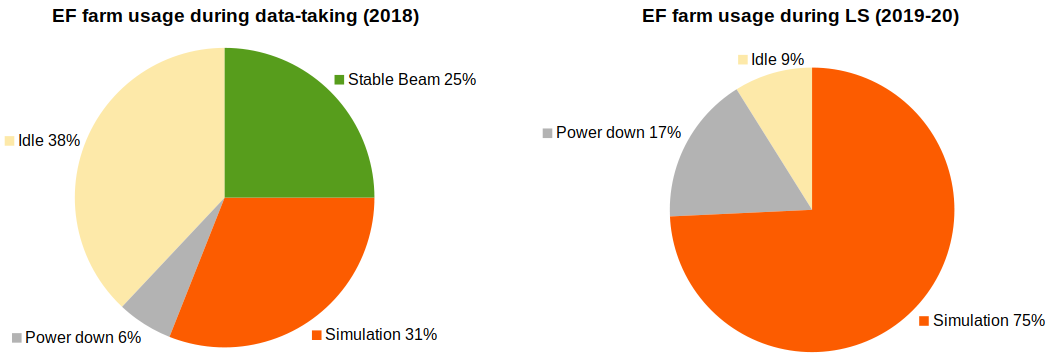

Figure 2.3: *Time sharing of the EF farm during data-taking (left) and the first two years of LS2 (right). When idle, power consumption is down to 30%. The farm was switched to simulation mode only when a break of at least 24 hours in data-taking was foreseen.*

optimized for tasks such as tensor manipulation, dataflow algorithms and graph processing. This trend is currently seen in High Performance Computing (HPC) systems, which provide large compute capacity via the use of accelerators. Unlike these external resources, the choice of EF architecture is under the control of ATLAS and an ideal environment to develop and deploy reconstruction algorithms that are able to optimally exploit heterogeneous computing resources. The invested effort will directly benefit and boost the R&D program described in the ATLAS Phase-II Computing CDR [13].

2.4 Demonstrators

This section describes two demonstrators that have been developed since the publication of the Phase-II TDAQ TDR [1] and a reminder of the previous work done to study the use of GPUs in the Run-3 trigger system. Section 2.4.1 describes a system based on an improved fast track reconstruction implemented entirely in software running on commodity CPUs and Section 2.4.2 shows the implementation of an FPGA-based demonstrator using commodity PCIe accelerator cards. Section 2.4.3 gives a brief recap on the prior work on track reconstruction on GPUs.

2.4.1 Fast-tracking CPU-based Demonstrator

Technical description

Since the Phase-II TDAQ TDR [1], a fully functional ITk fast trigger reconstruction prototype has been developed. In the following an overview of the algorithmic approach is given, a detailed description of the default and fast ITk reconstruction strategy can be

found in an ATLAS note [3]. The studies performed for this section used the ITk geometry 22-02-00 shown in Figure 1.1.

The largest contributor to the overall ITk reconstruction time is the Ambiguity Resolution. In the default ITk reconstruction about 60% of the input track candidates pass the Ambiguity Resolution algorithm. For the purpose of the fast ITk trigger reconstruction, the Ambiguity Resolution algorithm is omitted from the fast reconstruction chain. Instead, a tighter track selection is implemented, including limiting the number of shared hits between output tracks, in the Silicon Track Finding to remove duplicate and fake tracks already at this stage. For the track parameter estimation the fast Kalman Filter track fit is used directly, as applied by the Combinatorial Kalman Filter in the Silicon Track Finding algorithm. The fast Kalman Filter makes use of the offline cluster calibrations and corrections, but uses some approximations in particular in the material model. Therefore the fast Kalman Filter is expected to yield a slight loss in resolution compared to the full track fit.

The Silicon Track Finding in the default ITk reconstruction was tuned for robustness and physics performance, primarily to allow for studies of different detector layout options. The fraction of the CPU spent in Seed Finding can be as large as 50 % for the ITk for a sample with an average pile-up $\langle\mu\rangle = 200$. Speeding up the Seed Finding and improving the purity of the output seed collection by tuning the selection cuts leads to a significant reduction in the reconstruction time, as it also reduces the subsequent time spent in the road building and Combinatorial Kalman Filter. The five layer Pixel Detector covers the full range of $-4 < \eta < 4$. Hence only Pixel hit combinations are used for the fast ITk track reconstruction version of the Seed Finding, leaving out the strip seed iteration. In addition, the cuts on the final number of hits and non-shared hits, that were previously applied in the Ambiguity Resolution, are moved to the Silicon Track Finding to remove additional fakes and duplicate track candidates from the output track collection. The p_{T} cut in the central region was raised from 900 MeV to 1 GeV and the impact parameter cut was set to $|z_0| < 150$ mm, to match the length of the beam interaction region for HL-LHC. The full list of changes in Silicon Track Finding are listed in Table 2.2.

Table 2.2: *Changes in tracking cuts used in Silicon Track Finding for fast and default (in parentheses) track reconstruction, depending on the pseudo-rapidity interval. Here z_0 is defined with respect to the mean position of the beam spot.*

Requirement	Pseudorapidity interval		
	$\lvert\eta\rvert < 2.0$	$2.0 < \lvert\eta\rvert < 2.6$	$2.6 < \lvert\eta\rvert < 4.0$
Pixel+Strip hits	≥ 9 (7)	≥ 8 (7)	≥ 7 (7)
unique hits	≥ 7 (1)	≥ 6 (1)	≥ 5 (1)
shared hits	≤ 2 (no cut)	≤ 2 (no cut)	≤ 2 (no cut)
p_{T} [MeV]	> 1000 (900)	> 400 (400)	> 400 (400)
$\lvert z_0\rvert$ [cm]	≤ 15 (20)	≤ 15 (20)	≤ 15 (20)

For the purpose of this study the optional recovery for Bremsstrahlung is disabled in the Silicon Track Finding for primary track reconstruction. The electron reconstruction results with recovery for Bremsstrahlung are discussed separately below.

The Pixel Space Point Formation, as well as the Strip and Pixel Clustering, has also been optimised for timing performance. The only technical difference being that default pixel clustering looks at 8 neighbouring pixels, while fast clustering looks only at 4. Since the only use of the space points in reconstruction is for finding track seeds, and because the seeding using space points in the Strip Detector is turned off for the fast ITk track reconstruction, there is no need to run the space point finding on the strip clusters, further reducing the CPU time needed for fast ITk track reconstruction.

The new Pixel Detector is foreseen to use the RD53 readout chip [14] that features an improved data compression. The new pixel readout chip achieves for a sample of $t\bar{t}$ events with $\langle\mu\rangle = 200$ an average event size of 1.6 MB(compared to the 0.5MB for the Strip Detector). It should be expected that unpacking of the ITk pixel and strip raw data will take about the same time per kB as unpacking the current raw data information. The CPU time spent on pixel and SCT raw data decoding is measured for the current detector on a sample of high pile-up events taken in 2017. The results are then scaled by the ITk event size for the pixels and strips raw data for $\langle\mu\rangle = 140$ and 200. On MC simulation, the decoding of byte stream raw data has to be replaced by the decoding of the file containing the simulated hits.

Table 2.3: *The CPU required in* HS06 × sec *to reconstruct* $t\bar{t}$ *MC events with* $\langle\mu\rangle = 140$ *and 200 in the ITk. Listed are the results for the different reconstruction steps using the default and the fast ITk track reconstruction. See the text for details.*

$\langle\mu\rangle$	Tracking	Release	Byte Stream Decoding	Cluster Finding	Space Points	Si Track Finding	Ambiguity Resolution	Total ITk
140	default	21.9	2.2	6.4	3.5	31.6	43.4	87.1
	fast			6.1	1.0	13.4	-	22.7
200	default	21.9	3.2	8.3	4.9	66.1	64.1	146.6
	fast			8.1	1.2	23.2	-	35.7

To study the CPU performance of the prototype fast ITk track reconstruction, $t\bar{t}$ samples were used with $\langle\mu\rangle = 140$ and 200. Shown in Table 2.3 are the results using release 21.9. The timing numbers for ITk byte stream decoding are estimated as outlined above. The results for release 21.9 were measured on a node with two Intel Xeon E5-2620 v4, 2.1 GHz CPUs with a total of 16 cores and a performance of approximately 16.8 HS06 for single core running. For comparison, the samples are also reconstructed on this machine using the default ITk software.

Table 2.3 shows that the fast version of Silicon Track Finding is approximately six times faster than the sum of the default Silicon Track Finding plus the Ambiguity Resolution for

$\langle\mu\rangle = 140$ and 200, respectively. A significant part of the CPU time is spent in byte stream decoding, clustering and space point formation. The fast track reconstruction performance is close to the full offline tracking performance, as will be shown later in this section. The offline tracking performance meets or exceeds the performance requirements outlined in Section 2.2.2. In Figure 1.2 the CPU requirements for the fast ITk track reconstruction prototype is compared to the CPU time estimated for ITk reconstruction at the time of the Phase-II TDAQ TDR [1].

Table 2.4: *The CPU required in* HS06 × sec *to reconstruct* $t\bar{t}$ *MC events with* $\langle\mu\rangle = 140$ *and 200 in the ITk, using full-scan and regional (5% η-ϕ coverage) track selection cuts.*

$\langle\mu\rangle$	Tracking	Byte Stream Decoding	Cluster Finding	Space Points	Si Track Finding	Total ITk
140	full-scan	2.2	6.1	1.0	13.4	22.7
	regional	0.33	0.90	0.15	1.11	2.49
200	full-scan	3.2	8.1	1.2	23.2	35.7
	regional	0.48	1.23	0.18	1.92	3.81

The results presented in Table 2.3 are used to estimate the CPU resource requirements for the regional and full-scan trigger reconstruction. Table 2.4 shows CPU requirements for the two tracking setups. Based on the 150 kHz full-scan rate requirement, the CPU resources can be calculated in a straight forward way. For regional track reconstruction, the p_{T} threshold is increased to 2 GeV (800 MeV in the forward), which further reduces the CPU required for reconstruction. The regional tracking is executed for an η-ϕ coverage of 5%, which corresponds to 15% of the ITk detector elements [4]. The CPU required for LRT is small compared to the primary track reconstruction (see Table 2.7). The resulting resource requirements for running regional LRT on pre-selected triggers is within the uncertainties of the overall CPU estimate and hence it is not included separately.

Using the per-event CPU requirements from Table 2.4 we calculate the total required CPU, using a p_{T} thresholds of 2(0.8) GeV for regional tracking and 1(0.4) GeV for full-scan tracking in the central (forward) region, by multiplying these numbers by the input rate of 1 MHz (150 kHz) for the regional (full-scan) tracking, respectively. The required power is calculated assuming 0.25 W/HS06 (0.20 W/HS06) for Run-4 (Run-5). The resulting total resource requirements are shown in Table 2.5.

[4] This is calculated with a dedicated Monte Carlo simulation of the trigger menu. The simulation assumes the RoI sizes given in Section 2.2.1 and the menu and object multiplicities given in the original TDR. Object distributions are assumed to be flat in ϕ and in η over a range appropriate for the trigger. The set ITK modules needed to find the tracks in that set of $\eta - \phi$ regions is then determined based on tracks from the Extrapolator Engine in the full ITK geometry. The fraction of data is then calculated from the average fraction of the events in which each module is needed for regional tracking. A second calculation has also been made using the "RegionSelector" on the ITK geometry for the relevant RoI sizes and then calculating a total using the TDR menu.

Table 2.5: *Estimate of required CPU resources for running the fast track reconstruction in full-scan and regional mode.*

	Pile-up 140		Pile-up 200	
	full-scan	regional	full-scan	regional
Rate [MHz]	0.15	1.0	0.15	1.0
CPU Resource requirement [MHS06]	3.41	2.49	5.36	3.81
Tracking resource requirement [MHS06]	5.90		9.17	
Tracking power requirement [MW]	1.47		1.83	

Preliminary studies of the sensitivity of track reconstruction in the ITk to potential detector defects and dead modules was studied in Ref. [15] as part of an analysis of staged installation scenarios. These studies will be continued and extended in the future.

Large Radius Tracking for Searches for New Physics

Non standard track reconstruction for specific trigger signatures are of key interest for the HL-LHC physics programme. The LRT stands for a dedicated reconstruction pass that is optimised for events with late decays of (mostly heavy) particles predicted in certain classes of new physics models. Using the similar fast track reconstruction approaches for LRT, as for reconstructing primary particles, allows to significantly reduce the CPU required for reconstructing such signatures in the trigger.

The resource requirements for LRT strongly depend on the maximal impact parameter in $R\phi$ and z of candidate tracks and on the minimal p_{T} required for the reconstruction of primary particles. At the same time, many models of new physics predict heavy long lived particles that decay in the ITk volume and lead to secondaries with large opening angles w.r.t. the primary particle direction of flight. The range of opening angles to be covered is an additional parameter that strongly influences the resource needs for track reconstruction.

The LRT is running after the primary track reconstruction as a second pass, after removing the clusters used in primary tracks from the event [5]. The track seeding needs to be modified in the fast reconstruction approach to allow for the topology of late decays where most of the hits are expected in the Strip detector. Therefore, the LRT is seeded on space points formed from the stereo hits from both sides of the staves/petals in the Strip detector, that needs to run in addition to the Pixel space point formation required for the primary fast track reconstruction chain. The track finding step itself for LRT follows the same track seeding, road building and combinatorial track finder approach as the standard (fast) ITk track reconstruction, but with selection cuts adapted to the specialised use-case and with additional seed cleaning to further reduce the combinatorial overhead.

[5] Without running the primary track reconstruction first, the LRT would also find all primary tracks using a less optimal tracking strategy and looser selection cuts, resulting in much increased CPU needs.

Table 2.6: *Tracking cuts for LRT.*

Requirement	Cut
Pixel+Strip hits	≥ 8
unique hits	≥ 6
shared hits	≤ 2
$\lvert d_0 \rvert$	< 150 mm
$\lvert z_0 \rvert$	< 200 mm
$\lvert p_T \rvert$	> 5 GeV
$\lvert \eta \rvert$	< 2.4

Table 2.6 summarises the selection cuts for LRT as used for this study. The seed cleaning uses approximate estimates of the track parameters, built from the space point positions, to discard as soon as possible space point triplets incompatible with the previous requirements. Due to the lower precision of the strip measurements with respect to pixels, more precise position estimates are computed for the strips in offline reconstruction, combining the positions of the three space points used to form a seed, before applying p_{T} selections in particular. In order to reduce the combinatorial overhead early, the implementation used for this study instead directly relies on the less precise p_{T} estimate built from doublets of strip space-points already, like in the formation of pixel seeds, which brings a sizeable timing reduction with limited impact on the efficiency. The cuts used for this study represent a first tuning of the LRT reconstruction for specific physics models. A more elaborate tuning of the tracking cuts will be done in the future to adapt the settings for LRT to a specific trigger selection strategy. The software based LRT approach would even allow to optimise the tuning specifically to the needs of different LRT trigger chains, e.g. extending the reach in d_0 or z_0, or by reducing the p_{T} for certain channels.

Table 2.7: *The CPU required for large radius tracking in* HS06 $\times$ sec *to reconstruct* $t\bar{t}$ *MC events with* $\langle\mu\rangle = 140$ *and 200 in the ITk. Listed are the results for the different reconstruction steps needed to reconstruct secondary tracks from late decaying particles, that are in addition to the primary fast reconstruction. See the text for details.*

$\langle\mu\rangle$	Primary Tracking	Strip Sp.Points	LRT Track Finding	Total LRT Tracking
140	22.7	2.4	2.5	4.9
200	35.7	3.5	5.9	9.3

Table 2.7 summarises the CPU required for running the additional LRT tracking steps after the fast ITk reconstruction for finding the primary tracks. The results are listed for full event reconstruction on $t\bar{t}$ samples with an average of 140 and 200 pile-up, respectively. Running this configuration on e.g. 20 kHz of events at pile-up of 200 would increase the CPU cost for tracking only marginally, namely by less than 4%.

Performance of the Fast ITk Track Reconstruction

The tracking performance of the fast ITk reconstruction prototype has been documented in an ATLAS note [3] and in the ATLAS Phase-II Computing Conceptional Design Report [13]. For this document, the results have been updated using the latest Phase-II upgrade release, an updated detector geometry and further tuning of the track reconstruction algorithms. The MC simulation samples used in this report are using pixel sensors with a pixel size of $50 \times 50\ \mu\text{m}^2$. [6]

For this study, single muon samples were used to illustrate the technical performance and $t\bar{t}$ samples with $\langle\mu\rangle = 140$ and 200 are used to estimate tracking performance in hadronic jet events with Phase-II levels of pile-up. The $t\bar{t}$ sample plots are shown with a p_{T} cut of 2 GeV on the truth particle (except for the p_{T} plots themselves) to avoid turn-on effects due to the efficiency variation at the reconstruction p_{T} cut.

Figure 2.4 shows the efficiency as a function of η for the fast and the default ITk reconstruction on single muon samples. Efficiency losses for the fast track reconstruction are visible for $|\eta|$ close to 4 and in the barrel/end-cap transition region of the Strip Detector (see Figure 1.1), attributed to the preliminary nature of the tuning of the fast track reconstruction software.

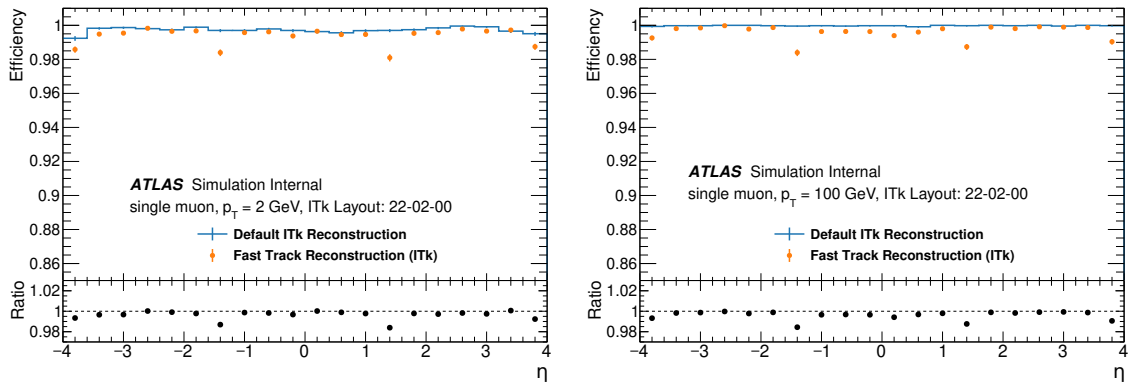

Figure 2.4: *Track reconstruction efficiency for single muon samples with p_{T} of 2 GeV (**left**) and 100 GeV (**right**) as a function of η for the fast and the default ITk reconstruction. The ratio is given by the efficiency for the fast reconstruction divided by the efficiency for the default reconstruction.*

Figure 2.5 shows the efficiency on $t\bar{t}$ events with $\langle\mu\rangle = 140$ and 200 as a function of η and p_{T} for the fast and the default ITk reconstruction. The fast reconstruction has a slightly lower efficiency at large η. Increasing the pile-up to 200 does not add visible track reconstruction inefficiencies. The by far dominating inefficiency is caused by hadronic interactions in the detector material in both cases. The efficiency as a function of p_{T} shows a more pronounced

[6] The final ITk design is using for the barrel staves in pixel layer-0 3D sensors with a pitch of $25 \times 100\ \mu\text{m}^2$ to further improve the b-tagging performance, see Reference [2] for details.

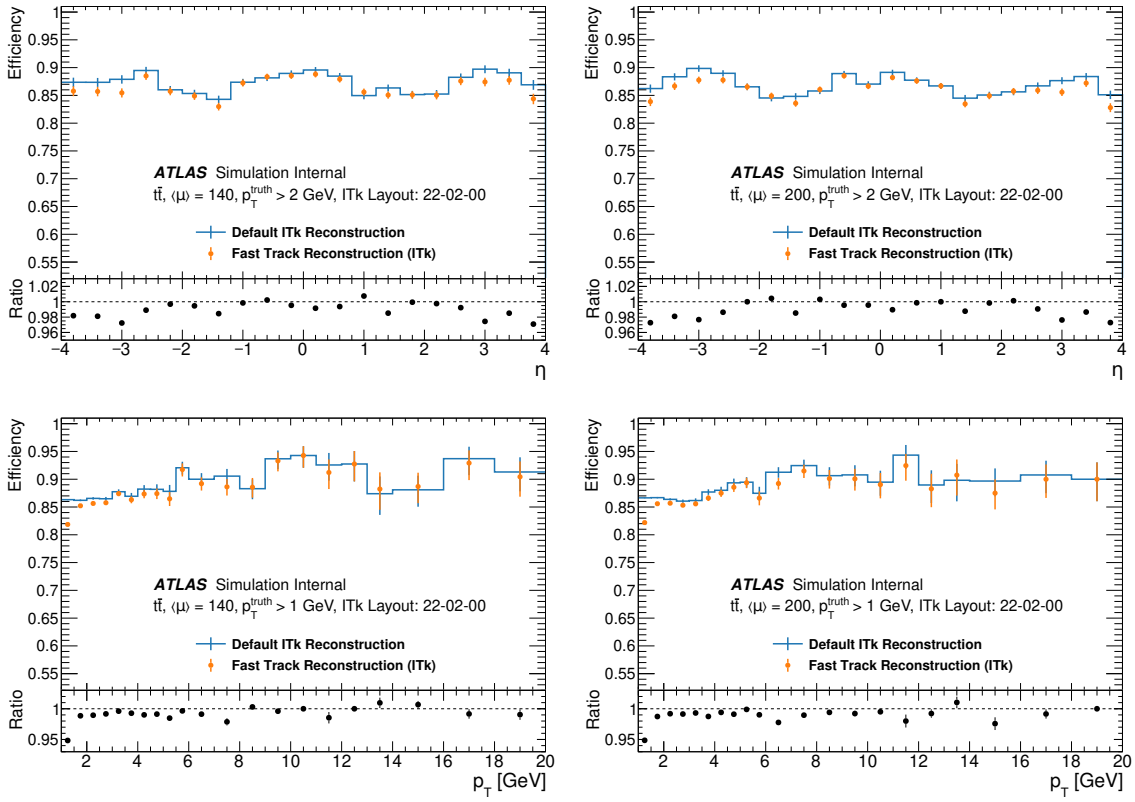

Figure 2.5: *Tracking efficiency as a function of η* (**top**) *and p_{T}* (**bottom**) *for the fast and the default ITk reconstruction. Samples of $t\bar{t}$ events are used with and $\langle\mu\rangle = 140$* (**left column**) *and 200* (**right column**). *A p_{T} cut of 2 GeV is used for the generated particles, to avoid turn-on effects. The ratio is given by the efficiency for the fast reconstruction divided by the efficiency for the default reconstruction.*

turn-on effect at 1 GeV, compared to the default reconstruction. This is attributed mostly to the increased p_{T} cut from 900(400) MeV to 1(0.4) GeV for the central(forward) part of the detector.

The number of tracks per event as a function of η is shown in Figure 2.6. This is an inclusive measure of the rate of additional tracks, fakes or duplicates, compared to the default reconstruction, which itself has a very low fake rate of around 10^{-4} for $\langle\mu\rangle = 200$ [2]. The fast reconstruction finds very similar track rates for most of the η coverage, illustrating the low rate of fakes for fast reconstruction, with no significant dependence when changing from 140 to 200 pile-up. Only in the very forward region a small increase in the number of found tracks is observed.

The number of pixel and strip hits associated to reconstructed tracks is a measure of the hit association efficiency of the reconstruction. Figure 2.7 shows a comparison between fast

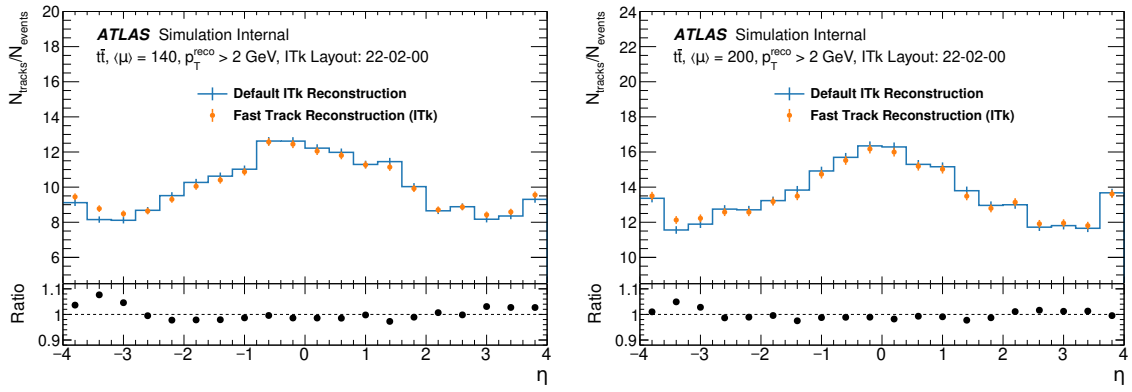

Figure 2.6: *Inclusive track rate as a function of η for the fast and the default ITk reconstruction. Samples of $t\bar{t}$ events are used with $\langle\mu\rangle = 140$ (**left**) and 200 (**right**). A p_{T} cut of 2 GeV is used to avoid turn-on effects. The ratio is given by the number of tracks for the fast reconstruction divided by the number of tracks for the default reconstruction.*

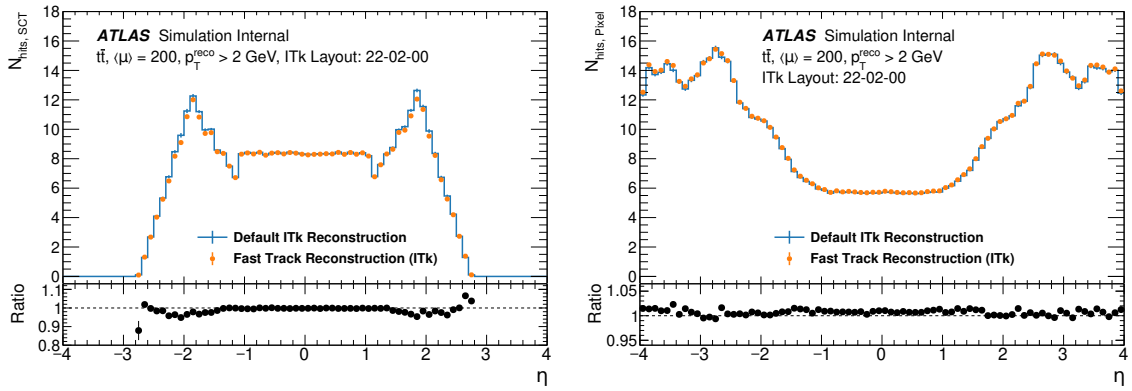

Figure 2.7: *Number of strip (**left**) and pixel (**right**) hits as a function of η for the fast and the default ITk reconstruction. A sample of $t\bar{t}$ events is used with $\langle\mu\rangle = 200$. A p_{T} cut of 2 GeV is used to avoid turn-on effects. The ratio is given by the number of hits on tracks for the fast reconstruction divided by the number of hits on tracks for the default reconstruction.*

and the default ITk reconstruction. The hit association to tracks is very similar for both versions of the reconstruction.

Figure 2.8 shows the d_0 and z_0 impact parameter resolutions as a function of η, as well as the relative resolution of the inverse transverse momentum for single muon samples of 2 and 100 GeV. The fast ITk reconstruction almost matches the resolution of the default reconstruction. At low p_{T} some residual differences are visible in the very forward region, illustrating remaining approximations in the material model in the fast Kalman filter track fit. Also at high p_{T} the fast track fit does not fully reproduce the default reconstruction performance.

