

MEASUREMENTS OF ATLAS,
MEASUREMENTS WITH ATLAS

Construction and characterisation of ITk Pixel detector
structures, and a search for leptoquarks in events with
di-tau final states



Simon Florian Koch
Brasenose College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2025

Copyright © Simon Florian Koch, 2025.



This thesis is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this work are included in the work's Creative Commons licence, unless indicated otherwise.

The illustrations labelled Illus. 1 and Illus. 3 have been included with the permission of the artists (Samuel Koch and Melanie Schmid, respectively), whom have agreed to license these under a Creative Commons Attribution 4.0 International License. The illustration on page 5 (Ref. [1]) is licensed under the Creative Commons Attribution Non-Commercial 2.5 License and reproduced under these conditions, and retains the original license and copyright.

Figures 1.1b, 3.1 to 3.3, 3.4a, 3.4b, 4.1, 4.5, 4.6, 6.1, 8.1, 8.4a, 8.4b, 8.6a and 8.6b have been reproduced from open-access articles or equivalent sources and are included here following the applicable attribution requirements. These figures retain their original licensing and copyright.

Figures 7.9a, 7.9b, 7.15a, 7.22a and 7.22b have been previously published by the author and collaborators in Ref. [2], and are subject to a Creative Commons Attribution 4.0 International License.

Figures 9.7, 9.11, 9.12, 9.14 to 9.17 and 9.19 have been previously submitted for publication by the ATLAS Collaboration in Ref. [3], and are subject to a Creative Commons Attribution 4.0 International License.

The author has received support for the completion of the work described in this dissertation from the Science and Technologies Funding Council [Doctoral Training Grant ST/W507726/1], the Clarendon Fund, and Brasenose College, Oxford.

To Ben

Acknowledgements

Just as it takes a village to raise a child, it has taken three laboratories, two testbeam campaigns, many advisers and friends, and an inordinate amount of coffee to educate this DPhil student. There are many people I must thank — I couldn't have done it without you.

First and foremost, my gratitude goes to my supervisors Daniela Bortoletto and Chris Hays. You have both guided me in very complementary ways, and I am grateful to have been able to learn from your respective perspectives. Daniela, thanks for bringing me into such a vibrant group, for your guidance and advice, and especially for the particular care you've taken for my wellbeing over these years. From difficult times to exciting times, you've always been there with understanding and a smile. Chris, thank you for introducing me to the wonderful world of physics analysis, for every engaging discussion from signal scale systematics to forays into anime, and for pushing me to explore new projects and research avenues.

Special thanks go to Brian Moser — who would have thought that a chat about a potential project for a summer student could turn into two successful testbeam campaigns, advice and support for many shenanigans, and most importantly, a new friendship. You always treated me as a peer and an equal, and for that I am forever grateful.

To the OPMD lab group, thanks for the collaboration, the chats, the pub trips. Thanks to Daniel for always making time for a chat, whether physics or personal. To Richard, thanks for running such a wonderful lab. Umberto and Jonathan, it's been a pleasure to mix glue and crash gantries with you both. Particular thanks to Alex for putting up with my occasional strange machining requests, and to Ashley for always being up for a good yarn. Thanks also to Nina, Annabelle, Alik, Sam, Martin, and Dan; you make OPMD all the more vibrant and cheerful.

I was very lucky to join an incredibly proficient and welcoming analysis group — it has been a pleasure to work with you. In particular, thanks to Chris Pollard and Chris Gutschow for your leadership and advice, and to Tomoya, Giovanni, and Federico for the many conversations and great collaboration. I will very much miss our hackathons and impromptu meetings at CERN.

My time at CERN was incredibly formative, and I owe a large debt of gratitude to Benedikt Vormwald and Susanne Kuehn for welcoming me so warmly into the SR1 Pixel group. Thanks Hans for your friendship, and for the many discussions on DCS and beyond. To Ismet and Leyre, thanks for teaching me some of your endless knowledge on the FELIX system.

The two material testbeam experiments were only a success due to the many people who jumped headlong into the chaos to make them happen. To our summer students Antonin and Ruixiang, you were both great and I wish you both a very successful future! A great thanks is owed to the MALTA group at CERN and beyond for lending equipment and providing support, and in particular Valerio Dao, Lucian Fasselt, Milou van Rijnbach, and Carlos Solans. Thanks also to the ETH Zurich CMS group for joining on our second sortie — Bane, Sophie, Malte, I had a blast working with you.

I would also like to acknowledge the Pixel Outer Endcap groups, and the ITk community more generally. In particular, thanks to Marta Sabate Gilarte for the many shared debugging sessions, Maria Mironova for your advice on topics from ITkPix to life, and Lingxin Meng for always being up for a chat when we bump into each other.

To Ian Shipsey — you were my first point of contact in Oxford; thank you for introducing me to this wonderful department. I miss you terribly, and will always remember your ever-positive outlook and infectious smile. To Elisabeth and Eleonora, thank you for organising the Bortoletto/Shipsey group meetings, with so many opportunities for fruitful discussions, and apologies for always replying late to your emails.

Every walk to OPMD needed a coffee stop at Horsebox on the way — thanks to Alessandro for listening to my unfounded opinions on machine learning, Maggie for sharing your well-founded ones, Gregor for being very much a kindred spirit on our respective hardware projects, and Sparshita for your endless positivity. Ynyr, Iza and Eimear, thanks for the companionship and the (at times terrifying) shared ski trips. Beth, Anna, and Seb, thanks for the fun times on LTA and friendship beyond. Baris, thanks for trusting me when I told you hardware would be fun, and I apologise for the inconvenience that I'm sure has followed. Tom, thanks for always being up for a coffee and a walk, and the many deep physics discussions. Thanks also to everyone else in our vibrant and talented department — I couldn't have asked for a better group of colleagues.

Tough times come for us all, and I am particularly grateful to Sonia, Earl, and Matt back in Australia, as well as our neighbours and community back home, for being there and helping my family and I pick ourselves up when our world suddenly shattered.

My family has been here for me through thick and thin, and I am incredibly grateful to all of you. Thanks to my aunt Ria and uncle Wolfgang for showing such an interest in what I do. To my family in Germany and Austria, thanks for having become my home since leaving Australia, and in particular *danke Oma, dass ich immer bei Dir willkommen bin, und dass du Eningen zu meinem zweiten Zuhause gemacht hast*. I am especially grateful to my parents, Sabine and Bernhard, and my brother Samuel — you have supported me unconditionally over the years, and have very much made me who I am now. A special thanks also to Sam and Melanie for providing beautiful illustrations for this work. Finally, to my partner Katharina — we started the PhD journey together, and you've been with me through all the twists and turn that came along with it. Thanks for your eternal patience during stressful testbeam weeks, your willingness to discuss any inane topic my mind gets stuck on, and your love and care over these past years.

Abstract

The Large Hadron Collider represents the energy frontier of the field of collider physics, and in recent years has produced some of the most precise measurements of the behaviour of fundamental particles. This thesis spans the lifecycle of a general-purpose particle detector: from a measurement and interpretation of data collected with the ATLAS detector, to the construction and characterisation of the ITk, a new inner tracker for ATLAS to address the challenges of the High-Luminosity era.

The ITk is one of the largest and most ambitious tracking detector projects to date, and faces a challenging new environment. The pixel detector comprises lightweight support structures populated with modules, and the construction and quality control of these is discussed with reference to half-rings for the Outer Endcap subdetector. Readout, control, monitoring, and safety systems have been developed that mirror the final detector services, and preliminary results demonstrate preparedness for the upcoming pre-production phase.

Precise knowledge of detector material is crucial both during the R&D phase and as an input to simulations. The first direct measurement of the radiation length of an ITk pixel module is presented, producing a 2D resolved map with $\mathcal{O}(10\%)$ uncertainty and $0.5\text{ mm} \times 0.5\text{ mm}$ resolution via a measurement of the multiple scattering of positrons at a testbeam. The understanding of our detector plays a key role in enabling the analysis of complex detector signatures, often in a search for new phenomena.

Probing off-shell production of new particles or tail effects encoded in effective field theory operators allows for searches for new phenomena beyond the naïve energy threshold of the LHC. A measurement of the Drell-Yan cross-section in a final state with two τ leptons in association with b -jets at high di-tau mass is combined with a search for leptoquark and Z' boson models, using data collected by the ATLAS detector during Run 2 of the LHC. In addition, model-independent constraints are extracted on new physics contributions to operators involving τ leptons. The unfolded cross-section matches Standard Model predictions, and stringent constraints are set on new physics models which begin to exclude the mass-coupling regions preferred by recent hints of lepton flavour universality violating anomalies.

Contents

List of Abbreviations	xi
Preface	1
I ATLAS now and ATLAS future	4
1 Toward a universal theory of particles and interactions	5
1.1 A physics of particles	7
1.1.1 Noether’s theorem and the role of symmetries	7
1.1.2 The fields of QFT	8
1.1.3 Interactions	10
1.2 The particles of physics	12
1.2.1 The Higgs field and spontaneously broken SU(2) symmetry	14
1.3 Effective Field Theories	16
1.3.1 The Standard Model Effective Field Theory	17
1.4 The Drell-Yan spectrum at colliders	18
2 Charged particle tracking with silicon detectors	20
2.1 Interactions of charged particles with matter	20
2.1.1 Radiation length and electromagnetic processes	21
2.2 Operating principles of silicon detectors	23
2.2.1 Detector readout — hybrid and monolithic	25
2.2.2 Radiation damage	27
2.2.3 Detector resolution and performance considerations	28
2.3 Summary	31
3 From collisions to ATLAS data	32
3.1 Probing rare phenomena in proton collisions at the LHC	32
3.1.1 Kinematics of proton-proton collisions	34
3.1.2 Luminosity and pileup	35
3.1.3 The role of triggers	36
3.2 The observables of interest to physics	37
3.2.1 Momentum and charge — the tracking detectors	38
3.2.2 Energy — calorimetry	40
3.2.3 The muon system	42
3.3 From observables to physics objects	43
3.3.1 Hadrons and other jets	43
3.3.2 Jet structure for b -quark identification	44
3.3.3 Identifying hadronically decaying τ -leptons	45

4	The HL-LHC and ATLAS ITk upgrades	47
4.1	The case for higher luminosity	48
4.1.1	Technical challenges	49
4.2	The ATLAS Phase-2 upgrade	51
4.3	ITkPix: readout for a new pixel detector	56
4.3.1	The ITkPix readout chip	56
4.3.2	On-chip power distribution	58
4.3.3	Injection circuitry and electrical testing	59
4.4	OEC support structures and services	62
4.4.1	Monitoring	62
4.4.2	Interlock	63
4.4.3	The optosystem	65
4.4.4	Outer endcap support structures	67
5	Large-scale tracking structures for the ATLAS ITk	69
5.1	Half-ring loading	70
5.1.1	Loading programs and the GiGo framework	72
5.1.2	Half-ring fiducial fitting algorithm	74
5.1.3	Module fiducial fitting algorithm	76
5.1.4	The thermal interface adhesive	76
5.1.5	Adhesive preparation	78
5.1.6	Adhesive deposition	80
5.1.7	Loading tests and progress	81
5.2	Half-ring QC testing	84
5.2.1	QC testing requirements and readout scheme	84
5.2.2	The test enclosure	87
5.2.3	Interlock design	89
5.2.4	Detector control system	92
5.2.5	Similarities to the SR1 RD53A demonstrator and slice test	96
5.2.6	Preliminary results	96
5.3	Conclusions and outlook	101
II	Measurements of ATLAS	102
6	Multiple scattering and material content	103
6.1	Material distribution in ATLAS	103
6.1.1	Impact of material on physics results	104
6.2	Multiple scattering of charged particles	105
6.2.1	The single Coulomb scattering distribution	106
6.2.2	Molière's theory of multiple scattering	109
6.2.3	The Lynch-Dahl-Highland approximation	110
6.2.4	The Frühwirth-Regler formalism	111
6.2.5	Scattering models in GEANT4	114

7	Measuring detector material via multiple scattering at testbeams	115
7.1	Experimental design considerations	116
7.2	The MONSTAR telescope	118
7.2.1	Thermal considerations	122
7.2.2	Datasets collected	123
7.3	Analysis methodology	124
7.3.1	Data preparation, alignment, and data quality checks	124
7.3.2	Tracking and angle extraction	125
7.3.3	MALTA half-chip efficiency correction	126
7.3.4	Telescope acceptance checks	127
7.3.5	Fitting angle distribution widths	128
7.3.6	Mechanics and air contribution subtractions	131
7.4	Material estimate	133
7.5	GEANT4 simulations of air and telescope scattering	134
7.5.1	Comparison of scattering distribution width extraction methods	137
7.5.2	Telescope response	138
7.5.3	Air and telescope scattering validation	140
7.5.4	Variation of air and telescope contributions with plane spacing	142
7.6	Results	142
7.6.1	Resolution and beam energy uncertainties	143
7.6.2	Comparison of fit methods and scattering models	145
7.6.3	Analysis without DUT data	146
7.6.4	Comparison to the 700 MeV campaign 1 dataset	146
7.7	Implications and conclusions	147
8	Lepton flavour universality violation in BSM models and EFT	150
8.1	Experimental evidence	151
8.1.1	$R(D^{(*)})$	152
8.1.2	$R(K^{(*)})$	153
8.2	Relevant effective field theory couplings	155
8.2.1	$\mathcal{C}_9^{\text{NP}}$ and other $b \rightarrow s\ell^+\ell^-$ operators	155
8.2.2	$b \rightarrow c\ell\bar{\nu}$ and leptoquark-equivalent couplings	157
8.2.3	Flavour-dependent SMEFT couplings	157
8.2.4	Electric and magnetic dipole moments	157
8.3	Leptoquark models	158
8.3.1	Scalar leptoquarks	160
8.3.2	Vector leptoquarks	161
8.3.3	Leptoquark searches at collider experiments	162
8.4	Z' boson models	165

III Measurements with ATLAS 150

9 Measurement of the $pp \rightarrow \tau\tau$ cross-section at high $m_{\tau\tau}$ and search for leptoquarks	166
9.1 Analysis design and observable	167
9.2 Data and simulation samples	169
9.2.1 MC-simulated background samples	169
9.2.2 Signal models	172
9.3 Object reconstruction and event selection	173
9.3.1 Trigger conditions	173
9.3.2 Tau candidate reconstruction	174
9.3.3 Jets	175
9.3.4 b -tagging	175
9.3.5 Overlap removal	175
9.4 Background estimation for jets faking taus	176
9.4.1 The Universal Fake Factor method	178
9.4.2 Application to the $e\tau_{\text{had}}$ and $\mu\tau_{\text{had}}$ channels	182
9.4.3 Application to the $\tau_{\text{had}}\tau_{\text{had}}$ channels	184
9.5 Systematic uncertainties	186
9.5.1 Experimental uncertainties	187
9.5.2 Modelling uncertainties on background and Drell-Yan samples	188
9.5.3 Modelling uncertainties on BSM samples	189
9.6 Modelling validation and detector-level cross-section	189
9.7 Measurement of the inclusive $\tau_{\text{had}}\tau_{\text{had}}$ cross-section	192
9.7.1 Accounting for detector effects	192
9.8 Statistical model	195
9.9 Search for non-resonant leptoquark and Z' interactions	198
9.9.1 Signal distribution correction for the ≥ 2 b -jet category	198
9.9.2 Results	201
9.10 SMEFT interpretation and limits on Wilson coefficients	205
9.10.1 Parametrisation of the cross-section in terms of couplings	207
9.10.2 Signal samples	209
9.10.3 Results	209
9.11 Conclusions and outlook	212

Conclusions and Outlook 214

Appendices

A Literature and calculated material X_0 values 217

B Further results from telescope simulations 219

References 222

List of Abbreviations

ASIC	Application-specific integrated circuit, typically using CMOS logic.
ATLAS	A Toroidal LHC ApparatuS, one of four main detectors at the LHC.
BDT	Boosted decision tree classifier.
CAN	Controller Area Network, low-level communication protocol used in the MOPS system and implemented in CAN interface cards (CICs).
CERN	<i>Conseil Européen pour la Recherche Nucléaire [FR]</i> . European Organization for Nuclear Research, Geneva, Switzerland.
CL	Confidence level.
cRIO	National Instruments compact reprogrammable I/O system.
CMOS	Complementary metal-oxide semiconductor logic.
DAC, ADC . . .	Digital-to-analog converter, analog-to-digital converter.
DSCB, SSCB .	Double- or single-sided Crystal Ball function.
DCS, DSS . . .	Detector control system, detector safety system.
DUT	Device or subject under test in a testbeam measurement.
EFT	Effective field theory. Specific prescriptions include the standard model EFT (SMEFT) and low-energy EFT (LEFT).
FELIX	Front-end Link Exchange, a new optical readout system for ATLAS.
FPGA	Field-programmable gate array.
FSM	Finite state machine.
GBCR	Gigabit Cable Receiver ASIC.
Geant4	Geometry and tracking toolkit for particle simulations.
HL-LHC	High-Luminosity upgrade program of the LHC.
HV, LV	High-voltage (bias) and low-voltage (power) provision to a detector.
IBL	Insertable B-layer, a later addition to the ATLAS Pixel detector.
ID	ATLAS Inner Detector, encompassing the Pixel detector, Semiconductor Tracker (SCT) and Transition Radiation Tracker (TRT).
ITk	ATLAS Inner Tracker upgrade for the HL-LHC.
ITkPix	Readout ASIC designed for the ITk Pixel detector.
JES, JER	Jet energy scale and resolution.
JVT	Jet vertex tagger algorithm to veto pile-up jets.
LISSY	Local Interlock and Safety System for the ATLAS ITk.

LFUV	Lepton flavour universality violation.
LHC	Large Hadron Collider at CERN.
LO, NLO,	Leading order, next-to-leading order, ...
lpGBT	Low-power gigabit transceiver ASIC.
MALTA	A depleted monolithic active pixel sensor developed for high-energy physics experiments.
MC	Monte-Carlo event generation.
ME	Matrix element for a given scattering process.
MONSTAR	Malta-based testbeam telescope designed for material measurements.
MOPS	Monitoring Of Pixel System ASIC, and readout system (MOPShub).
MH4B	MOPShub for Beginners, a small-scale MOPS readout system.
NIEL	Non-ionising energy losses.
NN	Neural network, including recurrent (RNN) and graph-based (GNN).
NTC	Negative-temperature-coefficient thermistor.
OEC, OB, IS	Outer endcap, outer barrel, and inner subsystems of the ITk Pixel detector.
OS, SS	Opposite-charge (sign) and same-charge (sign), referring to the leptons in a di-leptonic final state.
OPMD	Oxford Physics Microstructure Detector Laboratory.
PCB	Printed circuit board.
pdf	Probability density function, not to be confused with parton distribution functions.
PDF	Parton distribution function (for partons within a hadron). A very specific type of probability density function.
PS	Parton shower simulation.
POI	Parameter of interest in a functional or template fit.
PP0, PP1,	The four layers of patch panels forming the interfaces between the ITk detector and the services cavern.
QC	Quality control, referring to electrical or mechanical verification.
QFT	Quantum field theory.
QCD	Quantum chromodynamics, referring to the strong nuclear force.
RD53	Collaboration that designed the RD53A and ITkPix readout ASICs.
SLDO	Shunt low-dropout regulator.
SM, BSM	Standard Model of particle physics, and physics beyond the SM.
SMD	Surface-mount device.
SP-chain	Serial powering chain of pixel modules.
TDAQ	Trigger and data acquisition system.
TES	Energy scale for jets originating from hadronic tau decays.
TRACI	Transportable Refrigeration Apparatus for CO ₂ Investigation.
ToT	Time-over-threshold for a pixel hit, a measure of deposited charge.
VTRx+	Versatile link plus transceiver for electrical ↔ optical conversion.