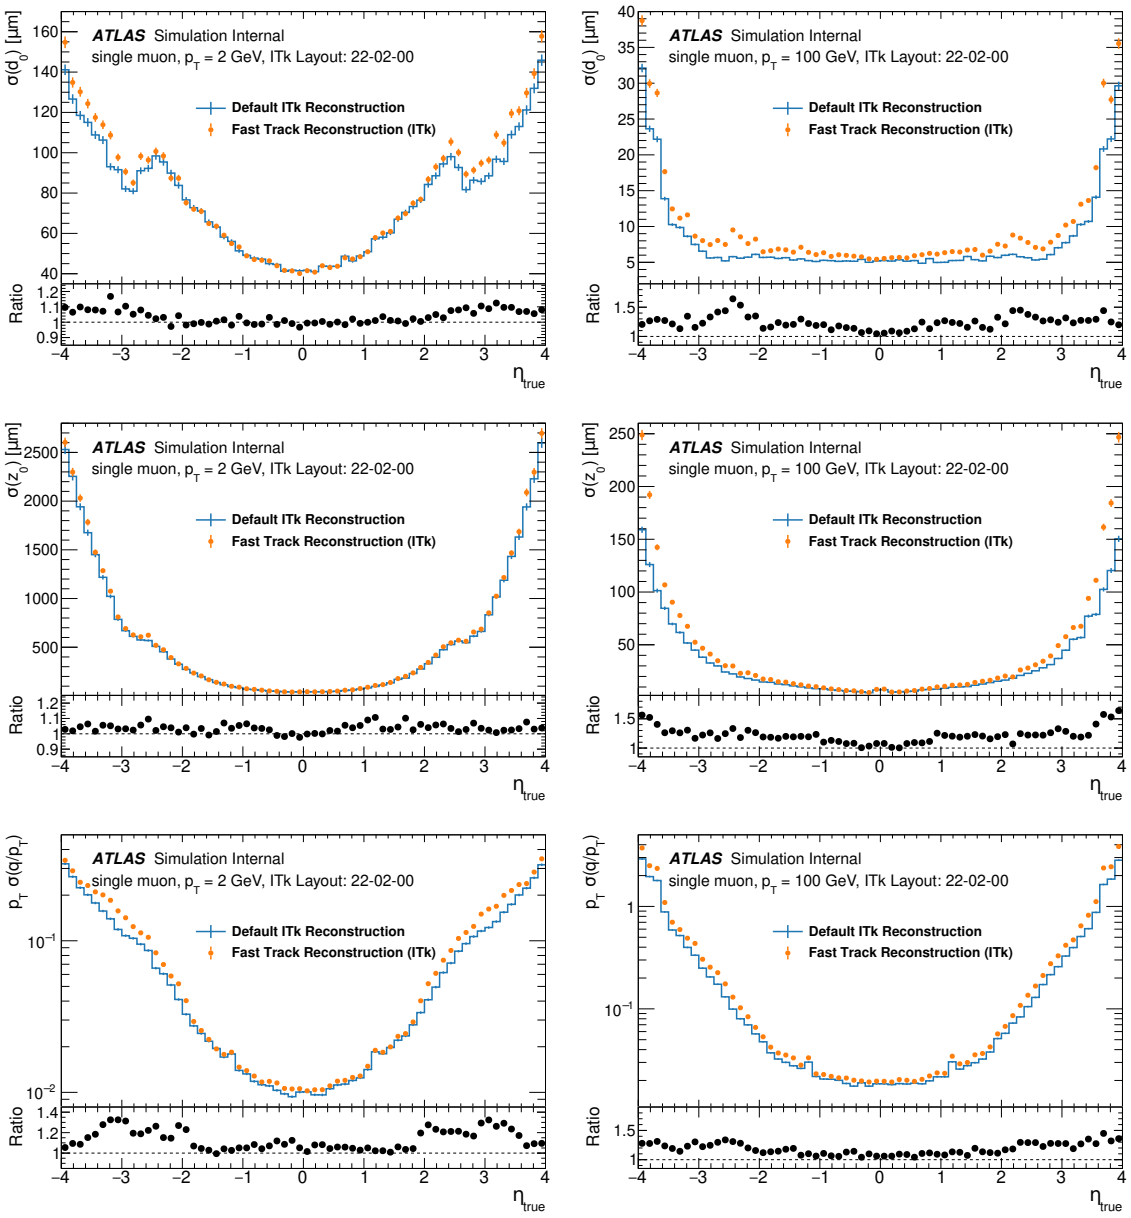

Figure 2.8: *Track parameter resolutions for a single muon sample of 2 GeV* (**left**) *and 100 GeV* (**right**) *for the fast and default ITk reconstruction. Shown are the* d_0 (**top**) *and* z_0 (**middle**) *impact parameter and the relative inverse transverse momentum* (**bottom**) *resolutions, as a function of* η*. The ratio is given by the resolution for the fast reconstruction divided by the resolution for the default reconstruction.*

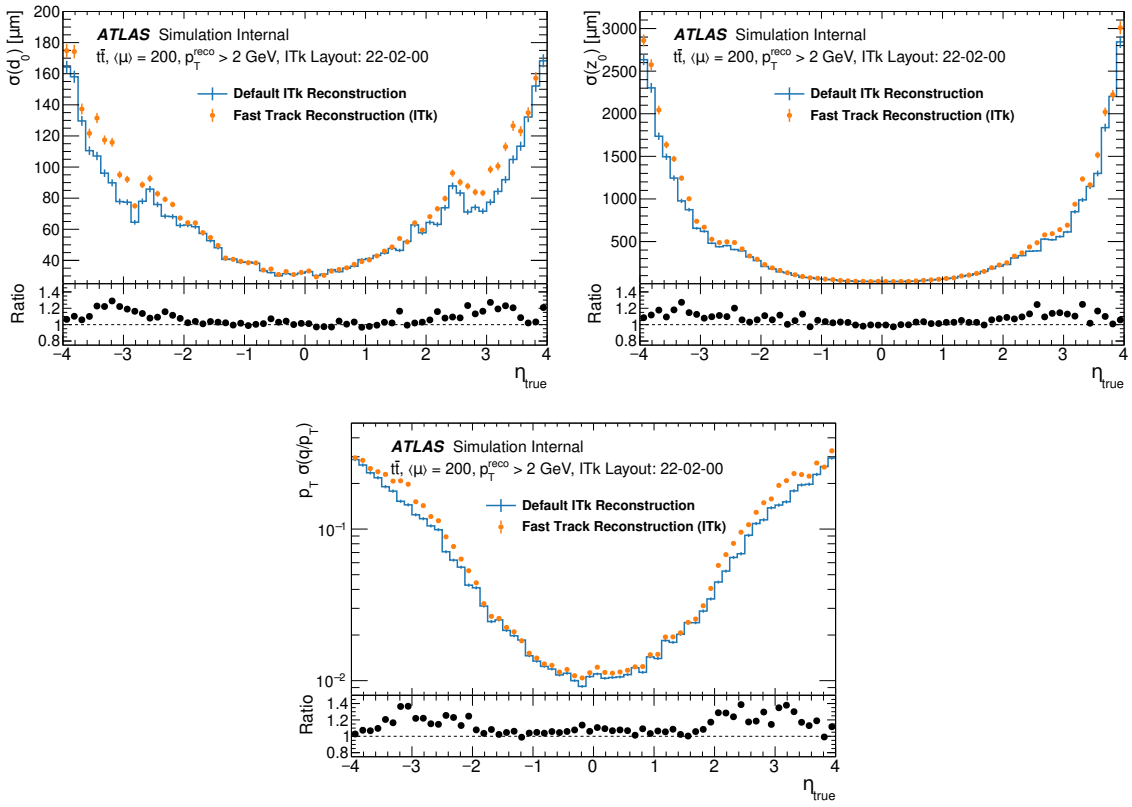

Figure 2.9: *Track parameter resolution for transverse impact parameter d_0 (**left**), longitudinal impact parameter z_0 (**right**) and relative transverse momentum p_T (**bottom**) as a function of the true pseudorapidity. Shown are results for the fast and the default ITk reconstruction. A sample of $t\bar{t}$ events is used with $\langle\mu\rangle = 200$. A p_T cut of 2 GeV is used to avoid turn-on effects. The ratio is given by the resolution for the fast reconstruction divided by the resolution for the default reconstruction.*

Figure 2.9 shows the d_0 and z_0 impact parameter resolutions, and the relative resolution of the inverse transverse momentum, as a function of η, for a $t\bar{t}$ sample with $\langle\mu\rangle = 200$. The ratios between the fast and default ITk reconstruction are at a similar level as seen in Figure 2.8 for low-p_T muons.

The efficiency for electron track reconstruction in the ITk is limited by the effect of Bremsstrahlung in the detector material. The default ITk reconstruction aims at recovering part of the inefficiency due to the forward radiation of a Brem.-photon in the combinatorial track finder in the Silicon Track Finding and in the full track fit in the Ambiguity Resolution. In both cases, the track model in the fit is altered to allow for the effect of the Brem.-photon (see [16] for details). Figure 2.10 shows the efficiency for electrons, muons and pions of $p_T = 10$ GeV as a function of η for both, the default and the fast ITk track reconstruction. As expected, the efficiency for muons is close to 1 for all η values. The inefficiency for pions is driven by the rate of inelastic hadronic interactions in the detector material. The efficiency

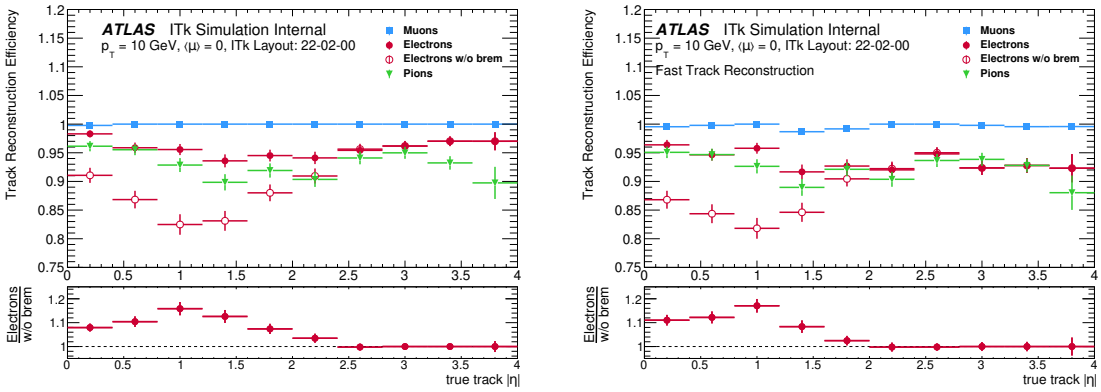

Figure 2.10: *Tracking efficiency for* $p_{\mathrm{T}} = 10$ GeV *particles as a function of* $|\eta|$ *for the standard ITk track reconstruction (***left***) and for the fast ITk track reconstruction (***right***). Compared are the results for muons, pions and for electrons with and without Brem.-recovery. See text for details. The ratio illustrates the relative gain from the Brem.-recovery for electron tracks.*

for electrons also follows the trend of the detector material, but part of the inefficiency can be recovered thanks to the Brem.-recovery in the fitting, in particular in the central region, where the material induced inefficiency before Brem.-recovery is the largest. It is visible that the fast reconstruction, before Brem.-recovery, has a slightly lower electron efficiency compared to the default reconstruction. For the trigger reconstruction, the Brem.-recovery would only be applied inside electron and di-electron RoIs, reducing the rate at which this technique is applied.

Performance of fast Large Radius Tracking

The track reconstruction efficiency for the fast LRT is estimated using two SUSY samples corresponding to the production of a pair of gluinos. Those gluinos are forced to decay into a neutralino and two quarks. The neutralinos are afterwards decaying through an R-parity violating coupling into three quarks, with a lifetime of 0.1 ns. In the first (second) sample, the mass of the gluino is set to 1600 (2000) GeV, while the neutralino mass is set to 50 (1950) GeV. The samples differ in the neutralino decay length and in the observed spectrum and angles of secondary decay products. In the 1950 GeV neutralino sample, the long-lived particles decay products are typically produced at a radius below 100 mm, while in the 50 GeV neutralino sample those can be produced beyond the last pixel layer due to the larger boost of the neutralino. In that case, the decay products are produced with a relatively small angle with respect to the neutralino direction of flight, leading to tracks typically pointing towards the hard-scatter vertex, while the heavy neutralino decays can produce a very large fraction of non-pointing tracks. Those samples have been simulated without pile-up. The efficiency is estimated for stable charged particles produced in the decay chain, with $p_{\mathrm{T}} > 5$ GeV and $|\eta| < 2.4$. The performance obtained with the fast LRT

configuration is shown in the following figures, together with the performance of the current offline LRT configuration, still under development, for reference. The larger efficiency observed for the offline LRT reconstruction for $R < 300$ mm is due to the additional pixel seeds used in association with the strip seeds.

Figure 2.11 shows the tracking efficiency for secondary tracks with $p_{\mathrm{T}} > 5$ GeV and $|\eta| < 2.4$ as a function of the neutralino decay radius R, η and p_{T} of its decay products. The efficiency of the fast LRT tracking is slightly lower than the results for the offline default LRT reconstruction, mainly due to a slightly more stringent set of selection cuts applied. The flexibility of the LRT reconstruction software allows to tune the fast LRT strategy to a particular trigger signature, further optimising tracking and/or CPU performance.

Conclusion

The CPU-based demonstrator described in this section provides confidence that the final EF Tracking system will meet all necessary specifications, including the tracking requirements detailed in Section 2.2.2. This demonstrator provides near offline-quality tracking performance in the phase space covered by these studies with CPU requirements that are compatible with the available power and space at Point-1. Further optimizations towards a final system are described in Section 2.5.

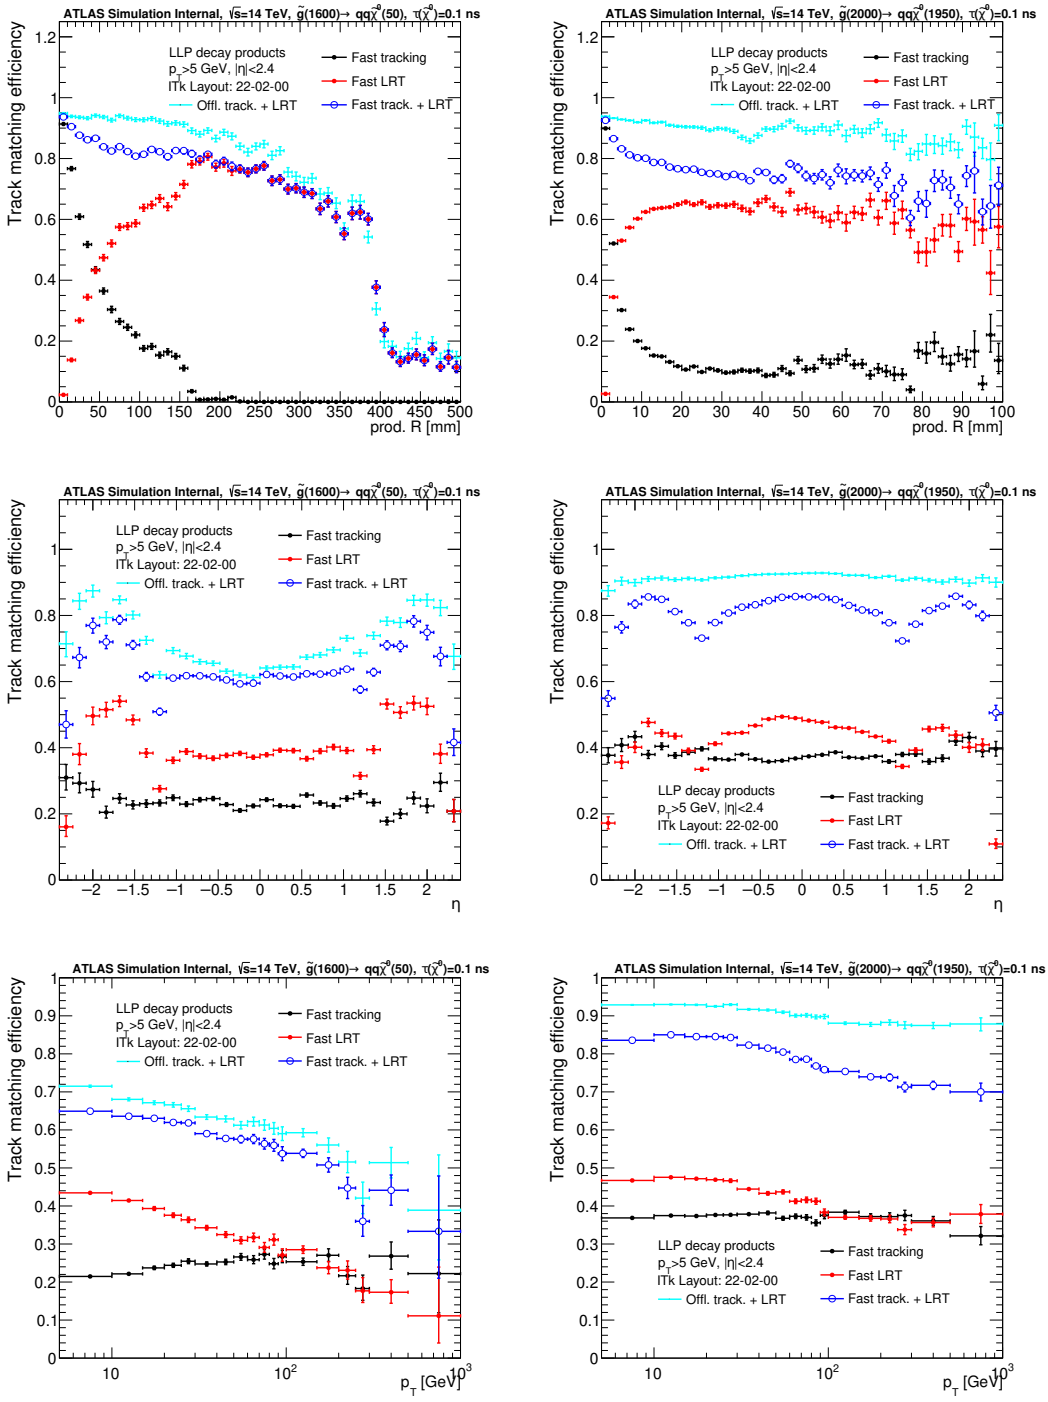

Figure 2.11: *Track reconstruction efficiency as a function of the production radius, η and p_T for the fast LRT reconstruction. The efficiency is estimated for long-lived neutralino decay products with p_T >5 GeV and $|\eta| < 2.4$ in two SUSY RPV samples at pile-up of 0. The efficiency is compared to the one of the current offline LRT reconstruction for reference.*

2.4.2 FPGA-based Demonstrator

Technical description

In establishing the viability of a heterogeneous commodity system, a tentative reconstruction flow for ITk data has been developed to demonstrate its feasibility. Lightweight CPU-based load-balancing software would route a particular event to an available FPGA. The following functionality would then be implemented in firmware in each FPGA: 1) decoding of the ITk data and clustering the hits in each ITk layer; 2) the option to prefilter and/or gang the ITk hit clusters to form "spacepoints" and/or "stubs"; 3) a pattern recognition step that analyzes a subset of hit clusters or spacepoints to identify hit combinations that would correspond to likely tracks; 4) and finally a series of duplicate and fake removal steps followed by an initial track fit. The most resource-intensive functionality is the pattern recognition stage; in the prototype example explored here, the assumed implementation is a Hough transform, but additional possibilities are discussed in Section 2.5.3 and will be considered in the future. Depending on findings of the optimization, the chosen configuration may not have all four of these pieces implemented in accelerators, and the demonstrator discussed in this section does not imply a specific implementation choice for the future. In the model studied here, the output from the FPGA is a set of track candidates and their associated hit clusters, along with any additional associated hit clusters, that would be passed to a precision Kalman filter that has been developed for the fast ITk track reconstruction for a final precision fit. The studies performed for this section use the geometry for the ITk (22-02-00) shown in Figure 1.1.

The border between the FPGA duplicate removal, fake rejection, and initial fitting functionality and the precision fit in software may still evolve with further study. In particular, the possibility of implementing an extrapolation and/or a second-stage track fit on the FPGA has not been studied in detail. These functionalities would require additional FPGA resources but would improve the quality of the tracks passed to the CPU. Additional study (including a software interface) would be needed to determine the overall optimization of a true heterogeneous system.

Hough Transform for Pattern Recognition The Hough transform [17, 18] has been investigated in the context of reconstructing charged particles that originate from the detector origin in the ATLAS ITk [19, 20]. The basic concept is that circular arcs traced out in the transverse $(x - y)$ plane can be transformed onto the qA/p_{T} versus ϕ_t plane, where q is the charge of the particle, A is the curvature constant for a 2 T magnetic field ($A = 3 \times 10^{-4}$ GeV mm^{-1}), p_{T} is the transverse momentum of the particle, and ϕ_t is the azimuthal angle of the particle at the origin. Defining r_h and ϕ_h as the radius and azimuthal

angle of the hit cluster in the detector, each hit cluster on a curved track in the transverse plane will form a straight line in the qA/p_{T}-ϕ_t plane:

$$\frac{qA}{p_{\mathrm{T}}} = \frac{\sin(\phi_t - \phi_h)}{r_h} \approx \frac{(\phi_t - \phi_h)}{r_h}. \tag{2.1}$$

The intersection of N lines, where N is a configurable threshold, indicates that the given set of hit clusters are consistent with coming from the same charged particle track. This qA/p_{T}-ϕ_t space is frequently referred to as an "accumulator". This approach is intrinsically two-dimensional (2D) and assumes that the charged particles come from the origin; in order to include additional z-information due to the spread of the beamspot, one can repeat the 2D scan in “slices" in the z-direction, as described below.

The nominal strategy for the Hough transform firmware implementation is to define the accumulator for an $\eta \times \phi = 0.2 \times 0.2$ region with a given binning in qA/p_{T} and ϕ_t. Thus, for full $\eta - \phi$ coverage ($|\eta| < 4.0$), 1280 regions can be defined across the detector. The nominal number of tracker layers used is eight, though this can be optimized in the future. In the central barrel, these layers are comprised of the outermost pixel layer, the inner side of the first strip layer, and both sides of the remaining strip layers.

The nominal choice to use the outer layers is based on the relative occupancy of the ITk. In order to handle the inner layers, a Hough-like scan in both p_{T} and d_0 would be required, as the requirement of 2 mm impact parameter coverage causes significant ϕ shifts in the inner pixels.

An example Hough transform image for single muons, without and with pileup, is shown in Figure 2.12. This figure illustrates the challenge presented by the HL-LHC environment.

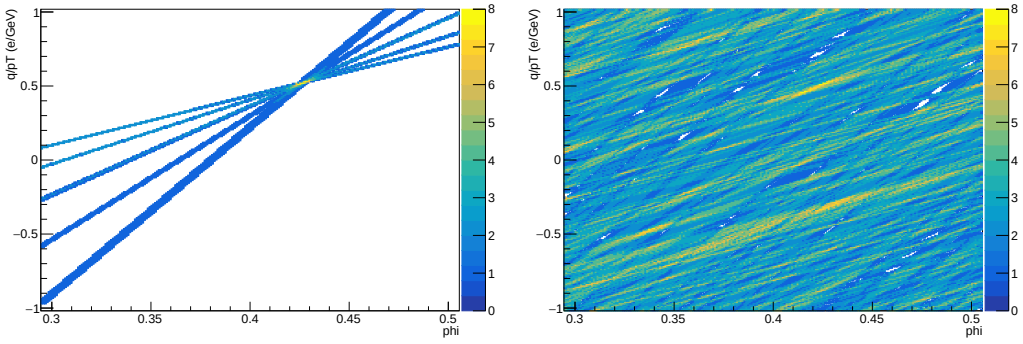

(a) (b)

Figure 2.12: *Hough transform images of a single muon at true $\phi_t = 0.43, q/p_{\mathrm{T}} = 0.52$ in a single z-slice, (a) without and (b) with pileup. The single muon candidate is reconstructed by the Hough transform pattern recognition; additional hit combinations from pileup can also be observed.*

A first step in tackling the large detector occupancy of the HL-LHC is dividing an $\eta - \phi$ region into so-called "slices"in the z-direction [21]. It was determined that slicing in z at a radius of 562 mm (a so-called "key layer") reduced the duplication compared to defining slices at $r = 0$ at central η; the nominal number of key-layer slices was chosen to be 6 in the central η region. The next steps in the optimization of the Hough transform algorithm include: qA/p_T binning, ϕ binning, number of key-layer slices, padding, and hit extension. Here the hit padding refers to the inclusion of additional bins outside of the nominal $p_T - \phi$ ranges of the accumulator. In other words, it extends the accumulator at the edges beyond the nominal region. The hit extension refers to the number of additional, adjacent bins filled in the accumulator in the ϕ direction to account for resolution effects and tracks with non-zero d_0 values. A hit extension of zero implies only the bin in the accumulator from Equation 2.1 is filled. A hit extension of n means that this bin is filled, and in addition, so are n bins on each side in the ϕ direction. The hit extension can be chosen separately for each accumulator so that it can differ per silicon layer. Figure 2.13 shows a pictorial representation of the studied Hough transform options. The baseline scenario uses a hit extension setting of 2 in the single used pixel layer, 1 in the first used strip layer, and 0 elsewhere, as well as paddings of 6 (2) in the ϕ (q/p_T) direction of the accumulator image.

A configurable threshold is applied to the Hough transform accumulator, requiring hit clusters or spacepoints/stubs in a certain number of layers; each bin that passes this threshold is referred to as a road. Additional fake rejection, duplicate removal, and fitting is required to bring the number of roads (and thus track candidates) down to a manageable level for the precision fit. The performance of these strategies is described in Section 2.4.2.

Technical overview of system The following assumptions are made in order to estimate the proposed size and layout of a heterogeneous Hough transform-based commodity system:

- FPGA-based hardware accelerator cards are assumed.
- The FPGA target used for firmware prototyping and resource estimates is the commercially available Xilinx Alveo U250, wherever possible. The outlook for potential future FPGA targets is discussed in Section 2.5.
- A full ITk event (1.6 MB (pixel) + 0.5 MB (strip) = 2.1 MB) can be processed by each FPGA accelerator. Summing the estimated FPGA resource usage indicates this is a reasonable, but perhaps not conservative, assumption for the three firmware implementation options.
- The Hough transform has been adopted as the strategy for the pattern recognition functionality on the FPGA. Advanced algorithms for pattern recognition that take advantage of machine learning algorithms such as Graph Neural Networks (GNNs) will be considered in the future; the potential for these strategies are discussed in Section 2.5. However, machine learning techniques are considered for fake track rejection here.

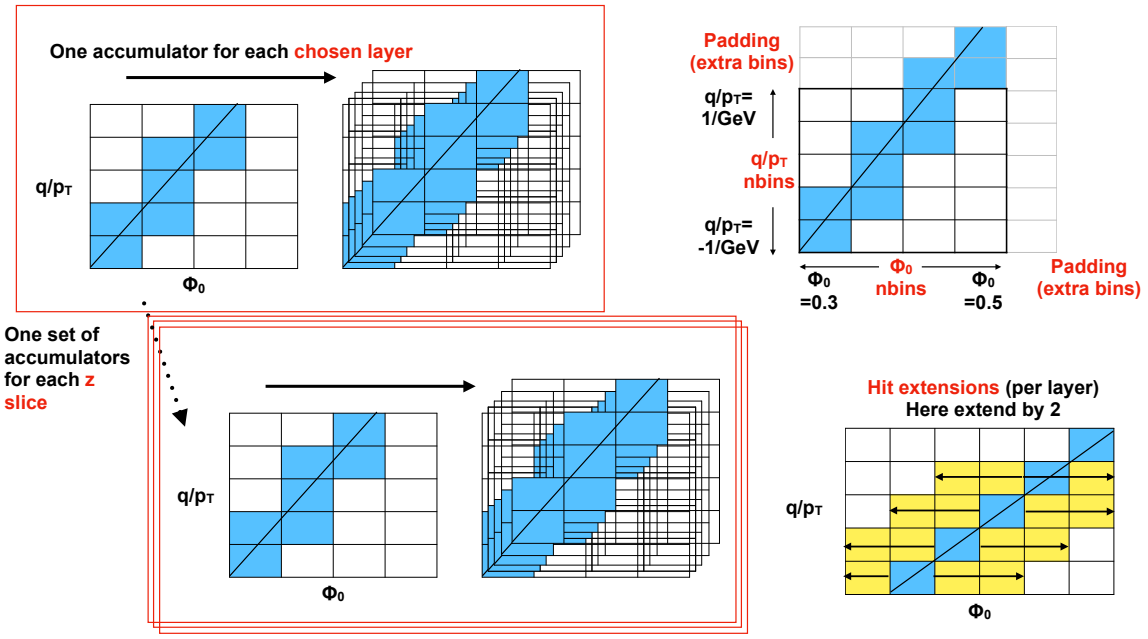

(a)

Figure 2.13: *Example options available for the Hough transform optimization. The Hough transform runs on a set of chosen silicon layers. Accumulators are filled for each of these layers, and evaluated for a given number of slices along the z direction. Each accumulator can have an arbitrary number of bins in the ϕ and q/p_T dimensions, and can be "padded" with extra bins beyond the nominal values to not lose efficiency near region edges. For each q/p_T bin, all ϕ values consistent with a point in the q/p_T range are filled. In addition, to maintain high efficiency, particularly for non-prompt tracks, an extra "hit extension" can be added such that n extra bins on either side in ϕ are filled. The hit extensions can be set separately for the accumulators of each layer.*

- Techniques such as pre-filtering hit clusters (“stub-finding") and combining hit clusters on opposite sides of a stave (“spacepoint formation") are promising as they could reduce the number of input points to the pattern recognition. Studies indicate that these techniques can be reasonably robust against radiation damage.
- For each track candidate, track parameters as well as identifiers for and positions of the clusters associated with that candidate, are sent to the EF CPU to be used in an extension to layers not considered in the FPGA (parameters) and as part of the final precision track fits (associated hit clusters).

The number of track candidates per event exiting the FPGA accelerator determines the overall CPU usage for the precision fit. An important area of system optimization is the precise determination of this boundary between functionality performed in hardware versus software (including track extrapolation, duplicate and fake track removal, and the track fitting strategy). The default scenario considered here is one where the full track extension

and 13-layer fit is done in the FPGA, followed by only the final precision fit in CPU for the final rejection. Other options are discussed in Section 2.5.4.

FPGA Firmware Development Resource Estimate Summary This section provides a high-level summary of the firmware blocks developed or referenced for the heterogeneous commodity system. These blocks are shown in Fig. 2.14 and described briefly below. The resource estimates are based on the Xilinx Alveo U250 FPGA.

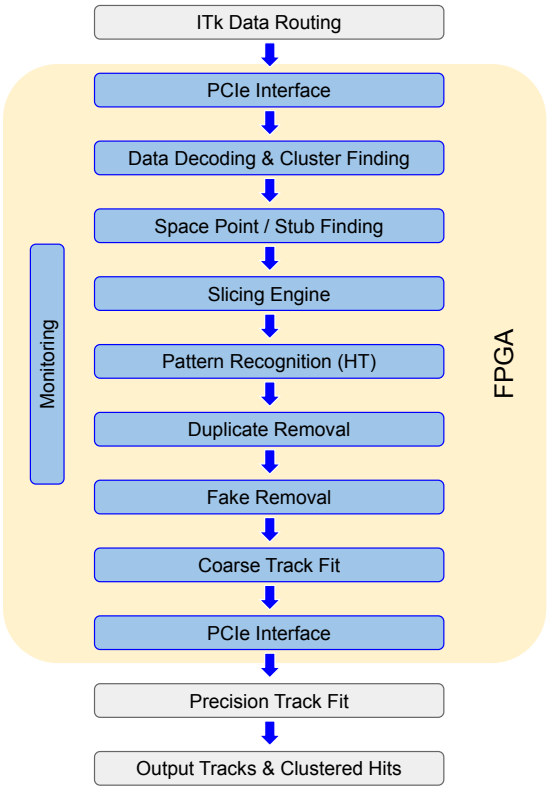

Figure 2.14: *Diagram illustrating the firmware blocks in the heterogeneous commodity system.*

PCIe: The PCIe module has been designed for the ATLAS Front-End LInk eXchange (FELIX) project and provides a simple Direct Memory Access (DMA) interface.

Clustering: The clustering module provides the cluster position and size information from the raw front-end data ITk pixel and strip data. The resource estimate is based on the Xilinx VU37P FPGA, which is similar in size/resources to the Alveo U250.

Stub-Finding: The stub-finding module inputs the data stream containing ITk strip event data, extracts the cluster data, and performs one or both of stub-filtering and space-point formation. It then replaces the cluster data with the spacepoint data and passes the resulting event data to the next stage.

Slicing Engine: The slicing engine sorts the input data based on the z-slicing configuration and $\eta - \phi$ region size.

Coordinate Conversion: The local-to-global coordinate conversion converts local coordinates from each ITk cluster and converts them to global hit positions.

Pattern Recognition: Two complementary approaches to the firmware implementation of the Hough transform accumulator have been investigated. Though the full event is loaded on the FPGA, data corresponding to a single $\eta - \phi$ region are loaded into the accumulator at a given time; individual regions and z-slices are processed sequentially:

- A nominal 2D Hough transform implementation as described above.
- The 1D Hough transform implementation, which is significantly different and utilizes the fact that within a slice the radii of the layers are close enough to constant that for a coarse ϕ binning ($\mathcal{O}(64 \text{ per } 0.2\,\text{rad})$) the bin migration due to the track angle for tracks with $p_{\text{T}} \geq 1$ GeV is of order one bin. Only ϕ bins are constructed and the p_{T} binning is achieved by bit shifts This produces more candidate tracks, which are then filtered using more z information and finally full-granularity ϕ information.

Duplicate Removal: The duplicate removal module is estimated based on a Xilinx Kintex-7 implementation from the Fast TracKer (FTK) project [22]. The resources are scaled to the number of track inputs expected from $\langle\mu\rangle \approx 200$ implementation and quoted with respect to the available resources on the Alveo U250.

Fake Removal (Neural Network (NN)): The architecture of the NN used to reject fake tracks has been developed with a firmware implementation in mind; the fake rejection module is implemented using HLS4ML [23]. This NN is described in further detail later in this section.

Coarse Track Fit: The track fitting module extracts the candidate tracks from the roads found by the pattern matching. A set of possible hit combinations is produced with one hit per layer. The χ^2 of each candidate track is computed and if a selected threshold is passed, the track parameters are then evaluated. The resource usage is determined from implementation on an Intel Stratix 10 MX FPGA.

2nd-Stage Track Extension and Fitting: The 2nd-stage module performs an extension of the 8- or 9-layer track formed in the pattern recognition stage through to the innermost pixel layers to form a full 13-layer track. The hits from the 1st stage and matrix constants are used to extrapolate to the new layers. A fit similar to the first-stage fit is performed based on all 13 layers.

Monitoring (IPBus): The monitoring module handles the communication between the FPGA and PC; the communication occurs via ethernet using IPbus protocol over User Datagram Protocol (UDP). The resource usage is determined from implementation on an Intel Stratix 10 MX FPGA. Another option being considered is to take advantage of the experience in developing the FELIX firmware for this functionality.

The estimated resource usage for each block is summarized in Table 2.8. Care should be taken when interpreting these numbers and the total resource usage in each column. The purpose of this table is to demonstrate the feasibility of each block and demonstrate that a bottom-up estimate of the overall resource usage fits within the resource budget on the Alveo U250. It is expected that the total resources would change after integrating the various firmware modules.

An illustration of the pipelined dataflow through the firmware is presented in Fig. 2.15.

Processing Time and System Size The event processing time is dominated by the input and output of the pattern recognition firmware block. Two quantities are of interest: the time it takes to load hit clusters into the Hough transform accumulator (t_{loading}) and the time it takes to read out the accumulator (t_{readout}). The resulting event processing times (and rates) are shown in Table 2.9. Given these rates, and assuming an EF tracking rate of 150 kHz full-scan tracking + 5% × 1 MHz regional tracking, one can determine the number of FPGA accelerators required by dividing the EF tracking rate by the event processing rate for a given FPGA implementation and scenario. The resulting numbers of FPGA accelerators (N_{accel}) required for each implementation and scenario are also summarized in Table 2.9.

Byte Stream Decoding, Clustering, and Spacepoints The byte stream decoding, clustering, and spacepoint formation of the full event will be performed on the FPGA and the clusters are expected to be suitable for use in the software[7].

FPGA Output Unpacking Because the FPGA firmware can be very flexible in the construction of the data structure, it is possible to design the output data format from the FPGA such that the corresponding CPU usage to read this data would be comparable to the CPU usage to read the software data model input in the track extension and fitter step. These output data would include the hit cluster positions, identifiers, and the parameters (p_{T}, quality, etc) and hit clusters associated with each track candidate. As a conservative estimate, we scale the ITk byte stream decoding estimate (3.2HS06 × sec/2.1 MB) according to data volume estimates. We then assign a range of 50% to 100% of the conservative estimate. The offline track Event Data Model (EDM) format would be created at the end of the software functionality.

From HTT studies, it is estimated that the average number of bits per 8- or 9-layer track to be 640 bits and the average number of bits per 13-layer track to be 896 bits (including the associated hit clusters). We assume a size of 32 bits/cluster for the pixel clusters, which

[7] Detailed comparisons of the performance for the current algorithm show 1-2% cluster splitting/merging differences. Further development work to get closer to the offline performance is anticipated. For example, the calculation of the actual cluster centroid can follow the offline implementation.

Table 2.8: *Estimates of firmware resource usage for a Xilinx Alveo U250 FPGA. The Hough transform accumulator settings can be tuned to optimize the FPGA usage and requirements. The range for the clustering indicates how the load will depend on the number of FPGAs (a smaller system means a larger fraction of the FPGA needs to do clustering). Since the track fitting module has been estimated on a different FPGA target, only an indication of the relevant resource usage has been provided. The 1D and 2D Hough firmwares correspond to two potential pattern recognition blocks, and have different minimally efficient $\eta \times \phi$ regions, as indicated. The range of 2D firmware resource usage corresponds to different preliminary implementations from two groups. These regions (and z-slices) are processed sequentially through the same pattern recognition block. The difference is accounted for in determining the event processing rate shown in Table 2.9. In the Xilinx architecture there is three times more Ultra Random Access Memory (URAM) than Block Random Access Memory (BRAM); for many applications the URAM can be used in place of BRAM.*

Firmware Block	LUT (%)	flip-flop (FF) (%) Logic Functions	BRAM/ URAM (%)	DSP (%)
PCIe	0.6	0.6	0.3	–
Clustering	1−4	0.14−0.51	1.3−5.4	–
Stub-Finding	0.2	0.05	0.1	–
Slicing Engine	0.1	0.07	13	–
Hough (2D, 0.2×0.2)	39-59	10-30	1-5	1.8-21
Hough (1D, 0.2×0.8)	12	7	27	1
Fake Rejection (NN)	8	1	0.02	29
Duplicate Removal	1	1	–	–
Track Fitting		∼ 10	–	∼ 10
Monitoring (IPbus)		∼ 1	–	–
2nd-Stage Fitting		∼ 10	∼ 30	∼ 15
Total	**44−94**	**32−55**	**33−41**	**55−75**

Table 2.9: *Estimates of the event processing time and system size for a Xilinx Alveo U250 FPGA. Event processing time for two implementations of a pattern recognition block are provided; the range of 2D firmware resource usage corresponds to different preliminary implementations from two groups. The number of accelerators (N_{accel}) is determined using effective EF tracking rates of 150 kHz (full-scan) + 50 kHz (regional) as described in the text.*

Per Event	Firmware Implementation & Scenario Hough (2D)	Hough (1D)
Loading Time (ms)	1.9-2.8	0.7
Readout Time (ms)	2.7-3.4	1.3
Total Time (ms)	max(loading, readout) 2.7-3.4	1.3
Processing Rate (Hz)	294-534	741
N_{accel}	374-680	270

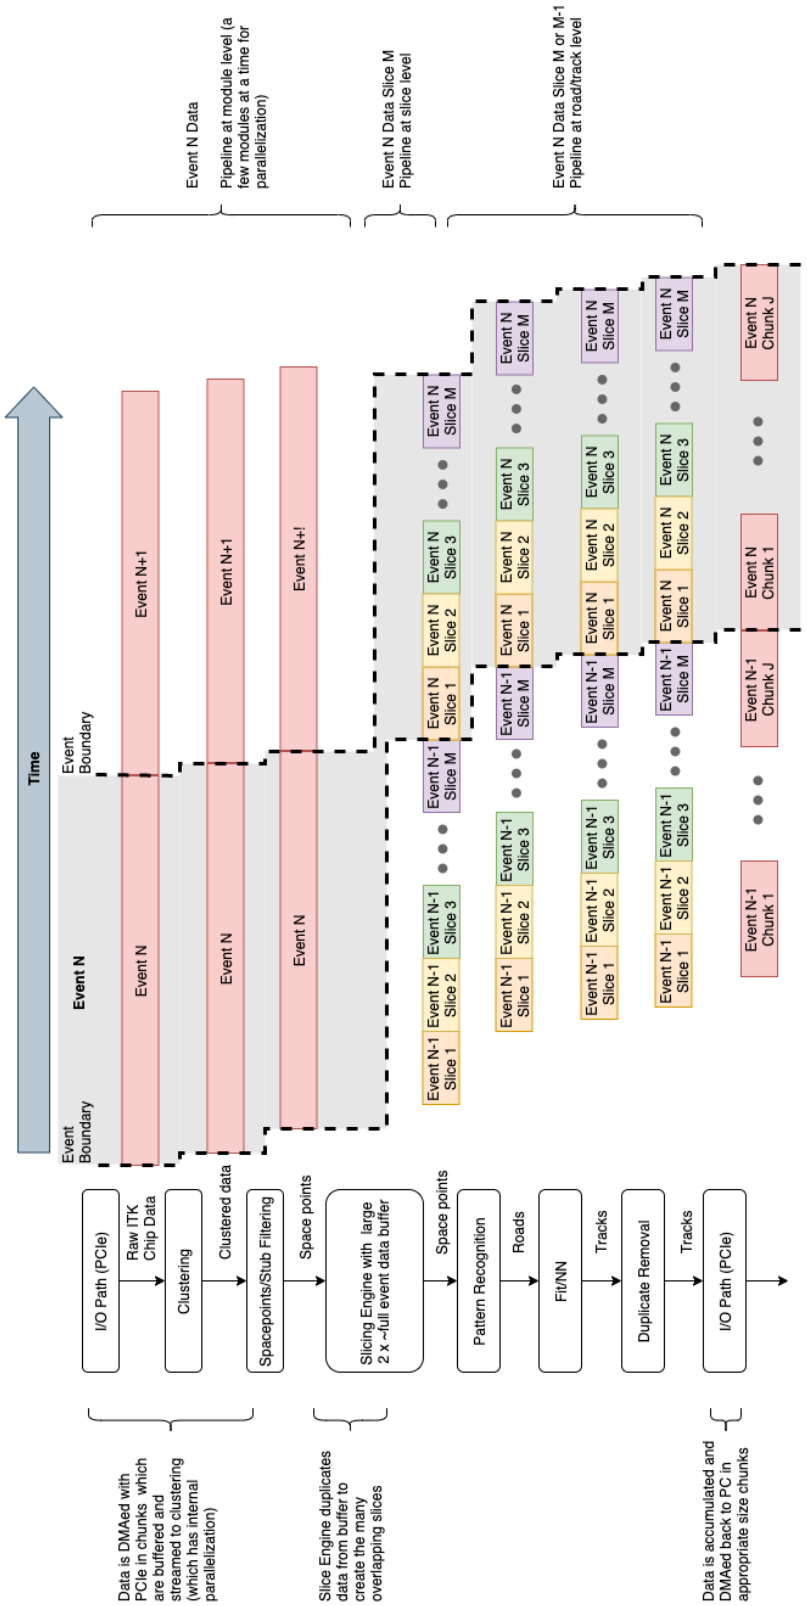

Figure 2.15: *Illustration of the pipelined dataflow through the FPGA firmware in the heterogeneous commodity system.*

is enough information to provide a position resolution of approximately 1/10th of a pixel width.

CPU Estimate and CPU System Size The heterogeneous system can be optimized to take advantage of the relative strengths of FPGAs and CPUs for high-quality EF tracking. The software capabilities here have been estimated based on extensive experience with the Kalman filter developed for the fast ITk reconstruction prototype documented in Refs. [20, 3, 9], and updated in the software-only task force report and this document. A critical element of a successful heterogeneous commodity system will be the handshake between the FPGA firmware described in previous sections and the offline fast tracking software.

As input to our calculations, we assume 10 track candidates per $\eta - \phi$ region, or 12,800 track candidates per event for the 8 or 9-layer input to the extension step. The fake track rejection factor for the Hough transform plus NN fake rejection is 9.1, 15.5, and 23.6 for $0.1 < \eta < 0.3$, $0.7 < \eta < 0.9$, and $2.0 < \eta < 2.2$. We expect the implementation of a fitter and other optimizations to reduce the fake rate further (see Section 2.5.4). Experience with the offline fast tracking software suggests that the improvement on the outer-layer pattern recognition plus fit could potentially reduce this to as low as order(5) tracks per region.

Track Extension and Fitting For the silicon track finding, it is estimated that 30% of the CPU usage is seeding, 50% is the combinatorial Kalman filter, and 20% is the final (precision) fit. The seeding is effectively replaced by the FPGA functionality.

Given that 50% of the silicon track finding is estimated to be used by the Kalman filter, we can then apply a 4/13 factor to extend the 8-layer/9-layer tracks provided by the FPGA into the inner four pixel layers. This number is then scaled by the number of seeds per event used by the offline as an input to track-finding.

Rejection from the extension has not been simulated, so we consider two estimates. The first is based on the expected gHTT performance. This gives 2.9 tracks per 0.2×0.2 region, corresponding to $2.9 \times 1280 = 3712$ total tracks. The main difference between the gHTT and proposed system is that the fitter would need to be constructed to use fewer constants because all the constants need to fit in a single FPGA as opposed to being distributed across many FPGAs. There are concepts for how to achieve comparable results with fewer constants using the Hough output as a pre-fit to correct the geometry to an idealized geometry, but they have not yet been demonstrated. CMS has demonstrated good performance with a similar pre-fit scheme [24]. The other estimate for the rate after extension is based on offline experience, which suggests that if input to the extension is reduced to $\mathcal{O}(5)$ then the extension might further reduce the number of candidates by a factor of two giving a final number of tracks entering the precision fit of 2.5, similar to the HTT-based estimate.

All the Hough simulation has been done with $\langle\mu\rangle \approx 200$, so for $\langle\mu\rangle \approx 140$, it is estimated that the CPU usage scales approximately linearly with luminosity. This is appropriate for the final fit where the track purity is high, but for the input to the extension, this is conservative since the combinatorial effects leading to fake tracks will be decreased.

Estimated Total System Size An estimate of the overall system size based on the demonstrator includes the total number of FPGA accelerators and the needed CPU resources for an effective tracking rate of 150 (full-scan)+50 (regional) kHz; this is shown in Table 2.10. The physical footprint of the accelerator part of the system is estimated to be 26-64 servers in 3-7 racks, where the ranges correspond to the envelope of the three different preliminary firmware implementations. Alveo U250 cards each take up 2 PCIe slots, and we assume 200 W of power usage per card (the cards are rated to 225 W) and 350 W per server of non-accelerator power. Assuming 48U racks and 20kW of usable power per rack, 8 cards per server and 10 4U servers / rack would be populated.

To estimate the power for the CPU farm 0.25 W/HS06 (0.20W/HS06) for Run 4 (Run 5) are assumed.

Table 2.10: *The estimated size and total power of an FPGA-based heterogeneous system for effective tracking rate of* 150 *kHz (full-scan)* + 50 *kHz (regional). The ranges in the number of accelerators and power correspond to the envelope of the three different preliminary firmware implementations. The CPU is estimated based on release 21.9 software with the range indicated a best estimate as the low end and the high value includes a +50% uncertainty on the rates input to the calculation. The Run 4 CPU estimate is derived from the Run 5 estimate by scaling down by 140/200. This a minor over-estimate of the Run 4 requirement, since the number of fake tracks will go down more than linearly with luminosity.*

	Run 4	Run 5
# of Accelerator Cards	270-680	
CPU resource requirement [MHS06]	1.1-1.7	1.6-2.4
Accelerator Power [MW]	0.08-0.18	
CPU Power [MW]	0.28-0.42	0.32-0.48
Total Power [MW]	0.4-0.6	0.4-0.7

The addition of LRT at 20 kHz using the straight track approximation, described below, would add 27% additional accelerator cards and an additional 0.10 MHS06 (0.19 MHS06) of computing in Run 4 (Run 5).

Pattern Recognition Performance

The efficiency versus p_T is given in Figure 2.16 for muons, pions and electrons in three regions of η for the ϕ range $0.3 < \phi < 0.5$. The performance for the regions in the central part of the detector is comparable to the offline software, as is the muon performance in

the endcap region. In the forward region the performance for pions and electrons is not as good as the offline performance. This region has been studied less than the other regions. The performance deficit here may be addressed with further configuration optimization or potentially modifying the underlying algorithm.

Fake Rejection and Duplicate Removal The samples used in this section are single muons overlaid on top of $t\bar{t}$ events at $\langle\mu\rangle \approx 200$. This allows for the study of busy $t\bar{t}$ events added on top of the complexity of the pileup environment, all while providing an adequate number of tracks to study. A variety of duplicate/overlap removal and fake track rejection strategies have been investigated, using the same definitions as the CPU demonstrator. The described sequence is demonstrating promising performance, but has not yet been fully optimized.

The first duplicate removal stage is a neighbor-based filter implemented on the output of the Hough transform. The procedure is as follows. For each bin in the Hough transform accumulator (histogram), look at $n \times n$ neighbors. If any neighbor:

1. has more hit layers,
2. has equal hit layers but more hits, or
3. has equal hit layers and hits, but is in a lower-left bin,

then roads are not created from this bin. This procedure thus removes “duplicates" arising from multiple bins in the accumulator. This procedure (with $n = 2$) is applied to all of the plots in this section with pileup.

The next stage is a NN trained to provide fake track rejection against fake hit combinations. Other future options discussed in Section 2.5.4 include χ^2 fits that could be used in addition to or instead of the NN. The detailed NN architecture used here was chosen with FPGA implementation in mind. The pre-processing optimization and NN architecture were chosen based on studies performed using fake hit combinations independent of the samples used in the testing, while the final training is performed by discriminating hit combinations as found by the default offline ATLAS reconstruction vs hit combinations from the Hough transform step with a truth probability score of less than 50%. Finally, the NN output scores are used to implement an overlap removal step inspired by the Hit Warrior algorithm [1]: If two or more output track candidates share more than 5 hits with one another, the track candidate with the highest NN score is selected. The NN input variables are the x/y/z values of all hits on the track candidate, pre-processed to improve performance:

1. First, the ϕ angle of the first hit is calculated, and all hits are rotated by the negative of that to remove the ϕ degree of freedom
2. The x/y/z hit positions are scaled such that max position is O(1)
3. Finally, the hits are ordered by R

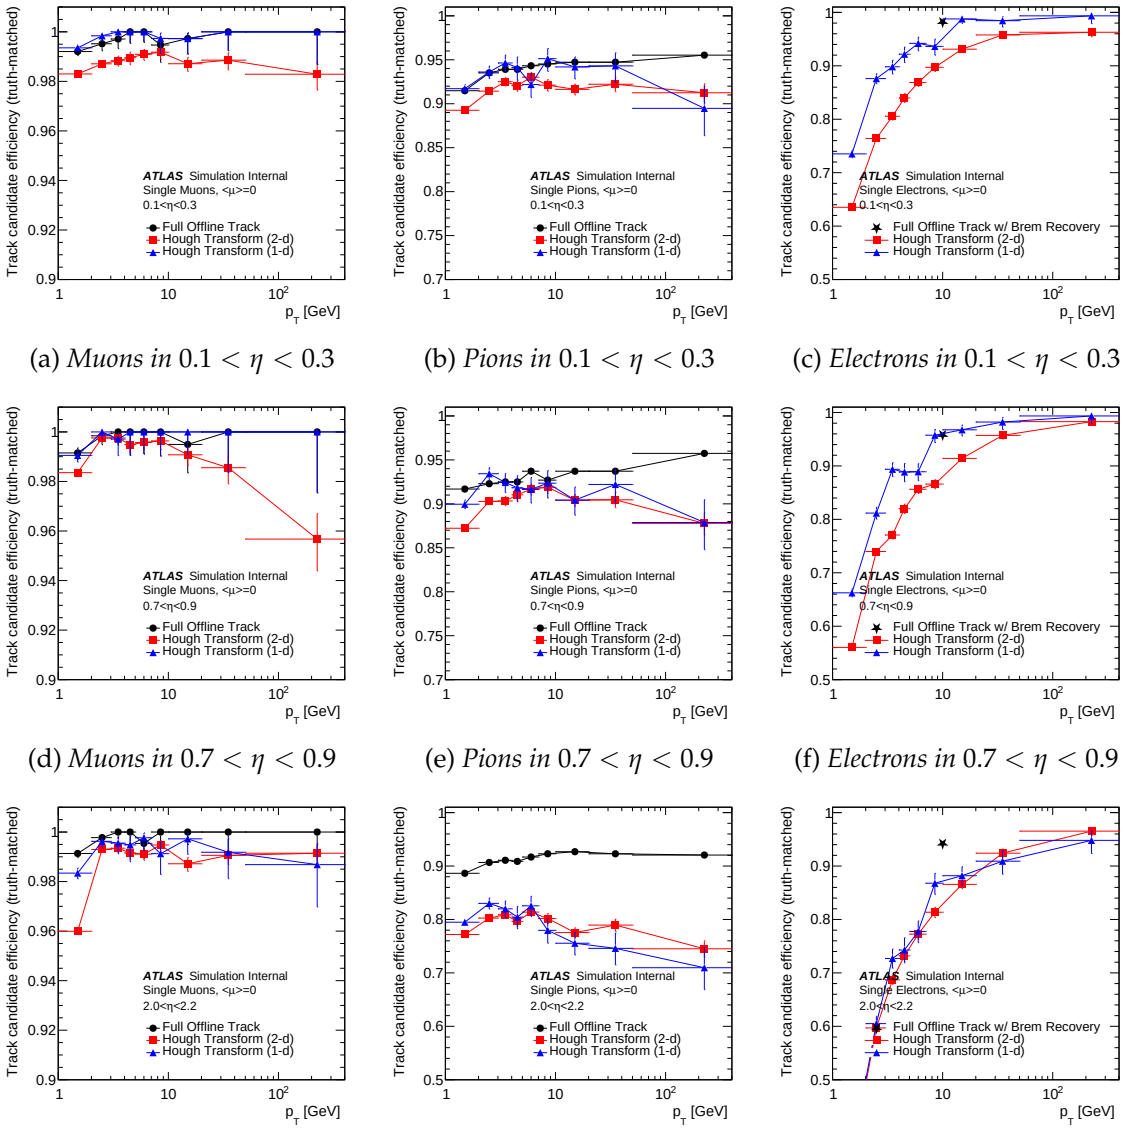

(a) *Muons in* $0.1 < \eta < 0.3$ (b) *Pions in* $0.1 < \eta < 0.3$ (c) *Electrons in* $0.1 < \eta < 0.3$

(d) *Muons in* $0.7 < \eta < 0.9$ (e) *Pions in* $0.7 < \eta < 0.9$ (f) *Electrons in* $0.7 < \eta < 0.9$

(g) *Muons in* $2.0 < \eta < 2.2$ (h) *Pions in* $2.0 < \eta < 2.2$ (i) *Electrons in* $2.0 < \eta < 2.2$

Figure 2.16: *Pattern Recognition efficiencies for a traditional 2-d Hough transform approach (8-layer) and the 1-d "Bit-Shift" Hough transform approach (9-layer), for various* η *ranges and* $0.3 < \phi < 0.5$. *These are compared to the efficiency for full offline tracks which are extended to all layers and have fake reduction already applied. At this stage, the Hough candidates have a very large fake rate which is reduced using the methods described in Section 2.4.2 where the corresponding additional efficiency loss is estimated for the current algorithms. For electrons, the offline value is based on the software including the Bremsstrahlung recovery algorithm.*