Preface

This dissertation walks a somewhat meandering path between two topics at the heart of collider physics: from a discussion of how we build the detectors we plan to use, to an exploration of how we use the detectors that we have built. To cover these topics in a comprehensive manner, brief reviews of the relevant subject areas are interwoven with novel contributions, which this preface aims to disambiguate.

PART I begins with a current snapshot of the Standard Model of particle physics, laying the theoretical foundation for goals of the Large Hadron Collider programme. Chapter 2 follows, relating high-energy particles to low-energy phenomena in the context of silicon tracking detectors. A brief review of practical design considerations provides the basis for the discussions of pixel detectors throughout the remainder of the work. The ATLAS experiment is introduced in Chapters 3 and 4, first focusing on the detector as currently installed and operating, and later moving to the future plans for the detector, and specifically the Inner Tracker (ITk) upgrade project.

Chapter 5 describes efforts at Oxford Physics Microstructure Detector (OPMD) forming part of a larger construction effort for the ITk detector. Two aspects of the work at Oxford are explored: the production of loaded half-ring structures for the new detector, and the quality control (QC) testing of these in an electrical test setup. The methods and results discussed in this chapter are my work, including the development of the GiGo software framework and the gantry loading programs, the refinement of the method used for adhesive deposition during loading, and the loading tests presented. I led the development of the half-ring QC test setup, including choices on equipment to be used, designs for cabling and interconnect, the design of an interlock system for larger structures on the basis of the module-testing interlock, and the setup of the associated detector control system (DCS). As part of this work, I have collaborated with the ITk

Pixel system tests team at CERN to expand their common DCS software to our use case, and have made substantial contributions to this effort.

PART II describes an effort to measure the material content of ITk detector structures through testbeam measurements of multiple scattering. Following a review of relevant scattering theory in Chapter 6, a measurement using a new testbeam telescope is presented in Chapter 7, which has been published in Refs. [2, 4]¹. The measurement itself was performed in a small team, using the T9 beamline at the CERN Proton Synchrotron. I designed and constructed the experimental apparatus, led the data-taking, performed the analysis of the resulting datasets, and co-wrote the manuscript for the paper. I also played a role in the managerial aspects, including applications for beam time and the training and supervision of shifters who assisted during the data-taking period. The results presented led to a more recent measurement which has not been described in this work, and for which a future publication is planned.

PART III introduces the principle of lepton flavour universality violation (LFUV) and implications in terms of potential physics beyond the Standard Model of particle physics (SM). Chapter 8 provides a brief review of current results on LFUV in b -physics, and introduces the concepts of leptoquarks, Z' bosons, and discusses relevant effective field theory (EFT) couplings.

The measurement and search for new physics presented in Chapter 9 have been internally reviewed by the ATLAS collaboration and submitted to the Journal of High Energy Physics, and a preprint has been made public as Ref. [3]². This chapter describes a measurement of the Drell-Yan cross-section at high energies in final states with two tau leptons, and an associated search for beyond-SM (BSM) behaviour, including leptoquark and Z' models, in ATLAS data collected between 2015 and 2018. The work was performed as part of a larger analysis group, and I contributed to the analysis throughout its duration. I worked on the development of the analysis software framework used, and developed and maintained post-processing pipelines that collated the outputs and prepared them for


¹Figures similar to those in this publication have been included here, but have been modified to match the style and colour scheme of this document. In many cases, figures have been augmented with further results or more detailed explanations compared to the published versions.

²Following ATLAS Collaboration rules, result plots from the published work have been left unaltered, and are marked with an “ATLAS” label. Figures without this label are my own work.

plotting and further interpretation through profile likelihood fits. I also contributed to the plotting and fit code used, as well as the implementation and validation of the data-driven estimate in the $\tau_{\text{lep}}\tau_{\text{had}}$ channel of the analysis, and the development and implementation of the forward-folding method used to estimate signal yields in the $\geq 2b$ -jet category.


The development processes for both detectors and data analyses involve overcoming various technical challenges, and generally benefit from a wide viewpoint on the respective fields. Alongside the work described here, I have made contributions to the Local Interlock and Safety System for the final ITk detector, including the integration of this interlock system into the common DCS software described above. I have also contributed to the testing of pixel modules produced in OPMD, and to the validation of safe transport methods for modules post-construction. I have been an active developer for the ICICLE instrument control framework [5], and have integrated interfaces for many devices used for the work presented in Parts I and II. As part of the Drell-Yan di-tau analysis, I also performed comparisons of τ lepton identification efficiencies between different Monte-Carlo generators, and developed signal region optimisations in an attempt to improve the sensitivity of the analysis. These efforts are not documented here.

PART I



ATLAS now and ATLAS future

Illus. 1: *The LHC in Geneva*, by Samuel Koch, 2025.



— Randall Munroe, *xkcd* #2933 [1]

1

Toward a universal theory of particles and interactions

Contents

1.1 A physics of particles	7
1.1.1 Noether's theorem and the role of symmetries	7
1.1.2 The fields of QFT	8
1.1.3 Interactions	10
1.2 The particles of physics	12
1.2.1 The Higgs field and spontaneously broken $SU(2)$ symmetry . .	14
1.3 Effective Field Theories	16
1.3.1 The Standard Model Effective Field Theory	17
1.4 The Drell-Yan spectrum at colliders	18

So far, every well-accepted theory within physics has eventually shown itself to be an *effective theory*, valid only within a particular parameter space. Newtonian gravity provides near-perfect predictive power for very massive objects moving slowly, Einstein's theory of special relativity disregards gravitational effects, and general relativity breaks down at high energies and small length scales. The Standard Cosmological Model (Λ CDM) shows great predictive power for the current scale and energy spectrum of the universe, but is likely a low-energy limit of a more complete theory that could account for phenomena such as dark energy and dark matter [6]. Theorems from thermodynamics, condensed matter physics, or quantum theory are generally limited to a much stricter specific set of boundary conditions.

It has long been a goal of physicists to derive a single universally applicable theory, a

theory of everything. A natural starting point for this search is to build up from the smallest constituents of the universe, and the Standard Model of particle physics (SM) represents the current state of this endeavour [7]. The SM combines two theoretical frameworks, quantum field theory (QFT) and special relativity, and shows very good agreement with experimental results and strong predictive power. Experimental support for the model has come both from observations of natural phenomena including cosmic rays and radioactivity, and from dedicated experiments involving particle accelerators or colliders.

The fundamental limitations on the SM are given by the energy scales accessible to experiment. The development of quantum physics and formulation of the wave-particle duality unearthed a principle that has had strong implications on the field of particle physics since the turn of the 20th century — the *de Broglie relation* [8], given as

$$\lambda = \frac{h}{p} \quad (1.1)$$

where λ is the wavelength of a quantum-mechanical wavefunction, p is the momentum, and h is Planck's constant. Wave-particle duality implies this holds true for a particle represented by a given wavefunction, which through Einstein's mass-energy equivalence also implies massive particles at rest carry a characteristic wavelength known as the *Compton wavelength*, which is given as $\lambda = \frac{h}{mc}$ [9].

Applied at face value, the de Broglie relation implies that the characteristic length of a wavefunction — which dictates its propensity for interaction with other wavefunctions of a similar wavelength¹ — is inversely proportional to the momentum of the particle. This relationship has long been used in optics and microscopy, where it is well known that the resolving power of photons in the form of light is directly proportional to the frequency (and hence energy) of the incident beam [10]. Used as a probe for physical phenomena, this implies the energy of a particle is proportional to its resolving power — the higher the total energy, the smaller the subjects and phenomena it can resolve.

¹This can be shown by rearranging the Heisenberg uncertainty principle $\Delta x \Delta p \geq \hbar/2$ and using the relativistic relation between mass and energy $E^2 = (mc^2)^2 + (pc)^2$ to obtain $\Delta x \geq \frac{1}{2}(\frac{\hbar}{mc}) = \frac{1}{2}\lambda$, demonstrating that the de Broglie wavelength dictates a lower limit on the position uncertainty and hence resolving power of the wavefunction.

1.1 A physics of particles

The SM has a long and storied history, with individual components of the theory often motivated by several experimental, theoretical, or historical considerations — good introductory texts covering the topic include Cottingham and Greenwood [11] and Griffiths [12], whereas a more detailed review of the phenomenology can be found in Rubbia [13]. Peskin and Schroeder [14] and Schwartz [15] give detailed descriptions of the QFT foundations underlying the SM. As it is not possible to provide here a complete summary from the historical perspective, a current snapshot of the structure of the SM will be presented in this chapter, followed by the particular practical applications and implications discussed in the body of this thesis.

Taken together, quantum mechanics and special relativity introduce a set of fundamental principles that guide the structure of allowable QFTs. Special relativity gives a clear case for a four-dimensional theory with three spatial dimensions and one time dimension². The inherent causality implies locality of individual interactions. Quantum theory provides an existing framework within which interactions between objects of the theory can be described whilst maintaining the unitarity of the overall state of the system³.

QFT generalises Hamiltonian dynamics and quantum-mechanical scattering amplitudes to a quantised four-dimensional field theory [16]. The dynamics of a set of fields $\Phi(x)$, where x is a vector encoding a location within four-dimensional spacetime, is governed by the Lagrangian density $\mathcal{L}[\Phi]$ of the theory. Equivalently to the energy minimisation of the Hamiltonian in quantum theory, the minimisation of the action, defined as

$$S[\Phi] = \int d^4x \mathcal{L}[\Phi], \quad (1.2)$$

is used to derive the evolution over time of the fields $\Phi(x)$.

1.1.1 Noether's theorem and the role of symmetries

The allowable terms of \mathcal{L} are governed by symmetries imposed due to physical or theoretical considerations. These are generally formalised as groups or algebras within the framework

²Higher-dimensional theories are possible, but require complex mechanisms to reduce the impacts of additional dimensions at observable scales, and will not be considered here.

³Non-unitary probabilities invalidate the physical applicability of the theory.

of group theory, and restrict the allowable forms of the fields $\Phi[x]$ themselves, and interactions between these. Noëther's theorem states that every continuous symmetry of the action generates a conserved quantity [17].

The most fundamental symmetry is imposed by the special relativistic requirement for physics to behave the same in all inertial reference frames — *Lorentz symmetry*. This is formalised in the Poincaré algebra, and requires all fields $\Phi[x]$ to transform as irreducible representations of the Lorentz group [18]. Invariance under spatial translations, time translations, and rotations generate conservation laws for energy, momentum, and angular momentum. In addition, the Poincaré algebra allows for several non-mandatory discrete symmetries to be introduced — symmetry under charge conjugation (C), parity (P), or time reversal (T) operations.

Further symmetries can be imposed by assuming particular group structures for the objects and interactions within the theory, and are generally related to unphysical degrees of freedom obtained by embedding physical particles into field representations. These symmetries produce a series of *currents* coupling different fields within the theory.

1.1.2 The fields of QFT

At this point, it becomes instructive to consider what is meant by the term *particle* in the context of a field theory. A QFT consists of a set of universally-permeating fields satisfying one of two relativistic generalisations of the Schrödinger equation, that are then quantised. Particles arise as localised excitations within these fields. The propagation through spacetime and interactions between these are governed by the *equations of motion* that can be obtained through the minimisation of the action. The two generalisations of Schrödinger's equation are the Klein-Gordon equation [19, 20],

$$(\partial^2 + m^2)\phi = 0, \tag{1.3}$$

solutions to which describe fields of integer spin with mass m , and the Dirac equation [16],

$$(i\partial\!\!\!/ - m)\psi = 0, \tag{1.4}$$

describing fields of half-integer spin. Here, ∂^2 refers to the d'Alembert operator ($\partial_\mu \partial^\mu = \partial_t^2 - \nabla^2$)⁴, and $\not{\partial} = \gamma_\mu \partial^\mu$ where γ_μ are Gamma matrices forming a four-dimensional representation of the Clifford algebra $\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu}$, with the Minkowski metric $\eta^{\mu\nu}$.

Of the possible field configurations, three are reflected in the SM: spin-0 fields transforming as *scalars* under the Lorentz group, spin-1 fields transforming as *vectors*, and spin- $\frac{1}{2}$ fields transforming as *spinors*. In general, fields with spin greater than 2 require higher-order conserved currents to be introduced that would violate unitarity. Spin-2 fields (gravitons) can be introduced as a consequence of the stress-energy tensor of general relativity — in fact, it can be shown that at most one interacting type of spin-2 particle can exist, and that if introduced, it must be the graviton⁵ [21]. Spin-3/2 (Rarita-Schwinger) fields are relevant for Supersymmetry models introducing a spin- $\frac{3}{2}$ conserved supercurrent, but require careful treatment to maintain unitarity [22].

For each field, the non-interacting components of the Lagrangian comprise a kinetic term, and in the case of massive fields, a mass term. The forms of these are given in Table 1.1. In the case of a vector field, a bare mass term violates Lorentz invariance, and massive vector fields can only arise as consequences of another field (or more specifically, spontaneously broken gauge symmetry) introduced to the Lagrangian. From dimensional analysis of the kinetic terms, the scalar and vector fields are determined to have dimensions⁶ of $[\text{length}]^{-1}$, or equivalently $[\text{mass}]$, and the spinor field has $[\text{length}]^{-\frac{3}{2}} = [\text{mass}]^{\frac{3}{2}}$.

The Lorentz group $\text{SO}^+(1, 3)$ has a double-cover $\text{Spin}(1, 3) \cong \text{SL}(2, \mathbb{C})$, which allows irreducible representations (j_L, j_R) under left-handed and right-handed Lorentz rotations, respectively. Due to their integer spin, scalar and vector fields will necessarily have

Table 1.1: Kinetic and mass terms for scalar, vector, and spinor fields. The form of the field-strength tensor $F_{\mu\nu}$ depends on the corresponding gauge symmetry.

Scalar field ϕ		Vector field A^μ		Spinor field ψ	
kinetic term	mass term	kinetic term	mass term	kinetic term	mass term
$\frac{1}{2} \partial_\mu \phi \partial^\mu \phi$	$m_\phi^2 \phi^2$	$\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$	$m_A^2 A_\mu A^\mu$	$i \bar{\psi} \not{\partial} \psi$	$m_\psi \bar{\psi} \psi$

⁴Einstein summation over pairs of raised and lowered indices is used throughout this thesis.

⁵The Coleman-Mandula theorem [21] posits that with the exception of geometric currents (such as supersymmetric currents or the stress-energy tensor), currents within a Lorentz-invariant QFT must not carry any space-time indices other than the vector index. Since non-interacting fields are non-observable, this sets constraints on the set of useful fields within a realistic model.

⁶In natural units, dimensions of length, mass, and energy are related by Einstein's relation $E^2 = m^2 - p^2$.

representations $(0, 0)$ and $(\frac{1}{2}, \frac{1}{2})$. Spinor fields carry half-integer spin, and as a result can satisfy one of two possible representations, $\psi_L : (\frac{1}{2}, 0)$ or $\psi_R : (0, \frac{1}{2})$. This inherent *handedness* is termed *chirality* [23], and any spinor can be separated into left- and right-chiral components using the projection operators

$$P_{L/R} = \frac{1 \mp \gamma^5}{2}, \quad \psi = P_L \psi_L + P_R \psi_R \quad (1.5)$$

where $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$, and the chiral components form a complete basis.

An interesting effect occurs for vector fields, which carry three independent degrees of freedom⁷. Spin considerations require massless vector fields to carry exactly two possible polarisations⁸. The excess degree of freedom can be reinterpreted as free choices of a *gauge* [24, 25]. The Lagrangian must be invariant under the set of gauge transformations related to the choice of gauge, which make up the *gauge group*. In practical terms, computations typically rely on *gauge fixing*, in which a particular gauge is chosen, and non-physical *ghost fields*⁹ are introduced to explicitly cancel the excess degrees of freedom. Together with spin, the conserved quantities introduced by requiring invariance under gauge groups are typically referred to as *quantum numbers* and take discrete values as allowed by the form of the gauge group in question.

1.1.3 Interactions

Interactions between fields are generated by Lagrangian terms containing several fields, or derivatives of these. A full set of possible interaction terms can be derived by considering all possible terms that retain a mass dimension of 4 [26, 27]. However, due to the quantum mechanical nature of QFTs, interaction terms cannot be considered individually. Rather, information about interactions between particles is encoded in scattering amplitudes, which can be computed through the path integral [16, 28],

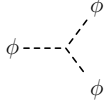
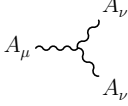
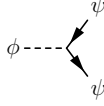
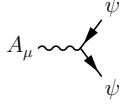
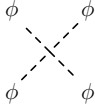
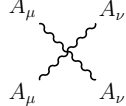
$$\mathcal{M} \propto \int D\Phi \exp(iS[\Phi]), \quad (1.6)$$

⁷Despite being four-component vectors, one degree of freedom is removed since the equations of motion require $\partial_\mu A^\mu = 0$.

⁸As they must travel at the speed of light, they can only be polarised transverse to their direction of motion, as longitudinal polarisation would imply a speed variation.

⁹These are generally of scalar type, but are chosen to satisfy the Grassman algebra and hence are anticommuting. It is important to note that ghost fields encode no physical phenomena, and are purely a mathematical tool.

Table 1.2: Tree-level interaction vertices with dimensionless coupling constants (except $\mu\phi^3$, where the dimensionful μ can naturally arise from a non-zero *vacuum expectation value* for the field ϕ). The propagators are associated with self-interaction terms including the kinetic term and mass term. Constant and combinatorial factors arising on these terms have been neglected.