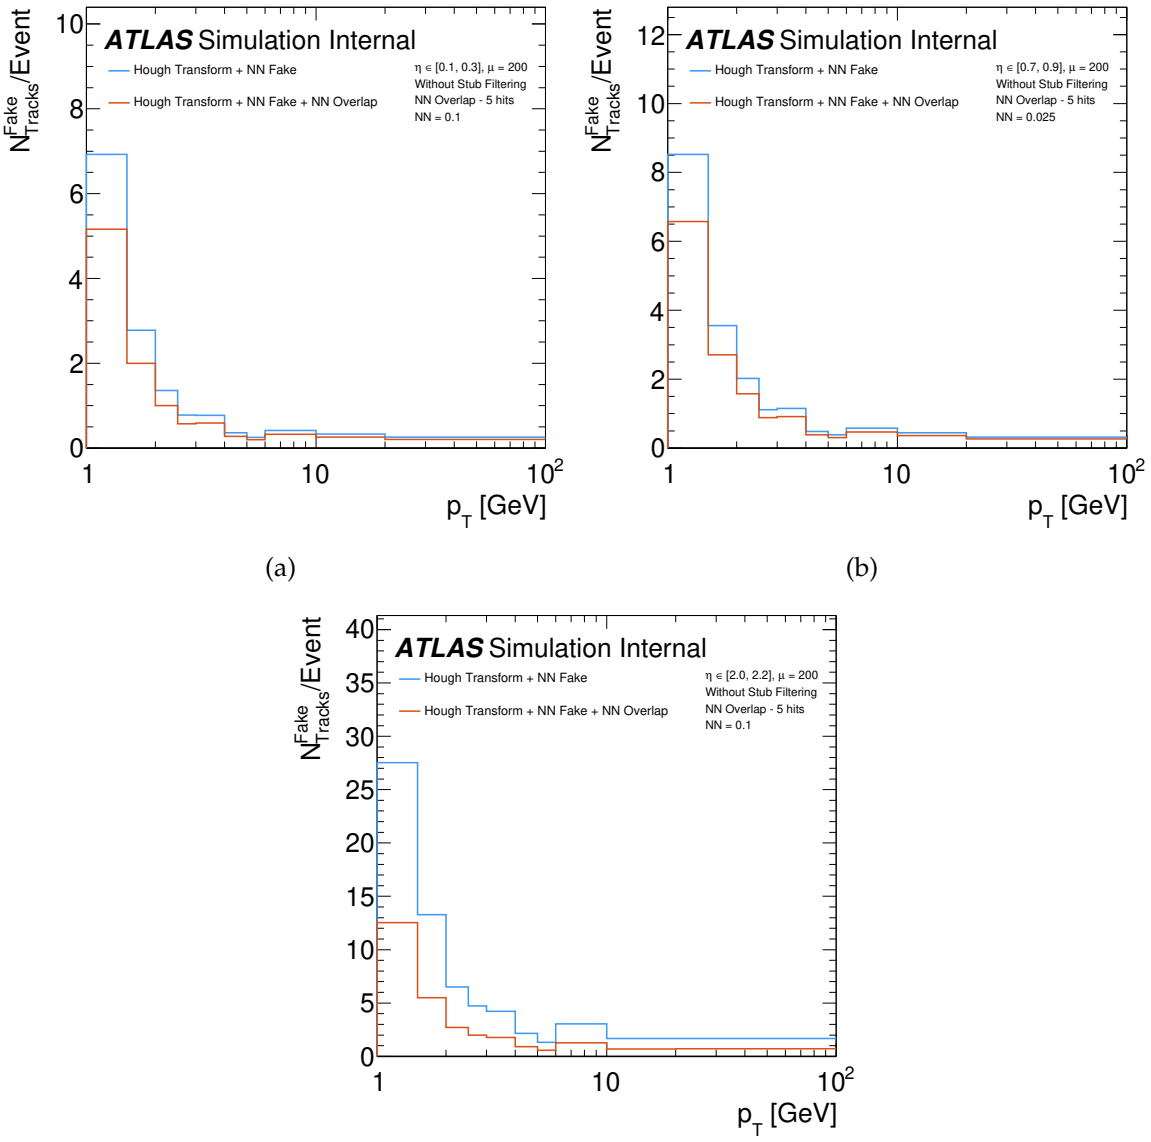

Figure 2.17: *The distribution of the number of fake track candidates per event in each p_T bin after all duplicate removal and NN selection for $0.1 < \eta < 0.3$ (a), $0.7 < \eta < 0.9$ (b) and $2.0 < \eta < 2.2$ (c). The x axis is the Hough transform p_T bin, and corresponds to the full considered range such that there is no overflow. The samples are muons on top of $t\dot{t}$ events at $\langle\mu\rangle \approx 200$. Hits from $t\dot{t}$ and pile-up are the source of combinatorial fakes. The blue curve is the number of fake tracks after the Hough transform pattern recognition and the NN fake rejection; the fake rejection reduces the number of tracks by more than 2 orders of magnitude. The red curve is after an additional overlap removal on the number of shared hits.*

To improve fake rejection and efficiency, p_T-dependent cut values from the NN are used. The resulting NN efficiencies are typically $\sim$ 97% above 2 GeV, with a total reduction in the number of tracks after application of the NN cut and overlap removal of a factor of 200 in the barrel and over a factor of 540 in the endcap. About 50% of fake track combinations have a p_T in range of 1.0 - 1.5 GeV, while an additional 10% of fake combinations have a p_T in range of 1.5 - 2.0 GeV. These two regions dominate the optimization for a single cut value, biasing the results for higher p_T regions. A p_T-dependent cut allows to optimize each regime separately, extracting the best possible results from this fake removal step. Typically, the NN cut becomes tighter at high p_T, signifying a better rejection of fakes tracks, while ensuring high true efficiency. However, in a few p_T bins, due to limited numbers of events available, the training does not perform as well and a looser NN cut with lower background rejection needs to be used. These limitations also cause the p_T dependence and fluctuations seen in the efficiency plots. It is expected that additional numbers of events at higher p_T for both NN training and performance evaluation will remove these trends. If the low-p_T regime continues to dominates, the flexibility of the NN allows to implement 2 separate modules with each trained in the specific p_T regime, subject to available FPGA resources.

Figure 2.17 shows the number of fake track candidates per event in each p_T bin after the full selection described above. Approximately, 40%-50% of the remaining fake track combinations are in the p_T = 1 - 2 GeV region, with a large fractions of that in the 1 - 1.5 GeV region. The current level of remaning fakes reflects the preliminary nature of the optimisation of the NN training strategy that is expected to improve in the future.

Options for Large-Radius Tracking Large-radius tracking (LRT) is challenging because without the restriction that a track should originate from the beamspot, the combinatoric possibilities for a group of hits increase. A successful LRT approach needs to be both fast and not overly prone to combinatoric fakes. Three methods are being considered: fast LRT in software (see Section 2.4.1) and two modified Hough transform approaches that could potentially be implemented in an FPGA. All three methods run as a second pass on the hit clusters remaining after the first pass of the prompt reconstruction. The following strategy is employed to efficiently remove the hit clusters associated with prompt track candidates. In the firmware, the clustered data is stored in the slicing engine during pattern recognition and fitting. The data will be organized such that as the complete tracks are produced, the hits in the slicing engine storage are marked as used. This will not take significant resources, but the exact indexing method needs to be developed to ensure that the look-up is efficient. The slicing engine would be required to retain an event until the last track is complete; however, the effect of this is small compared to the full event processing time. The unmarked clusters would then be passed to software for fast LRT or passed a second time through the firmware to run pattern recognition tuned for large-radius tracks.

For the FPGA-based approaches, the standard Hough transform is based on the assumption that the track passes through the origin. Without this restriction, a full Hough transform would be three-dimensional and would require accumulator binning in q/p_T , ϕ_t, and d_0.

The two modified Hough transform approaches described below explore different possibilities for reducing this problem to two dimensions without enforcing an assumption on the d_0 of the tracks. In the straight track approximation, one can take the limit of high-p_T tracks such that $q/p_T = 0$ and then populate an accumulator in the space of d_0 versus ϕ_t. The result is a version of the Hough transform that scales linearly with the number of hits and is sensitive to tracks with large d_0 (essentially up to the position of the smallest-radius detector elements used in the algorithm), at the expense of efficiency for low-p_T tracks. Thus, such an algorithm can provide an efficient way to provide trigger sensitivity to some, though not all, of the phase space of interest involving long-lived particles.

The hit positions and the track parameters are related by:

$$d_0 = (y_h - x_h \tan\phi_t)/\cos\phi_t.$$

This can be simplified by using the r, ϕ coordinates of the hit and taking the linear approximation

$$d_0 = r_h(\phi_h - \phi_t)$$

In contrast to the above, the displaced track Hough transform uses pairs of hits only. For each pair of hits, a collection of possible tracks can be defined in terms of only two of the three variables, with the third being implicitly fixed by the positional relationship of the hits. In this case, the implicitly fixed variable is ϕ and the accumulator is filled in the p_T and d_0 of the candidate tracks.

This "doublet" method uses a simple calculation to obtain the d_0 of a candidate track from the hit positions and a selected q/p_T. In this approach, no simplifying assumptions are made as to the minimum value of the track p_T and thus no performance degradation is expected at low track momenta.

Conclusion

The FPGA-based demonstrator described in this section provides confidence that the final EF Tracking system will meet all necessary specifications, including the tracking requirements detailed in Section 2.2.2. This demonstrator includes pattern recognition firmware for a baseline Hough transform with firmware estimates shown to fit within the target FPGA, software aspects that take advantage of novel developments in fast tracking software for ATLAS, and reduced power consumption. Further optimizations towards a final system are described in Section 2.5.

2.4.3 GPU-based Demonstrator

A GPU-based demonstrator for the use in Run-3 has already been described in the Phase-II TDAQ TDR [1] and for completeness the most relevant aspects are reproduced here. While that study was done for the current Inner Detector (ID), using a lower pile-up of 46 and by now outdated GPUs, the overall reconstruction strategy for silicon track finding in the ID also applies to the ITk.

Selected trigger reconstruction algorithms have been re-written as Compute Unified Device Architecture (CUDA) [25] modules to be executed on an NVIDIA GPU that are steered by the CPU code. Processing modules for GPU have been implemented for compute-intensive parts of the inner detector tracking [26], calorimeter topological clustering [11], cluster-splitting, jet reconstruction and a conformal-transform based algorithm for muon tracking [11].

Specifically for the track reconstruction, the compute intensive data preparation and track-seeding stage were implemented as CUDA modules. The algorithm for the data preparation assigns fragments from the bytestream to separate thread blocks and each word within the fragment is assigned to a separate thread, allowing all words to be decoded in parallel. A global output buffer is created so that the hits for a module can be written to a location that can be accessed by the clustering routine. Atomic integer markers are used to ensure that module hits are only recorded once to global memory, reducing write access and decoding time. After the decoding, a clustering kernel based on a cellular automaton is run to find all hits belonging to the same cluster. The track seeding is implemented as three separate kernels: the first calculates the number of doublets with middle and inner or outer spacepoints and allocates the global storage space for doublets accordingly; the second module forms the doublets and stores them as 32-bit indices of the corresponding outer/inner spacepoint; and the last module forms the triplets and applies kinematic and quality cuts.

At each step in the data preparation process, some knowledge of the detector geometry is required, and hence any decoding, clustering, or spacepoint formation algorithm requires use of the ATLAS detector geometry. To solve this problem, a minimal version of the detector geometry was developed for the GPU which utilizes a hash table for efficient lookup of geometry information for the several thousand detector modules.

Throughput measurements have been made for a system comprising two 14-core Intel Xeon E5-2695 CPU (2.3 GHz clock-speed) and a system with the same CPU plus a GPU serving as an accelerator. For the GPU two scenarios were tested:

i) up to four Nvidia GK210GL Kepler architecture GPU (in two K80 PCI cards) on the same PCIe bus as the CPU,
ii) a GTX1080 PCIe card with a Pascal architecture GPU (GP104) in a separate unit with a local CPU acting as a server and connected to the client node via a 10 GbE network.

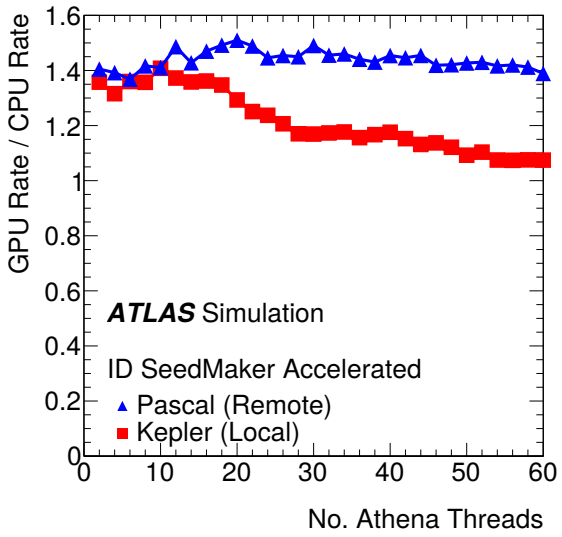

Figure 2.18: *Throughput increase as a function of the number of Athena worker threads when executing only ID reconstruction algorithms.*

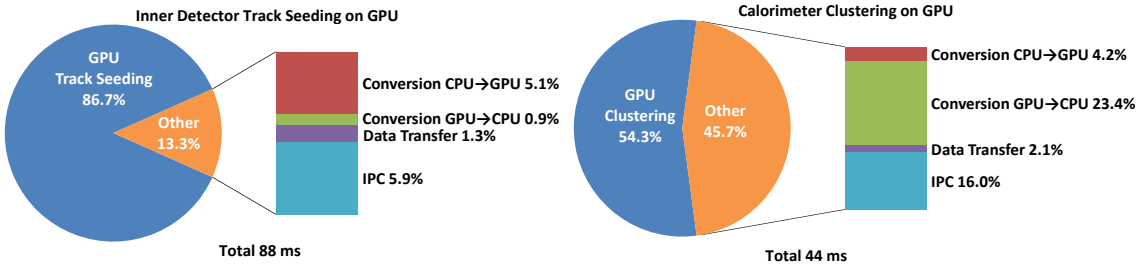

Figure 2.19: *Breakdown of the time per event (measured on the Kepler system) for (a) inner detector track-seeding and (b) calorimeter clustering offloaded to a GPU for the kernels running on the GPU and the overhead associated with offloading the work (other). The overhead comprises the time to convert data-structures between CPU and GPU data-formats, the data transfer time between CPU and GPU and the Interprocess Communication (IPC) time that accounts for the transfer of data between the Athena processes and the process handling communication with the GPU.*

For the ID track-seeding algorithm a speed-up factor of 28 (5.8) is obtained on the Pascal (Kepler) GPU relative to the execution on a CPU. Overheads for data conversion and inter-process communication reduce this to an effective speed-up of 15 (5) for the Pascal (Kepler) GPU cards, respectively. Figure 2.18 shows the throughput increase factor (defined as the ratio of the event rate with GPU to the CPU-only event rate) as a function of the number of Athena clients for full-event ID track reconstruction accelerated by a Pascal-architecture and a Kepler-architecture GPU. The input was a simulated Phase-I $t\bar{t}$ dataset converted to raw detector output format with an average of 46 minimum bias events superimposed. In the CPU-only case, the track-seeding algorithm takes 28% of the total event processing time, which limits the maximum possible throughput increase to a factor of about 1.4. The Kepler GPU shows evidence of saturation for 20 or more clients, while no saturation is seen for the Pascal GPU with 60 clients due to the shorter algorithm execution time on this GPU.

The speed-up factor obtained for the calorimeter algorithm is 3.6 on the Kepler GPU, but in this case the shorter algorithm execution time on the GPU is completely offset by the data-conversion and inter-process communication overheads, meaning no throughput increase was obtained. The latter can be observed in Figures 2.19a and 2.19b showing the break-down of the time per event for ID track-seeding and calorimeter clustering, respectively. This illustrates the importance of implementing a suitable event data format in the offline and trigger code to avoid expensive data-format conversions. More details can be found in [11].

Based on the measurements made with the demonstrator at the time, it was estimated that using the tested hardware it would cost approximately the same to increase the farm throughput by adding GPUs or CPUs. However, the cost-effectiveness of adding GPUs to the EF for Run-4 depends on the evolution of CPU and GPU in terms of price, performance and packaging. The above work should be taken as an indication that the software can be successfully adapted to other architectures in order to perform a full cost/benefit evaluation with current and future GPU models. Further developments required to achieve this goal are discussed in Section 2.5.

2.5 Optimization of Event Filter Tracking System

The input from EF Tracking to the choice of the final system architecture of the EF will follow the process described in Section 3.3. This will take all the system requirements into account as detailed in Section 2.2, with the goal to maximize the trigger performance and to choose the best processing technology (CPUs and accelerators) to achieve it. The working prototypes in Section 2.4 provide confidence that the final system will meet all necessary specifications:

- The CPU-based demonstrator provides near offline-quality tracking performance with CPU requirements that are compatible with the available power and space at Point-1.
- The FPGA-based demonstrator includes pattern recognition firmware for a baseline Hough transform with firmware estimates shown to fit within the target FPGA, software aspects that take advantage of novel developments in fast tracking software for ATLAS, and reduced power consumption.
- The GPU-based demonstrator has shown that the tracking software can be adapted successfully to run on such hardware.

The exact system is unlikely to be one of these examples, and there are other promising avenues for exploration. The following subsections describe the needed optimizations of the Software Framework and Technology, followed by a discussion of the optimization considerations for each functional aspect related to EF Tracking: ITk Data Preparation, Track Seeding & Pattern Recognition, and Track Extension, Fitting, & Ambiguity Resolution. Finally, the strategy for the optimization of the System Design is described.

2.5.1 Software Framework & Technology

The CPU code to be deployed in the final system will be based on the tracking chain that is being developed in collaboration with the offline tracking group. The tracking performance reported in Section 2.4.1 has been obtained with a prototype of this fast ITk reconstruction. ATLAS has initiated the Acts open source tracking project [7] to modernise its track reconstruction software. The project is carried out in collaboration with tracking experts from other experiments, e.g. Belle-II, sPHENIX and FCC. The Acts project also has the goal of making tracking software more portable to exploit co-processors like GPUs or FPGAs, using compiler extensions like e.g. CUDA or oneAPI/SYCL [6]. High-level Synthesis (HLS) [27] tools for FPGAs greatly facilitate the translation of new algorithm ideas into digital implementations and are increasingly popular for the design of high-performance and energy-efficient heterogeneous systems. These tools are, by design, accessible not only to professionally-trained engineers but also technically-inclined scientists, which makes it easier to collaborate in addressing challenges such as fast tracking of charged particles for use in ATLAS trigger decisions.

It is the ATLAS goal in the coming years to port the full offline ITk reconstruction to Acts. Using the faster Acts combinatorial Kalman Filter implementation to avoid any approximations in the fitting process will recover the full tracking performance of the ITk without losing the CPU efficiency of the current fast ITk reconstruction prototype. More details of this development program are outlined in the Phase-II Computing CDR [13].

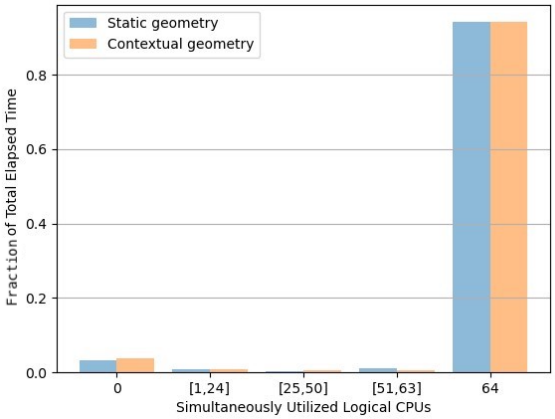

Figure 2.20: *Difference in CPU utilization for track propagation between an ideal detector and a per-event randomly misaligned detector in Acts. The x-axis shows the number of utilized threads. Optimal utilization of the given multi-core CPU is when all threads are constantly occupied, and neither of them are idle. In this example, the alignment has been pre-computed and loaded before execution, in which case the Acts code can give a practically optimal utilization of all available (64) threads, regardless whether the detector undergoes condition changes.*

One of the starting points of the Acts project is to establish a fully multi-threading ready code base for track reconstruction, with a practically cost-free handling of contextual data (conditions data, such as alignment) during the data flow and an optimized EDM. The EDM is fully based on the Eigen math library [28] and profits from its demonstrated great CPU performance when using compile-time fixed matrix operations. For this reason, all EDM objects in Acts are fixed-size at compile time in their dimension, and polymorphism is highly restricted. Figure 2.20 shows a demonstrator of the CPU utilization with 64 threads comparing a static detector without alignment and a randomly misaligned detector. While general code optimizations, such as a vast simplification of the EDM, the strict avoiding of virtual function calls where possible and general code cleanup have already led to sizeable speed-up compared to the modules available in ATLAS Tracking software[8] [30], bigger gains can be made by rationalizing certain operations.

While the choice of accelerator hardware to be deployed at Point-1 is entirely under the control of TDAQ, attention needs to be given to simulating the trigger (tracking) also on shared Grid and HPC resources that may or may not have access to the same accelerator

[8] The Acts vertex reconstruction will run as part of the ATLAS Run-3 track reconstruction and has proven a roughly 40% CPU improvement without any loss of physics performance [29].

hardware. Techniques need to be developed to either transparently compile the code for multiple platforms (hopefully with the support of future compilers), or dedicated simulation code needs to be written to efficiently simulate the trigger reconstruction on standard CPU hardware.

2.5.2 ITk Data Preparation

The ITk data preparation functionality described in Section 2.1 (data decoding, hit clustering, spacepoint formation, and stub filtering) can be resource-intensive and needs to be studied and optimized for the final architecture. Prototype firmware to perform these tasks exists (as described in Section 2.4.2) and has been demonstrated to utilize a modest amount of FPGA resources, but the firmware needs to be optimized and possibly extended to provide as close to the offline clustering performance as possible. Software studies are needed to understand how to exploit the pre-clustering that is done in the front-end chips, and studies are also needed for possible GPU implementations. The investigation of multiple options is needed as input to the full system design.

2.5.3 Track Seeding & Pattern Recognition

A broad suite of track seeding and pattern recognition strategies are available for exploration. Within the context of the Acts project, ATLAS is pursuing several development initiatives, from optimising classical tracking strategies as used in the fast ITk track reconstruction, to studies using machine learning (ML) techniques.

Approaches that build on Hough or conformal transforms require detailed optimization studies in order to maximize their performance on a given architecture. For example, the nominal strategy for the Hough transform firmware implementation described in Section 2.4.2 is to define the accumulator for an $\eta \times \phi = 0.2 \times 0.2$ region with a given binning in qA/p_{T} and ϕ and to use eight tracker layers. In order to handle the inner layers, a Hough-like scan in both p_{T} and d_0 would be required, as the nominal 2 mm impact parameter coverage causes significant ϕ shifts in the inner pixels. Additional tuning of the accumulator may prove to yield better performance. The Hough transform stage was found to be the most resource-intensive functionality on the FPGA; thus, an important aspect of a track pattern recognition stage implemented in firmware will be the optimization of FPGA resource usage.

Novel reconstruction algorithms beyond Hough or conformal transforms that find track candidates are being developed that are suited to exploit the potential of accelerators. While the potential for advanced algorithms is certainly not limited to ML methods, these methods are ubiquitous and have proven to be very powerful in improving the ATLAS detector capabilities and advancing the ATLAS physics program. Popular open-source software

such as TENSORFLOW and PYTORCH have helped democratize machine learning for scientists by providing powerful and flexible libraries to develop and test ML models.

One such advanced algorithm is based on a type of *geometric learning* of cluster patterns produced by charged particles traversing the ITk. Graph Neural Networks (GNNs) [31, 32] are a type of geometric deep learning algorithm that is well suited for capturing spatial dependencies via message passing over graphs. Tracking detector hits can be represented as graph nodes and the relationships between them can be represented by graph edges that represent possible track segments. Particle tracking applications of GNNs were first studied in [33]. Recent work by the Exa.TrkX project and other collaborations has demonstrated that edge-classifying GNNs are well-suited to particle tracking applications [33, 34, 35, 36, 37] and that considerable speedup over CPU-based execution is possible on GPUs or FPGAs [38].

The most advanced studies in Acts using GPUs have been done for the seed-finding stage of the track reconstruction. This algorithm exists in several flavours using either GPU backends (with native CUDA or high level OneAPI/sycl implementation) compared to a standard CPU implementation. All of the implementations achieve almost identical results, and allow for the exploitation of future heterogeneous architectures. Figure 2.21 shows the speedup achieved running the seed finding on an NVIDIA GTX 1070 compared to running on a single CPU core. Although in general speed improvements have been observed, several bottlenecks concerning memory transfer between host and device became evident, and longer chain demonstrators are needed for the relevant event data model and algorithm R&D. Further focus has been drawn on a chained track reconstruction demonstrator consisting of clustering, spacepoint formation and seeding (as part of the traccc [39] project of Acts). Finally, an attempt to create a track reconstruction geometry description for GPU usage is currently underway [40].

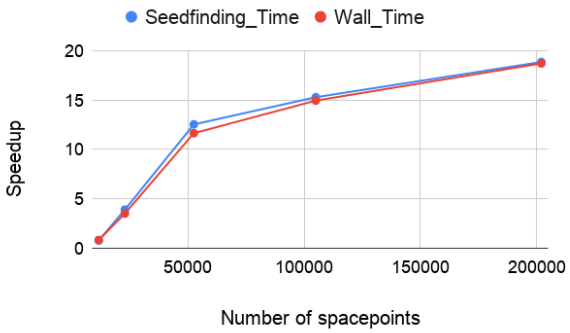

Figure 2.21: *Seed finding speed-up on a NVIDIA GTX 1070 GPU compared to a single core of an Intel i7-5820K (3.3 GHz) CPU.*

Individual algorithms will also need to be optimized to take into account the overall system performance, including the track extension, fitting, and ambiguity resolution stages

discussed in Section 2.5.4 required to meet the defined tracking requirements. Tunable parameters such as the choice of layers included in this stage of track finding and the variables considered by the algorithm will also require detailed study.

Machine learning-based end-to-end solutions for track finding are being currently exploited using Acts and/or Acts simulated data. Most prominently, big advances have been achieved by the Exa.TrkX project [41] that demonstrated a first feasible track finding setup with high track reconstruction efficiency on the TrackML data set [42, 43]; an adaption to simulated data for the ITk is currently underway.

An optimization program will be needed to evaluate and compare the full suite of track seeding and pattern recognition algorithms given the available technologies for implementation. The chosen algorithm will need to meet the tracking performance requirements defined by the project as well as be suitable for implementation on the final technology chosen for the EF processor farm.

2.5.4 Track Extension, Fitting, & Ambiguity Resolution

Control of the duplicate/fake rate out of the track seed finding & pattern recognition stage is an important consideration in the optimization of the system.

In an FPGA-based approach, a high-quality fitting method could replace or augment the NN filtering of roads, and/or could be used for a second-stage extension of tracks and corresponding filtering. One candidate method is to use the p_{T} and ϕ values from the Hough transform to translate the geometry into an idealized ϕ-symmetric geometry [44]. The transformation is simple because the layers are close to ϕ-symmetric, so the correction is small and errors on the correction are even smaller. The η and z_0 are restricted by the slicing so they also have a known crude estimate. A linearized track fit would then find the small correction from the pre-fit to the full track fit. The number of necessary fit constants (matrices) for the idealized geometry are small because of the ϕ-symmetry that is introduced. Additional optimizations could also be made to the NN itself; the slicing and Hough transform information could be used to compute pre-estimates of track parameters that could be fed into the network along with the Cartesian coordinates of each hit. In addition, the network could be made more aware of holes on tracks. Finally, it is clear that additional training samples are a key component of optimization of the network performance.