Propagators	ϕ ----- ϕ	A_μ ----- A_μ	ψ ----- ψ	
3-point interactions	 $\mu \phi^3$	 $g \partial_\mu A^\mu A_\nu A^\nu$	 $y \bar{\psi} \psi \phi$	 $g \bar{\psi} \psi \gamma_\mu A^\mu$
4-point interactions	 $\lambda \phi^4$	 $g^2 A_\mu A^\mu A_\nu A^\nu$		

where the matrix element (ME) \mathcal{M} allows the computation of the transition probability from an initial state $|\Phi_1(x_1^i)\Phi_2(x_2^i)\dots\rangle$ to a final state $\langle\Phi_1(x_1^f)\Phi_2(x_2^f)\dots|$, which is proportional to $\left|\langle\Phi_1(x_1^f)\Phi_2(x_2^f)\dots|\mathcal{M}|\Phi_1(x_1^i)\Phi_2(x_2^i)\dots\rangle\right|^2$. Once expanded, \mathcal{M} includes all possible combinations of interaction terms, which are typically expressed as a series of *Feynman diagrams* [26, 28]. For interactions with a coupling strength much less than unity, higher-order diagrams including more couplings will be successively suppressed, allowing computations to low orders of the expansion to produce relatively accurate results. First-order expansion is typically referred to as *tree-level*, and further orders as one-loop-level, two-loop-level, and so on. The full set of basic tree-level interactions with dimensionless coupling constants and their contributions to the matrix element are shown in Table 1.2.

One interesting corollary is the form of the propagator, which appears as internal lines in Feynman diagrams. Consider the case of a pair of incoming scalars with four-momenta p_1, p_2 interacting to produce an intermediate state I , which then decays into two final states with momenta p_3, p_4 . The propagators are Green's functions for the Klein-Gordon equation, satisfying $(p^2 - m^2)\Psi = 0$ in momentum space. The naïve solution is given by

$$\frac{i}{p^2 - m^2}, \quad (1.7)$$

which develops a pole for the cases where $s = (p_1 + p_2)^2 = m_I^2$, the residue of which must be the produce of two corresponding three-particle amplitudes for the production

and decay of the intermediate state [29]. Equivalent poles exist for t -channel and u -channel states, where $t = (p_1 + p_3)^2 = m_I^2$ or $u = (p_1 + p_4)^2 = m_I^2$. The resulting form demonstrates the viability of *virtual* particles, for which the propagator is heavily suppressed but still remains non-zero for s , t , or u far away from these poles.

At one-loop-level and above, a significant challenge can arise when attempting to compute integrals over the loop propagators [30]. Since the particles in the loop are not required to have uniquely defined momenta, such integrals will often diverge, as they must integrate over all possible combinations of energy and momentum that could travel in the loop. These divergences do not represent non-physical behaviour, but rather limitations of the mathematical framework used. Several *renormalisation* schemes exist which allow the effects of these divergences to be mitigated by including the effects within the tree-level mass terms and three- or four-point interactions by adjusting the relevant mass or coupling parameter. Upon expansion of the renormalised parameters, this introduces *counter-terms* into the theory which cancel the divergences¹⁰. A consequence is that *renormalised* couplings are generally no longer truly constant, and will vary with the momentum transferred in the interaction.

1.2 The particles of physics

The SM is defined by its invariance under the gauge group $SU(3) \times SU(2)_L \times U(1)_Y$, with the first gauge group associated with quantum chromodynamics (QCD), and the remaining two groups describing the electroweak interaction [11, 13]. The $SU(2)_L$ group couples only to left-chiral fermions¹¹. The field content is given in Table 1.3 together with the relevant quantum numbers for each. Four species of spinor fields (the associated particles are termed *fermions*) are introduced — up-type quarks, down-type quarks, charged leptons, and neutrinos — and for each, three generations have been observed. Three types of vector field (*bosons*) are associated to the individual subgroups, and act as force carriers for these. The

¹⁰To make this cancellation mathematically precise, a form of *regularisation* needs to be applied, in which an arbitrary modification to the loop integrals (a *regulator*) suppresses the integrals at high momenta, and hence acts as a form of momentum cutoff. The limit as the regulator tends to infinity then recovers the original integrals. Several forms of regularisation exist, but will not be discussed here.

¹¹This behaviour is considered *maximally parity violating* due to its non-conservation under the parity transformation.

Table 1.3: Field content of the SM and associated quantum numbers. For fermions, the baryon number B and lepton number L are given, as are the electric charge Q , the third component of the weak isospin I_3 , and the weak hypercharge Y . The rightmost column shows the corresponding gauge group representation for $SU(3) \times SU(2)_L \times U(1)_Y$. Colour triplets carrying colour charge are denoted **3**, and colour singlets carrying no colour charge as **1**. The gluon octet (**8**) forms a basis of $SU(3)$. Weak isospin doublets are denoted **2** and singlets **1**, with the basis of $SU(2)_L$ denoted **3**. The third component of the representation shows the value of the weak hypercharge. Fields and couplings before and after symmetry breaking are separated by an arrow.

Spin-1 gauge fields								
	Symbol	Associated charge			Coupling		Gauge group representation	
Gluons	G_μ^α	colour			g_s	$(\mathbf{8}, \mathbf{1}, 0)$		
Electroweak bosons	$W_\mu^I \rightarrow W_\mu^\pm$	weak isospin I			$g \rightarrow g$	$(\mathbf{1}, \mathbf{3}, 0)$		
	$B_\mu \rightarrow Z, \gamma$	weak hypercharge Y			$g' \rightarrow \theta_{W,e}$	$(\mathbf{1}, \mathbf{1}, 0)$		
Spin- $\frac{1}{2}$ fermion fields								
	Symbol	B	L	Q	I_3	Y	colour	Gauge group representation
left-chiral quarks	$q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	$\frac{1}{3}$	0	$\begin{pmatrix} +\frac{2}{3} \\ -\frac{1}{3} \end{pmatrix}$	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$\frac{1}{6}$	$\mathbf{3}$	$(\mathbf{3}, \mathbf{2}, \frac{1}{6})$
right-chiral u -type quarks	u_R	$\frac{1}{3}$	0	$+\frac{2}{3}$	0	$\frac{2}{3}$	$\mathbf{3}$	$(\mathbf{3}, \mathbf{1}, \frac{2}{3})$
right-chiral d -type quarks	d_R	$\frac{1}{3}$	0	$+\frac{2}{3}$	0	$-\frac{1}{3}$	$\mathbf{3}$	$(\mathbf{3}, \mathbf{1}, -\frac{1}{3})$
left-chiral leptons	$\ell_L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}$	0	1	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$	$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$-\frac{1}{2}$	$\mathbf{1}$	$(\mathbf{1}, \mathbf{2}, -\frac{1}{2})$
right-chiral leptons	e_R	0	1	-1	0	-1	$\mathbf{1}$	$(\mathbf{1}, \mathbf{1}, -1)$
Spin-0 scalar fields								
	Symbol	Q			I_3	Y	colour	Gauge group representation
Higgs boson	$H = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 \\ h \end{pmatrix}$	$\begin{pmatrix} +1 \\ 0 \end{pmatrix}$			$\begin{pmatrix} +\frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$+\frac{1}{2}$	$\mathbf{1}$	$(\mathbf{1}, \mathbf{2}, \frac{1}{2})$

full set of gauge-invariant interactions under the gauge group is introduced by expanding the derivative within the fermion kinetic terms into the gauge-covariant derivative,

$$D_\mu = \partial_\mu + i\frac{g_s}{2}\lambda_\alpha G_\mu^\alpha + i\frac{g}{2}\sigma_I W_\mu^I + i\frac{g'}{2}Y B_\mu, \quad (1.8)$$

where the coupling constants corresponding to the strong force, weak isospin, and weak hypercharge are given as g_s , g and g' , respectively. The matrices λ_α and σ_I correspond to the generators of $SU(3)$ and $SU(2)_L$, respectively, with the latter taking the form of the Pauli matrices, and Y represents the weak hypercharge. In a similar way, the full set of couplings between gauge bosons can be obtained by defining the field strength tensors $G_{\mu\nu}^\alpha$, $W_{\mu\nu}^I$ and $B_{\mu\nu}$ in terms of the covariant derivative, as $F_{\mu\nu}^a = D_\mu F_\nu^a - D_\nu F_\mu^a + \tilde{g} f^{abc} F_\mu^b F_\nu^c$ where F_μ is the vector field and f^{abc} is the set of structure constants of the associated gauge group with charge \tilde{g} . The SM Lagrangian is hence given as

$$\mathcal{L}_{\text{SM}} = -\frac{1}{4}G_{\mu\nu}^\alpha G^{\mu\nu\alpha} - \frac{1}{4}W_{\mu\nu}^I W^{\mu\nu I} + \frac{1}{4}B_{\mu\nu}B^{\mu\nu} + \sum_{\psi=q_L, u_R, d_R, \ell_L, e_R} \bar{\psi} i \not{D} \psi + \mathcal{L}_{\text{Higgs}}, \quad (1.9)$$

where α iterates over the eight generators of the $SU(3)$ colour charge, and I iterates over three weak isospin configurations. The SM Lagrangian leads to several accidental global symmetries, including conservation of lepton number L and baryon number B as part of a five-fold flavour symmetry discussed further in Chapter 8.

1.2.1 The Higgs field and spontaneously broken $SU(2)$ symmetry

Neglecting $\mathcal{L}_{\text{Higgs}}$ gives the form prior to symmetry breaking, in which all bosons and fermions remain massless. To introduce mass terms, the mechanism of *spontaneous symmetry breaking* is introduced via a new complex scalar doublet, $\phi = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}$ [31–33]. The self-interaction potential for ϕ is chosen such that the minimum of the potential is non-unique, and occurs at a *vacuum expectation value* v defined by $\phi^\dagger \phi = v^2/2$. The vacuum state can then be chosen where $\phi_1 = \phi_2 = \phi_4 = 0$ and $\phi_3 = v$, and the potential expanded about this minimum. After adding a full set of possible couplings to the SM particles, a particular gauge can be chosen such that the *Goldstone bosons* ϕ_1 , ϕ_2 , and ϕ_4 are reinterpreted as longitudinal polarisation modes for the vector fields, effectively giving these mass. The remaining scalar field becomes a massive scalar field h termed the *Higgs field*, with $\phi \rightarrow H = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix}$. The additional terms to the Lagrangian are

$$\mathcal{L}_{\text{Higgs}} = (D_\mu H)^\dagger (D_\mu H) - \lambda \left(H^\dagger H - \frac{v^2}{2} \right)^2 - \left(\tilde{H}^\dagger \bar{q}_L Y_d u_R + H^\dagger \bar{q}_L Y_d d_R + \tilde{H}^\dagger \bar{\ell}_L Y_d e_R + \text{h.c.} \right). \quad (1.10)$$

where $\tilde{H} = i\sigma_2 H^\dagger$, and Y_u , Y_d and Y_e are 3×3 matrices of Yukawa couplings. Upon expansion of the Yukawa terms, fermions receive a mass term proportional to $Y_\psi^{ii} v$ for a fermion of type ψ and generation i . The Higgs mass is given as $m_H = \sqrt{2\lambda v^2}$.

After symmetry breaking, an $SU(3) \times U(1)_Q$ symmetry is retained, with Q denoting the electric charge defined by the relation $Q = I_3 + Y$. The interaction basis of the electroweak bosons is given as a mixture of B_μ and W_μ^i , with the W^\pm bosons chosen to such as to obtain unitary electric charge, and the electrically neutral Z boson and photon γ chosen such that the latter is massless:

$$W_\mu^\pm = \frac{1}{\sqrt{2}} (W_\mu^1 \pm iW_\mu^2), \quad \begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} B_\mu \\ W_\mu^3 \end{pmatrix}. \quad (1.11)$$

The off-diagonal terms of the Yukawa coupling matrices allow the mass basis to differ from the interaction basis. This behaviour is encoded in the unitary Cabibbo-Kobayashi-Maskawa (CKM) matrix [34],

$$V = \begin{pmatrix} V_{ud} & V_{cd} & V_{td} \\ V_{us} & V_{cs} & V_{ts} \\ V_{ub} & V_{cb} & V_{tb} \end{pmatrix}, \quad (1.12)$$

where the off-diagonal elements of V allow couplings between fermion generations for *charged-current interactions* mediated by the W^\pm bosons.

The strong coupling constant α_s is defined as $\alpha_s = g_s^2/(4\pi)$, and post-symmetry-breaking, the electroweak couplings are related by

$$g_W = g \quad g_Z = \frac{g}{\cos \theta_W} \quad e = \frac{g}{\sin \theta_W} \quad \tan \theta_W = \frac{g'}{g} \quad (1.13)$$

Due to renormalisation, the strong coupling weakens at higher energy scales¹², and the weak and electromagnetic couplings strengthen.

The current formulation of the SM contains 19 free parameters that must be determined experimentally. These include the three coupling strengths of the strong, and electroweak interactions, the nine Yukawa coupling parameters generating masses for the quarks and charged leptons, the four degrees of freedom of the unitary CKM matrix¹³, the mass and vacuum expectation value of the Higgs field, and additional parameters governing CP -violation in strong interactions and neutrino oscillations which are not discussed here.

¹²This property generates the spatial *confinement* and short range experienced by the strong force.

¹³These are typically parameterised by three mixing angles and one complex CP -violating phase.

1.3 Effective Field Theories

The SM in its current form has been highly successful, predicting the weak bosons, Higgs boson, and top quark, and showing great agreement with measurements. Despite this, it still exhibits several experimental and theoretical shortcomings, many of which could be alleviated through extensions of the theory (BSM). Experimental evidence including observations of dark matter, indications of lepton flavour universality violation, and the baryon asymmetry of the universe could be indicative of new fields only accessible at higher energy scales. The large number of experimentally-determined parameters of the model could be generated by a higher-order symmetry group that collapses to the SM at low energies — for example, Grand Unified Theories attempt to unify the strong and electroweak interactions at a unification scale well above what is currently accessible at particle colliders [35], and Supersymmetry can solve several theoretical qualms by the introduction of new particle content and a new symmetry linking the lepton and quark sectors [36].

The form of the propagator in Equation (1.7) indicates that very massive fields can appear within internal lines of processes, but will be suppressed by a factor proportional to p^2/m_X^2 for each internal line, where the new field X has mass m_X . Since the internal momenta will generally scale with the overall energy scale of the process E , this allows a low-energy effective description of the high-energy theory to collapse the internal line, and instead introduce a new vertex with mass dimension greater than 4, balanced by a coupling suppressed by $E^d/\Lambda^{d-4} \sim E^d/m_X^d$, where d is the dimension of the vertex. A classic example of such a theory is the Fermi four-point interaction [37], which acts as an effective theory for the electroweak interaction at energy scales well below the W and Z masses.

Formalising this notion leads to effective field theories (EFTs), in which couplings with arbitrary mass dimension are introduced together with appropriate suppression scales. A generic EFT Lagrangian may be written as

$$\mathcal{L} = \sum_d \sum_i \frac{c_i^{(d)}}{\Lambda^{d-4}} \mathcal{Q}_i^{(d)}, \quad (1.14)$$

forming a sum over arbitrary operators $\mathcal{Q}_i^{(d)}$ of mass dimension d with associated dimensionless couplings $c_i^{(d)}$, which are termed *Wilson coefficients*. It is immediately clear that since all SM operators are of mass dimension 4, they survive at all mass scales.

1.3.1 The Standard Model Effective Field Theory

A logical starting point for building EFT operators applicable to scales beyond those already excluded is the existing field content and symmetry structure of the SM. The Standard Model Effective Field Theory (SMEFT) comprises the set of all gauge invariant operators that can be constructed from SM fields [38, 39]. The SMEFT Lagrangian takes the form

$$\mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}} + \sum_i \frac{c_i^{(5)}}{\Lambda} \mathcal{Q}_i^{(5)} + \sum_j \frac{c_j^{(6)}}{\Lambda^2} \mathcal{Q}_j^{(6)} + \sum_k \frac{c_k^{(7)}}{\Lambda^3} \mathcal{Q}_k^{(7)} + \sum_l \frac{c_l^{(8)}}{\Lambda^4} \mathcal{Q}_l^{(8)} + \dots, \quad (1.15)$$

and is typically ordered in powers of E/Λ , since higher-dimension terms are increasingly suppressed for $\Lambda \gg E$. The odd-dimension terms are generally related to the Weinberg operator $\mathcal{Q}^{(5)} = (\ell^c \tilde{H}^*)(\tilde{H}^\dagger \ell)$ generating Majorana mass terms for neutrinos, or higher order expansions including further fermions or Higgs fields. Such operators necessarily violate either baryon number (B) conservation or lepton number (L) conservation, or both. The number of even-dimension operators grows exponentially for each increase in mass dimension due to the combinatorics of the field permutations.

Table 1.4 summarises the classes of operators in the dimension-6 SMEFT as parametrised in the Warsaw basis. A total of 2499 independent operators can be categorised

Table 1.4: The eight operator classes of the dimension-6 SMEFT in the Warsaw basis. Each class is annotated with the total number of possible operators within this class, and the number retained in the SMEFTSIM `top` prescription [39]. The SM fermion fields are denoted by ψ , the Higgs field is denoted H , and X represents any of the gauge bosons G_μ^α , W_μ^I or B_μ .

$X^3 — 4/4$	$H^6 — 1/1$	$H^4 D^2 — 2/2$	$\psi^2 H^3 — 8/8$
$X^2 H^2 — 54/14$	$\psi^2 X H — 144/36$	$H^2 \psi^2 D — 81/21$	$\psi^4 — 2205/189$

into eight classes, as shown. The largest class concerns four-fermion interactions, due to the combinatorics of flavour, generation, and chirality configurations.

Assuming certain symmetries of the higher-order theory generating these couplings allows the number of independent operators to be significantly reduced. One common assumption on the fermion sector is the $U(3)^5$ flavour symmetry, for which only flavour-conserving couplings (and by extension, chirality-conserving couplings) are allowed.

A symmetry model developed for top-quark measurements generalises the flavour structure to a $U(2)^3 = U(2)_q \times U(2)_u \times U(2)_d$ flavour symmetry on the first two generations of quarks, allowing the third to couple freely. In the lepton sector, a $U(1)^3 = U(1)_e \times U(1)_\mu \times U(1)_\tau$ flavour symmetry is introduced, imposing a $U(1)$ symmetry for each lepton generation separately, and hence allowing for lepton flavour universality violating couplings. This scheme is implemented as the `top` model in the SMEFTSIM package often used to calculate SMEFT predictions [39]. An alternative `topU31` model replaces the $U(1)^3$ symmetry of the lepton sector by the original $U(3)^2 = U(3)_\ell \times U(3)_e$ lepton flavour symmetry of the $U(3)^5$ flavour symmetry group.

1.4 The Drell-Yan spectrum at colliders

Investigations of SM behaviour have historically focused on four types of particle collisions: fixed-target experiments, lepton-antilepton collisions, hadron-(anti)hadron collisions and lepton-hadron collisions. The first and last have generally focused on probing the properties of the nuclear forces through inelastic scattering off nuclei or hadrons. Lepton colliders provide a very clean method of producing pairs of new particles on a resonant peak through lepton-antilepton annihilation, with peaks typically isolated by scanning the centre-of-mass energy s across the accessible range of the collider. Hadron colliders are arguably the most versatile category, allowing a plethora of different production processes involving the valence quarks, sea quarks, and gluons of the accelerated hadrons, but as a result tend to suffer from very complex final states with additional remnants of the initial protons.