The Acts Kalman Fitter (as will any future track fitter implementation within Acts) includes in particular the hole search directly in the track fitting procedure, a relatively CPU-intensive operation that currently has to be performed in the Ambiguity Resolution after track fitting in the default ITk track reconstruction, which in turn requires yet another round of magnetic field integration. Using the Acts fast Kalman Filter as the combinatorial Kalman filter in Track Finding will produce track candidates with full offline resolutions, removing the residual differences between the full and fast ITk track reconstruction as shown in Section 2.4.1. In this scenario the Ambiguity Solution will only act on ambiguous tracks

with shared hits or merged clusters to recover some residual inefficiencies in the core of high-p_T jets

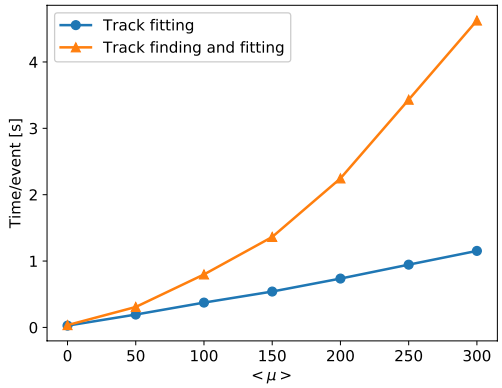

Figure 2.22: *CPU time spent in track finding with the combinatorial Kalman filter, such as for track fitting itself above 1 GeV in a typical HL-LHC pile-up condition range. The hole search, one pivotal part for flagging and excluding fake tracks is already included in this timing measurement. The additional time in the track finding comes from the combinatorial search and can be reduced by implementing smart pre-selection of measurements provided to the track search.*

While Acts has focused so far on feature completion and most CPU optimizations are simply a result of best practice programming, competitive timing results for track finding and fitting can already be achieved. Figure 2.22 shows the track fitting and track finding time for HL-LHC environment which reaches the order of the fast tracking setup. Further optimisations to the baseline Acts reconstruction are to be expected as a result of a dedicated future profiling and optimization campaigns. Recently the Combinatorial Kalman Filter has also been demonstrated on a GPU [45].

2.5.5 System Design & Integration

Several system architecture options will be considered, such as: different choices of CPUs with or without embedded accelerators, PCIe-based accelerator cards in each EF node, and separate accelerators in their own servers. The architecture choice will be driven by the available technology on the market, the required ratio of CPU-to-accelerator compute power and operational simplicity.

In the case of a separate accelerator farm, multiple accelerator cards would be housed in servers that can be horizontally cooled and are readily available commercially as of the writing of this document; PCIe Gen 3.0 is likely sufficient, and PCIe Gen 4.0 is readily available. The event processing time determines the number of accelerators needed and can be further optimized. To simplify the dataflow, one event would be processed per accelerator; based on the demonstrators described in Sections 2.4.2 and 2.4.3 this is a reasonable

assumption. Additional load-balancing software would be required to distribute the workload to each card.

In all cases, the required detector data is retrieved by the EF nodes from the readout system and forwarded to the accelerator devices if needed. No limitations are expected from the dataflow system, which is expected to be able to deliver up to 1 MHz of fully built events to the EF nodes.

A program of investigation will be required to track industry developments including: how data is transferred to and from accelerators, how the accelerator workloads are managed, how the system is configured and monitored, and the physical power, cooling, and networking. Finally, a critical element of a successful heterogeneous commodity system will be the accelerator-software handshake, including the development of an efficient interface.

2.5.6 System-level Optimization and Outlook

The flexible nature of a heterogeneous system provides an opportunity for innovation at all stages of the EF tracking pipeline. A system-level optimization that capitalizes on the most performant functionalities of the individual components is needed, taking into account the constraints and opportunities provided by the available technologies, to deliver an optimized system that meets the requirements for EF tracking. Should the final system exploit accelerators, optimizing the CPU-accelerator interface while maintaining the expected tracking performance will be an important consideration.

The roadmap towards an optimized, heterogeneous, commodity system for EF tracking described in this section does not contain an exhaustive list of avenues to be pursued. Advancements in software, firmware, and hardware technology will be studied throughout the process. The final choice of system architecture will be determined following the schedule in Section 3.3, including decision points and ATLAS reviews. The criteria for the technology choice will be listed in details as part of the requirements and specifications phase that will be reviewed at the Specification Validation Review. They will include an assessment of whether the proposed solution fits within the cost envelope of the system, the expected tracking performance, the power requirements, the operational procedures, the resource requirements and an evaluation of risks.

3 Project Management and Organization

3.1 Project Management and Organization

This section focuses on the organizational structure of EF Tracking within the TDAQ Phase-II Upgrade Project (UPR). The full TDAQ Phase-II project structure and reporting lines to the TDAQ system and ATLAS are described in the Phase-II TDAQ TDR [1].

The TDAQ Phase-II Upgrade project comprises three systems at Level-2 of the organisation structure: the Level-0 (L0) Trigger system, the Data Acquisition (DAQ) system, and the Event Filter (EF) system. Each L2 system coordinator reports to the TDAQ Phase-II UPL.

The EF system is organised in the following sub-systems: the EF Farm, EF Tracking, EF Calo, EF Muon, and EF Core Software. The EF Farm sub-system coordinators (L3) coordinate the overall EF farm requirements, the technology decision process, and procurement. The EF Tracking, EF Calo and EF Muon L3 coordinators are responsible for the reconstruction of objects primarily from tracking, calorimeter, and muon detectors, respectively. The EF Core Software L3 coordinators' remit includes the software framework, common accelerator-related infrastructure and the interface to the DAQ Dataflow subsystem (Dataflow) sub-system. The L2 manager will organise regular coordination between the managers and technical coordinators of the L3 sub-systems to ensure that studies converge on a common, optimal EF farm specification.

Figure 3.1 shows the organizational structure of the EF Tracking subsystem comprising five L4 areas, four coordinators, a project officer and an AI/ML liaison. Ultimately the L3 managers are responsible for the overall delivery of the EF Tracking sub-system. Following from the description of the system in Section 2.3.1, the general strategy of the EF Tracking organization is centered on the tracking algorithms rather than along technology architectures. The goal of each of the L4 areas will be to choose an implementation which best achieve the requirements. Although for some system implementations the boundaries between the first three areas might be fluid, categorizing the tracking steps as shown will be applicable to most cases and a breakdown helps to coordinate the work. The five L4 areas are:

ITk Data Preparation This area covers ITk data preparation, including input data organization, decoding, cluster-finding, stub filtering, and spacepoint formation.

Track Seeding & Pattern Recognition This area covers the development of track seeding and pattern recognition algorithms, potentially including but not limited to Hough or

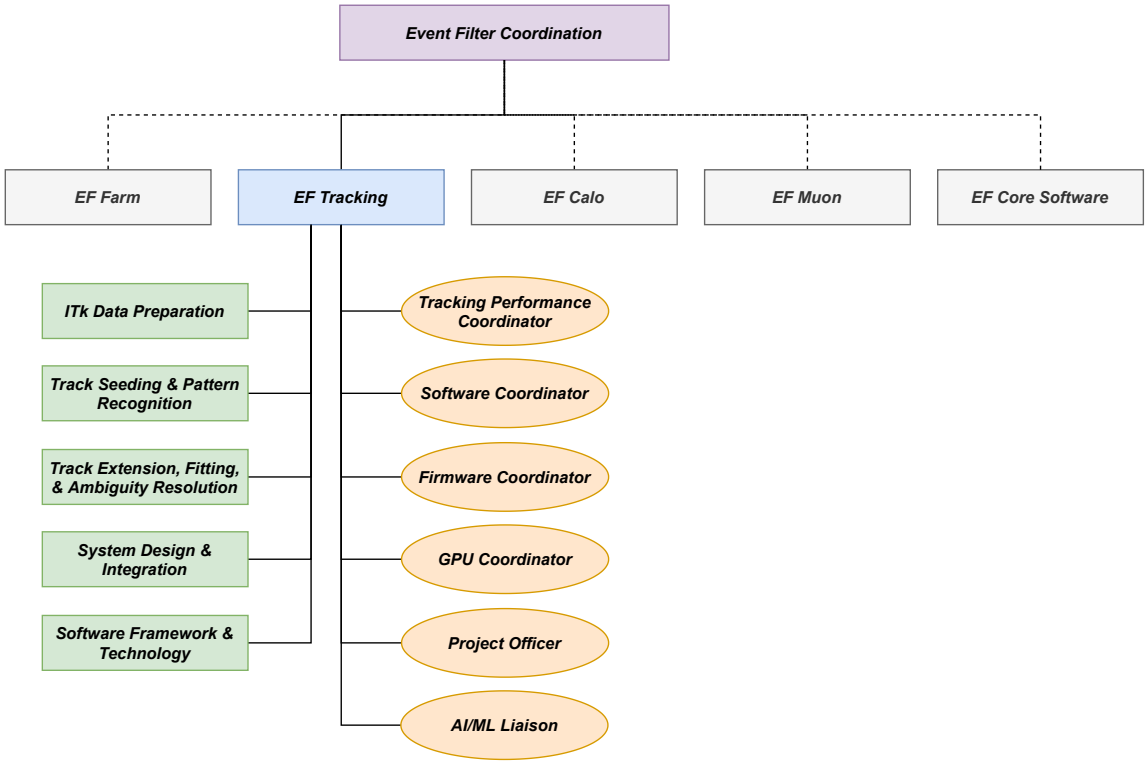

Figure 3.1: *Organization of the EF Tracking subsystem in the context of the EF system. The blue and gray boxes are the L3 areas, the green boxes are the L4 managers responsible for deliverables, and the orange ovals indicate coordinators that span the L4 areas.*

conformal transforms and advanced algorithms based on artificial intelligence (AI)/ML strategies.

Track Extension, Fitting, & Ambiguity Resolution This area is responsible for track extension (if needed) after the pattern recognition stage, coarse or precision fits and track parameter determination, and ambiguity resolution due to duplicate or fake track candidates.

System Design & Integration This L4 area is responsible for system engineering, infrastructure and dataflow considerations (including interfaces between the CPUs and accelerators), vertical integration design and tests, and new technology evaluation in conjunction with the relevant L3 areas. This area also provides a liaison with the EF Farm for tendering and procurement of commercial hardware. It is foreseen to have two L4 managers for this area, one of which would be an engineer who would be the System Engineer for EF Tracking.

Software Framework & Technology This area is responsible for the tracking software framework at the basis of the EF track reconstruction implementation. Members of this L4 area are foreseen to contribute to the development of the Acts tracking software for ATLAS and its integration into ATHENA, as well as to the system-level optimization for accelerator-based tracking implementations. The software framework needed for

the emulation of accelerators and the simulation of EF tracking in MC production are also included in this area. The L4 manager(s) will closely collaborate with the other EF L3 areas and the offline software community.

In addition to the above L4 areas, a set of coordinators will also be appointed. While the L4 managers will be responsible for the deliverables and milestones associated with their respective L4 area, the coordinators will ensure continuity of tools and approaches across the EF Tracking project. These will include tracking performance, software, firmware, and GPU. The tracking performance coordinator will establish tracking performance requirements and ensure that the L4 areas meet those requirements, as well as interface with the Performance, Physics, and Event Selection coordinators within the TDAQ Phase-II Upgrade Project and the Upgrade Tracking and Upgrade Physics coordinators within the ATLAS Upgrade Project. The Software, Firmware and GPU coordinators will ensure coherent standards and tools used across the L4 areas, oversee the validation, verification and automation of testing and will manage the relevant releases and version control. The Firmware and GPU coordinators will also interact with the EF Calo, EF Muon sub-systems in order to coordinate studies and harmonize any accelerator specifications. A project officer will be appointed to assist the L3s with the monitoring of EF Tracking tasks, milestones and deliverables, risk monitoring and mitigation, and preparation of quarterly statusing and baseline change proposals. An AI/ML Liaison will also be appointed to follow current developments in ML tracking and liaise with the ATLAS and Acts ML fora for example.

3.2 Project Cost Estimates

The strategy for the cost estimate of the EF Farm follows the cost management plan and cost policy and methodology described in the Phase-II TDAQ TDR [1]; some aspects of the costing policy and methodology are summarized here in Section 3.2.1.

The project is organised in a set of deliverables that are summarised in the Product Breakdown Structure table in Section 3.2.2. The cost estimate and expected spending profile is presented in Section 3.2.3. Though this document focuses on the revised plans for EF Tracking, the cost of EF Tracking will be included in the total cost of the EF Farm presented in this section.

3.2.1 Summary of CORE Costing Policy and Methodology

For reference, the relevant aspects of the CORE costing policy and methodology are summarized here. At the time of the original construction, the LHCC established a COsting REview (CORE) committee charged with reviewing the experiment's costs. A costing policy was established to define a metric that assigns cost values to each deliverable. Since then CORE equivalent costs are used to estimate and evaluate the effective costs of upgrade

projects. CORE cost does not represent the entire cost of a project; there are non-CORE costs that are required to complete successful production and delivery of the components. It is the responsibility of each institution participating in the UPR to obtain financial support from its Funding Agency for both CORE and non-CORE expenditures for the deliverables where the institution is responsible.

CORE costing is based on the Current Best Estimate (CBE) concept, i.e. on the information available at the time of the estimate, and it is naturally associated with an uncertainty. Several factors determine the level of cost uncertainty of an item: (i) the maturity of the technical development and design of a particular item, (ii) recent experience in other construction projects of items with similar technical complexity, (iii) the availability of vendor quotes through tendering processes or standard catalogues for commercial-off-the-shelf (COTS) components, (iv) understanding the procurement processes by the Institutions responsible for the deliverable item. To describe the level of uncertainty of an estimate, a quality factor ranging from 1 to 5 is used. The quality factors and the criteria are shown in Table 3.1: **QF1** has the highest certainty and is based on a vendor quote for the final item or a catalogue price; **QF5** has the lowest amount of certainty and is based on a rough estimate without any detailed design.

Table 3.1: *Quality Factor (QF) definitions used to estimate uncertainties on the CORE values of the project elements.* ***QF1*** *has the lowest uncertainty, while* ***QF5*** *has the highest one.*

Factor	Definition of the criteria based upon
QF1	(i) Items for which there is a recent (1 year max.) quote or catalogue price, based on a nearly completed design and for which there is more than one potential vendor (ii) Items that are a copy or are almost identical to an existing design for which there is a recent catalogue price or quote and for which there is more than on potential vendor.
QF2	Items that just fall short of satisfying the QF1 criteria: (i) Items that have only one potential vendor. (ii) Estimates based on a detailed, but not completed, design. (iii) Items adapted from an existing design with minor modifications. (iv) Items having quotes >1 year old, but deemed still to be sufficiently reliable based on experience.
QF3	(i) Items with quotes > 2 years old. (ii) Items whose cost estimates are based on a conceptual design or adapted from an existing design with extensive modifications. (iii) Items whose costs are estimated using physicist or engineering experience regardless of the maturity of the design.

continued ...

Factor	Definition of the criteria based upon
QF4	Items that have unproven fabrication yields or for which there are unique issues, e.g. a special-order item and/or a single preferred vendor
QF5	Items that are still in a conceptual stage with no detailed specifications or design

3.2.2 EF Tracking Product Breakdown Structure (PBS)

The EF Farm already exists in the PBS as item 1.3.2.1, commodity PC servers on which the EF selection software is executed. The compute power is to be provided by CPUs with the option of including accelerators, so this is consistent with the heterogeneous commodity solutions being considered for the EF Tracking. The scope of this EF Farm will therefore be expanded to meet the full requirements for EF Tracking. Consequently there is no addition to the PBS under 1.3.4.

3.2.3 EF Farm Cost Estimate

The required CPU and CORE cost to achieve the projected EF farm sizes in Run-4 and Run-5 of the HL-LHC are presented in Table 3.2.

The CPU requirements for non-tracking trigger processing have been estimated for the TDAQ Phase-II TDR [1] to be 1.86 MHS06 for the level of pile-up of 140 for Run-4 and to 2.27 MHS06 for a pile-up of 200 for Run-5.

The required resources for tracking (Table 2.5) are based on the CPU-based fast-tracking demonstrator described in Section 2.4.1. The reasons for this are two-fold: first, estimating CPU requirements is less uncertain than estimating the cost of a potential heterogeneous system, and second, based on the outcomes of the Task Force reports described in Section 1.4, it is expected that a heterogeneous system could fit into the cost envelope presented below. The associated risk linked to this assumption is discussed in Section 3.5.

Two factors contribute to the CORE cost estimates: (i) the CPU cost extrapolations, together with the projected hardware configuration of the commodity servers; and (ii) the continuation of the M&O rolling replacement foreseen to continue throughout the lifetime of the experiment.

CPU cost extrapolations: The EF farm purchases will be done in two batches to provide the required CPU for the different levels of pile-up in Run-4 and Run-5 according to the expected luminosity profile of the HL-LHC (Figure 2.1). Using the cost projection of CERN IT shown in Figure 1.3 with a 15% cost reduction per year, we assume a compute cost of

1.3 CHF/HS06 for Run-4 (2027) and 0.6 CHF/HS06 for Run-5 (2032), respectively. Further delay of the LHC schedule would result in further price reductions.

M&O rolling replacement: It is assumed that old servers are subject to periodic replacement, which is typically covered by M&O Common funds and is not considered part of project costs. The rolling replacement of the existing Run-3 EF farm would provide an estimated 2.85 MHS06 at the start of Run-4 and will reflect the EF farm technology decision. The excess farm capacity required to accommodate EF Tracking is considered as CORE cost.

Table 3.2: *Required CPU and CORE cost to achieve the projected* **total** *EF farm sizes in Run-4 and Run-5 of the HL-LHC. The reported costs are the costs of the farm extension required in addition to the available CPU from the previous run period. The total CORE cost also includes a 3% increase to cover additional services required to operate the farm.*

	Run-4 (2027)	Run-5 (2032)
CPU for tracking [MHS06]	5.90	9.17
Other reconstruction [MHS06]	1.86	2.27
Required CPU [MHS06]	7.76	11.44
Available CPU [MHS06]	2.85	7.76
Compute cost [CHF/HS06]	1.3	0.6
Farm extension cost [MCHF]	6.4	2.2
Total CORE cost [MCHF]	8.8	

The uncertainty on the compute cost is driven by the CHF/HS06 extrapolation method, which is updated annually by CERN IT (Fig. 1.3) and has varied considerably from year to year in the past. This variability reflects the market-driven nature of the commodity compute technology sector. Since this extrapolation method (based on physicist experience) drives the cost uncertainty, we assign an overall quality factor of QF3.

3.3 Planning and Schedule

The EF Tracking subsystem will adhere to the Schedule Management Plan summarized in the Phase-II TDAQ TDR [1] and detailed in the references therein.

Additionally, the EF Tracking subsystem will follow the ATLAS review process in alignment with the EF Farm subsystem. These reviews are done in agreement with ATLAS Management and are organised by the ATLAS Electronics Coordinator, in his/her function of chair of the ATLAS Electronics Review Office. There are four types of design reviews planned:

- the Specification Validation Review (SVR),

- the Preliminary Design Review (PDR),
- the Final Design Review (Final Design Review (FDR)),
- the Production Readiness Review (Production Readiness Review (PRR)).

Annual software and firmware reviews will be carried out throughout the project as described below. As is already done for the trigger and offline software, the software will be developed with standard ATLAS tools and processes: continuous integration builds and tests and peer review of software before approval of each new request.

Detailed bottom-up plans of the activities of the EF Tracking sub-system are used as the basis for the overall project plan, schedule and milestones. The schedule includes all the tasks required until the delivery of the project to ATLAS, ready for installation, and all their interdependencies. The plans will evolve as the project moves forward, including the addition of more detailed breakdown of tasks as they approach. This will enhance the ability of the TDAQ Upgrade Project management to track progress, detect delays as early as possible, and take preemptive actions.

The high-level development phases for the EF Tracking sub-system are presented in Fig. 3.2; a milestone is associated with the end of each phase as summarized in Table 3.3. The initial phase allows for the definition of the requirements and specifications of the system, and ends with the SVR. The first and second demonstrator development phases apply to software, firmware, and GPU development. Each phase has a fixed duration of one year and will begin with the specification of the goals for that development phase, followed by a high-level design period, a technical implementation period, an integration period, a documentation period. Each phase will then conclude with an EF review of the demonstrator, including software and firmware. The goals of these and subsequent development phases will be synchronized with the offline upgrade tracking software development effort. Critical aspects impacting both the online and offline tracking for Run 4 include the integration of Acts into ATHENA, the development of the EDM and heterogeneous computing support in Acts and ATHENA, and reconstruction optimizations.

The development of test stands (both remote and at CERN) will be critical to the technology evaluation of the project as well as to understand load-balancing and internal dataflow considerations.

The software framework and test stand results are required as input to the “Preparation for Technology Choice" phase, during which the direct comparison of a variety of algorithm and technology choices will be carried out. This phase will culminate in the PDR at which time a technology choice will be made and the baseline system will be fully defined. The criteria for the technology choice are listed in Section 2.5.6.

A “Prototype Development" phase follows (involving test stands at remote sites and at CERN), at the end of which a working prototype of the needed software (and firmware) for the system will be delivered for the FDR. A “Final System Preparation" phase will allow for

any updates to the prototype as well as documentation of the specific system requirements that will be provided to the EF Farm sub-system to begin the procurement process.

There are two floats associated to the EF Tracking sub-system. The first is defined as the number of working days between the “SW/FW Ready for Commissioning" milestone and the “SW/FW Needed by ATLAS for Commissioning" date (the same date as for the rest of the EF software), while the second float is defined as the number of working days between the “Requirements Ready for EF PRR" milestone and the “Requirements Needed for EF PRR" date provided by WBS 1.3.2. Each float is approximately seven months and three months respectively. The installation and commissioning milestones are held by ATLAS Technical Coordination.

The PRR will be conducted within the EF Farm sub-system, since that sub-system is responsible for the procurement of the EF commodity system as a whole, and is planned for Q2 2027.

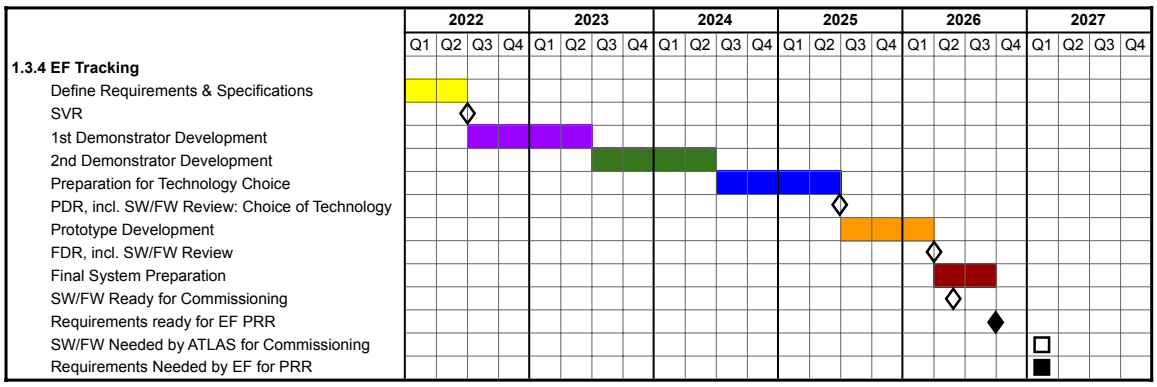

Figure 3.2: *Development phases of the EF Tracking sub-system. Each phase culminates in a milestone; the high-level milestones for WBS 1.3.4 are listed in Table 3.3.*

Table 3.3: *Principal milestones for the EF Tracking sub-system (WBS 1.3.4). The milestones indicated with a * are review stages defined by ATLAS, as described above. The PRR is conducted under the EF Farm sub-system (WBS 1.3.2), and thus is not listed here.*

WBS	Descr.	Milestone	Date
1.3.4	EF Tracking		
		SVR*	01.07.2022
		First Demonstrator SW/FW Review	01.07.2023
		Second Demonstrator SW/FW Review	01.07.2024
		PDR*, incl. SW/FW Review: Choice of Technology	01.07.2025
		FDR*, incl. SW/FW Review	01.04.2026
		SW/FW Ready for Commissioning	01.06.2026
		Requirements Ready for EF PRR	01.10.2026
		SW/FW Needed by ATLAS for Commissioning	05.01.2027
		Requirements Needed for EF PRR	05.01.2027

3.4 Resource Requirements and Institutional Responsibilities

This chapter documents the tasks and the associated required human resources for the EF Tracking sub-system. The UPL and the TDAQ Institute Board chair are surveying the aspirations of the Institutes, and investigating with their Representatives whether the required resources are available locally among those Institutes that have expressed interest in the EF Tracking sub-system. This is an iterative process that will ultimately conclude with an updated "Memorandum of Understanding" (MoU) document codifying the agreement between ATLAS and the Institutes after approval by the CERN Resource Review Board. The MoU will be signed by ATLAS and the Funding Agencies of the participating Institutes. Formal commitments are made only after the MoU has been signed.

The Resource Management Plan is documented in the Phase-II TDAQ TDR [1]. A summary of the required effort for EF Tracking is presented in Table 3.4 and Fig. 3.3 for the case of a heterogeneous commodity system. A total of 89.7 FTEs are needed over a five-year period, for an average of 18 FTEs per year. In the table and figure the effort is divided into the following professional categories based on the required expertise: scientist, software/firmware (SW/FW) engineer, and student. GPU experts are counted under the SW/FW engineer category.

The required effort was estimated for each L4 area and skill set with the development phases described in Fig. 3.2 in mind. The required effort for each development phase is based on analogous development phases in similar projects within the Phase-I and Phase-II TDAQ upgrades, experience from the task forces and Acts development efforts, and expert opinion. It is expected that the two demonstrator development phases and the preparation for the technology choice include developments on a variety of platforms that require FPGA, GPU, and software expertise. After the technology choice in mid-2025, the expert-level effort will focus on the prototype development for the chosen platform. During the prototype development and final system preparation phases, the balance between firmware and software engineering effort will depend on the relevant expertise for the technology choice, but the overall SW/FW FTE level is not expected to exceed the estimated envelope. There is also a ramp-up of student and scientist effort during this phase to focus on the optimization of the tracking for the relevant trigger signatures in conjunction with the trigger signature experts. An estimated effort level is also included for the coordination of the project based on the organizational structure presented in Fig. 3.1.

The studies of trigger signatures and menus are shared with the TDAQ Performance, Physics, and Event Selection, ATLAS Upgrade Physics, and ATLAS Trigger activity groups. The project also benefits from close collaboration with Upgrade Tracking within the ATLAS Tracking Combined Performance group and the Heterogeneous Computing and Accelerator Forum (HCAF). This is taken into account in the effort estimates.

Additional detailed information is being prepared for the Upgrade Cost Group confidential material in the form of Excel spreadsheets.

Table 3.4: *Required effort in the EF Tracking sub-system expressed in FTEs and divided by professional category for the duration of the project construction (2022-2026).*

1.3.4: EF Tracking	2022	2023	2024	2025	2026	Total
Scientist	6.5	6.5	6.2	5.4	4.7	29.3
SW/FW Engineer	8.3	8.3	7.5	5.6	3.6	33.3
Student	6.0	6.0	5.2	4.5	5.5	27.2
Total	20.8	20.8	18.9	15.5	13.7	89.7

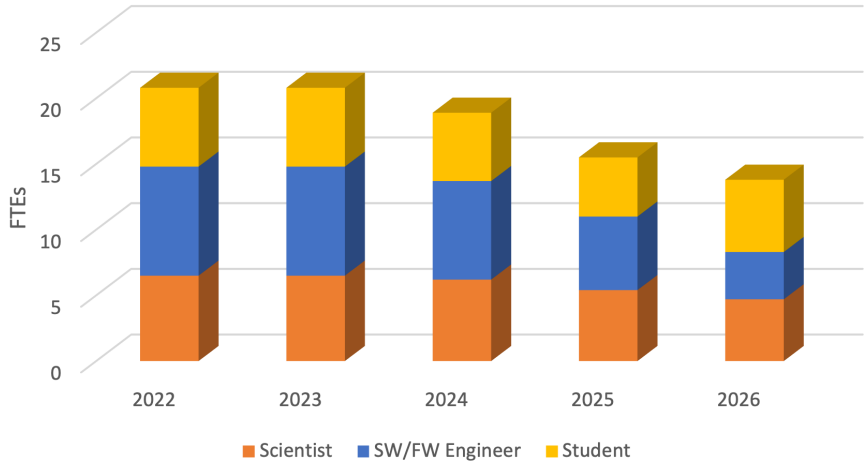

Figure 3.3: *Estimated resources required in the EF Tracking subsystem, expressed as FTEs for each year during the project construction. Three professional categories are stacked on top of each other: Scientist (orange), SW/FW Engineer (blue), and Student (yellow).*

The interests and aspirations of Institutes (within and outside the current ATLAS TDAQ collaboration) to participate in the EF Tracking subsystem are presented in Table 3.5.

Table 3.5: *List of participating Institutes in the Event Filter Tracking subsystem.*

Country / Institute	Event Filter Tracking
CERN	
CERN	✓
Denmark	
Copenhagen NBI	✓
Germany	
Heidelberg PI	✓
Greece	
Thessaloniki	✓
Italy	
Bologna	✓
Genova	✓
Pisa	✓
Netherlands	
Nijmegen	✓
Nikhef	✓
Poland	
Cracow AGH/UJ cluster	✓
Switzerland	
Geneva	✓
United Kingdom	
Edinburgh	✓
London RHBNC	✓
London UC	✓
Manchester	✓
RAL	✓
Sussex	✓
USA	
Argonne	✓
Arizona	✓
Brookhaven	✓

continue . . .

. . . continued

Country / Institute	Event Filter Tracking
California (Irvine)	✓
Chicago	✓
Duke	✓
Illinois	✓
Northern Illinois	✓
Ohio State	✓
Pennsylvania	✓
SLAC	✓
Wisconsin	✓

3.5 Risk Analysis and Mitigation Strategies

The EF Tracking subsystem will follow the risk management plan described in the Phase-II TDAQ TDR [1] and the TDAQ Upgrade Project Risk Management Plan [46]. Risks are managed by a structured and integrated process for identifying, evaluating, tracking, mitigating, responding to, and managing project risks in terms of three risk categories: cost, schedule, and scope/technical performance.

The Upgrade Cost Group Confidential Material package contains a full description of the Risk Register, with risk identification, impact analysis, and mitigation strategy. Only the elements identified at this stage with the potential to have a significant impact on the project are presented in this section.

The EF Tracking sub-system and the overall EF system share a common technology, so they have several risks in common as described below.

3.5.1 Detector simulation software schedule

The EF Tracking software depends on a timely delivery of simulation of the ITk detector data. EF Tracking software components require realistic input data, both in terms of test samples and large physics datasets, in order to develop and test them, and make realistic measurements of the resources. Such a delay would have an impact on the schedule and/or performance of the EF Tracking software. This is mitigated by connecting the plans of the EF Tracking software with those for offline tracking simulation and ensuring good communication across the communities. This will be the responsibility of the Software Coordinator and of the Software Framework & Technology L4 managers, which are described in Section 3.1

3.5.2 Delay to external tracking software

The ATLAS track reconstruction will be developed through a close collaboration between the TDAQ upgrade EF Tracking and offline upgrade tracking groups, and will use the Acts project with the expectation that both TDAQ and offline will benefit from shared expertise, code and libraries. This does mean that the EF Tracking will depend on tracking software that is not all under its direct management. If such a part of the software is delayed, it could affect the schedule for EF Tracking. This will be controlled by working closely with the Upgrade Tracking group conveners to reconcile the priorities of the two groups and by carefully following progress with respect to the EF Tracking schedule. If the risk is realized, the response would be to adapt the EF Tracking schedule and seek additional effort to recover lost time.

3.5.3 Loss of key personnel

If a key person ceases to be available, or is available at a much lower than planned FTE level, then the consequences are that there will be a delay to the project. This is mitigated by identifying key tasks that need to be spread over several persons, by ensuring adequate documentation is produced, and by choosing software and firmware tools that are used by the majority of the community. Key people are more likely to be in long-term positions at their institutes and the L3 managers will have regular communication with each key person.

3.5.4 Size and power of the system

If the assumed number of servers that can fit into a rack is too optimistic or more auxiliary items are needed in a rack, then more racks will need to be purchased. Similarly, the power requirement of the system might exceed the available power. This is mitigated by optimizing the rack layout during the vertical integration tests, by scaling the number and type of accelerators as part of the system or by choosing more expensive CPUs with lower power needs.

3.5.5 Cost and availability uncertainties

There are cost and availability uncertainties due to market-driven prices for commodity components. This can be considered both a risk and an opportunity. For example regarding CPUs, the server market in the last decade was dominated almost exclusively by Intel. With AMD becoming competitive again, one would expect a further decline of the prices. New competitors like ARM are entering the market. However, price fluctuations due to other components (e.g. RAM) are not excluded and have impacted cost in the recent past. Similar availability and cost issues may affect the FPGA and GPU markets. The availability risk is mitigated by monitoring the markets. It may be possible to adjust the procurement schedule if helpful.

3.5.6 Insufficient throughput

If the EF Tracking is not fast enough to achieve the throughput required during data taking, this would require a temporary action to reduce the L0 trigger rate or change the EF menu to reduce the use of tracking, which would impact the physics performance of lower-priority triggers. To recover the physics performance, either an increase in the size of the EF processing farm or more effort to improve the performance would be needed. This will be controlled by regular performance checks during the development and integration phases.

3.5.7 New physics requirements

It is possible that during the course of the EF Tracking construction or once data-taking has started that new physics requirements could come along that present a significant additional tracking requirement. An example of this is the LRT use case discussed in Section 2.2 for which tracking resources are dependent on the specific LLP model simulated. A different model could require much more tracking resources than have been allowed for. This risk is mitigated by the Performance Coordinator ensuring good communication with the Performance, Physics, and Event Selection group to keep abreast of such developments. Early warning gives time to adapt: develop an optimal tracking working point and triggering strategy taking into account the EF farm capacity and if necessary re-prioritise the overall trigger menu.

3.5.8 Technical or performance problems during prototyping

Although the planning of the project allows for a careful evaluation of the suitability of each technology (CPU, FPGA, GPU) prior to the PDR, there remains a small risk that during the Prototype Development phase some technical difficulties may still be encountered that could delay the project or increase its cost, or that unexpected performance issues arise. This risk is mitigated by a thorough PDR review process with clear criteria and benchmarks based on complete and realistic integration test results.

3.5.9 Replacement of accelerators

Accelerators may be replaced when they reach the end of their warranty, become unreliable, or the servers that host them are replaced. This is likely to happen a few times during the lifetime of the HL-LHC.

This presents the risk that there might be a need for substantial changes to the firmware when accelerators are replaced, beyond the usual maintenance performed during a Run. In order to mitigate this risk accelerator firmware design should make use of standard tools, which are portable between devices, whenever possible. In addition, operations plans should include sufficient effort to make these updates. This risk only applies to the system during its operations period and leads only to an increase in the amount of engineering effort required.

The replacement of accelerators during the period of HL-LHC running may also present an opportunity by offering an upgrade in computing power which would enable the EF to do more and/or higher-quality tracking. This would also provide additional flexibility to the EF system by allowing new tracking algorithms to be developed for more performant accelerators.

3.5.10 Other opportunities

In addition to the risks described above that are mainly threats, the EF Tracking project also presents some interesting opportunities. For example, in the event of a new emergent technology applicable to the project, the performance of the tracking algorithms could be improved or the cost of the EF farm be reduced. Similarly, the cost of the farm would also be reduced if the performance of the tracking algorithms ends up being much better than anticipated. If there is a significant delay to the start of the HL-LHC running period, the cost of the EF farm is also likely to be reduced, due to the policy of purchasing commodity components as late as possible to benefit from the trend of a decrease in cost with time.

A Glossary

Acts A Common Tracking Software 12, 53–58, 62, 63, 67, 69, 73

AI Artificial Intelligence 62, 63

ASIC Application-Specific Integrated Circuits 3, 5

BRAM Block Random Access Memory 41

CBE Current Best Estimate: the value of a resource or of a parameter extracted through simulations and/or extrapolation of existing data to the conditions foreseen in the HL-LHC phase and/or best current knowledge of what will be the ATLAS experimental layout 64

combination A specific track candidate with one hit per logical layer. Used as the input to a track fitter with a fixed number of layers. 39

CORE COsting REview; costing policy established by the LHCC at the time of the original construction 1, 63–66

COTS Commercial Off-The-Shelf electronics component 64

CPU Central Processing Unit 1, 5, 6, 19, 23, 32, 34, 37, 38, 40, 43, 44, 53, 56

CUDA Compute Unified Device Architecture 50, 53, 56

DAQ Data Acquisition System 9, 61

Dataflow Dataflow subsystem 61

DMA Direct Memory Access 38

DSP Digital Signal Processor 5, 41

EDM Event Data Model 40, 54, 67

EF Event Filter . 1–3, 7–9, 11–13, 16–20, 32, 37, 40, 41, 43, 49, 52, 53, 57–59, 61–63, 65–70, 73–76

eTDSG Extended TDAQ Steering Group 7, 8

FDR Final Design Review 2, 67, 68

FELIX Front-End LInk eXchange 38, 39

FF Flip-Flop 41

FPGA Field Programmable Gate Array . 1, 3, 5–7, 17–20, 34, 36–45, 48, 49, 53, 55–57, 69

FTE Full Time Equivalent 69, 70

FTK Fast TracKer 39

gHTT Full-scan Global Hardware-based Tracking for the Trigger 43

GNN Graph Neural Network 36, 56

GPU Graphics Processing Unit. Usable for general purpose calculations (also known as GPGPU) 1, 5, 6, 17–20, 50–53, 55, 56, 58, 63, 67, 69, 74, 75

HL-LHC High Luminosity LHC 1, 12, 13, 17, 21, 24, 35, 36, 58, 65, 66, 77

HLS High-level Synthesis 53

HPC High Performance Computing 20, 54

HTT Hardware-based Tracking for the Trigger 3, 40, 43

ID Inner Detector 50–52

IPC Interprocess Communication 51

ITk Inner Tracker . . . 1, 3–5, 7, 11–13, 18, 20–31, 34–36, 38–40, 43, 50, 53–57, 61, 73

L0 Level-0 Trigger 15, 16, 61

LHCC LHC Experiments Committee 3, 8, 63

LLP Long-Lived Particle 1, 12, 16, 17, 75

LRT Large-Radius Tracking: reconstruction of tracks with large impact parameter that do not originate from the primary vertex . 1, 16, 17, 23–25, 31–33, 44, 48, 75

LUT Lookup Table. An array that replaces runtime computation with a simpler array indexing operation. 41

MC Monte-Carlo 13, 17, 19, 22, 23, 25, 26, 63

ML Machine Learning 55, 56, 62, 63

MoU Memorandum of Understanding 69

NN Neural Network 39, 41, 43, 45, 47, 48, 57

P2UG Phase-II Upgrade Cost Group 3

PBS Product Breakdown Structure 65

PCIe Peripheral Component Interconnect Express 6, 20, 38, 41, 50, 58

PDR Preliminary Design Review 67, 68

Point-1 ATLAS experiment site on the LHC ring 18, 32, 53

PRR Production Readiness Review 2, 67, 68

road A loose track candidate output generally defined by a set of track parameters (p_T , ϕ) and a z slicing. It may have several hits per layer (i.e. hits consistent with parameters). 36, 39, 45

RoI Region of Interest 13, 15, 23, 31

SVR Specification Validation Review 66–68

TDAQ ATLAS Trigger and Data Acquisition System 1, 3–5, 7–9, 11–13, 15–18, 20, 23, 50, 54, 61, 63, 66, 67, 69, 70, 73, 77, 79

TDIB TDAQ Institutional Board 8

TDR Technical Design Report 1, 3–5, 7, 8, 11–13, 15, 16, 18, 20, 23, 50, 61, 63, 66, 69, 73

TLA Trigger-object Level Analysis 1

UDP User Datagram Protocol 39

UPL Upgrade Project Leader 61, 69

UPR Upgrade Project 61, 64

URAM Ultra Random Access Memory 41

B References

[1] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*, Tech. Rep. CERN-LHCC-2017-020. ATLAS-TDR-029, CERN, Geneva, Sept., 2017. https://cds.cern.ch/record/2285584.

[2] ATLAS Collaboration, *Expected tracking and related performance with the updated ATLAS Inner Tracker layout at the High-Luminosity LHC*, tech. rep., CERN, Geneva, Jul, 2021. http://cds.cern.ch/record/2776651. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-024.

[3] ATLAS Collaboration, *Fast Track Reconstruction for HL-LHC*, tech. rep., CERN, Geneva, Oct., 2019. https://cds.cern.ch/record/2693670. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2019-041.

[4] M. Michelotto et al., *A comparison of HEP code with SPEC benchmarks on multi-core worker nodes*, J. Phys. Conf. Ser. **219** (2010) 052009.

[5] B. Panzer-Steindel. Private communication, Jan., 2021.

[6] (online). https://www.oneapi.com/.

[7] X. Ai, C. Allaire, N. Calace, P. Gessinger, H. Grasland, H. Gray, C. Gumpert, J. Hrdinka, M. Kiehn, F. Klimpel, R. Langenberg, A. Salzburger, and B. Schlag, *Acts Project*, (online), Nov., 2020. https://doi.org/10.5281/zenodo.3937454.

[8] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, Tech. Rep. CERN-LHCC-2017-005, CERN, Geneva, Apr, 2017. https://cds.cern.ch/record/2257755.

[9] P. Calafiura, J. Catmore, D. Costanzo, and A. Di Girolamo, *ATLAS HL-LHC Computing Conceptual Design Report*, tech. rep., CERN, Geneva, Sept., 2020. https://cds.cern.ch/record/2729668.

[10] M. Michelotto, M. Alef, A. Iribarren, H. Meinhard, P. Wegner, M. Bly, G. Benelli, F. Brasolin, H. Degaudenzi, A. D. Salvo, I. Gable, A. Hirstius, and P. Hristov, *A comparison of HEP code with SPEC1benchmarks on multi-core worker nodes*, Journal of Physics: Conference Series **219** (2010) no. 5, 052009. https://doi.org/10.1088/1742-6596/219/5/052009.

[11] ATLAS Collaboration, P. Conde Muíño, *Multi-threaded algorithms for GPGPU in the ATLAS High Level Trigger*, J. Phys. Conf. Ser. **898** (2017) no. 3, 032003.

[12] F. Berghaus et al., *ATLAS Sim@P1 upgrades during long shutdown two*, EPJ Web Conf. **245** (2020) 07044.

[13] P. Calafiura, J. Catmore, D. Costanzo, and A. Di Girolamo, *ATLAS HL-LHC Computing Conceptual Design Report*, Tech. Rep. CERN-LHCC-2020-015. LHCC-G-178, CERN, Geneva, Sept., 2020. https://cds.cern.ch/record/2729668.

[14] RD53 Collaboration, *RD53B Design Requirements*, Tech. Rep. CERN-RD53-PUB-19-001, CERN, Geneva, Feb., 2019. https://cds.cern.ch/record/2663161.

[15] ATLAS Collaboration, *Study of tracking and software performance of the ATLAS Inner Tracker in staged-installation scenarios at the High-luminosity LHC*, tech. rep., CERN, Geneva, Nov, 2021. https://cds.cern.ch/record/2790800. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-040.

[16] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, Tech. Rep. CERN-LHCC-2017-021. ATLAS-TDR-030, CERN, Geneva, Sept., 2017. https://cds.cern.ch/record/2285585.

[17] P. V. Hough, *Method and means for recognizing complex patterns*, Dec., 1962. US Patent 3,069,654.

[18] R. O. Duda and P. E. Hart, *Use of the Hough Transformation to Detect Lines and Curves in Pictures*, Commun. ACM **15** (1972) no. 1, 11–15. https://doi.org/10.1145/361237.361242.

[19] M. Martensson, *Fast pattern recognition of ATLAS L1 track trigger for HL-LHC*, PoS **Vertex 2016** (2017) 069.

[20] J. Baines et al., *Report of the ATLAS Alternative Event Filter Tracking Working Group*, Tech. Rep. ATL-COM-DAQ-2019-173, CERN, Geneva, Oct., 2019. https://cds.cern.ch/record/2693094.

[21] M. Mårtensson, *A search for leptoquarks with the ATLAS detector and hardware tracking at the High-Luminosity LHC*. PhD thesis, Uppsala U., 2019.

[22] M. Shochet, L. Tompkins, V. Cavaliere, P. Giannetti, A. Annovi, and G. Volpi, *Fast TracKer (FTK) Technical Design Report*, tech. rep., Jun, 2013. https://cds.cern.ch/record/1552953. ATLAS Fast Tracker Technical Design Report.

[23] J. Duarte, S. Han, P. Harris, S. Jindariani, E. Kreinar, B. Kreis, J. Ngadiuba, M. Pierini, R. Rivera, N. Tran, and et al., *Fast inference of deep neural networks in FPGAs for particle physics*, Journal of Instrumentation **13** (2018) no. 07, P07027–P07027. http://dx.doi.org/10.1088/1748-0221/13/07/P07027.

[24] E. Clement, M. De Mattia, S. Dutta, R. Eusebi, K. Hahn, Z. Hu, S. Jindariani, J. Konigsberg, T. Liu, J. Low, and et al., *A high-performance track fitter for use in ultra-fast electronics*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **935** (2019) 95–102. http://dx.doi.org/10.1016/j.nima.2019.05.018.

[25] NVIDIA, *CUDA*, (online), 2020. https://developer.nvidia.com/cuda-toolkit.

[26] 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.

[27] R. Nane, V.-M. Sima, C. Pilato, J. Choi, B. Fort, A. Canis, Y. T. Chen, H. Hsiao, S. Brown, F. Ferrandi, J. Anderson, and K. Bertels, *A Survey and Evaluation of FPGA High-Level Synthesis Tools*, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **35** (2016) no. 10, 1591–1604.

[28] (online). https://eigen.tuxfamily.org.

[29] ATLAS Collaboration, *Development of ATLAS Primary Vertex Reconstruction for LHC Run 3*, tech. rep., CERN, Geneva, Apr, 2019. https://cds.cern.ch/record/2670380. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2019-015.

[30] ATLAS Collaboration, *Software Performance of the ATLAS Track Reconstruction for LHC Run 3*, tech. rep., CERN, Geneva, May, 2021. http://cds.cern.ch/record/2766886. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-012.

[31] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, *A Comprehensive Survey on Graph Neural Networks*, IEEE Transactions on Neural Networks and Learning Systems **32** (2021) no. 1, 4–24.

[32] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, *The Graph Neural Network Model*, IEEE Transactions on Neural Networks **20** (2009) no. 1, 61–80.

[33] S. Farrell, P. Calafiura, M. Mudigonda, Prabhat, D. Anderson, J.-R. Vlimant, S. Zheng, J. Bendavid, M. Spiropulu, G. Cerati, L. Gray, J. Kowalkowski, P. Spentzouris, and A. Tsaris, *Novel deep learning methods for track reconstruction*, 2018. arXiv:1810.06111 [hep-ex].

[34] X. Ju, S. Farrell, P. Calafiura, D. Murnane, Prabhat, L. Gray, T. Klijnsma, K. Pedro, G. Cerati, J. Kowalkowski, G. Perdue, P. Spentzouris, N. Tran, J.-R. Vlimant, A. Zlokapa, J. Pata, M. Spiropulu, S. An, A. Aurisano, J. Hewes, A. Tsaris, K. Terao, and T. Usher, *Graph Neural Networks for Particle Reconstruction in High Energy Physics detectors*, 2020. arXiv:2003.11603 [physics.ins-det].

[35] X. Ju, D. Murnane, P. Calafiura, N. Choma, S. Conlon, S. Farrell, Y. Xu, M. Spiropulu, J.-R. Vlimant, A. Aurisano, J. Hewes, G. Cerati, L. Gray, T. Klijnsma, J. Kowalkowski, M. Atkinson, M. Neubauer, G. DeZoort, S. Thais, A. Chauhan, A. Schuy, S.-C. Hsu, A. Ballow, , and A. Lazar, *Physics and Computing Performance of the Exa.TrkX TrackML Pipeline*, 2021. arXiv:2103.06995 [hep-ex].

[36] G. Dezoort, S. Thais, I. Ojalvo, P. Elmer, V. Razavimaleki, J. Duarte, M. Atkinson, and M. Neubauer, *Charged particle tracking via edge-classifying interaction networks*, arXiv:2103.16701 [hep-ex].

[37] C. Biscarat, S. Caillou, C. Rougier, J. Stark, and J. Zahreddine, *Towards a realistic track reconstruction algorithm based on graph neural networks for the HL-LHC*, 2021. arXiv:2103.00916 [physics.ins-det].

[38] A. Heintz et al., *Accelerated Charged Particle Tracking with Graph Neural Networks on FPGAs*, in *34th Conference on Neural Information Processing Systems*. 11, 2020. arXiv:2012.01563 [physics.ins-det].

[39] (online). https://github.com/acts-project/traccc.

[40] (online). https://github.com/acts-project/detray.

[41] X. Ju et al., *Physics and Computing Performance of the Exa.TrkX TrackML Pipeline*, arXiv:2103.06995 [hep-ex].

[42] S. Amrouche, L. Basara, P. Calafiura, V. Estrade, S. Farrell, D. R. Ferreira, L. Finnie, N. Finnie, C. Germain, V. V. Gligorov, and et al., *The Tracking Machine Learning Challenge: Accuracy Phase*, p. 231–264. Springer International Publishing, Nov, 2019. http://dx.doi.org/10.1007/978-3-030-29135-8_9.

[43] S. Amrouche, L. Basara, P. Calafiura, D. Emeliyanov, V. Estrade, S. Farrell, C. Germain, V. V. Gligorov, T. Golling, S. Gorbunov, H. Gray, I. Guyon, M. Hushchyn, V. Innocente, M. Kiehn, M. Kunze, E. Moyse, D. Rousseau, A. Salzburger,

A. Ustyuzhanin, and J.-R. Vlimant, *The Tracking Machine Learning challenge : Throughput phase*, arXiv:2105.01160 [cs.LG].

[44] E. Clement et al., *A High-performance Track Fitter for Use in Ultra-fast Electronics*, Nucl. Instrum. Meth. A **935** (2019) 95–102, arXiv:1809.01467 [physics.ins-det].

[45] X. Ai, G. Mania, H. M. Gray, M. Kuhn, and N. Styles, *A GPU-based Kalman Filter for Track Fitting*, arXiv:2105.01796 [physics.ins-det].

[46] The TDAQ Group, *TDAQ Phase-II Upgrade Risk Management Plan*, Tech. Rep. ATL2-D-MG-0004, EDMS Id 1874388 v.1, CERN, Geneva, Nov, 2017. https://edms.cern.ch/document/1874388/1.

The ATLAS Collaboration

The ATLAS Collaboration

Argentina
Buenos Aires
M.F. Daneri, G. Otero y Garzon, R. Piegaia, M. Toscani

La Plata
M.J. Alconada Verzini, F. Alonso, M.T. Dova, J. Hoya, F. Monticelli, S.J. Noacco Rosende, G.E. Orellana, F. Sili, H. Wahlberg

Australia
Adelaide
P. Jackson, A.X.Y. Kong, H. Potti, T.A. Ruggeri, A.S. Sharma, E.X.L. Ting, M.J. White

Melbourne
E.L. Barberio, S. Goldfarb, E.F. McDonald, P.C. McNamara, T. Pham, F. Scutti, J. Shojaii, G.N. Taylor, P. Urquijo

Sydney
S. Balaji, K.E. Varvell, B. Yabsley

Austria
Innsbruck
E. Kneringer, A. Manousos

Azerbaijan
Baku
N. Huseynov

Belarus
Minsk AC
S. Harkusha, Y. Kulchitsky, Y.A. Kurochkin, P.V. Tsiareshka

Minsk NC
A. Hrynevich

Brazil
Juiz de Fora UF
V. Araujo Ferraz, A.S. Cerqueira, L. Manhaes de Andrade Filho, D. Oliveira Goncalves, B.S. Peralva

Sao Paulo
M. Donadelli, M.A.L. Leite

Rio de Janeiro UF
Y. Amaral Coutinho, M. Begalli, L.P. Caloba, J.V. Da Fonseca Pinto, E. Egidio Purcino De Souza, W.S. Freund, P. Gaspar, J. Lieber Marin, C. Maidantchik, J.M. Seixas, M. Verissimo De Araujo

Sao Joao del Rei UF
M.A.B. Do Vale

Canada
Alberta
J.P. Biswal, D.M. Gingrich, J.H. Lindon, N. Nishu, J.L. Pinfold

Carleton
A. Bachiu, A. Bellerive, B. Davis-Purcell, D. Gillberg, K. Graham, J. Heilman, S. Kaur, J.S. Keller, C. Klein, T. Koffas, A. Laurier, L.S. Miller, M. Naseri, F.G. Oakham, D.A. Pizzi, M.G. Vincter, S.A. Weber, N. Zakharchuk

McGill
W.S. Ahmed, A. Ambler, A. Canesse, F. Corriveau, T. Kwan, Z. Li, J.P. Mc Gowan, C. Nelson, S.H. Robertson, S. Saha, S. Stärz, B. Vachon, A. Warburton

Montreal
J-F. Arguin, G. Azuelos, H.M. Borecka-Bielska, D. Godin, C. Leroy, K. Mochizuki, C. Papadatos, B. Pascual Dias, M. Usman

Simon Fraser
J.C. Burzynski, M. Danninger, E. Drechsler, B.P. Jaeger, K. Lehmann, B.X. Liu, J. Llorente Merino, D. Mori, D.C. O'Neil, A. Poley, S. Singh, B. Stelzer, M.C. Vetterli

Toronto
L. Adamek, O. Arnaez, M.J. Basso, J.W.S. Carter, A.F. Casha, B.M. Ciungu, K. Dette, J.H. Foo, C.A. Garner, N. Giangiacomi, Y.F. Han, N. Ilic, S. Ketabchi Haghighat,

D.P. Kisliuk, T.M. Knight, P. Krieger, H.Y. Meng, A. Milic, O. Miu, R.S. Orr, T.H. Park, E. Rossi, P. Savard, P. Sinervo, S. Singh, R.J. Teuscher, W. Trischuk, L.M. Veloce

TRIUMF
S.V. Chekulaev, N.P. Hessey, L.L. Kurchaninov, B. Lefebvre, E. Perez Codina, D.M. Portillo Quintero, O. Stelzer-Chilton, M. Swiatlowski, R. Tafirout, I.M. Trigger, M. Valente

York
C. David, M.J. Kareem, A.M. Rodríguez Vera, W.Y. Song, D.M. Starko, W. Taylor

Vancouver UBC
F. Cormier, C. Gay, C. Gubbels, R.L. Hayes, A. Lister, R. Newhouse, Z. Tao, D.A. Trischuk, V.W.S. Wong

Victoria
J. Albert, C. Camincher, E.M. Carlson, B. Chen, Y.H. Chiu, M. Ghasemi Bostanabad, K. Hamano, E.F. Kay, R. Keeler, R. Kowalewski, M. Lefebvre, D.M. Mac Donell, K.D. McLean, R.A. McPherson, J.C. Rivera Vergara, R.A. Rojas, H.L. Russell, M.J. Shroff, R. Sobie, G.A. Vasquez

Chile
ULS
P.A. Ulloa Poblete

UTA
S. Kabana, T. Lagouri, S.A. Olivares Pino

UNAB
S. Kuleshov

Valparaiso UTFSM
W.K. Brooks, E. Carquin, R. Pezoa, C.M. Robles Gajardo, N. Viaux Maira

Santiago
A.C. Abusleme Hoffman, M.A. Diaz, F.M. Garay Walls, R.Y. González Andana, J.P. Ochoa-Ricoux

China
Beijing Tsinghua
W. Ding, M. Feng, B. Li, M. Xia, Y. Xu, G. Zhang

Nanjing

M.K. Ayoub, H. Chen, A. De Maria, A. D'onofrio, L. Han, X. Huang, Z. Jia, Y. Liu, F.L. Lucio Alves, T. Min, X. Wang, L. Xia, H. Ye, B. Zhang, L. Zhang

Beijing IHEP
X. Chu, H. Cui, Y. Fan, Y. Fang, F. Guo, Y.F. Hu, M. Li, B. Liu, P. Liu, G. Lu, A.F. Mohammed, K. Ran, J. Wu, X. Yang, K. Zhang, C. Zhu