Figure 1.1a shows an example Drell-Yan production process at hadron colliders [40], which acts as an analogue to the annihilation process at lepton colliders. A strict definition of Drell-Yan production assumes a vector current in the initial state, which enforces

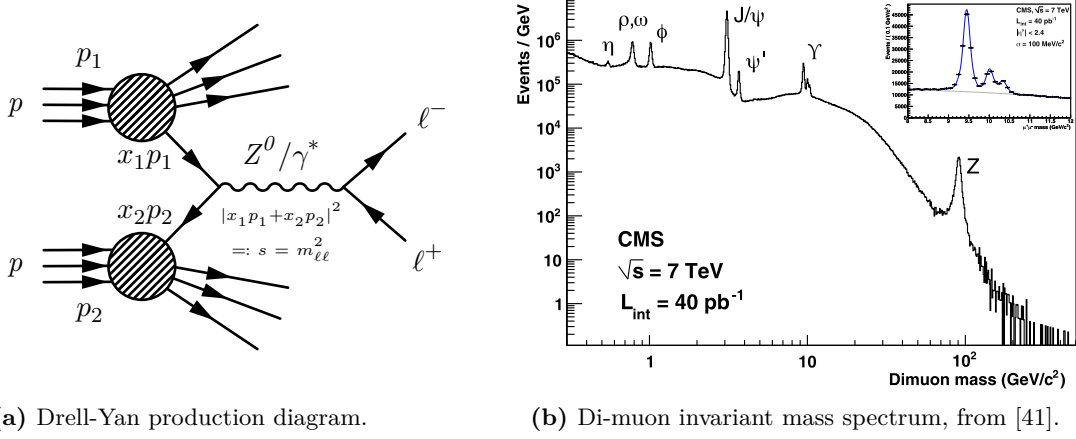


Figure 1.1: A diagram of Drell-Yan production at proton colliders, and the di-muon mass spectrum observed at the CMS experiment. The inset in (b) shows the four $\Upsilon(b\bar{b})$ resonances.

characteristic quantum numbers $J^{PC} = 1^{--}$ for the intermediate state, corresponding to a spin-1 bosonic entity that must accrue a sign flip under parity or charge-conjugation transformations. The former property implies non-trivial angular distributions in the final state, and the latter requires the intermediate state to carry no electric charge ($Q = 0$).

Drell-Yan production is predominantly measured in di-lepton final states, in which the mass of the intermediate vector state can be reconstructed as the invariant mass of the di-lepton system, an example of which is shown in Figure 1.1b. Since the initial partons will each carry a variable fraction of the total proton momentum (denoted as x_1 and x_2 in Figure 1.1a), Drell-Yan production probes both the parton distribution functions (PDFs) of the protons, and the spectrum of available vector states. The $m_{\ell\ell}$ spectrum consists of a continuum of states mediated by off-shell Z^0 and γ^* bosons, with a resonant peak at the Z^0 mass. Overlaid on this spectrum are resonant peaks for other $J^{PC} = 1^{--}$ states, namely neutral vector mesons with $q\bar{q}$ content¹⁴.

Drell-Yan s -channel production closely follows the relationship between energy and resolving power outlined by the de Broglie relation: probing shorter-range forces requires greater collision energies. Chapters 8 and 9 explore an alternative to direct resonance searches that allows some leverage against this effect: investigating interference between s -channel Drell-Yan and t -channel BSM diagrams.

¹⁴One may notice that the η meson shows in Figure 1.1b, but does not carry the correct quantum numbers to act as the intermediate vector in the Drell-Yan diagram. In this case, the resonant peak at the η mass arises from background processes.

For a $50\mu\text{m} \times 50\mu\text{m}$ segmented pixel detector of thickness $100\mu\text{m}$, a signal-to-background ratio of deposited charge carriers to thermally produced carriers of 10^{-6} can be achieved by cooling the detector below about 205K^a — a substitute to depletion?

^aIncidentally, solid CO_2 ($T_{\text{crit}} = 194.7\text{K}$ at atmospheric pressure) could be an easily sourced, safe, suitable method of cooling.

2

Charged particle tracking with silicon detectors

Contents

2.1 Interactions of charged particles with matter	20
2.1.1 Radiation length and electromagnetic processes	21
2.2 Operating principles of silicon detectors	23
2.2.1 Detector readout — hybrid and monolithic	25
2.2.2 Radiation damage	27
2.2.3 Detector resolution and performance considerations	28
2.3 Summary	31

2.1 Interactions of charged particles with matter

Particle detectors provide position, timing, or calorimetric information via interactions between incident particles and the detector material. These interactions can be of many forms, ranging from ionisation and scattering to the emission of Čerenkov and transition radiation for charged particles. Photons will generally interact via the photoelectric effect, Compton scattering, or the production of electron-positron pairs [42].

Silicon tracking detectors are designed to provide precise positional information on charged particles. This is often coupled with a strong magnetic field which induces curvature in the trajectory of particles and hence allows their momenta to be determined after passing through multiple layers of detector. They rely on ionisation within a depleted silicon substrate, in which charged particles interact with electrons, resulting in

the liberation of charge carriers within the substrate. The mean rate of energy loss from ionisation with respect to path length is well described by the Bethe-Bloch formula, given as

$$-\frac{1}{\rho} \left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}^2}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right], \quad (2.1)$$

where z , $\beta = v/c$, and γ are the charge number, fractional speed, and Lorentz factor of the interacting particle. Z , A , I , and ρ are the atomic number, atomic mass, mean excitation energy, and density of the material, respectively [43, 44]. The constant m_e is the electron mass, and $K = 4\pi N_A r_e^2 m_e c^2$ defines a constant coefficient equal to $0.307\,075 \text{ MeV cm}^2 \text{ mol}^{-1}$, where N_A is Avogadro's number, r_e is the classical electron radius, and c is the speed of light. A density correction factor $\delta(\beta\gamma)$ accounts for the flattening and extension of the particle electric field at higher energies, leading to an increase in ionisation.

For pure (undoped) silicon, the mean energy loss per path length for a minimally ionising particle (MIP, $\beta\gamma \approx 3.5$) may be calculated as $\langle \frac{dE}{dx} \rangle = 3.87 \text{ MeV cm}^{-1}$ [45]. Since the mean energy required to produce an electron-hole pair in silicon is $I_0 = 3.62 \text{ eV}$, this corresponds to sufficient energy to liberate 107 electron-hole pairs per μm path length in the medium. However, it is important to note that the processes inducing energy loss are of a statistical nature and exhibit fluctuations in either the number of ionisation processes, or the lost energy during the process. The probability density function (pdf) for the energy loss over a macroscopic distance $\Delta E/\Delta x$ follows the *Landau-Vavilov* distribution [46, 47], for which the *most probable* value is typically lower than the mean energy loss, with a tail to very large energy depositions. In practice the most probable number of charge carriers¹ produced in silicon is typically much closer to 50 to 80 per μm path length, and is dependent on the thickness of the sensor layer.

2.1.1 Radiation length and electromagnetic processes

Most purely electromagnetic processes invoking energy loss of highly energetic charged particles (including bremsstrahlung and pair production by photons) can be parametrised in terms of the radiation length of the traversed material, where the radiation

¹For the purposes of modelling or simulating detector behaviour, the most probable value is typically used rather than the mean. Fluctuations are often modelled as a Poisson process moderated by a *Fano factor* F encoding the underlying statistical fluctuations in charge deposition. For Si, typical values would be $\sigma_{e/h} = \sqrt{F N_{Eh}} \approx 0.3 \sqrt{N_{e/h}}$, with $F_{\text{Si}} \approx 0.8$ [48].

length is a characteristic path length dependent only on the material properties, which may be approximated as

$$X_0 = (716.4 \text{ g cm}^{-2}) \cdot \frac{A}{Z(Z+1) \ln \frac{287}{Z}}. \quad (2.2)$$

In the case of bremsstrahlung, X_0 defines the mean distance over which a high-energy electron will lose all but $1/e$ of its energy, and for pair production the mean free path is $7/9 \cdot X_0$. Although X_0 is often quoted in g cm^{-2} as in Equation (2.2), it may also be divided by the density of the material ρ to give the quantity in cm.

In the context of material budget for tracking detectors, the radiation length is a natural measure with which to quantify the impact on particles of electromagnetic interactions within detector components. The main electromagnetic effects of energy loss via ionisation and position resolution smearing from multiple scattering become dominant for particles at sub-GeV energies, and hence have a substantial impact on measurements of lower-energy final-state particles in ATLAS and other detectors at hadron colliders. Photon conversions in which photons pair-produce within detector material also significantly contribute to particle identification uncertainties at and above these energies. Since detectors are generally highly heterogeneous in terms of material content, the material is often quoted as a relative thickness for each component (with respect to a reference particle trajectory) in terms of radiation lengths, x/X_0 . This *fractional radiation length* is additive between discrete layers of a composite structure:

$$(x/X_0)_{\text{composite}} = \sum_{\text{components}} (x/X_0)_{\text{component}} \quad (2.3)$$

It is useful to note that a second quantity, the *nuclear interaction length* λ_I is often used as a similar measure for interactions of hadrons mediated by nuclear forces. This quantity is material dependent, but exhibits a loose relationship with the radiation length, and is used to parametrise secondary vertex production through interaction of hadrons with the detector material, which become the dominant uncertainty contribution for high-momentum particles.

2.2 Operating principles of silicon detectors

The ionisation of static lattice atoms mentioned in Section 2.1 provides a robust method for the non-destructive detection of charged particles. Charges are deposited by ionisation on traversal of the sensor layer, which can then be read out and digitised. However, thermal production of charge carrier pairs in appropriate unmodified (*intrinsic*) semiconductors (such as silicon) liberate numbers of charge carriers² ($\sim 10^{10}/\text{cm}^3$) several orders of magnitude larger than those released by ionisation ($\sim 10^4/\text{cm}$ path length) for a detection layer thin enough to enable charge traversal and readout ($\sim 100\text{ }\mu\text{m}$) [45]. For deposited charge carriers to be successfully read out, an electric field gradient must be present to induce a drift of the charge carriers through the bulk toward a collection electrode, and the bulk must be sufficiently free from thermally produced charge carriers.

Small quantities of controlled impurities can be introduced into the bulk to change the valence and conduction band structure of the semiconductor (a process generally referred to as *doping*, with the resultant material termed an *extrinsic semiconductor*). *Electron donors* (P, As, Sb; *n*-type) produce an excess of conduction electrons compared to free holes, whilst *hole donors* (B, Al, Ga, *p*-type) lead to an excess of holes. Bringing together *p*- and *n*-type layers, a thin region depleted of thermally produced charge carriers will form at the boundary. Semiconductor detectors consist of such *p-n* junctions, typically in a construction such as that shown in Figure 2.1a. The depletion region can be extended by the application of a reverse bias voltage³ across the junction, pulling the majority charge carriers out of this region and toward the implants. On passage of a charged particle, electron-hole pairs will be produced in the depleted region, which then drift along the electric field lines toward a collection electrode, as shown in Figure 2.1b. Despite the suppression of free holes in the *p*-type depleted region and electrons in the *n*-type depleted region, a small residual *leakage current* will be produced by the remaining

²Calculable from the intrinsic carrier concentration n_i predicted by the thermal equilibrium between excitation and recombination charges given the Si band gap and Boltzmann distribution: $n_i^2 = N_C N_V \cdot \exp(-E_G/k_B T)$, where $N_C = 3.05 \times 10^{19} \text{ cm}^{-3}$ is the effective density of states in the conduction band, and $N_V = 2.55 \times 10^{19} \text{ cm}^{-3}$ the effective density of states in the valence band, each at a temperature of 300 K. E_G is the Si band gap of 1.11 eV, and k_B the Boltzmann constant [45].

³Electrically, a *p-n* junction acts as a *diode*, allowing current flow in a single direction on application of a *forward voltage*. Application of a voltage against the natural direction of charge flow expands the depletion region, and to first order prevents current flow, with only small leakage currents of thermally produced charge carriers remaining.




Figure 2.1: Principle of operation of an n -in- p silicon detector. These diagrams are not to scale.

thermally-generated carrier density⁴. Auxiliary processes exciting charge carriers within the semiconductor will also influence the leakage current, such as electrons freed by the photoelectric effect from ambient light sources.

Pixel detectors segment the charge collection implants into a two-dimensional grid (as opposed to strip-type detectors, which only segment in one direction), providing positional information for the particle in three dimensions, the resolution of which is dictated by the pixel pitch and sensor thickness. The bias voltage applied to planar sensors induces a magnetic field oriented perpendicular to the sensor surface to allow charge collection at surface implants, as shown in Figure 2.1. In 3D sensors, the collection and backplane implants are constructed as vertical columns through the silicon perpendicular to the surface, and the electric field is oriented horizontally between these. Typical designs have a collection column at the centre of each pixel, and backplane implants at each corner.

ATLAS ITk pixel modules utilise n -in- p designs for both planar and 3D pixel sensors. In both cases, n^+ -type (highly-doped) implants are embedded in a weakly doped p -type substrate. This is then connected to a p^+ -type backplane, providing two dopant boundaries. The region between these can generally be fully depleted with a sufficient bias voltage, as shown in Figure 2.1.

⁴For example, pure Si doped with 10^{16} As atoms has residual carrier densities of $10^{16} e^-/\text{cm}^3$ and 10^4 holes/ cm^3 [45].

2.2.1 Detector readout — hybrid and monolithic

Traditionally, silicon-based tracking detectors have adopted a *hybrid* design, in which the implants on a sensor (built from high-resistivity silicon bulk) are connected to pads on a dedicated readout application-specific integrated circuit (ASIC)⁵ via bond wires or solder bumps⁶ [45]. This allows the sensor and readout chip to be independently optimised, with the caveat for pixel detectors that bump-bonding can be a very expensive and potentially low-yield process. The current ATLAS detector adopted such an approach for its original sensors, and this design will also be used for the ITk upgrade, as is detailed in Chapter 4.

Developments in camera and imaging technology in the past three decades have led to the widespread use of CMOS for silicon-based photosensors, allowing the integration of sensor and readout in the same chip. These have motivated the development of *monolithic* designs for pixel detectors, in which the CMOS readout circuitry is typically nestled between collection implants. Such sensors have seen success at the Solenoidal Tracker at RHIC (STAR) and A Large Ion Collider Experiment (ALICE) [49, 50]. In the latter case, a high-resistivity epitaxial layer below the implants allowed partial depletion




Figure 2.2: Schematic cross-section of an ALPIDE-style monolithic pixel detector, showing the deep p -well structures shielding the n -wells from the sensor, allowing full complementary metal-oxide semiconductor (CMOS) in-sensor. In the standard process, it is difficult to deplete the epitaxial layer over its full width. Not to scale.

⁵Usually developed in CMOS circuitry on low-resistivity silicon.

⁶Bump or *flip-chip* bonding is a back-end process for semiconductor fabrication, in which small solder bumps are reflowed between pairs of metallised pads on two silicon chips, forming an electrical interconnect.




Figure 2.3: Process modification applied to the MALTA readout chip to allow for full depletion at high bias voltages [52]. A low-dose n -type layer implant is embedded into the epitaxial layer below the n -well and deep p -well implants. This diagram is not to scale.

of this region, and deep p -wells below the n -wells allowed for full CMOS circuitry in-pixel, as is shown diagrammatically in Figure 2.2. However, the lack of full depletion reduces the radiation tolerance of these chips substantially over hybrid designs, initially limiting their use to lower-dose applications. Figure 2.3 shows a process modification allowing for full depletion developed by the monolithic from ALICE to ATLAS (MALTA) project on top of the ALICE ALPIDE sensor design [51, 52]. MALTA has produced several generations of prototypes, originally as candidate designs for the outer layers of the ATLAS ITk Pixel upgrade.

Digitisation is generally performed via amplification and discrimination stages. In the case of the RD53-type readout chips used for the ATLAS ITk and CMS tracker upgrades, a charge-sensitive amplifier (CSA) is coupled to a Krummenacher feedback circuit to discharge the sensor at a constant rate [53, 54]. This allows a discriminator to measure the *time-over-threshold* (ToT) via digital counting circuitry, which will be proportional to the deposited charge. In sophisticated readout chips such as the RD53 series, the digitised data is zero-suppressed, packaged, and serialised to high-speed electrical links, allowing for very fast efficient readout via low-voltage differential signalling (LVDS), and conversion to optical signals elsewhere within the detector, as is detailed for the case of the ITk in Section 4.3. Designs less constrained by a limited number of output links often adopt simpler systems; the original generation MALTA sensors opted against full serialisation logic in favour of reading out partially packaged data over a wide parallel bus.

2.2.2 Radiation damage

Particle detectors are subjected to large doses of both ionising and non-ionising radiation in the form of charged and neutral particles, as well as γ -rays and X-rays. Depending on the radiation flux and detector composition, this can lead to degradation of the detector performance over time [45, 55, 56].

Compared to many alternatives, silicon detectors provide outstanding radiation tolerance; however, they still deteriorate with prolonged exposure to high doses [57]. In addition, they are often placed very close to an interaction or collision point, and therefore generally experience much greater fluxes than detector layers placed further away⁷. The mechanisms for damage can generally be separated into two categories:

Damage to the silicon crystal through **non-ionising energy losses (NIEL)** is caused by interactions between incoming charged or neutral particles and the nuclei within the crystal lattice. This damage generally results in impurities in the lattice structures, which results in localised changes to the band structure of the semiconductor. Donor or acceptor atoms are removed, or new impurities that act as donors or acceptors created, thus inducing changes to the overall effective doping concentration. A common consequence is an increase in leakage currents, which can be counteracted by cooling to reduce the intrinsic carrier concentration and suppress thermal creation of free charge carriers. Measurements have shown that every reduction in temperature by 7 K approximately halves the current [58]. The effects of radiation damage can also be mitigated by increasing the bias voltage of the sensor to improve depletion and increase the drift velocity of charge carriers. For *p-in-n* sensors and extremely high radiation fluences, the changes to the effective doping concentration in the *n*-type bulk can become strong enough to invoke a *type inversion*, in which the bulk effectively becomes *p*-type, since radiation damage typically generates more *p*-type defects than *n*-type. Electrons and holes can also become trapped by localised defects and released after some time in a process known as *charge trapping*, which results in decreased signal amplitudes if the intrinsic time constant for this process is much longer than the signal formation time interval.

⁷As governed by the inverse square law.

Ionising radiation will generally induce damage at the sensor surface or at boundaries between interfaces. Charge carrier production from X- or γ -ray interactions at these boundaries can lead to excessive trapping of holes on the boundary in question, and the positive associated charge can affect the electric field structure, and can also affect the operating characteristics of transistors in CMOS circuitry. In addition, charge deposition via ionisation or pair production in CMOS logic used for digitisation and readout can lead to single-event upset (SEU) events, in which data stored within a register or memory structure or traversing a transistor are modified by the introduction of the additional charge⁸. Logic triplication is a common technique to combat such events, and involves the introduction of three identical but spatially separated copies of each logic path, with a majority voting system to decide the output. In this way, SEUs are almost entirely suppressed, as it is very improbable that two paths will be affected simultaneously.