UCAS

TDLI
N. Brahimi, C. Kato, S. Li, D. Liu, K. Liu, Q. Liu, D. Sampsonidou, W. Su, Z. Wang

Shanghai
J. Chen, X. Chen, B. Dong, M. Faraj, J. Guo, J. Hong, S. Hu, C-Q. Li, J. Li, L. Li, J. Tang, C. Wang, X. Wang, J. Yan, H.J. Yang, X. Zhang, N. Zhou

Shandong
H. Atmani, C. Feng, B. Li, H. Li, H. Li, T. Li, Z. Li, L.L. Ma, K. Tariq, S. Wang, Z. Xu, R. Yuan, X. Zhang, T. Zhao, C.G. Zhu

Hefei
A. Baroncelli, C. Chen, Y. Chen, M.J. Da Cunha Sargedas De Sousa, D. Du, Y. Fu, J. Gao, A. Giannini, K. Han, L. Han, F. He, Y. Huang, Y. Jiang, D. Krasnopevtsev, C. Li, H. Li, Q.Y. Li, J.B. Liu, M. Liu, M.Y. Liu, X. Liu, Y.W. Liu, R. Ospanov, H. Peng, S. Su, X. Su, C. Wang, T. Wang, W.X. Wang, Y. Wang, C. Wei, Y. Wu, M. Xie, X. Xie, H. Xu, H. Xu, L. Xu, S. Yang, X. Yang, Z. Zhao, H.L. Zhu, Y. Zhu

Hong Kong HKUST
K. Lie, K. Prokofiev, J. Xiang, T. Yang

Hong Kong HKU
D. Paredes Hernandez, C. Peng, K.C. Tam, Y. Tu

Hong Kong CUHK
H.C. Cheng, M.C. Chu, L.R. Flores Castillo, J.M. Iturbe Ponce, K.W. Kwok, T.S. Lau, J. Wang

Colombia
Bogota UNAL
C. Sandoval

Bogota UAN
G. Navarro, Y. Rodriguez Garcia

Czech Republic
Olomouc
P. Baron, K. Cerny, M. Hrabovsky, T. Komarek, J. Kvita, L. Nozka, J. Pacalt

Prague AS
J. Chudoba, J. Hejbal, O. Hladik, P. Jacka, O. Kepka, J. Kroll, A. Kupco, V. Latonova, M. Lokajicek, R. Lysak, M. Marcisovsky, M. Mikestikova, S. Nemecek, O. Penc, P. Sicho, P. Staroba, M. Svatos, M. Tasevsky

Prague CTU
B. Ali, K. Augsten, B. Bergmann, T.R.V. Billoud, M. Havranek, Z. Hubacek, S. Mondal, M. Myska, L. Novotny, V. Petousis, R. Polifka, S. Pospisil, K. Smolek, A. Sopczak, V. Vacek, P. Vokac, V. Vrba, O. Zaplatilek

Prague CU
P. Balek, P. Berta, T. Davidek, J. Dolejsi, Z. Dolezal, J. Faltova, P. Kodys, M. Krivos, R. Leitner, K. Petukhova, V. Pleskot, P. Reznicek, M. Rybar, D. Scheirich, J. Smiesko, M. Spousta, M. Sykora, T. Sykora, P. Tas, S. Todorova-Nova, V. Vorobel

Denmark
Copenhagen NBI
A. Camplani, M. Dam, J.B. Hansen, J.D. Hansen, P.H. Hansen, D.S. Nielsen, T.C. Petersen, C. Wiglesworth, S. Xella

France
Annecy LAPP
C. Adam Bourdarios, M. Belfkir, N. Berger, F. Costanza, M. Delmastro, L. Di Ciaccio, L. Franco, C. Goy, T. Guillemin, T. Hryn'ova, S. Jézéquel, I. Koletsou, J. Levêque, N. Lorenzo Martinez, L. Portales, E. Rossi, A. Sanchez Pineda, E. Sauvan, L.P. Selem

Clermont-Ferrand
D. Boumediene, D. Calvet, S. Calvet, J. Donini, R. Madar, T. Megy, L. Rustige, C. Santoni, A. Tnourji, L. Vaslin, F. Vazeille

Grenoble LPSC
J. Collot, S. Crépé-Renaudin, J.B. De Vivie De Regie, P.A. Delsart, M.H. Genest, R. Hulsken, M. Kuna, N. Lalloue, A. Lucotte, F. Malek, B. Trocmé

IJCLab
C. Agapopoulou, A. Bassalat, L. Duflot, M. Escalier, L. Fayard, D. Fournier, J.-F. Grivaz, L. Guo, L. Iconomidou-Fayard, A. Kotsokechagia, A. Lounis, N. Makovec, N. Morange, P. Puzo, D. Rousseau, A.C. Schaffer, L. Serin, S. Simion, R. Tanaka, A. Trofymov, D. Varouchas, D. Zerwas, Z. Zhang

Marseille CPPM
G. Aad, M. Barbero, G. Bartolini, T.P. Calvet, Y. Coadou, C. Diaconu, F. Djama, A. Duperrin, L. Feligioni, E. Fortin, G.D. Hallewell, N. Kumari, E. Monnier, S. Muanza, E. Nagy, E. Petit, P. Pralavorio, A. Rozanov, J. Sardain, J. Stark, T. Strebler, M. Talby, S. Tisserant, A. Vallier, N.K. Vu

Paris LPNHE
T. Beau, G. Bernardi, M. Bomben, R. Bouquet, G. Calderini, R. Camacho Toro, F. Crescioli, F. Derue, L. Ginabat, Y. He, M. Khandoga, M.W. Krasny, D. Lacour, B. Laforge, B. Malaescu, G. Marchiori, I. Nikolic-Audit, I. Nomidis, J. Ocariz, L. Poggioli, M. Ridel, L. Roos, R.H.M. Taibah, S. Trincaz-Duvoid

Saclay CEA
N. Andari, H. Bachacou, F. Balli, F. Bauer, M. Boonekamp, L. Chevalier, F. Deliot, A. Formica, P.F. Giraud, C. Guyot, F. Jeanneau, T. Kawamoto, M. Kolb, W. Kozanecki, J.F. Laporte, M. Manisha, B. Mansoulie, J-P. Meyer, R. Nicolaidou, A. Ouraou, A. Peyaud, M. Saimpert, L. Schoeffel, Ph. Schune, Ph. Schwemling, M. Vandenbroucke, Z. Wu

Georgia
Tbilisi SU
B. Chargeishvili, T. Djobava, A. Durglishvili, J. Khubua, G. Mchedlidze, I.A. Minashvili, M. Mosidze, T. Zakareishvili

Tbilisi IP
J. Jejelava, E.G. Tskhadadze

Germany
Berlin HU
N.B. Atlay, M. Bahmani, D. Berge, S. Grancagnolo, D. Guest, C. Issever, T.J. Khoo, K. Kreul, H. Lacker, T. Lohse, M. Michetti, Y.S. Ng, C. Scharf, F. Schenck, P. Seema, T. Theveneaux-Pelzer

Bonn
A. Bandyopadhyay, S. Bansal, P. Bauer, P. Bechtle, F. Beisiegel, F.U. Bernlochner, I. Brock, K. Desch, C. Deutsch, F.G. Diaz Capriles, C. Dimitriadi, J. Dingfelder, P.J. Falke, C. Grefe, S. Gurbuz, M. Hamer, F. Hinterkeuser, T. Holm, M. Huebner, F. Huegging, C. Kirfel, O. Kivernyk, T. Klingl, H. Krüger, K. Lantzsch, T. Lenz, A. Melzer, C. Nass, Ö.O. Öncel, D. Pohl, M. Standke, C. Vergis, E. Von Toerne, N. Wermes

DESY
X. Ai, S. Amoroso, L. Aperio Bella, J.-H. Arling, A. Basalaev, M. Baselga, C. Becot, J.K. Behr, I. Bloch, P. Bokan, F. Braren, K. Brendlinger, B. Brüers, C.D. Burgard, S. Carrá,

E. Cheremushkina, J.M. Clavijo Columbie, S. Díez Cornell, B. Dutta, J. Ferrando, P. Gadow, G. Gaycken, N.E.K. Gillwald, V. Goumarre, I.M. Gregor, K. Grevtsov, A. Guida, R. Gupta, M. Habedank, S. Heim, B. Heinemann, L. Helary, J. Hofer, R.M. Jacobs, J. Katzy, D. Kuechler, J.T. Kuechler, T. Kuhl, W.A. Leight, C.E. Leitgeb, X. Li, M. Liberatore, A. Linss, Y. Liu, E.M. Lobodzinska, A. Lopez Solis, F. Meloni, P. Moder, K. Mönig, F. Nechansky, J. Neundorf, T. Novak, M.L. Ojeda, P. Pani, K. Peters, T. Poulsen, G. Pownall, A. Renardi, S. Richter, M. Rimoldi, P. Rivadeneira, M. Robin, L. Rossini, C.O. Sander, S. Schmitt, C. Seitz, S. Sharma, S. Sinha, D. South, M.M. Stanitzki, B. Stapf, L.R. Strom, N.A. Styles, K. Tackmann, E.A. Thompson, J. Von Ahnen, A.F. Wongel, S.D. Worm, Y.C. Yap

Dortmund
T. Dado, S. Dungs, J. Erdmann, E.M. Freundlich, K. Kroeninger, T. Kupfer, O. Nackenhorst, K. Sedlaczek, J. Weingarten, B. Wendland, S.V. Zeißner

Dresden
T. Herrmann, D. Kirchmeier, W.F. Mader, M. Maerker, J. Manjarres Ramos, A.Nag Nag, O. Novgorodova, F. Siegert, S. Todt, H. Torres, C. Wiel

Freiburg
S. Argyropoulos, F. Becherer, M. Boehler, J.A. Fernandez Pretel, S. Gargiulo, C. Heidegger, K.K. Heidegger, G. Herten, D. Hohn, J.C. Honig, K. Jakobs, P. Jenni, E. Karentzos, A. Knue, K. Köneke, O. Kuprash, U. Landgraf, V.S. Lang, A. Lösle, V. Magerl, T. Moskalets, U. Parzefall, V. Plesanovs, D. Rafanoharana, A. Rodriguez Rodriguez, B. Roland, B. Rottler, F. Rühr, Z. Rurikova, D. Sammel, F. Sauerburger, K.E. Schleicher, P.G. Scholer, M. Schumacher, S. Solomon, K. Solovieva, D. Sperlich, C. Weiser, L.A.M. Wiik-Fuchs, B.T. Winter, C.J.S. Young, D. Zanzi, S. Zimmermann

G"ottingen
K. Abeling, B. Achkar, J.F. Beirer, M. Bindi, L.O. Gerlach, R. Goncalves Gama, J. Grosse-Knetter, Y. Hong, G. Kawamura, A. Kirchhoff, S. Korn, S. Lai, J.C. Lange, A.H. Melo, S. Mobius, M. Niemeyer, I. Pokharel, A. Quadt, E. Shabalina, S. Sindhu, A. Skaf, Y. Tian, J. Veatch

Giessen
D. Caforio, M. Düren, H. Stenzel

Heidelberg PI
M.M. Czurylo, S.J. Dittmeier, A. Krishnan, C. Sauer, A. Schoening, L. Vigani, J. Zinsser

Heidelberg KIP
L.M. Baltes, F. Bartels, F.L. Castillo, F. Del Rio, M. Dunford, S. Franchino, M. Klassen, T. Mkrtchyan, F. Napolitano, P.S. Ott, D.F. Rassloff, S. Rodriguez Bosca, H-C. Schultz-Coulon, R. Stamen, P. Starovoitov, S.M. Weber, M. Wessels, X. Yue

Mainz
L. Adam, M. Akbiyik, J. Balz, A. Basan, J. Blumenthal, B. Brickwedde, V. Büscher, J. Damp, F. Fiedler, M. Geisen, S. Groh, R. Gugel, A. Hadef, H. Herr, C. Kahra, P. Kramer, J.A. Kremer, A. Laudrain, L. Masetti, F. Neuhaus, S. Rave, A. Reiss, M. Robles Manzano, U. Schäfer, B. Schlag, K. Schmieden, C. Schmitt, M. Schott, D. Ta, S. Tapprogge, E. Tzovara, O.V. Vujinovic, R.-J. Wang, M. Weirich, N. Wieseotte, A. Wolf

Munich LMU
Z.P. Arrubarrena Tame, O. Biebel, G. Duckeck, N.M. Hartmann, M.G. Herrmann, R. Hertenberger, C.S. Jagfeld, F.F. Klitzner, F. Krieter, J. Lorenz, A.M. Lory, A. Mann, A. Matic, S. Mehlhase, P. Mogg, M.P. Rinnagel, D. Schaile, E. Schanet, C. Valderanis, R. Walker

Munich MPI
T. Barillari, S. Bethke, D. Britzger, D. Cieri, D. Duda, F. Guescini, S. Hadzic, M. Holzbock, A. Hönle, J. Jimenez Pena, A.E. Kiryunin, F. Klimpel, S. Kluth, O. Kortner, S. Kortner, H. Kroha, S.R. Maschek, T.G. McCarthy, S. Menke, R. Nisius, M.B. Rendel, R. Richter, M. Sahinsoy, P. Schacht, M. Spalla, S. Stonjek, A. Verbytskyi, E. Voevodina, V.M. Walbrecht

Siegen
B. Batool, P. Buchholz, M. Cristinziani, C. Diez Pardos, F. Fischer, I. Fleck, I. Ibragimov, V.V. Kostyukhin, J.K.R. Meshreki, B. Mondal, A. Petrukhin, A. Rej, W. Walkowiak, M. Ziolkowski

W"urzburg
D.S. Bhattacharya, M. Haleem, G. Khoriauli, L. Klein, R. Ströhmer, T. Swirski, T. Trefzger

Wuppertal
V.A. Austrup, O. Bessidskaia Bylund, G. Brandt, C. Dülsen, F. Ellinghaus, M. Errenst, T. Flick, T. Harenberg, D. Hirschbuehl, G. Jäkel, S. Kersten, M. Rezaei Estabragh, J. Roggel, A. Sahu, M. Sandhoff, F. Schroeder, O. Thielmann, M. Vogel, W. Wagner, S. Wahdan, M. Wensing, C. Zeitnitz

Greece
Athens NKUA
S. Angelidakis, D. Fassouliotis, L. Fountas, I. Gkialas, C. Kourkoumelis

Athens NTU
T. Alexopoulos, C. Bakalis, E.N. Gazis, P. Gkountoumis, C. Kitsaki, S. Maltezos, C. Paraskevopoulos, M. Perganti, P. Tzanis, G. Zacharis

Demokritos
G. Fanourakis, T. Geralis, M.M. Prapa, G. Stavropoulos, O. Zormpa

Thessaloniki
K. Bachas, I. Karkanias, E. Kasimi, K. Kordas, C. Lampoudis, A. Leisos, I.M. Maniatis, A. Marantis, I. Maznas, C. Petridou, D. Sampsonidis, A. Tsiamis, A. Tsirigotis, M. Tsopoulou, S. Tzamarias, A. Vgenopoulos

Israel
Technion Haifa
H. Abreu, G. Avner, E. Kajomovitz, D. Melini, Y. Rozen, S. Tarem, O. Tsur

Tel Aviv
H. Abramowicz, L. Barak, U. Barron, G. Bella, Y. Benhammou, B. Chen, H. Cohen, E. Etzion, G. Koren, M.S. Lutz, R. Nayak, L. Pascual Dominguez, D. Reikher, A. Soffer, N.M. Tamir, D. Vannicola

Weizmann Rehovot
I. Aizenberg, M. Birman, R. Brener, S. Bressler, Z.H. Citron, E. Dreyer, E. Duchovni, E. Gross, N. Hod, A. Ivina, T. Jakoubek, N. Kakati, L.J. Levinson, D.C. Mankad, G. Mikenberg, A. Milov, B. Parida, I. Ravinovich, A. Santra, D. Shaked Renous, J. Shlomi, E. Shulga, V. Smakhtin

Italy
Bologna
G.L. Alberghi, F. Alfonsi, L. Bellagamba, S. Biondi, D. Boscherini, A. Bruni, G. Bruni, M. Bruschi, G. Cabras, G. Carratta, N. Cavalli, A. Cervelli, L. Clissa, S. De Castro, L. Fabbri, M. Franchini, A. Gabrielli, B. Giacobbe, F. Lasagni Manghi, L. Massa, M. Negrini, A. Polini, L. Rinaldi, M. Romano, C. Sbarra, A. Sbrizzi, N. Semprini-Cesari, A. Sidoti, M. Sioli, K. Todome, S. Valentinetti, M. Villa, C. Vittori, A. Zoccoli

Cosenza
M. Capua, G. Carducci, G. Crosetti, F. Curcio, I. Gnesi, D. Malito, A. Mastroberardino, E. Meoni, D. Salvatore, M. Schioppa, E. Tassi

Frascati
P. Albicocco, M. Antonelli, C. Arcangeletti, M. Beretta, V. Chiarella, G. Maccarrone, G. Mancini, A. Sansoni, M. Testa, E. Vilucchi

Genova
D. Barberis, A. Coccaro, G. Darbo, G. Gagliardi, C. Gemme, A. Lapertosa, P. Morettini, F. Parodi, S. Passaggio, M. Ressegotti, L.P. Rossi, C. Schiavi, F. Sforza, L. Vannoli

Trieste ICTP
B.S. Acharya, L. Serkin

Udine
M. Cobal, M.P. Giordani, G. Giugliarelli, J. Magro, G. Panizzo, M. Pinamonti

Lecce
M.S. Centonze, G. Chiodini, E. Gorini, F.G. Gravili, M. Primavera, E.J. Schioppa,
S. Spagnolo, A. Ventura

Milano
G. Alimonti, A. Andreazza, E. Ballabene, A. Carbone, L. Carminati, M. Citterio, S. Coelli,
S. D'Auria, L. Dell'Asta, M. Fanti, D. Giugni, T. Lari, M. Lazzaroni, C. Meroni, D.P. Mungo,
A. Murrone, L. Perini, F. Piazza, F. Ragusa, S. Resconi, A. Stabile, G.F. Tartarelli,
C. Troncon, R. Turra

Napoli
A. Aloisio, M.G. Alviggi, S. Auricchio, V. Canale, G. Carlino, F. Cirotto, F. Conventi,
R. De Asmundis, M. Della Pietra, C. Di Donato, A. Doria, V. Izzo, M. Lavorgna,
P. Massarotti, L. Merola, E. Rossi, G. Sekhniaidze

Pavia
J. Agarwala, E.M. Farina, R. Ferrari, G. Gaudio, G. Introzzi,
A. Kourkoumeli-Charalampidi, A. Lanza, M. Livan, A. Negri, G. Polesello, D.M. Rebuzzi,
A. Rimoldi, G. Rovelli, S. Sottocornola

Pisa
A. Annovi, N.V. Biesuz, M. Calvetti, V. Cavasinni, G. Chiarelli, G. Di Gregorio,
P. Francavilla, P. Giannetti, S. Leone, P. Mastrandrea, C. Roda, F. Scuri, M. Verducci

Roma I
F. Anulli, P. Bagnaia, M. Bauce, C. Bini, N. Bruscino, M. Carnesale, A.R. Chomont,
M. Corradi, A. De Salvo, F.A. Di Bello, A. Di Domenico, S. Falciano, G. Frattari, P. Gauzzi,
S. Gentile, S. Giagu, V. Ippolito, M. Kado, F. Lacava, I. Longarini, C. Luci, L. Martinelli,
A. Nisati, E. Pasqualucci, S. Rosati, L. Sabetta, F. Safai Tehrani, R. Vari, S. Veneziano

Roma II
G. Aielli, E. Alunno Camelia, P. Camarri, R. Cardarelli, L. Cerrito, U. De Sanctis,
M. De Santis, A. Di Ciaccio, M. Faucci Giannelli, F. Giuli, B. Liberti, S. Loffredo,
L. Marcoccia, L. Pizzimento, G. Proto, A. Rocchi, M. Vanadia

Roma III
I. Bashta, M. Biglietti, M.T. Camerlingo, V. D'Amico, B. Di Micco, R. Di Nardo, A. Farilla,
M. Iodice, D. Orestano, G. Salamanna, M. Sessa

Trento

M. Cristoforetti, A. Di Luca, F.M. Follega, R. Iuppa, E. Ricci, D.M.S. Sultan

Japan
KEK
M. Aoki, K. Hanagaki, S. Higashino, Y. Ikegami, A. Mizukami, K. Nagano, Y. Nakahama, K. Nakamura, S. Odaka, O. Sasaki, Y. Takubo, M. Togawa, K. Tokushuku, M. Tomoto, S. Tsuno

Kobe
H. Hibi, H. Kurashige, J. Maeda, A. Ochi, K. Takeda, Y. Yamazaki

Kyoto
S. Akatsuka, R. Kobayashi, Y. Mino, Y. Noguchi, Y. Okazaki, T. Sumida, Y. Tsujikawa

Kyoto UE
R. Takashima

Kyushu
K. Kawagoe, D. Kobayashi, Y. Miyazaki, S. Oda, H. Otono, J. Tojo, N. Yamaguchi

Nagoya
H. Asada, S. Hayashida, Y. Horii, Y. Kano, T. Kawaguchi, M. Wakida

Ochanomizu
K. Asai, M. Fujimoto, T. Kono

Osaka
M. Hirose, H. Nanjo, L. Wickremasinghe

Shinshu
Y. Hasegawa, K. Kawade

Tokyo ICEPP
S. Asai, Y. Enari, S. Ganguly, R. Iguchi, M. Ishino, T. Kishimoto, T. Kodama, T. Mashimo, T. Masubuchi, N. Matsuzawa, Y. Minegishi, M. Morinaga, T. Nobe, R. Oishi, Y. Okumura, M. Saito, T. Saito, R. Sawada, K. Sugizaki, A. Tanaka, J. Tanaka, G. Tateno, K. Terashi, K. Uchida, K. Uno, Y. Yang

Tokyo Tech
Y. He, O. Jinnouchi, E. Kim, A. Kubota, M. Kuze, H. Oide, R. Ushioda, Y. Yamaguchi

Tsukuba
K. Hara, S. Hirose, K. Sato, F. Ukegawa, S. Wada, R. Wakasa, H. Yamauchi

Waseda
H. Ito, T. Kaji, T. Mitani, T. Nitta, Y. Shimogama, K. Yorita

Morocco
Kenitra
Y. El Ghazali, M. Gouighri

Oujda
J. Assahsah, M. Ouchrif

Rabat
A. Aboulhorma, S. Batlamous, R. Cherkaoui El Moursli, H. El Jarrari, F. Fassi, H. Hamdaoui, B. Ngair, Z. Soumaimi, Y. Tayalati, M. Zaazoua

Casablanca
D. Benchekroun, K. Bouaouda, Z. Chadi, A. El Moussaouy, S. Ezzarqtouni, H. Imam, S. Zerradi

UM6P

Marrakesh

Netherlands
Nijmegen
A. Aggarwal, S. Caron, N. De Groot, F. Filthaut, C.A. Gottardo, P. Moskvitina, C. Nellist, A. Perrevoort

Nikhef
A. Alfonsi, P.J. Bakker, R. Balasubramanian, S. Bentvelsen, G.J. Bobbink, W.S. Chan, Y.S. Chow, M. De Beurs, P. de Jong, J. Degens, F.A. Dias, T.A. du Pree, P. Ferrari, G. Gilles, P. Kluit, B. Kortman, A.E. McDougall, B. Moser, F. Pasquali, A. Pizzini, H.L. Snoek, M. Stamenkovic, J.J. Teoh, I. Van Vulpen, M.J. Veen, W. Verkerke, J.C. Vermeulen, M. Vreeswijk

Norway
Bergen
T. Buanes, J.I. Djuvsland, G. Eigen, N. Fomin, S.K. Huiberts, G.R. Lee, A. Lipniacka, B. Martin dit Latour, B. Stugu, A. Traeet

Oslo
V. Ananiev, M.K. Bugge, D. Cameron, J.R. Catmore, E. Gramstad, E.S. Haaland, A.L. Heggelund, V. Morisbak, F. Ould-Saada, A.L. Read, O. Røhne, E.B. Rye, H. Sandaker, K.O.H. Vadla

Poland
Krakow Jagiellonian
D.T. Gil, M. Palka, E. Richter-Was

Krakow AGH-UST
L. Adamczyk, T. Bold, W. Dabrowski, M. Dyndal, I. Grabowska-Bold, K.W. Janas, S. Koperny, I.K. Lakomiec, K. Maj, B. Mindur, A. Ogrodnik, M. Przybycien, K.S. Swientek, A.W. Wysocki

Krakow IFJ PAN
E. Banas, P.A. Bruckman de Renstrom, J.J. Chwastowski, K.M. Ciesla, D. Derendarz, B.S. Dziedzic, P.A. Erland, A. Kaczmarska, K. Korcyl, J.A. Krzysiak, Pa. Malecki, A. Olszewski, J. Olszowska, M. Slawinska, A. Smykiewicz, E. Stanecka, R. Staszewski, M. Trzebinski, A. Trzupek, M.W. Wolter, B.K. Wosiek, K.W. Woźniak

Portugal
Portugal 8- IST

Portugal 4-CFNU Lisboa

Portugal 5-Minho

Portugal 6-Granada
J.A. Aguilar-Saavedra, M. Chala

Portugal 2-Lisboa

Portugal 3-Coimbra
J.C. Macario Ribeiro

Portugal 1-LIP
S.P. Amor Dos Santos, N.F. Castro, R.F. Coelho Barrue, P. Conde Muiño, T. Dias Do Vale, G. Evans, M.C.N. Fiolhais, A. Gomes, R. Gonçalo, A. Maio, E.D. Mendes Gouveia, A.L. Moreira De Carvalho, I. Ochoa, L.F. Oleiro Seabra, A. Onofre, R. Pedro, A.P. Pereira Peixoto, H. Santos, J.G. Saraiva, R.P. Serrano Fernandez, M. Stolarski, F. Veloso, H. Wolters

Romania
Iasi UAIC
C. Agheorghiesei, P. Postolache

Brasov TU
S. Popa

ITIM, Cluj Napoca

G.A. Popeneciu

Politehnica Bucharest
R. Hobincu

Timisoara WU
P.M. Gravila

Bucharest IFIN-HH
C. Alexa, A. Chitan, I.A. Cioară, D.A. Ciubotaru, I-M. Dinu, A.E. Dumitriu, A.A. Geanta, A. Jinaru, V.S. Martoiu, J. Maurer, A. Olariu, D. Pietreanu, M. Renda, M. Rotaru, G. Stoicea, G. Tarna, I.S. Trandafir, A. Tudorache, V. Tudorache, M.E. Vasile

Russia
JINR Dubna
F. Ahmadov, I.N. Aleksandrov, V.A. Bednyakov, I.R. Boyko, I.A. Budagov, G.A. Chelkov, A. Cheplakov, E. Cherepanova, M.V. Chizhov, D.V. Dedovich, M. Demichev, A. Gongadze, M.I. Gostkin, S.N. Karpov, Z.M. Karpova, E. Khramov, U. Kruchonak, V. Kukhtin, E. Ladygin, V. Lyubushkin, T. Lyubushkina, S. Malyukov, M. Mineev, E. Plotnikova, I.N. Potrap, F. Prokoshin, N.A. Rusakovich, R. Sadykov, A. Sapronov, M. Shiyakova, A. Soloshenko, S. Turchikhin, I. Yeletskikh, A. Zhemchugov, N.I. Zimine

Moscow FIAN
A.V. Akimov, F. Dubinin, I.L. Gavrilenko, O. Meshkov, P.Y. Nechaeva, A.A. Snesarev, V.V. Sulin, L. Sultanaliyeva, V.O. Tikhomirov, K. Zhukov

Moscow MEPhI
K. Belotskiy, N.L. Belyaev, O. Bulekov, A. Kurova, D. Ponomarenko, N. Proklova, D. Pyatiizbyantseva, A. Romaniouk, S.Yu. Smirnov, Y. Smirnov, E.Yu. Soldatov, S. Timoshenko, K. Vorobev

Moscow SU
L.K. Gladilin, N. Korotkova, V.A. Kramarenko, V. Sinetckii, S.Yu. Sivoklokov, L.N. Smirnova