For damage from non-ionising sources, the NIEL hypothesis posits that the damage to the sensor depends only on the total deposited energy. As a result, damage caused by protons, pions, neutrons, and electron fluences can be scaled to each other based on existing theory and experimental observations⁹. As a result, NIEL fluences are generally quoted in the damage equivalent of a flux of 1 MeV neutrons per unit area, $n_{\text{eq}}/\text{cm}^2$.

In the case of the ATLAS ITk, the choice of *n-in-p* over a more traditional *p-in-n* design helps reduce charge carrier trapping, since electrons are collected at the implants rather than the less mobile holes, and also prevents type inversion of the larger substrate from *n*-type to *p*-type. However, *n-in-p* sensors tend to be more expensive and complex to manufacture.

2.2.3 Detector resolution and performance considerations

Tracking detectors are generally designed to meet specific spatial and timing resolution requirements, which are driven by the desired vertexing precision and momentum resolution of the detector [45]. The spatial resolution of a pixel detector in a given direction is a

⁸Extreme cases can lead to single-event latch-up (SEL) behaviour, in which such events induce a low-impedance path between the power rails of the CMOS logic and can lead to a persistent or even permanent change of state.

⁹Assuming no change in the composition of this radiation flux over time.

function of the pixel pitch in this direction, and the probability and proportion of charge sharing with neighbouring pixels. Second-order dependencies exist on the resolution of the charge measurement in the case of proportional sharing, and the threshold set on the readout circuitry as this will influence charge sharing. By neglecting all effects other than pixel pitch, the resolution can be estimated as the standard deviation of the occupancy distribution [45], which is given as

$$o(x) = \begin{cases} \epsilon(x) & \text{if } x \in \left[-\frac{d}{2}, \frac{d}{2}\right], \\ 0 & \text{otherwise.} \end{cases} \quad (2.4)$$

Here $\epsilon(x)$ is the efficiency of the binary charge readout as a function of the location within the pixel. Assuming a uniform flux of particles over the pixel as well as uniform readout efficiency for all parts of the pixel, the occupancy is constant across the pixel, and the standard deviation may be given as

$$\begin{aligned} \sigma &= \sqrt{\frac{\int_{-\infty}^{\infty} x^2 \cdot o(x) \cdot dx}{\int_{-\infty}^{\infty} o(x) \cdot dx}} = \sqrt{\frac{\int_{-d/2}^{d/2} x^2 \cdot \epsilon(x)}{\int_{-d/2}^{d/2} \epsilon(x)}} \\ &= \frac{d}{\sqrt{12}} \quad \text{if } \epsilon(x) = \epsilon \text{ constant.} \end{aligned} \quad (2.5)$$

Thus, for a rectangular pixel the expected resolution in a given direction is $d/\sqrt{12}$, where d is the pixel pitch in that direction. For the purposes of deriving track resolutions and vertex resolutions, the resolution of each pixel hit can be combined in a linear or curved track fit, and then propagated to intersection points to derive the vertexing resolution.

Material interactions in layers along the track will add an additional contribution to the tracking and vertexing uncertainty. At low energies, multiple scattering will be the dominant interaction type, as is discussed in detail in Chapter 6. Consider a simple model of a two-layer cylindrical pixel detector surrounding a beam pipe at radius r_0 , with the pixel layers positioned at radii r_1 and r_2 , and uniform pixel pitches in each layer [59]. The traversal of a particle with momentum p , speed β and charge z through the beam pipe will generate a deflection due to multiple scattering approximately satisfying a Gaussian distribution with RMS $\Theta_0(p, \beta, z)$; this is traced visually in Figure 2.4a. Another deflection with RMS $\Theta_1(p, \beta, z)$ occurs through scattering in the first pixel layer, as shown

in Figure 2.4b. Now assuming Θ_0 and Θ_1 to be small, the spatial resolution at each layer σ_i can be modified to add in quadrature the contributions from scattering

$$\begin{aligned}\sigma_1 &= \sqrt{\left(\frac{d}{\sqrt{12}}\right)^2 + ((r_1 - r_0) \cdot \Theta_0)^2}, \\ \sigma_2 &= \sqrt{\left(\frac{d}{\sqrt{12}}\right)^2 + ((r_2 - r_1) \cdot \Theta_1)^2 + ((r_2 - r_0)^2 \cdot \Theta_0)^2}.\end{aligned}\quad (2.6)$$

The vertex resolution can be derived in this model by intersecting a straight track through the layers with the beam at $r = 0$, and computing the uncertainty in the vertex location. Consider a shift in each layer due to the effective resolution of the layer, as shown by the horizontal uncertainty brackets in Figure 2.4a for the first layer, and Figure 2.4c for the second. The contribution of each layer to the vertex resolution will be given by

$$\sigma_{\text{vtx},1} = \frac{r_2}{r_2 - r_1} \sigma_1, \quad \sigma_{\text{vtx},2} = \frac{r_1}{r_2 - r_1} \sigma_2. \quad (2.7)$$

Combining these in quadrature and accounting for the correlation of the terms

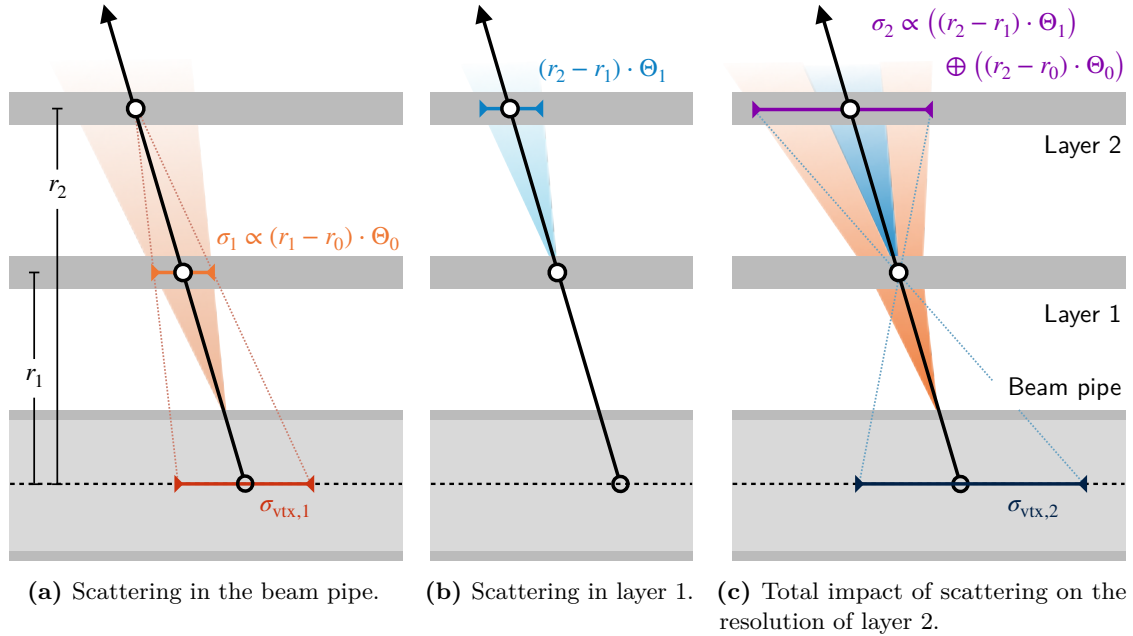


Figure 2.4: A geometric argument for contribution of scattering to the position resolution in each layer, and the impact of the resolution of each layer on the vertex resolution given in Equation (2.7).

originating from scattering at the beam pipe gives

$$\begin{aligned}
\sigma_{\text{vtx}}^2 &= \frac{r_2^2 \sigma_1^2 + r_1^2 \sigma_2^2 + 2r_2 r_1 (r_1 - r_0)(r_2 - r_0) \cdot \Theta_0^2}{(r_2 - r_1)^2}, \\
&= \frac{r_2^2 + r_1^2}{(r_2 - r_1)^2} \cdot \left(\frac{d}{\sqrt{12}} \right)^2 + \frac{(r_0(r_1 + r_2) - 2r_1 r_2)^2}{(r_2 - r_1)^2} \cdot \Theta_0 + r_1 \cdot \Theta_1, \\
&\approx \left(1 + \frac{r_1^2}{(r_2 - r_1)^2} \right) \cdot \left(\frac{d}{\sqrt{12}} \right)^2 + (2r_1 - r_0) \cdot \Theta_0 + r_1 \cdot \Theta_1.
\end{aligned} \tag{2.8}$$

The approximation in the last line has been made by assuming $r_2 \gg r_1$. The first term of the final line shows that alongside the intrinsic resolution derived from the pixel pitch, a second contribution is generated that depends on the detector geometry, and can be suppressed by increasing the lever arm $r_2 - r_1$. The contribution from scattering in the beam pipe is dependent on the distance between the beam pipe and first layer $r_1 - r_0$, motivating the placement of the first layer as close to the beampipe as possible. Scattering in the first pixel layer is coupled to a factor r_1 , and can also be suppressed by reducing the distance to the first tracking layer. However, in most cases a reduction of the absolute distance to the first layer from the beam axis will be more difficult to implement than changes to the lever arm, pixel pitch, or beam pipe placement. As a result, minimising the material within each layer (and hence the x/X_0 of each) is critical for optimising vertex resolution. The material-dependent terms in Equation (2.6) will also propagate to curvature measurements in a track fit, allowing similar arguments to be made for a dependence of momentum resolution on material content.

2.3 Summary

Modern tracking detectors have come to rely on semiconductor technology in the forms of pixel and strip-type detectors for the non-destructive detection of charged particles. These technologies offer unprecedented spatial resolution whilst achieving good radiation hardness, and allowing for low material content with careful design. In this particular context, five parameters are central to their performance: detection efficiency, spatial resolution, data readout rates, geometry (number of layers and inter-layer spacing), and material content.

It is well-known trivia around CERN that CMS would sink ($\rho_{\text{CMS}} \sim 3.9 \text{ g cm}^{-3}$) whilst ATLAS would float ($\rho_{\text{ATLAS}} \sim 0.31 \text{ g cm}^{-3}$), but one barely ever considers the implications of the medium. In the case of mercury flooding, both would sink, and in the somehow even more unlikely case of catastrophic liquid helium leaks at a large enough scale to fill the cavern, both would float. Luckily for us, both are entirely non-buoyant in their true ambient medium of air.

3

From collisions to ATLAS data

Contents

3.1 Probing rare phenomena in proton collisions at the LHC . .	32
3.1.1 Kinematics of proton-proton collisions	34
3.1.2 Luminosity and pileup	35
3.1.3 The role of triggers	36
3.2 The observables of interest to physics	37
3.2.1 Momentum and charge — the tracking detectors	38
3.2.2 Energy — calorimetry	40
3.2.3 The muon system	42
3.3 From observables to physics objects	43
3.3.1 Hadrons and other jets	43
3.3.2 Jet structure for b -quark identification	44
3.3.3 Identifying hadronically decaying τ -leptons	45

The Large Hadron Collider (LHC) represents the current energy frontier within the field of particle physics, regularly setting new records for collision energy during its runtime [60]. This chapter introduces the acceleration and collision of protons at the LHC in Section 3.1, and the detection of collision products in the general-purpose ATLAS detector is discussed in Section 3.2. Section 3.3 discusses the object reconstruction methods most relevant to the ATLAS analysis described in Chapter 9.

3.1 Probing rare phenomena in proton collisions at the LHC

The LHC is a circular hadron collider of circumference 26.7 km, consisting of two counterdirectional synchrotron rings sharing common magnetic elements [60]. Both

the dipole magnets used for beam bending and the quadrupole magnets used for focussing utilise superconducting NbTi coils operated at a temperature of 1.9 K and capable of producing magnetic fields of up to 8.33 T. Bunched proton beams are provided by a multi-stage injector chain shown in Figure 3.1, and are accelerated by a radiofrequency (RF) acceleration system up to a maximum design energy of 7 TeV in each counterrotating direction. Collisions occur at four interaction points, at which the clockwise and counterclockwise travelling beams are collimated and passed through each other to induce proton-proton (pp) collisions. ATLAS is located at *Point 1*, which together with Point 5 for CMS defines the two high-intensity interaction points available at the LHC capable of providing nominal instantaneous luminosities exceeding $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

Points 2 and 8 host the ALICE and LHCb experiments, respectively, and results from

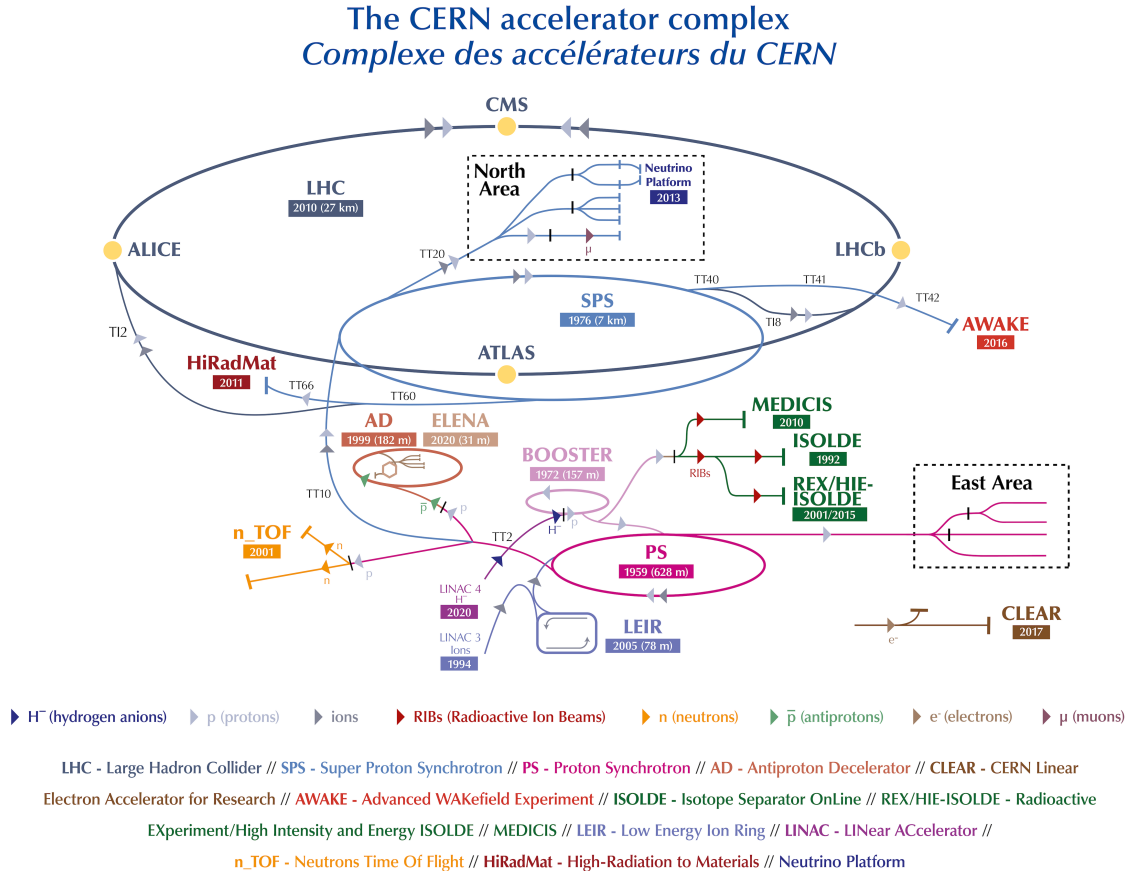


Figure 3.1: Layout of the accelerator complex at CERN. Proton beams from the Linear Accelerator (LINAC) 4 are accelerated in several circular stages including BOOSTER and the Proton Synchrotron (PS) and Super Proton Synchrotron (SPS) prior to injection in the LHC. Prior to 2020, the LINAC 2 initial acceleration stage was used in place of the current LINAC 4. Reproduced from [61].

the latter are relevant in Chapter 8. The East and North Area dedicated experimental halls contain secondary targets off the PS and SPS accelerators, respectively, which are often used for testbeam experiments. The T9 beamline in the East Area was used for the measurements discussed in Chapter 7.

The transverse structure of the LHC beam is controlled by quadrupole magnets, utilising strong focusing techniques to ensure beam stability and reduce the transverse phase-space of the bunches. The longitudinal structure is determined by the 400.789 MHz frequency of the RF cavities, which results in short bunches with a root-mean-square (RMS) bunch length of 0.25 ns, or about 75 mm. Only every tenth RF cycle is populated with a bunch, resulting in the characteristic *bunch-crossing frequency* of 40 MHz. The associated time constant of 25 ns defines the shortest inter-event time window for the LHC detectors.

3.1.1 Kinematics of proton-proton collisions

Objects in ATLAS events are typically defined relative to a right-handed cylindrical reference system, with the z -axis oriented along the beam pipe with origin at the interaction point. The azimuthal angle ϕ is defined counter-clockwise about the z -axis, with $\phi = 0$ defined as pointing to the centre of the LHC ring. The zenith angle θ is defined in the plane parallel to the z -axis¹.

Since protons are composite objects and collisions occur between constituent partons, the z -momentum of any given parton is not well-constrained, and proton remnants typically continue to travel along the beam axis and are not detected. As a result, observed momenta only sum to zero in the plane transverse to the z -axis. The transverse momentum p_T is hence defined as

$$\vec{p}_T = \vec{p} - p_z \hat{z} \qquad p_T = \sqrt{|\vec{p}|^2 - p_z^2}, \qquad (3.1)$$

and the p_T sum over all final state particles in a collision must give zero. This fact is often used to indicate the presence of one or more undetectable particles (typically neutrinos or very weakly interacting BSM particles) as deficits in the p_T balance — a

¹The Cartesian x -axis faces along $\phi = 0$, and y -axis points upward following the right-handedness of the system, but the Cartesian axes are typically only used for low-level objects or in special cases.

concept referred to as *missing transverse energy* (\cancel{E}_T). It is also useful to redefine the rapidity y in terms of the zenith angle θ , as

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \stackrel{m/E \rightarrow 0}{\approx} \eta = -\ln \left(\tan \frac{\theta}{2} \right), \quad (3.2)$$

where the rapidity y simplifies to the pseudorapidity η for massless or relativistic particles², and rapidity intervals are invariant under Lorentz boosts.

Since the momentum configurations of the colliding partons are not well known, it is typical to define pp interaction cross-sections in terms of parton distribution functions (PDFs) $f_i(x, \mu_F)$ that parametrise the probability of a parton of type i carrying a fraction x of the total proton momentum p , at a characteristic energy set by the factorisation scale μ_F . Typically, only one parton per proton will interact, and hence the inclusive cross-section for a process $\sigma_{pp \rightarrow X}$ may then be given in terms of the parton-level cross-sections $\sigma_{ij \rightarrow X}$ as

$$\sigma_{pp \rightarrow X} = \sum_{ij} \int dx_1 dx_2 f_i(x_1; \mu_F) f_j(x_2; \mu_F) \cdot \sigma_{ij \rightarrow X}(x_1 p_1, x_2 p_2, \mu_F, \mu_R). \quad (3.3)$$

The initial proton momenta are labelled p_1, p_2 , with each imparting a momentum fraction x_1, x_2 to the donated parton, respectively. The parton indices i, j are summed over all allowable combinations of initial partons, and the parton-level cross-section is renormalised at the renormalisation scale μ_R . The dependencies on μ_F and μ_R are determined by the renormalisation scheme used for the PDFs and cross-section, respectively.

As a result, the amount of available energy within the collision is variable, and is bounded from above by the centre-of-mass energy s :

$$s = |p_1 + p_2|^2 c^2 > |p_1 x_1 + p_2 x_2|^2 c^2 = E_{\text{avail}}^2. \quad (3.4)$$

3.1.2 Luminosity and pileup

The yield of a particular scattering reaction during collisions is defined in terms of the instantaneous luminosity L and the cross-section σ , which can be defined either as a rate or integrated over a data-taking period:

$$R = \frac{dN}{dt} = \sigma \cdot L \quad N = \sigma \cdot \mathcal{L}_{\text{int}}, \quad (3.5)$$

²Useful inclination angles to know are $\eta = 0, 1, 2, 3, 4 \leftrightarrow \theta = 90^\circ, 41^\circ, 15.4^\circ, 5.7^\circ, 2.1^\circ$.

where the rate R depends on the instantaneous luminosity L , and the total number N is integrated over a data taking period with integrated luminosity $\mathcal{L}_{\text{int}} = \int dt \cdot L$. The instantaneous luminosity is related to the collider parameters by

$$L = \frac{N_1 N_2 f N_b}{4\pi \sigma_x \sigma_y}, \quad (3.6)$$

in collisions between N_b bunches containing N_1 protons circulating at a frequency f in one direction, and N_2 in the other. Both beams are assumed to have equal widths of σ_x and σ_y in the x - and y -direction, respectively. Run 2 of the LHC achieved instantaneous luminosities up to $2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ with a centre-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$, and the total dataset collected from 2015 to 2018 corresponded to an integrated luminosity of $(140.0 \pm 1.2) \text{ fb}^{-1}$ [62].

A vast majority of pp collision events result in soft inelastic interactions. In order to maximise the proportion of the hard scattering events of interest, L is tuned to produce several interactions per bunch-crossing. A hard scatter vertex will hence be accompanied by several soft scatter vertices within the same event termed *pile-up*, the products of which must be rejected by reconstruction algorithms. Interaction vertices will typically occur as an approximately Gaussian spread in z with $\sigma_Z \approx 50 \text{ mm}$ [63], and are separated by precise tracking to identify the originating vertex for reconstructed objects. The mean pile-up $\langle \mu \rangle$ is proportional to L , and will vary during an LHC fill as the number of surviving protons per bunch (N_1, N_2) reduces over time. The mean pile-up has been steadily increased over the lifetime of ATLAS, with the initial ~ 20 collisions per bunch-crossing in Run 1 increasing to 40 to 50 in Run 2, and typically around 60 for Run 3 [64].

3.1.3 The role of triggers

The large proportion of uninteresting soft scatter events motivates careful selection of retained data. Additionally, for most subdetectors, the bunch-crossing rate of 40 MHz far exceeds the rate of events that could reasonably be read out and sent to long-term storage. As a result, ATLAS implements a two-level trigger system [65]. A subset of calorimeter and muon system data is sent to a hardware trigger system, where initial p_T and energy cuts are applied to reduce the event rate to below 100 kHz. These events are passed to the software-based High-Level Trigger (HLT), which performs online object

reconstruction, and applies identification and kinematic criteria to further reduce the event rate to approximately 1 kHz. After this, the selected events are sent to storage if the triggers belong to the *main physics* stream, or to various in-situ calibration and data quality tasks in the case of the calibration/monitoring and express streams.

3.2 The observables of interest to physics

Particles produced in collisions pass through several detector layers designed to measure direction, momentum, or energy. These measurements are performed by dedicated detectors using optimised technologies, and the data produced is used to reconstruct *observables* — quantities that can be used as estimators for the true kinematic properties of the particle.

Most of the subdetectors within ATLAS are structured as a series of concentric layers, each with one or more cylindrical *barrel* layers and a pair of *endcaps*, providing near-hermetic coverage surrounding the interaction point [66]. The layers of the detector are shown in Figure 3.2, and are discussed in the following sections.




Figure 3.2: Illustration of the ATLAS detector, with a cutaway showing the detector layers. Reproduced from Ref. [66].

3.2.1 Momentum and charge — the tracking detectors

The Inner Detector (ID) is designed to provide location information for traversing charged particles, using non-destructive sensing technologies, and is shown in Figure 3.3 [69, 70]. The innermost layers utilise silicon pixel detectors for optimal position resolution, followed by layers of microstrip detectors that allow for greater area coverage with a smaller channel count. The silicon layers provide a total of at least eight high-resolution measurement points per track for $|\eta| < 2.5$ across all polar angles. These layers are enclosed in a straw-tube detector termed the Transition Radiation Tracker (TRT), which provides approximately 36 additional lower-precision points per track. Surrounding this, a 4.5 cm thick Nb-Ti superconducting solenoidal magnet system produces a magnetic field of up to 2 T [66], which bends the tracks of charged particles passing through the ID according to the equation of cyclotron motion

$$R = \frac{p_T}{qB} = \frac{p \cos \theta}{qB}. \quad (3.7)$$

Here, R is the radius of the circular path taken by a particle of electric charge q traversing at an angle of inclination θ to a magnetic field of strength B , hence allowing a measurement of the momentum p (or alternatively, the transverse momentum p_T).




Figure 3.3: Cutaway of the layers of the ATLAS ID, showing the beam pipe, Pixel detector layers (including the IBL), SCT layers, and the TRT. The diagram is annotated with radial distances from the interaction region within the beam pipe. Reproduced from Refs. [67, 68].

The Pixel detector consists of four layers, the outermost three of which utilise sensors with a pixel pitch of $50\,\mu\text{m} \times 400\,\mu\text{m}$, read out via a dedicated ASIC termed the FEI [69–71]. In the central z -region, barrel layers hold modules perpendicular to the radial r direction, and in the forward regions circular endcap disks hold modules parallel to r and perpendicular to z . In both cases, the long edges of pixels are oriented parallel to the z -axis. The innermost layer was installed during the first long shutdown of the LHC, and is referred to as the Insertable B-Layer (IBL) [72]. It consists of a single barrel layer at a radius of 33 mm from the interaction point, and uses $50\,\mu\text{m} \times 250\,\mu\text{m}$ pixels and an updated ASIC called the FEI-4b.

The Semiconductor Tracker (SCT) utilises long strips with a pitch of $80\,\mu\text{m}$, and comprises eight barrel layers and nine disks per endcap [69, 70]. Barrel strips are oriented in an equivalent arrangement to the pixel layers, whilst endcap strips are oriented radially. Strips of neighbouring layers are offset by a stereo angle of 40 mrad to give spatial resolving power in both directions, resulting in resolutions of $17\,\mu\text{m}$ across the strip pitch, and $580\,\mu\text{m}$ along the strip length.

The Transition Radiation Tracker (TRT) is composed of 73 layers of straw drift tubes of radius 4 mm filled with either Argon- or Xenon-based gas mixtures [69, 70, 73]. Barrel straws are oriented along the z -axis within the region $|\eta| < 0.7$, and each endcap orients straws radially and extends coverage to $|\eta| < 2$. Low-density hydrocarbon materials are added to induce *transition radiation* in the form of X-rays as fast charged particles (especially electrons) traverse the boundary between media with different dielectric constants. This transition radiation is absorbed by the straw tubes, and generally deposits a larger signal. Since transition radiation is predominantly emitted by electrons rather than charged hadrons, the larger signals can be used for particle identification³.

Tracks are reconstructed from all available hits, typically using a Kalman filter algorithm [74]. Two variants exist: *inside-out tracking* seeds tracks from triplets of Pixel

³However, due to Xenon gas leaks, the gas mixture has recently been replaced by pure Argon, leading to a large decrease in identification efficiency.

and/or SCT hits, whilst *outside-in tracking* seeds with muon system or TRT hits [75]. A combinatorial Kalman filter is applied, testing potential compatible hits and branching where required. Once track candidates are built, a final fit is applied to determine track parameters, and tracks are ranked and filtered based on quality criteria to resolve shared hits. The curvature parameters of the track are used to determine the momentum [66], which in ATLAS typically achieves a resolution of

$$\frac{\sigma_{p_T}}{p_T} = 0.05\%/\text{GeV} \cdot p_T \oplus 1\%. \quad (3.8)$$

Important track parameters include the momentum (both in the transverse plane and globally), the transverse impact parameter d_0 defined as the minimum transverse distance between the track and the z -axis, and the longitudinal impact parameter z_0 defined as the z -distance of this closest approach to the detector centre. Together with the azimuthal (ϕ_0) and polar (θ_0) angles of the track at closest approach, these parameters form a complete description of any given track, and the impact parameters are particularly important for the identification of secondary vertices.

3.2.2 Energy — calorimetry

ATLAS integrates two types of sampling calorimeter for destructive energy measurements of charged and neutral particles, providing hermetic coverage in ϕ up to $|\eta| < 4.9$ [76]. Sampling calorimeters intersperse passive *showering* layers with active ionisation- or scintillation-based detectors which sample the shower depth, shape, and deposited energy [77]. Important metrics include the total depth of the calorimeter d (in radiation lengths X_0), the relative proportions of active sampling layers to inactive material f_{samp} , and the thickness of the sampling layers x . A general form for the energy resolution of a calorimeter can be given as

$$\frac{\sigma_E}{E} = \frac{A}{\sqrt{E}} \oplus \frac{B}{E} \oplus C, \quad (3.9)$$

where A is a stochastic term related to the statistical fluctuations of the measured signal and the sampling scheme, B includes noise effects from the readout system, and C includes any constant systematic effects on the measurement such as uncertainties related to digitisation, mechanical or density variations, or incorrect calibration. For

sampling calorimeters, A is approximately proportional to $\sqrt{x/(f_{\text{samp}}Z)}$ where Z is the atomic number of the inactive material, since the ratio f_{samp} predicts the fraction of the total energy which is actually deposited within the active elements⁴. Using showering layers with higher Z results in greater particle multiplicity during showering, hence improving the statistical power when sampling.

The Electromagnetic Calorimeter (ECal) combines lead showering layers with Liquid Argon (LAr) drift chamber layers in which ionisation charges are collected on electrodes placed within the chambers [78]. An electric field is applied between the electrodes and showering layers. The calorimeter barrel is arranged in an *accordion* geometry to ensure hermetic coverage in ϕ , whilst the endcaps use radially oriented plates. The total depth d is at least $22 X_0$ across the covered solid angle, and a separation into three (barrel, $|\eta| < 2.5$) or two (endcap, $|\eta| > 2.5$) layers allows shower direction information to be extracted⁵. The first layer is segmented for measurements of shower shape that help to discriminate between photons, neutral pions, and converted photons. For the electromagnetic part of the Forward Calorimeter (FCal) ($|\eta| > 3.1$), the lead absorber is replaced by copper for the innermost layer. The energy resolution is $\frac{\sigma_E}{E} \approx \frac{10\%}{\sqrt{E/\text{GeV}}} \oplus 0.7\%$ in the barrel and endcap regions, and $\frac{\sigma_E}{E} \approx \frac{28.5\%}{\sqrt{E/\text{GeV}}} \oplus 0.035\%$ in the forward region [79, 80].

The Hadronic Calorimeter (HCal) layers are designed to induced showers from charged and neutral hadrons that survive the ECal. For the endcaps, a similar structure to the ECal is used, with copper showering layers interspersed between active LAr elements. The outer two layers of the FCal use tungsten as the showering material. Instead of LAr, the barrel region ($|\eta| < 1.7$) uses a scintillating tile calorimeter, in which plastic scintillators are layered between steel absorbers and read out via photomultiplier tubes coupled via optic fibres [81]. The HCal exhibits an overall resolution of $\frac{\sigma_E}{E} \approx \frac{50\%}{\sqrt{E/\text{GeV}}} \oplus 3\%$ for the barrel and endcap, and $\frac{\sigma_E}{E} \approx \frac{100\%}{\sqrt{E/\text{GeV}}} \oplus 10\%$ in the forward region [79, 80].

⁴The sampling fraction f_{samp} is typically much smaller in hadronic calorimeters due to the larger fraction of invisible components in the hadronic shower compared to a purely electromagnetic shower, leading to much larger uncertainties arising from the sampling term.

⁵This information is useful to determine the propagation direction of photons.

3.2.3 The muon system

The calorimeter layers of the ATLAS detector are surrounded by an eight-coil superconducting toroid magnet in the barrel region [82], and two identically segmented smaller toroids in the forward regions [83]. Both magnets are air-core to reduce scattering, and produce a non-uniform ~ 0.5 T to 2 T magnetic field in the barrel region, and ~ 1 T to 3 T field for the endcaps.

Muon Drift Tube chambers (MDTs) provide precision tracking throughout the barrel and endcap regions through single-wire cylindrical drift tubes with a diameter of 30 mm and spatial resolution of $80\,\mu\text{m}$ ⁶, arranged in several layers [84]. In the forward region with very high particle rates ($> 1\,\text{kHz}/\text{cm}^2$), the drift time is significantly longer than the mean interval between hits, resulting in pile-up of MDT hits. To combat this, Cathode Strip Chambers (CSCs) are introduced, which are multi-wire proportional chambers with a time resolution of $\sim 7\,\text{ns}$ and per-wire spatial resolution of $60\,\mu\text{m}$ that can tolerate rates up to $10\,\text{kHz}/\text{cm}^2$.

Resistive Plate Chambers (RPCs) and **Thin-Gap Chambers (TGCs)** provide fast readout that is used for triggering ATLAS on muon detections. The barrel uses three layers of RPCs, which have a spatial resolution of $\sim 1\,\text{cm}$ and a time resolution of 1 ns. On a multi-layer coincidence of hits between layers, muon transverse momenta estimated from the track inclination are used to form p_{T} -dependent trigger signals. The endcaps include three to four layers of TGCs, which are thin multi-wire proportional chambers with a 2 mm to 3 mm spatial resolution and $\sim 2\,\text{ns}$ time resolution.

The muon momentum resolution achieved varies with muon p_{T} . For low-energy muons, energy loss fluctuations in the calorimeters can affect p_{T} resolution, but this issue can be mitigated by combining track information from both the ID and muon systems, if available. Typical p_{T} resolutions vary from 2% to 3% at 50 GeV and exceed 10% above 1 TeV [66, 85].

⁶The spatial resolution is generated by precise measurement of the drift time.

3.3 From observables to physics objects

ATLAS analyses utilise a wide array of algorithms to reconstruct particle-equivalent *physics objects* from detector signals. A small subset relevant to the analysis presented in Chapter 9 are briefly discussed here⁷.

3.3.1 Hadrons and other jets

Due to the confinement property of the strong force, non-colour-singlet states are only free at high energy scales where the strong coupling α_s is reduced due to the running of the coupling from renormalisation. In hadron collisions, this leads to free quarks *hadronising* and producing jets of approximately collinear energetic hadrons. Jets typically produce large numbers of tracks in the ID and one or more large clustered depositions in the HCal.

In ATLAS, jets are typically reconstructed using the anti- k_T algorithm [87, 88], which clusters calorimeter deposits from largest to smallest p_T using the distance parameters

$$d_{ij} = \min(d_i, d_j) \frac{\Delta R_{ij}^2}{R^2}, \quad d_i = \frac{1}{p_{Ti}^2}, \quad (3.10)$$

where the angular distance is defined as $\Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$. The jet radius R defines the maximal angular size of the jet cone. Objects i and j where $d_{ij} < d_i$ are clustered together in the order defined above. If $d_{ij} > d_i$, the object i is labelled as a complete jet, and removed from further consideration.

Jet reconstruction relies on several corrections and calibrations [89]. The *jet energy scale* (*JES*) is calibrated against dedicated in-situ measurements, using a well-calibrated object as a reference [90, 91]. For low-momentum jets ($20 \text{ GeV} < p_T < 500 \text{ GeV}$), Z +jet events are used in which the jet recoils against a Z boson decaying to leptons, allowing the jet momentum to be computed from the p_T balance. For intermediate momenta up to 900 GeV , γ +jet events are utilised. For very high p_T , the JES must be computed from multi-jet events containing a high- p_T jet recoiling against lower- p_T jets. The *jet energy resolution* (*JER*) is extracted from a Gaussian fit to the corrected p_T . Typical systematic uncertainties on the JES calibration are around 4.5% at 20 GeV , and reduce to $\mathcal{O}(1\%)$ above 300 GeV for multi-jet events with two jets in the final state.

⁷The algorithms here are described as implemented in *Release 21* of the ATLAS ATHENA framework [86].

Pile-up vertices produce large numbers of jets, which are identified and removed using a dedicated jet vertex tagger (JVT) algorithm [92, 93]. A discriminant is built from information related to tracks contained within the jet. This algorithm has been calibrated for several working points with particular efficiencies and purities.

3.3.2 Jet structure for b -quark identification

The classification of jets by their originating parton is referred to as *flavour tagging*, and in ATLAS jets are typically categorised into b -jets, c -jets, light-jets (u , d , s quarks, and gluons), and jets originating from hadronically decaying τ leptons [94]. Non-stable initial partons with moderate lifetimes (specifically, B meson and D meson states originating from b and c quarks, respectively, and hadronic decays of τ leptons) will produce secondary vertices displaced from the primary interaction as they decay. These can be identified by increased impact parameters z_0 and d_0 compared to light jets originating from the primary vertex, or through dedicated track coincidence algorithms to find secondary vertices. Since such jets will hadronise into mesons prior to the weak decays that induce flavour changes, jet properties including shape, depth, track density, and potential tertiary vertices can give useful information on the topological structure of the hadronic decay chain, which can be correlated to the jet categories.

In ATLAS, the outputs of a series of lower-level algorithms are fed to a single high-level classifier [95, 96]. These inputs include the jet kinematics (p_T and $|\eta|$), likelihoods computed from the impact parameters for each quark category, and predictions from a low-level recurrent neural network (RNN) fed with impact-parameter variables. In addition, combinations of kinematic variables of secondary vertices identified by a dedicated algorithm termed SV1 [97] and information on the displaced vertex chain computed using the JETFITTER algorithm [98] are used. A deep neural network (DNN) consumes these input variables to produce multi-class discriminants for quark-type jets, which are typically arranged into log-likelihood ratios. The b -tagging discriminant is built as

$$\text{DL1R} = \log \left(\frac{p_b}{f_c \cdot p_c + (1 - f_c) \cdot p_{\text{light}}} \right) \quad (3.11)$$

where p_b , p_c , and p_{light} are the normalised DNN output scores, and f_c is a tunable parameter affecting c -jet rejection.