Novosibirsk BINP and NSU
A.V. Anisenkov, E.M. Baldin, K. Beloborodov, V.S. Bobrovnikov, A.G. Bogdanchikov, A.R. Buzykaev, V.F. Kazanin, A.G. Kharlamov, T. Kharlamova, A.L. Maslennikov, D.A. Maximov, S.V. Peleganchuk, O.L. Rezanova, A.M. Soukharev, A.A. Talyshev, Yu.A. Tikhonov, V. Zhulanov

NRC KI - IHEP, Protvino

S.P. Denisov, R.M. Fakhrutdinov, A.B. Fenyuk, D. Golubkov, A. Kamenshchikov, A.N. Karyukhin, A.S. Kozhin, A.A. Minaenko, A.G. Myagkov, V. Nikolaenko, A. Ryzhov, A.A. Solodkov, O.V. Solovyanov, E.A. Starchenko, A.M. Zaitsev, O. Zenin

NRC Kurchatov Institute - ITEP
A. Gavrilyuk, P.A. Gorbounov, P.B. Shatalov, I.I. Tsukerman

Petersburg NPI
S. Barsov, A. Ezhilov, O.L. Fedin, G. Fedotov, V. Gratchev, M. Levchenko, V.P. Maleev, D. Pudzha, V.A. Schegelsky, V. Solovyev

Tomsk SU
M. Aliev, O. Kuchinskaia, M. Penzin

Serbia
Belgrade IP
E. Bakos, J. Jovicevic, V. Maksimovic, Dj. Sijacki, N. Vranjes, M. Vranjes Milosavljevic, L. Živković

Slovak Republic
Kosice
D. Bruncko, P. Murin, F. Sopkova, P. Strizenec, J. Urban

Bratislava
R. Astalos, D. Babal, P. Bartos, M. Dubovsky, B. Eckerova, S. Hyrych, O. Majersky, I. Sykora, S. Tokár, T. Ženiš

Slovenia
Ljubljana
V. Cindro, A. Filipčič, A. Gorišek, B. Hiti, B.P. Kerševan, G. Kramberger, B. Leban, I. Mandić, B. Maček, M. Mikuž, J. Novak

South Africa
NIP-UPD
M. Flores

Cape Town
R.J. Atkin, K.N. Barends, J.M. Keaveney, S. Yacoob

Witwatersrand

T. Chowdhury, L.D. Christopher, S. Dahbi, H. Jivan, D. Kar, M. Kumar, J.E. Mdhluli, B.R. Mellado Garcia, G. Mokgatitswane, X. Ruan, E.M. Shrif, E. Sideras Haddad, S. Sinha, S.H. Tlou

Johannesburg
D. Casadei, S.H. Connell, N. Govender, L.L. Leeuw, L. Truong

iThemba

UNISA

Spain
Barcelona
M.N. Agaras, D. Bogavac, M. Bosman, M.P. Casado, L. Castillo Garcia, F.A. Förster, S. Gonzalez Fernandez, C. Grieco, S. Grinstein, A. Juste Rozas, S. Kazakos, I. Korolkov, M. Martinez, L.M. Mir, C. Moreno Martinez, J.L. Munoz Martinez, N. Orlando, A. Pacheco Pages, C. Padilla Aranda, I. Riu, A. Salvador Salas, A. Sonay, S. Terzo, D. Vazquez Furelos

Madrid UA
M. Alvarez Estevez, F. Barreiro, D. Camarero Munoz, J. Cantero, J. Del Peso, C. Glasman, M.A. Principe Martin, J. Terron

Portugal 6-Granada
J.A. Aguilar-Saavedra, M. Chala

Valencia
K.R. Amos, J.A. Aparisi Pozo, A.J. Bailey, S. Cabrera Urbán, F. Cardillo, V. Castillo Gimenez, M.J. Costa, M. Didenko, C. Escobar, O. Estrada Pastor, L. Fiorini, E. Fullana Torregrosa, J. Fuster, C. García, J.E. García Navarro, S. González de la Hoz, G.R. Gonzalvo Rodriguez, J.G.R. Guerrero Rojas, E. Higón-Rodriguez, C. Lacasta, J.J. Lozano Bahilo, J. Mamuzic, S. Marti-Garcia, P. Martinez Agullo, M. Miralles Lopez, V.A. Mitsou, L. Monsonis Romero, M. Moreno Llácer, J. Navarro-Gonzalez, J. Poveda, A. Prades Ibanez, A. Ruiz-Martinez, P. Sabatini, J. Salt, J. Sánchez, V. Sanchez Sebastian, I. Sayago Galvan, U. Soldevila, E. Torró Pastor, A. Valero, J.A. Valls Ferrer, M. Villaplana Perez, M. Vos

Sweden
Stockholm
S.Y. Andrean, F. Backman, L. Barranco Navarro, C. Bohm, C. Clement, K. Dunne, S. Hellman, T. Ingebretsen Carlson, X. Lou, D.A. Milstead, T. Moa, P. Pasuwan, L. Pereira Sanchez, N.W. Shaikh, S.B. Silverstein, J. Sjölin, S. Strandberg, A. Strubig, E. Valdes Santurio

Lund
T.P.A. Åkesson, E.E. Corrigan, C. Doglioni, J. Geisen, E. Hansen, V. Hedberg, G. Jarlskog, B. Konya, E. Lytken, K.H. Mankinen, C. Marcon, J.U. Mjörnmark, G.A. Mullier, R. Poettgen, N.D. Simpson, E. Skorda, O. Smirnova

Stockholm KTH
A. Leopold, O. Lundberg, B. Lund-Jensen, C.C. Ohm, G. Ripellino, R. Shaheen, D.R. Shope, J. Strandberg

Uppsala
E.M. Asimakopoulou, E. Bergeaas Kuutmann, R. Brenner, T. Ekelof, V. Ellajosyula, M. Ellert, A. Ferrari, R. Gonzalez Suarez, T. Mathisen, S. Oerdek, O. Sunneborn Gudnadottir

Switzerland
Bern
J.K. Anders, H.P. Beck, M. Chatterjee, A. Ereditato, L. Franconi, L. Halser, S. Haug, A. Ilg, R. Mueller, A.P. O'neill, M.S. Weber

CERN
A. Abed Abud, Y. Afik, A. Ahmad, M. Aleksa, C. Allaire, C. Amelung, N. Aranzabal, A.J. Armbruster, H. Arnold, G. Avolio, M-S. Barisits, T.A. Beermann, C. Beirao Da Cruz E Silva, L. Beresford, T. Bisanz, J.D. Bossio Sola, J. Boyd, L. Brenner, N. Calace, S. Camarda, T. Carli, A. Catinaccio, A.E.C. Coimbra, L.D. Corpe, A. Cueto, P. Czodrowski, V. Dao, A. Dell'Acqua, P.O. Deviveiros, A. Di Girolamo, F. Dittus, O.A. Ducu, A. Dudarev, M. Dührssen, N. Ellis, M. Elsing, S. Falke, D. Francis, D. Froidevaux, A. Gabrielli, P. Gessinger-Befurt, L. Goossens, B. Gorini, S. Guindon, G. Gustavino, R.J. Hawkings, L. Heinrich, C. Helsens, A.M. Henriques Correia, L. Hervas, A. Hoecker, M. Huhtinen, P. Iengo, T. Javůrek, J.J. Junggeburth, M. Kiehn, P. Klimek, T. Klioutchnikova, N.M. Köhler, A. Koulouris, A. Krasznahorkay, E.S. Kuwertz, M. Lassnig, M. LeBlanc, G. Lehmann Miotto, L. Longo, I. Lopez Paz, S. Manzoni, A. Marzin, S. Meehan, M. Mentink, J. Montejo Berlingen, A.K. Morley, L. Morvaj, P. Moschovakos, A.M. Nairz, M. Nessi, S. Palestini, T. Pauly, H. Pernegger, S. Perrella, B.A. Petersen, N.E. Pettersson, L. Pezzotti, L. Pontecorvo, M.E. Pozo Astigarraga, M. Queitsch-Maitland, M. Raymond, C. Rembser, P. Riedler, S. Roe, A. Rummler, D. Salamani, A. Salzburger, S. Schlenker, J. Schovancova, A. Sharma, O. Sidiropoulou, M.V. Silva Oliveira, R. Simoniello, C.A. Solans Sanchez, G. Spigo, G.A. Stewart, M.C. Stockton, A.N. Tuna, G. Unal, T. Vafeiadis, W. Vandelli, T. Vazquez Schroeder, B. Vormwald, R. Vuillermet, T. Wengler, H.G. Wilkens, S. Zambito, L. Zwalinski

Geneva
S. Adorni, C.S. Amrouche, C. Antel, A. Clark, D. Della Volpe, L.F. Ehrke, D. Ferrere, T. Golling, S. Gonzalez-Sevilla, M. Guth, G. Iacobucci, T. Iizawa, P. Mermod,

H.R. Nindhito, C.E. Pandini, L. Paolozzi, R. Poggi, J.A. Raine, C. Rizzi,
J.A. Sabater Iglesias, S. Schramm, A. Sfyrla, S. Shirabe, X. Wu, K. Zoch

Taiwan
Hsinchu NTHU
K. Cheung, P.J. Hsu, Y.J. Lu

Taipei AS
S. Ali, S. Hou, S.C. Lee, R. Mazini, T. Varol, S.M. Wang

Turkey
TOBB
S. Sultansoy

Ankara
O. Cakir, H. Duran Yildiz, S. Kuday, I. Turk Cakir

Istanbul Aydin

Gaziantep
A. Bingul, Z. Uysal

Bogazici Istanbul
A. Adiguzel, A. Bayirli, A.J. Beddall, E. Celebi, S.A. Cetin, S. Istin, K.Y. Oyulmaz,
V.E. Ozcan, S. Ozturk, S. Simsek

Bahcesehir

Bilgi

United Arab Emirates
UoS
R. Soualah

UAEU

NYU Abu Dhabi

United Kingdom
Birmingham
P.P. Allport, P. Bellos, G.A. Bird, J. Bracinik, D.G. Charlton, A.S. Chisholm, H.G. Cooke,
T. Fitschen, P.M. Freeman, W.F. George, L. Gonella, F. Gonnella, N.A. Gorasia,
C.M. Hawkes, S.J. Hillier, J.J. Kempster, J. Kendrick, D.J. Lewis,
A.M. Mendes Jacques Da Costa, T.J. Neep, P.R. Newman, K. Nikolopoulos, J.M. Silva,

A. Stampekis, J.P. Thomas, P.D. Thompson, G.S. Virdee, R.J. Ward, A.T. Watson, M.F. Watson

Cambridge
W.K. Balunas, J.R. Batley, O. Brandt, J.T.P. Burr, J.D. Chapman, J.W. Cowley, W.J. Fawcett, L. Henkelmann, B.H. Hodkinson, L.B.A.H. Hommels, D.M. Jones, C.G. Lester, C. Malone, D.L. Noel, H.A. Pacey, M.A. Parker, C.J. Potter, D. Robinson, D. Rousso, R. Tombs, S. Williams

Edinburgh
S. Alderweireldt, T.M. Carter, P.J. Clark, S.M. Farrington, Y. Gao, G.N. Hamity, A. Hasib, M.P. Heath, S. Jiggins, C. Leonidopoulos, V.J. Martin, L. Mijović, S. Palazzo, V.A. Parrish, E.P. Takeva, N. Themistokleous, E.M. Villhauer, A. Vishwakarma, B.M. Wynne, E. Zaid

Glasgow
R.L. Bates, A. Blue, A.G. Borbély, D. Britton, A.G. Buckley, P.J. Bussey, C.M. Buttar, G. Callea, I.A. Connelly, W.R. Cunningham, A.T. Doyle, F. Fabbri, J. Howarth, J. Jamieson, L. Mince, V. O'Shea, M. Owen, B. Ravina, A.G. Rennie, A. Robson, D.P. Spiteri, N. Warrack, K. Wraight

Lancaster
A.E. Barton, I.A. Bertram, G. Borissov, E.V. Bouhova-Thacker, H. Fox, R.C.W. Henderson, R.W.L. Jones, V. Kartvelishvili, P.A. Love, L. Meng, D. Muenstermann, K. Rybacki, I.G. Sanderswood, M. Smizanska, S. Spinali, A.M. Wharton, M.R. Yexley

Liverpool
E. Arena, A. Biondini, S. Burdin, W.Y. Chan, P. Dervan, M. D'Onofrio, R. Gonzalez Lopez, C.B. Gwilliam, H.S. Hayward, A.E. Jaspan, T.J. Jones, M. Klein, U. Klein, J. Kretzschmar, Z. Li, H. Lyons, A. Mehta, M.W. O'Keefe, J. Patton, N. Rompotis, A.J. Ruby, C.D. Sebastiani, M.J. Sullivan, H. Teagle, J.H. Vossebeld

London QMUL
A.J. Bevan, U. Blumenschein, M. Bona, T.P. Charman, I. Dawson, A.A. Elliot, A.C. Freegard, R. Gamboa Goni, J.M. Hays, N.D. Hehir, M.P.J. Landon, S.L. Lloyd, J.P. Mandalia, P.S. Miyagawa, T. Qiu, E. Rizvi, E.J. Thorpe, L. Vozdecky, S. Zenz

London RHBNC
D.K. Abhayasinghe, T. Berry, V. Boisvert, C.D. Booth, A.J. Bozson, G. Cowan, S. George, S.M. Gibson, B. Green, S.D. Lawlor, H.P. Lefebvre, J.G. Panduro Vazquez, F. Pastore, J.D. Shinner, F. Spanò, P. Teixeira-Dias, D.W. Thomas, P.J. Verschuuren, S.D. Walker, L.J. Wilkins

London UC

G. Barbour, J.M. Butterworth, M. Campanelli, C. Gutschow, D.P. Huang, A.S. Kelly, N. Kimura, V. Konstantinides, N. Konstantinidis, A. Korn, N.N. Lad, A.C.A. Lee, Y. Ma, A.C. Martyniuk, D.J. Mclaughlin, A.P. Morris, R.R.B. Norisam, E. Nurse, C. Pitman Donaldson, S. Rettie, T. Scanlon, P. Sherwood, L. Shi, A.L. Sopio, S. Van Stroud, B.M. Waugh

Manchester
C.R. Basson, C.J. Birch-sykes, A. Bitadze, S.E. Clawson, C. Da Via, M. D'uffizi, S. Dysch, A.C. Forti, R. Hankache, T.G. Hitchings, B.P. Honan, Z. Lawrence, B. Le, J. Masik, F.J. Munoz Sanchez, M. Murin, A. Oh, J.R. Pater, R.F.Y. Peters, A.D. Pilkington, D. Price, Y. Qin, N. Scharmberg, S.M. Shaw, V. Vecchio, M. Vozak, T.R. Wyatt

Oxford
L. Ambroz, G. Artoni, A.J. Barr, D. Bortoletto, F. Celli, E.I. Conroy, A.M. Cooper-Sarkar, M.G. Foti, J.A. Frost, G.E. Gallardo, E.J. Gallas, J.C. Grundy, C. Gwenlan, Y.T. Harris, C.P. Hays, T.B. Huffman, K. Karava, Z. Li, C. Merlassino, M. Mironova, K. Nagai, R.B. Nickerson, C.S. Pollard, K. Potamianos, E. Schopf, I.P.J. Shipsey, H.A. Smith, M. Stankaityte, I. Veliscek, G.H.A. Viehhauser, Y. Wei, A.R. Weidberg, P.J. Windischhofer, R. Wölker, J. Wuerzinger, S. Yan

RAL
T. Adye, J.T. Baines, W. Buttinger, J. Dopke, D. Emeliyanov, B.J. Gallop, J. Kirk, S. Martin-Haugh, S.J. McMahon, R.P. Middleton, R.E. Owen, P.W. Phillips, D.P.C. Sankey, C. Sawyer, B.H. Smart, J. Walder, M. Wielers

Sheffield
C. Anastopoulos, M.T. Anthony, D. Costanzo, A. Fell, J.J. Hall, M.C. Hodgkinson, P. Johansson, K. Lohwasser, C.M. Macdonald, J.C. MacDonald, H. Mildner, M.B. Norfolk, T.D. Powell, N.P. Readioff, E.M. Rüttinger, K.A. Saoucha, P. Sommer, D.R. Tovey, T. Vickey, O.E. Vickey Boeriu, D.F. Zhang, T.G. Zorbas

Sussex
B.M.M. Allbrooke, M.A. Aparo, L. Asquith, A. Cerri, A. De Santo, M.O. Evans, M. Grandi, D. Kelsey, D.M. Koeck, O. Kovanda, J.A. Mcfayden, B. Safarzadeh Samani, F. Salvatore, K. Shaw, M. Spina, T.J. Stevenson, M.R. Sutton, F. Trovato, I. Vivarelli, I. Xiotidis

Warwick
K. Becker, G. Facini, P.F. Harrison, E. Jones, T.A. Martin, A. Mitra, S. Morgenstern, W.J. Murray, B.A. Roberts, M. Spangenberg, V. Vladimirov

United States of America
Argonne

V.S. Bhopatkar, S. Chekanov, S. Darmora, W.H. Hopkins, E. Kourlitis, T. LeCompte, J. Love, J. Metcalfe, A.S. Mete, A. Paramonov, J. Proudfoot, P. Van Gemmeren, R. Wang, T. Xu, J. Zhang

Arizona
S. Berlendis, E. Cheu, Z. Cui, C.M. Delitzsch, A. Ghosh, K.A. Johns, W. Lampl, R.E. Lindley, P. Loch, J.P. Rutherfoord, E.W. Varnes, H. Zhou, Y. Zhou

Arlington UT
D. Bakshi Gupta, B. Burghgrave, J.C.J. Cardenas, K. De, A. Farbin, H.K. Hadavand, J.D. Little, A.J. Myers, N. Ozturk, G. Usai, A. White

Austin
T. Andeen, C.D. Burton, K. Choi, N. Nikiforou, P.U.E. Onyisi, D.K. Panchal, A. Roy, M. Unal, A.F. Webb

Berkeley LBNL
D.J.A. Antrim, R.M. Barnett, J. Beringer, P. Calafiura, F. Cerutti, A. Ciocio, A. Dimitrievska, G.I. Dyckes, K. Einsweiler, I. Ene, L.G. Gagnon, M. Garcia-Sciveres, C. Gonzalez Renteria, H.M. Gray, C. Haber, S. Han, T. Heim, I. Hinchliffe, X. Ju, K. Krizka, C. Leggett, Z. Marshall, W.P. McCormack, M. Muškinja, B.P. Nachman, G.J. Ottino, S. Pagan Griso, V.R. Pascuzzi, M. Pettee, E. Pianori, E.D. Resseguie, E. Reynolds, B.R. Roberts, S.N. Santpur, M. Shapiro, B. Stanislaus, V. Tsulaia, C. Varni, H. Wang, T. Yamazaki, H.T. Yang, W-M. Yao

Boston
J.M. Butler, Z. Yan

Brandeis
S.V. Addepalli, J.R. Bensinger, L.J. Bergsten, P. Bhattarai, C. Blocker, F. Capocasa, J. Chen, D. Dodsworth, M. Goblirsch-Kolb, Z.M. Schillaci, G. Sciolla, D.T. Zenger Jr

Brookhaven BNL
S.H. Abidi, K. Assamagan, G. Barone, M. Begel, D.P. Benjamin, M. Benoit, R. Bi, D. Boye, E. Brost, V. Cavaliere, H. Chen, G. D'amen, J. Elmsheuser, V. Garonne, Y. Go, H.A. Gordon, G. Iakovidis, C.W. Kalderon, A. Klimentov, E. Lançon, F. Lanni, D. Lynn, H. Ma, T. Maeno, C. Mc Ginn, C. Mwewa, J.L. Nagle, P. Nilsson, M.A. Nomura, D. Oliveira Damazio, J. Ouellette, D.V. Perepelitsa, M.-A. Pleier, V. Polychronakos, S. Protopopescu, S. Rajagopalan, G. Redlinger, J. Roloff, B.D. Seidlitz, S. Snyder, P. Steinberg, S.A. Stucci, A. Tricoli, A. Undrus, C. Weber, T. Wenaus, S. Ye, E. Zhivun

Cal State
K. Grimm, J. Moss, A.J. Parker

Chicago
R.W. Gardner, M.D. Hank, L.A. Horyn, Y.K. Kim, J.K.K. Liu, D.W. Miller, M.J. Oreglia, D. Schaefer, M.J. Shochet, E.A. Smith, C. Tosciri, I. Vukotic

Columbia
K. Al Khoury, A. Angerami, G. Brooijmans, E.L. Busch, B. Cole, A. Emerman, B.J. Gilbert, J.L. Gonski, D.A. Hangal, Q. Hu, A. Kahn, K.E. Kennedy, D.J. Mahon, S. Mohapatra, J.A. Parsons, P.M. Tuts, D.M. Williams, P. Yin, W. Zou

Dallas SMU
A. Betti, A.M. Deiana, K.J.C. Leney, M.A. McKay, C.D. Milke, R. Narayan, S. Parajuli, S.J. Sekula, R. Stroynowski, J. Ye

Dallas UT
S.W. Ferguson, J.M. Izen, B. Meirose, K. Reeves

Duke
A.T.H. Arce, J.B. Beacham, D.R. Davis, M.G. Eggleston, A.T. Goshaw, A. Kotwal, M.C. Kruse, K. Pachal, S. Sen, P. Zhao

Harvard
N.A. Asbah, A. Badea, B.A. Bullard, A.W. Fortman, S. Francescato, M. Franklin, J. Huth, L. Lee, M. Morii, G. Rabanal Bolanos, A.M. Wang, R. Wang, A.S. White

Indiana
P. Calfayan, H. Evans, B.C. Forland, C.A. Johnson, R. Kopeliansky, S. Lammers, R.A. Linck, F. Luehring, C. Meyer, G. Myers, G. Palacino, C.P.A. Roland, D. Zieminska

Iowa
M. Cano Bret, D. Garg, M. Lu, U. Mallik, W.T. Wang

Iowa State
W.D. Heidorn, N. Krumnack, A. Lebedev, S. Tapia Araya

Louisiana Tech
L. Sawyer, M. Wobisch

Massachusetts
D.C. Abbott, B. Brau, Y. Chou, R. Coelho Lopes De Sa, C. Dallapiccola, M. Javurkova, S. Krishnamurthy, R.J. Langenberg, V.I. Martinez Outschoorn, Z.A. Meadows, E.J.W. Moyse, J.A. Sandesara, P. Tornambe, M. Vessella, C. Wagner, S. Willocq

Michigan

D. Amidei, P.A. Atmasiddha, A. Chen, T. Dai, E.B. Diehl, C. Ferretti, P. Fleischmann, C. Grud, L. Guan, Y. Guo, C. Hayes, M.H. Klein, D. Levin, Y.L. Liu, S.P. Mc Kee, G. Merz, K. Nelson, J. Qian, T.A. Schwarz, S. Sun, M. Tsai, Z. Wang, Z. Wang, X. Xiao, W. Xu, Z. Yang, M. Yuan, S. Zhang, B. Zhou, J. Zhu

Michigan SU
C.J. Buxo Vazquez, H. De la Torre, T. Farooque, W.C. Fisher, J.P. Gombas, R. Hauser, D. Hayden, J. Huston, K.S. Krowpman, R. Les, K. Lin, I. Pogrebnyak, R. Schwienhorst, A. Tarek Abouelfadl Mohamed

New Mexico
A. Grummer, E.A. Narayanan, R. Novotny, S.C. Seidel

Northern Illinois
J. Adelman, D. Chakraborty, L. D'Eramo, W.H. Kostecka, M. Mlynarikova, R.G. Oreamuno Madriz, E.W. Parrish, B.L. Stamas

NYU New York
Y. Abulaiti, K. Cranmer, A. Haas, A. Held, A.I. Mincer, P. Rieck, M. Ronzani, C.J. Treado, M.L. Vesterbacka

Ohio SU
B.M. Cote, K.K. Gan, H. Kagan, M. Montella, Z.B. Pollock, B. Reynolds, R. Rosten, S. Shrestha

Oklahoma
B. Abbott, M. Alhroob, A. Gekow, N.A. Grieser, P. Gutierrez, J.E. Lambert, M. Marjanovic, J.M. Muse, A.R. Roepe, H. Severini, P. Skubic, M. Strauss, D.J. Wilbern

Oklahoma SU
E. Antipov, A.M. Burger, J. Haley, A. Khanov, F. Rizatdinova, E.R. Vandewall

Oregon
J. Barkeloo, R. Bielski, J.E. Brau, A. Dattagupta, G.R. Gledhill, M. Glisic, G. Gonella, J.J. Heinrich, L. Jeanty, A. Kilgallon, N.A. Luongo, S. Majewski, I. Siral, A.L. Steinhebel, D.M. Strom, E. Torrence, K. Whalen, F. Winklmeier

Pennsylvania
S. Chen, J.R. Dandoy, L. Flores, L.F. Gutierrez Zagazeta, J.G. Heinlein, A. Kahn, J. Kroll, E. Lipeles, S. Lu, J.J. Mullin, B.J. Rosser, J.D. Shahinian, E. Thomson, H.H. Williams, R. Xu

Pittsburgh
M. Bandieramonte, R. Bi, R.M. Bianchi, J. Boudreau, H. Cai, B.T. Carlson, T.M. Hong, J. Mueller, A.A. Myers, C.W. Ng, H.J. Stelzer

Santa Cruz UC
A.A. Affolder, M. Battaglia, V. Fadeyev, C.M. Gee, M. Gignac, A.A. Grillo, M. Hance, C.M. Helling, N.J. Kang, S.M. Mazza, J. Nielsen, H.F-W. Sadrozinski, B.A. Schumm, A. Sciandra, A. Seiden, G.H. Stark, Y. Zhao

Seattle Washington
C. Alpigiani, Q. Buat, S.J. Gasiorowski, A.G. Goussiou, S.-C. Hsu, E.E. Khoda, A.K. Kvam, K. Li, H.J. Lubatti, M.L. Proffitt, J. Schaarschmidt, T.R. Van Daalen, G. Watts, H. Zhao

SLAC
T. Barklow, R. Bartoldus, C. Bernius, L. Bryngemark, V.M.M. Cairo, A.R. Cukierman, R.B. Garg, P. Grenier, N.M. Hartman, H. Herde, R. Hyneman, M. Kagan, M. Kocian, J. Pearkes, M. Safdari, A. Salnikov, A. Schwartzman, S. Sevova, D. Su, R. Teixeira De Lima, L. Tompkins, C. Vernieri, M. Wittgen, Z. Xu, C. Young, Z. Zheng

Stony Brook
C.P. Bee, A. Behera, S. Bhatta, Y. Hernández Jiménez, J. Hobbs, J. Jia, Y. Ke, I. Luise, G. Piacquadio, M. Rijssenbeek, R.D. Schamberger, F. Tsai, V. Tsiskaridze, D. Tsybychev

Tufts
P.H. Beauchemin, V. Croft, A.S. Drobac, F.I. Kaya, H. Son

UCI
A. Armstrong, T. Cuhadar Donszelmann, M.J. Fenton, A. Ghosh, A.J. Lankford, Y.W.Y. Ng, K. Ntekas, J.L. Oliver, M.J.R. Olsson, M. Schernau, P. Sundararajan, A. Taffard, G. Unel, D. Whiteson, J. Wollrath

Urbana UI
M. Atkinson, V.R. Bailey, Y. Cao, C. Fan, M. Feickert, B.H. Hooberman, T.M. Liss, J.D. Long, R. Longo, M.S. Neubauer, M.W. Phipps, T.T. Rinn, A.C. Romero Hernandez, A.M. Sickles, C.J. Smith, M. Tian, X. Wang, J.C. Zeng, M. Zhang, K. Zheng, D. Zhong

Wisconsin
D. Biswas, J. Chan, C.L. Cheng, W. Guan, W. Islam, A. Pathak, S. Sun, A.S. Tee, A.Z. Wang, W. Wiedenmann, S.L. Wu, R. Zhang, C. Zhou

Yale
O.K. Baker, E.G. Castiglia, S. Demers, J. Pan, G. Pezzullo, C.O. Shimmin, P. Tipton