For such algorithms, several *working points* are defined based on the discriminant, which result in particular efficiencies (ranging from 60% to 80%) and purities on a calibration dataset. At the 77% working point used in Chapter 9, a light-jet rejection factor of 170 and c -jet rejection factor of 5 have been measured in simulated $t\bar{t}$ events. The working points are calibrated in datasets enriched with each jet type: for b -jets, $t\bar{t}$ samples are used in a tag-and-probe method, light jets are sourced from Z +jets events, and c -jets from complex varied samples [99, 100]. Data-to-simulation scale factors are derived as a function of jet momentum, which are then applied to simulated events used in analyses.

3.3.3 Identifying hadronically decaying τ -leptons

Tau leptons decay via the weak nuclear force, either to light leptons (e, μ) with a branching fraction of 35%, or to quarks with a branching fraction of 65%. In the former case, two neutrino emissions carry away some fraction of the τ lepton momentum, resulting in a significant \cancel{E}_T . However, such decays are nearly indistinguishable from prompt lepton production, and are not generally separated but rather treated in an inclusive selection.

Hadronically decaying τ leptons produce a jet alongside a single neutrino emission. The comparatively low τ lepton mass combined with the available quantum numbers results in two possible decay topologies: *one-prong* ($\tau^\pm \rightarrow \nu_\tau \pi^\pm \dots$ or $\tau^\pm \rightarrow \nu_\tau K^\pm \dots$) or *three-prong* ($\tau^\pm \rightarrow \nu_\tau \pi^\pm \pi^\pm \pi^\mp \dots$)⁸, categorised by the number of charged daughter particles. Additional neutral mesons (typically pions) can occur in either category.

In ATLAS, hadronic tau candidates τ_{had} are seeded by anti- k_T jets with radius $R = 0.4$ [101, 102], which are required to have $p_T > 5$ GeV and $|\eta| < 2.5$ due to a reliance on accurate track information. Tracks are associated with a candidate jet if they sit within $\Delta R < 0.25$ of the jet axis and carry $p_T > 1$ GeV, and contain at least two pixel hits and seven hits total across the ID. The distance of closest approach to the associated vertex must satisfy $d_0^v < 1$ mm and $z_0^v \cdot \sin \theta < 1.5$ mm.

An RNN processes track-level, cluster-level and high-level inputs, combining these to form a single discriminant to reject hadronic jets [101]. Four working points are defined, which are shown alongside efficiencies and background rejection factors in Figure 3.4a. An


⁸The prevalence of these decay modes is linked to meson resonances matching the quantum numbers of the τ lepton, such as $\tau^- \rightarrow \nu_\tau \rho^- \rightarrow \nu_\tau \pi^- \pi^0$ and $\tau^- \rightarrow \nu_\tau a_1^- \rightarrow \nu_\tau \pi^- \pi^+ \pi^-$.

additional boosted decision tree (BDT) has been trained to improve the rejection rate of jets originating from electrons [102]. These are discriminated from 1-prong τ_{had} candidates using shower shape information and calorimetric variables for the jet, as well as tracking information. The nominal working point corresponds to a τ lepton selection efficiency of 95%, with the characteristic behaviour shown in Figure 3.4b. Another dedicated BDT takes as input kinematic variables for track and calorimeter objects associated to the τ_{had} candidate, and categorises each as one of five decay modes:

$$\underbrace{h^\pm, \quad h^\pm\pi^0, \quad h^\pm(\geq 2\pi^0)}_{1\text{-prong}}, \quad \underbrace{h^\pm h^\pm h^\mp, \quad \text{and } h^\pm h^\pm h^\mp(\geq 1\pi^0)}_{3\text{-prong}}, \quad (3.12)$$

where h^\pm stands for π^\pm or K^\pm .

The tau energy scale (TES) is computed from the corresponding JES for cluster depositions, with additional corrections to account for expected pile-up depositions within the tau jet cone using a calibration function extracted from simulated $Z/\gamma^* \rightarrow \tau^+\tau^-$ events [102]. The TES is refined using a boosted regression tree (BRT) trained on the calculated TES value as well as other variables [101]. The typical energy resolution ranges from 4.5% to 6% for 1-prong τ_{had} candidates, and 4% to 5% for 3-prong candidates.



(a) Identification RNN performance on 1/3-prong τ_{had} , from Ref. [101].

(b) Electron veto BDT performance, from Ref. [102].

Figure 3.4: Performance of τ_{had} reconstruction algorithms: the identification RNN (a), and the electron veto BDT (b). In both cases, the background rejection rate is shown as a function of signal selection efficiency, and in (a), the markers denote the four available working points: from left to right, *very loose*, *loose*, *medium*, *tight*. In (b), the response is given in ranges of $|\eta^{\text{trk}}|$, which refers to the core track of the 1-prong τ_{had} candidate. In (a), the background rejection is measured on simulated di-jet events, whilst (b) utilises simulated $Z \rightarrow e^+e^-$ events.

The ITk Pixel detector is guaranteed to survive at least 31 000 years of operation in geostationary orbit. At least in terms of radiation flux...

— The ITk Pixel TDR [103], *in as many words...*

4

The HL-LHC and ATLAS ITk upgrades

Contents

4.1 The case for higher luminosity	48
4.1.1 Technical challenges	49
4.2 The ATLAS Phase-2 upgrade	51
4.3 ITkPix: readout for a new pixel detector	56
4.3.1 The ITkPix readout chip	56
4.3.2 On-chip power distribution	58
4.3.3 Injection circuitry and electrical testing	59
4.4 OEC support structures and services	62
4.4.1 Monitoring	62
4.4.2 Interlock	63
4.4.3 The optosystem	65
4.4.4 Outer endcap support structures	67

Following the discovery of the Higgs boson at the LHC in 2012, both the ATLAS and CMS collaborations have pivoted from a focus on discovering the Higgs boson to an extensive programme of measurements of its properties. With the confirmation of a Higgs mass at 125 GeV, CERN has decided to prioritise luminosity upgrades for the LHC rather than targeting substantial increases in energy. The High-Luminosity LHC (HL-LHC) program will turn the LHC into a veritable Higgs factory, delivering 3000 fb^{-1} , $10\times$ more data than Runs 1–3 combined, whilst still operating largely within the existing LHC magnet and detector infrastructure. To achieve this, substantial upgrades will be required to parts of the beam infrastructure and the experiments to prepare for the increase in luminosity.

This chapter will begin by considering the case for higher luminosity in the context of the ATLAS detector, and the technical challenges involved. An overview of the Phase-2 upgrade plans will be presented, followed by a detailed description of the Inner Tracker (ITk) Pixel detector, with a particular focus on the Outer Endcap (OEC) subsystem and relevant support structures and services.

4.1 The case for higher luminosity

Precision measurements of the Higgs boson parameters will be a key part of the HL-LHC physics programme at both ATLAS and CMS. Extrapolations of current analysis results to the full HL-LHC dataset prepared for the 2026 European Strategy for Particle Physics have shown that the increased total luminosity will allow statistically limited precision measurements to profit greatly [104]. Several rare processes including the $H \rightarrow \mu^+ \mu^-$ and $H \rightarrow Z\gamma$ decay modes and di-Higgs production will reach a 5σ or greater significance¹ in combination with CMS results. Measurements of Higgs boson couplings at ATLAS are expected to reach a precision of $\mathcal{O}(20\%)$ for the $H \rightarrow Z\gamma$ coupling, and $\mathcal{O}(10\%)$ or better for the other accessible coupling terms. The combined ATLAS+CMS measurement precision on the Higgs boson trilinear coupling κ_3 is expected to reach 30%, and additionally will constrain the result to a single unique minimum within the Higgs potential².

Beyond Higgs physics, the HL-LHC programme represents a $\sim 10\times$ increase in yield for any statistically-limited SM measurement or BSM search, and will play a large role in terms of measuring parameters of and setting limits on TeV-scale phenomena within and beyond the SM. Specific targets include the charm Yukawa coupling and very rare processes such as four-top-quark production and longitudinally polarised vector boson scattering ($W_L W_L$).

Beyond the physics reach, the HL-LHC also represents a large step forward in terms of magnet and detector technology, and in the latter case the technical challenges associated with the high-luminosity environment have driven the development of higher-resolution sensor designs and more performant data readout.

¹A 5σ significance is the generally accepted threshold for *discovery* of a new particle or process, corresponding to a ~ 1 in 2 million chance that the result is due to random statistical fluctuations.

²Current measurements of κ_3 find several local minima in the likelihood function [105].

4.1.1 Technical challenges

The higher data rate will pose a very challenging environment for the detectors at the HL-LHC due to the increased peak instantaneous luminosity of $(5 \text{ to } 7.5) \times 10^{-34} \text{ cm}^{-2} \text{ s}^{-1}$, and average pile-up $\langle \mu \rangle$ of 140 to 200 pp interactions per bunch-crossing. These increases pose two main challenges: higher radiation doses, and more stringent requirements on tracking and vertexing resolution to ensure good pile-up rejection. Both are most critical very close to the interaction point, and the ITk upgrade program will build a completely new inner detector system designed to meet these requirements, which have been summarised in Table 4.1.

The increased radiation levels require the innermost tracking layers to be radiation hard to a NIEL fluence of $2 \times 10^{16} \text{ n}_{\text{eq}}/\text{cm}^2$, and to 10 MGy of total ionising dose (TID), exceeding the expected total dose for the ID by more than an order of magnitude³. This has necessitated the development of many new radiation-hard components for the current inner detector, both for the detector front-ends, and for the services and monitoring systems installed within the inner detector volume. The expected total dose corresponding to a total integrated luminosity of 3000 fb^{-1} is visualised in Figure 4.1.

The much higher pile-up necessitates both a greater density of spatial measurements to build clean tracks in a much busier tracking environment, and a better vertexing

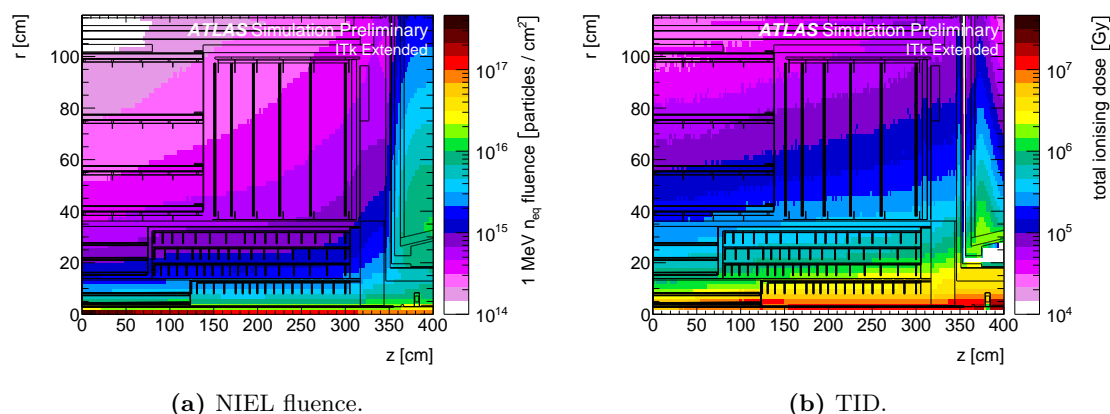


Figure 4.1: Map of the expected total radiation dose of the ITk detector in an $r-z$ quadrant, corresponding to an integrated luminosity of 3000 fb^{-1} , from Ref. [111].

³As of mid-2024, the ID has experienced a fluence of $8.5 \times 10^{14} \text{ n}_{\text{eq}}/\text{cm}^2$ [106]. This has corresponded to a reduction in signal response of about 25% at the end of Run 2, expected to increase to around 50% before the end of Run 3.

Table 4.1: Expected conditions for the ITk detector assuming 3000 fb^{-1} of integrated luminosity at $\langle\mu\rangle = 200$, and corresponding performance requirements [103, 107–109]. The material estimates quoted here are those given in the ITk Pixel and Strip technical design reports (TDRs) [103, 110].

<i>Expected conditions</i>			
Exp. radiation dose	Pixel	10 MGy	$2 \times 10^{16} \text{ n}_{\text{eq}}/\text{cm}^2$
	Strip	333 kGy	$8.1 \times 10^{14} \text{ n}_{\text{eq}}/\text{cm}^2$
	Optosystem	50 kGy	$1 \times 10^{14} \text{ n}_{\text{eq}}/\text{cm}^2$
Maximum occupancy (sim. $t\bar{t}$ events, $\langle\mu\rangle = 200$)	Pixel	0.16%	247 hits/chip
	Strip	0.23%	310 hits/chip
<i>Required performance</i>			
Track reconstruction efficiency (at 10 GeV)		$> 99\%$ (μ)	$> 85\%$ (e, π)
Fake rate			$< 10^{-5}$
Max. noise occupancy per channel	Pixel		1×10^{-3}
	Strip		1×10^{-3}
Max. acceptable channel loss	Pixel		15%
	Strip		10%
Readout rate			1 MHz
Material budget	Pixel	$< 2\% X_0 / \text{layer}$	$0.67\% X_0 / \text{module}$
	Strip	$< 2\% X_0 / \text{layer}$	0.55% to 0.65%
			X_0 / module

resolution to allow the discrimination of pile-up vertices from the main vertex within the $\sim 2 \text{ cm}$ long interaction region. The ID achieves a resolution on the transverse impact parameter d_0 of about $10 \mu\text{m}$, and a longitudinal impact parameter z_0 resolution of about $70 \mu\text{m}$. Recent simulations of the ITk layout performed with $\langle\mu\rangle = 200$ show a 60% improvement in d_0 resolution and a 200% improvement in z_0 resolution for 200 GeV muons, and similar performance to the ID for 2 GeV muons [112]. However, despite these improvements, the primary vertex reconstruction and selection efficiency exhibit an overall reduction from the $> 98\%$ achieved in Run 3 to about 92.5% expected for the ITk, largely driven by increased pile-up.

The improvements in vertexing resolution have been achieved by a substantial reduction in the pixel and strip pitch compared to the ID, and the replacement of the TRT. The

increase in the number of channels and the greater occupancy of the upgraded detector will necessitate updates to the Trigger and Data Acquisition (TDAQ) system, and carefully designed readout chains to handle the higher data rates. In addition, the higher luminosity will correspond to an increase in the rate of “interesting” events for any given statistically-limited trigger channel, and an increase in the event acceptance rate of the software trigger from 100 kHz to 1 MHz is planned.

4.2 The ATLAS Phase-2 upgrade

To meet the challenges posed by the HL-LHC, the ATLAS collaboration is undertaking several upgrades that will ensure the detector can meet the physics goals outlined above.

The current Pixel Detector, SCT and TRT will be replaced by a new all-silicon tracking system that has been termed the Inner Tracker (ITk) [107]. It will consist of two new detector subsystems: a silicon microstrip detector for the outer 4 layers, and a pixel detector for the inner 5 layers, in the configuration shown in Figure 4.2 [113, 115]. The ITk will also increase tracking coverage from its current limit of $|\eta| = 2.5$ to $|\eta| = 4$, and will ensure a minimum of 13 hits per track in the forward regions of the detector, and 9 in the barrel. The strip and pixel detectors will cover 168 m^2 and 13 m^2 of active area, respectively. Both detectors will improve considerably on their current counterparts with respect to tracking resolution and data rates. Alongside the tracking detectors, the ITk volume will also host an upgraded Beam Condition Monitor (BCM') [116], and a

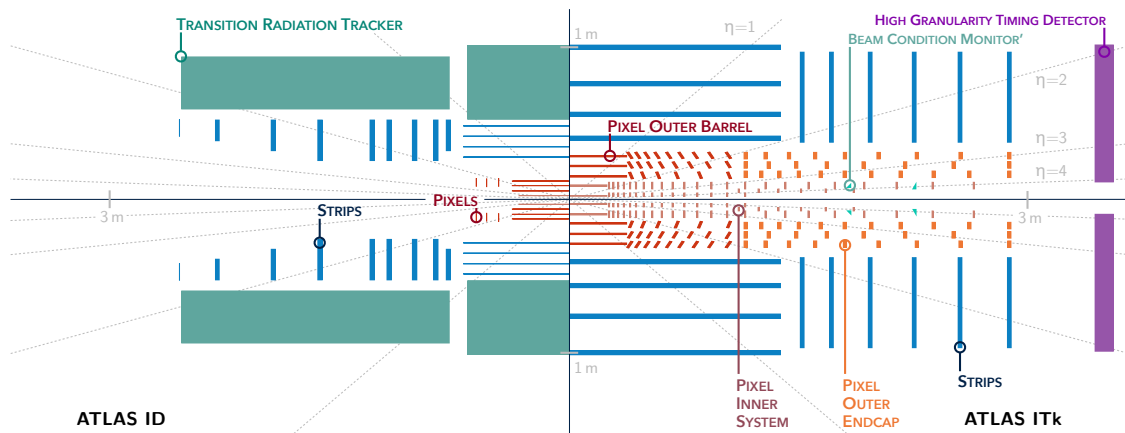


Figure 4.2: Cross-sectional layout of the ITk Pixel (red) and Strip (blue) detectors (right), and comparison to the ID (left) [113, 114]. The Pixel subsystems are labelled, and the Beam Condition Monitor (BCM') and High Granularity Timing Detector (HGTD) shown in relation to the ITk.

High Granularity Timing Detector (HGTD) at the edge of the volume in the forward regions to disentangle pile-up using timing information [117].

ITk Pixel detector

The new pixel detector is split into three sub-systems which utilise common front-end ASICs and readout chains, but differ in sensor choices, support structures, and services [103]. In all subsystems, ITkPix readout chips are bump-bonded to sensors, and then assembled into modules, as shown in Figure 4.3. The readout chips are powered in parallel within each module, and modules are combined into serial power (SP) chains in which they are powered in series. Dedicated circuitry has been added within the ITkPix chips to regularise their current draw and allow them to operate in this mode, which is discussed in detail in Section 4.3. Data are transferred electrically from all sub-systems to the optosystem positioned outside the ITk detector volume⁴, where they are converted to optical signals that are then sent to Front-End Link Exchange (FELIX) cards hosted in the service cavern. Bi-phase CO₂ cooling⁵ is provided from a cooling plant in the service cavern to all parts of the inner detector and the optosystem.




Figure 4.3: Annotated render of an ITkPix quad module. The cutout exposes the readout chips, bump bonds, and sensor behind the flexible PCB (flex) printed circuit board (PCB) used for interconnect. Charges deposited in the sensor are transferred to the readout chip via the bump bonds. After digitisation, packaged data are sent through the wire bonds and flex PCB, which also provide power and command to the readout chip.

⁴In the electromagnetic calorimeter gap, at $R = 1450$ mm and $z = 3500$ mm.

⁵A cooling system in which the pressure is adjusted to keep the coolant just below the liquid-gas phase transition at the desired temperature, at which point absorbed heat from the cooled components induces a phase transition rather than an increase in the coolant temperature.

The *Inner System (IS)* encompasses the innermost two layers of the pixel detector, which include both longerons as part of a cylindrical barrel, and rings that make up the endcaps. The rings and layer 0 longerons are populated with triplet modules consisting of three ITkPix chips bump-bonded to individual sensors and glued to a single flexible PCB acting as an interconnect for power and data. The layer 0 longeron uses a 150 μm -thick 3D sensor with pixels of size $100\,\mu\text{m} \times 25\,\mu\text{m}$, whilst the layer 0 and layer 0.5 rings use a similar sensor with $50\,\mu\text{m} \times 50\,\mu\text{m}$ pixels. The layer 1 longerons and rings are populated with quad modules consisting of 2×2 ITkPix chips bonded to a single thinned 100 μm n-in-p planar sensor. Modules are loaded on thin carbon supports, with integrated titanium piping allowing a flow of CO_2 for cooling. The inner system uses a complex data and power interconnect scheme, with the patch panel 0 (PP0) interconnect PCBs⁶ generally sitting near the outer radius of the system, and thin flex cables connecting to the data and power connectors on the modules.

The *Outer Barrel (OB)* and *Outer Endcap (OEC)* share a common quad module and sensor design, a render of which is shown in Figure 4.3. A 150 μm -thick n-in-p sensor containing a 800×768 matrix of $50\,\mu\text{m} \times 50\,\mu\text{m}$ pixels is bump-bonded to four ITkPix chips. In the OB longerons, individual modules are glued to a high-thermal-conductive graphite *cell*⁷, which is in turn attached with screws to a lightweight carbon frame. This improves the reworkability of the loaded structure during construction. A small circular metal implant in the centre of the cell is welded to a titanium cooling pipe, and the side of the carbon frame facing away from the interaction point hosts a PP0 PCB that provides both data and power to the modules. Lightweight flex cables connect the module data and power adapters to the PP0. The longerons are then assembled into a barrel structure. The OB also encompasses a series of *inclined half-rings*, which fill a transition region between

⁶The ATLAS services scheme consists of a series of *patch panels* placed at the edges of the local support, sub-detector, detector, and experimental cavern. Figure 4.8 shows the positions of these for the ITk pixel detector.

⁷In general, the process of attaching modules to supports is here termed *loading*, with the populated support then referred to as a *loaded support*.

the cylindrical barrel and vertical end-caps, and hold loaded cells at an angle perpendicular to a straight particle trajectory emanating from the interaction point.

In the OEC, modules are glued to a lightweight carbon fibre half-ring with an embedded titanium cooling pipe. The supports and services chain are described in detail in Section 4.4.4.

ITk Strip detector

The new strip system will consist of four barrel layers, and six endcap discs per endcap [110]. The barrel is composed of staves hosting 14 strip modules per side, and the endcap of petals supporting 9 modules each. Each module combines a 320 μm -thick n-in-p float-zone strip sensor with a number of ASICs — the ABCStar performs strip readout, whilst the HCCStar acts as a data concentrator and the AMACStar manages power and environmental monitoring. The sensors are produced in eight different shapes — two for the barrel with long and short strips, and six for the endcaps with roughly trapezoidal shapes. The ASICs are hosted on a flex that is glued to the sensor, and the conversion to optical data is performed at the stave- or petal-level by an auxiliary board hosting gigabit transceivers and optical drivers. In addition to the flex, a power board hosting a DC-DC converter and the AMACStar chip is glued to the sensor, with a flexible polyimide interposer recently added to this interface. Modules are then glued to a polyimide *Bus Tape*, which is co-cured with pre-impregnated partially-cured carbon fibres (pre-preg) that provide mechanical stiffness to the stave or petal.

Beam Condition Monitor (BCM')

An upgrade to the current Beam Condition Monitor (BCM), the BCM' system will monitor background beam activity within the ITk volume, and will be installed within the retractable IS of the pixel detector [116]. Alongside its protective function, BCM' will also act as an additional luminosity monitor for ATLAS.

High Granularity Timing Detector (HGTD)

A new set of timing layers will be installed at the border of the ITk volume in the endcap regions, which will provide < 70 ps timing resolution on tracks passing

through the forward region which is most susceptible to the increase in pile-up tracks ($2.4 < |\eta| < 4$) [117, 118].

Alongside the replacement of the inner detector, several upgrades to other detector systems are planned to ensure their continued performance in the high-luminosity environment.

TDAQ/FELIX

The much larger data rate from the sub-detectors expected during Run 4 and outlined in Figure 4.4 necessitates improvements to the trigger and data acquisition pipelines within ATLAS [119, 120]. By the beginning of Run 4, all subdetectors within ATLAS will have been upgraded to use the centrally-designed Front-End Link Exchange (FELIX) system for readout [121]. FELIX is a PCI-express card hosting a single field-programmable gate array (FPGA), up to 48 optical links for front-end communication, and a custom Trigger, Timing and Control (TTC) interface to relay these signals to each sub-detector. Front-end electronics will send data fragments via local FELIX boards over the network to the Data Handler, which performs data collection and may perform detector-specific preparation.

The hardware trigger will continue to rely on the identification of calorimeter objects and muons within localised regions known as regions of interest (ROIs). In addition, a new system known as the Global Trigger will be added to perform quasi-offline reconstruction to refine and combine the trigger object definitions, allowing exploitation of the event topology. The software trigger will be upgraded to run at an event-building rate of 1 MHz (~ 4.6 TB/s), and will no longer include regional readout and reconstruction components currently implemented in the HLT.




Figure 4.4: Run 3 ATLAS TDAQ operating conditions and requirements, and the expected changes for the HL-LHC upgrade [119].

It will make use of offline-like algorithms for reconstruction and selection, and will ensure that current trigger efficiencies and purities can be maintained or exceeded despite the higher pileup and event rates in the high-luminosity environment. The output event filter rate is expected to increase to 10 kHz (~ 46 GB/s), a $5\times$ increase over the current 2 kHz rate of events sent for data storage.

Calorimetry and Muon systems

New on-detector and off-detector electronics are being developed for the LAr and Tile calorimeters to enable full 40 MHz continuous readout and integration into the FELIX ecosystem, and to improve the radiation hardness of the electronics in use [122–125]. Similar upgrades are planned for the trigger and readout electronics of the Muon system, and additional layers of small-diameter Muon Drift Tube chambers (sMDTs), RPCs, and TGCs will be added to extend coverage to $|\eta| = 4$ [126, 127].

4.3 ITkPix: readout for a new pixel detector

In the early stages of R&D for the ITk upgrade, several potential readout schemes were developed, including a hybrid design led by the RD53 Collaboration [128, 129], and monolithic candidates such as the MALTA [130] and ATLASPix [131] designs that were initially considered for the OB and OEC. The hybrid design by the RD53 Collaboration has been adopted for the Phase 2 upgrades of both ATLAS and CMS, and the ATLAS variant of this chip is named ITkPix.

4.3.1 The ITkPix readout chip

The ITkPix ASIC supports a matrix of 400×384 square pixels of pitch $50\,\mu\text{m}$. The charge digitisation logic utilises a *differential* architecture, which couples a differential pre-amplification stage with continuous reset and adjustable gain to a differential pre-discrimination stage, as shown in Figure 4.5. In this second stage, two globally distributed voltages $V_{th,1}$ and $V_{th,2}$ introduce an offset on the two branches of the differential signal by producing additional gain before the final discriminator. The pre-amplified signal is discriminated in the third stage, with a 4-bit pixel-level DAC in form of a resistor ladder controlling the effective threshold. An additional bit chooses which branch in

the discriminator is adjusted, giving an effective 5-bit adjustment overall. A leakage current compensation (LCC) circuit has been added to the first amplification stage, which can be globally enabled or disabled. The differential architecture was chosen by ATLAS for the final design of the ITkPix readout chip.

The discriminated signals are consumed by clocked logic operating at the level of an 8×2 pixel *quarter core*, which measures ToT with respect to the beam clock, and packages data to be sent to the serialised output. Trigger commands are accepted within the chip periphery, and data streams are built from data sent by each core with a hit in the relevant bunch-crossing, including a quarter core address, an 8×2 pixel hit-map⁸, and ToT information for all hit pixels. Serialised data are output from the periphery via up to four 1.28 Gbit/s high-speed links using the proprietary Xilinx Aurora 64B/66B protocol [132]. Trigger commands and register read/write commands used to set and query DACs and operational registers are sent to the chip via a 160 Mbit/s differential command link with a custom DC-balanced encoding. Data frames are interspersed with regular *service frames* which return register information or monitoring values as requested via the command link.

The chip design implements several data inputs and a *data-merging* functionality that allow a master chip to merge serialised data from other (slave) chips into its data stream,




Figure 4.5: Schematic of the differential front-end in use in the ITkPix ASIC designs, from Ref. [128]. Global configuration digital-to-analog converters (DACs) are shown as **bold names**, whilst per-pixel configuration DACs are *italicised and underlined*.

⁸These are implemented as compressed hit-trees, as is described in detail in Ref. [129].

reducing the required number of output links for lower-occupancy parts of the detector.

Several submissions of the chip design were produced, including v1 (which showed problematic power draw), v1.1 (which showed some potential SEU sensitivities in simulations, corrupted/stuck ToT values, and issues with data-merging from other chips), and the final v2 that will be used for detector construction.

4.3.2 On-chip power distribution

Both the RD53A and ITkPix chips are designed for use in a serial powering chain, allowing for a reduction in cabling and detector material by providing power to a chain of modules connected in series from a single off-detector power supply channel. To ensure correct operation of modules in such an arrangement, the current draw of the readout chips must remain constant over time, to ensure consistent power balancing between readout chips. In order to achieve this, the ITkPix readout chips implement a specialised power distribution circuit termed a shunt low-dropout regulator (SLDO) [133], the operating concept of which is shown in Figure 4.6. The SLDO combines a linear low-dropout voltage regulator with a current shunt element that is activated in inverse proportionality to the current draw of the output load. This ensures a constant current consumption for each readout chip despite the typical fast variations in power draw due to switching currents in the digital logic. When modules are chained in series, this ensures the relative voltage drops across modules remain constant.

Focusing on the revised circuitry in the ITkPix versions, some improvements were made with respect to previous iterations to the power-on sequence. Rather than relying




Figure 4.6: Block diagram showing the operation of an SLDO as a linear regulator with constant current consumption, from Ref. [129]. As the current draw of the load I_L varies, the effective resistance of (and hence current I_s passed through) the shunt is adjusted to ensure $I = I_L + I_s$ remains constant.

on reference voltages generated by the SLDO itself, a secondary linear regulator is used to power the main reference current generator. This reference current I_{ref} has a nominal value of $20\text{ }\mu\text{A}$, and is used to derive the other reference currents and voltages required by the chip. Since the I_{ref} is generated by the internal band gap of the readout chip, wire bond settings can be used to make small adjustments to account for process variations. From this reference current, two reference voltages $V_{\text{ref,D}}$ and $V_{\text{ref,A}}$ are generated that are used to regulate the voltage output of the LDOs for the digital and analogue domains of the readout chip, respectively. The generation of these voltages can also be adjusted by on-chip DACs to account for process variations affecting the SLDO response.

4.3.3 Injection circuitry and electrical testing

The ITkPix design includes capacitor-based charge injection circuitry for self-testing and calibration of the pixel matrix. Two DACs generate boundary voltages $V_{\text{cal}}^{\text{hi}}$ and $V_{\text{cal}}^{\text{med}}$, which are used to charge a local capacitor for each pixel with a nominal capacitance of 8.02 fF . The injected charge is hence proportional to the difference between the two voltages, which is typically termed ΔV_{cal} . Calibration injections can be either *digital*, in which case the injection is performed directly into the hit processing logic of the pixel, bypassing the analogue front-end; or *analogue*, in which case the injection undergoes full amplification and discrimination.

Routine electrical testing of modules and larger-scale structures can be split into two general categories — measurements performed on the module and its components by probing from outside, and measurements that rely on the internal calibration circuitry and data readout. A very brief summary of the typical external tests used as part of module or loaded structure quality control checks is given below:

Sensor $I - V$ measurements ensure the correct response of the sensor under varying reverse biases, and are designed to catch early breakdown and high leakage currents.

SLDO $V - I$ measurements characterise the SLDO response and are used to determine the turn-on point and verify ohmic SLDO behaviour beyond this point.

Thermal cycling of modules or supports ensures there are no latent weaknesses or stresses in the bump-bonds and glue interfaces that may cause delamination or interface failure during operation.

The remaining tests rely on correct operation of and communication to the readout chips on the module in question. For single modules, this is generally performed in a relatively simple lab setup involving desk-top power supplies, water and/or Peltier cooling through a cold chuck or plate, and adapter boards from the module data and power flex cables to traditional lab cable standards⁹. Command encoding and data decoding is handled by an FPGA on a PCI-Express card¹⁰ hosting customised firmware, which then interfaces to the Yet Another Rapid Readout (YARR) software [134]. A range of calibration scans is performed with YARR:

Register read scans are used to verify basic command distribution and chip communication by sending and reading back register values using the service frames.

Eye diagram calibrations use the internal bit error rate test (BERT) infrastructure of the ITkPix chips to measure bit errors whilst varying the phase and pre-emphasis characteristics of the data drivers. An equivalent procedure can be performed using the BERT capabilities of optoboards when running with the full optical readout.

Digital scans perform digital injections in each pixel and measure the occupancy to test functionality of the hit readout. Pixels with non-unit occupancy are generally masked for further tests.

Analogue scans perform analogue injections in each pixel and measure occupancy to test the amplification and discrimination circuitry. Again, failing pixels are masked.

Noise scans send empty triggers without injections to characterise the latent noise distribution of the chip and sensor. Pixels above a certain threshold are masked.

⁹In this particular case, DisplayPort (DP) for data, ferrule-terminated banana leads for low/readout chip voltage (LV), and LEMO cables for high/bias voltage (HV).

¹⁰The typical lab setup uses a commercial Trenz TEF1001 FPGA with a custom “Ohio” breakout board that breaks out 20 of the gigabit transceivers on the FPGA from the on-board FPGA mezzanine card (FMC) connector to 4 DP connectors.




Figure 4.7: Idealised expected output of a threshold scan performed for a single pixel, with Gaussian noise. The value t gives the point at which V_{th} is applying an effective threshold equal to the injected charge I_{inj} , and n gives a measure of the discriminator noise¹¹.

Threshold measurements measure the per-pixel occupancy as a function of threshold for a given injected charge using analogue injections. As shown in Figure 4.7, this produces an s-shaped transfer function¹² for each pixel, the central value of which defines the point where the pixel threshold equals the injected charge, and the width of which gives some measure of the electrical noise at the discriminator.

Threshold equalisation involves a sequence of threshold measurements performed for varying per-pixel trim values, adjusting these to equalise pixel-level thresholds¹³.

Source scans or equivalently X-ray tests are often the first exposure modules have to real particles rather than injected charges. Generally, a β -emitter (such as ^{90}Sr or ^{60}Co) or keV-scale X-rays will be used, and either “random” triggering¹⁴ is employed, or the internal *hit bus* of the chip may be utilised for self-triggering. Occupancy substantially below the mean is usually an indicator of delaminated bump bonds for the affected pixels. Source scans are often combined with thermal cycling.

Crosstalk/Bump connectivity scans inject large charges in a pixel, and measure the occupancy of neighbouring pixels when doing so. The large injected charge will feed

¹¹In an ideal discriminator with no noise, the threshold is a δ -function and the occupancy forms a step function.

¹²Under the assumption of Gaussian noise in the discriminator response, this curve is the cumulative density function (CDF) of a Gaussian. Depending on the readout software, either a full fit may be performed on the response curve per-pixel (more computationally expensive) to extract the mean and width, or a crossing-point and gradient-based approximation may be used to estimate these quantities.

¹³Although a linear scan is perfectly suitable, faster implementations use a per-pixel binary search, reducing the maximal number of iterations from N to $\lceil \log_2(N) \rceil$, where N is the range of the trim setting.

¹⁴Barely ever truly random, but generally just implemented as evenly time-spaced.

back into the sensor volume of the injected pixel, and will induce small *cross-talk* charges in the sensor volume of neighbouring pixels. In the case of broken bump bonds in the injecting pixel no cross-talk will occur in any other pixels, whereas neighbouring pixels with broken bump bonds will not show cross-talk in any case.¹⁵

Module components, modules, and loaded structures will in general be electrically tested after each stage of construction (such as module assembly or local support loading), and also before if a transport has occurred since the last tests were performed (this is generally termed *reception testing*).

4.4 OEC support structures and services

The Pixel detector services are shared between the three subsystems, with some differences in support structures and routing for each, as has been summarised previously in Section 4.2. A general overview is shown in Figure 4.8.

4.4.1 Monitoring

The Monitoring Of Pixel System (MOPS) is a network of radiation-hard, low-power ASICs performing localised 12-bit analogue-to-digital conversion for up to 34 voltage

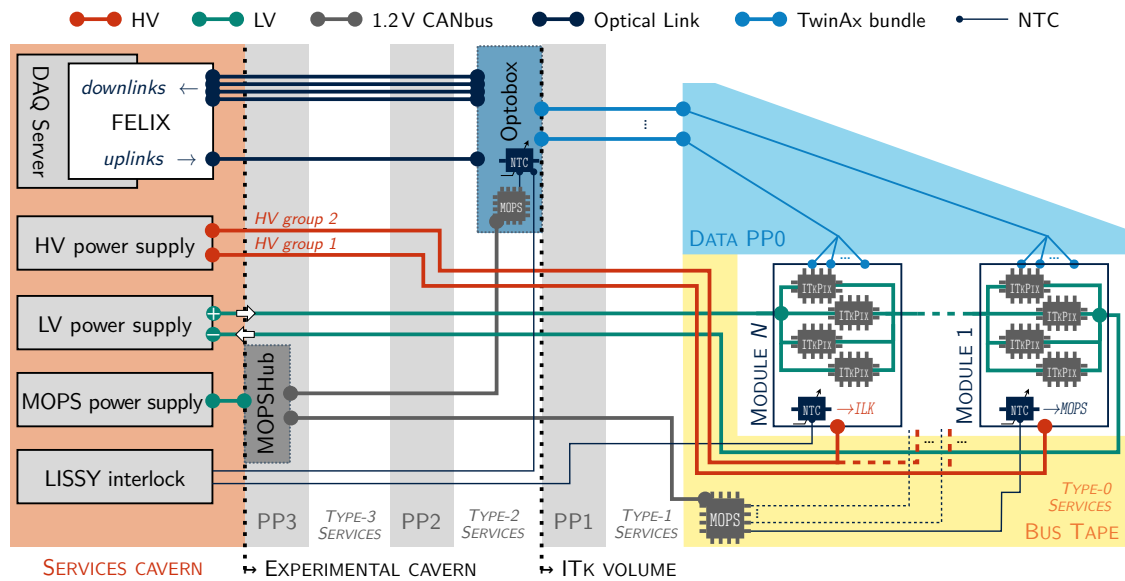


Figure 4.8: Diagram showing the services and interfaces for the ITk OEC subdetector.

¹⁵ An alternative form of bump connectivity scan has been proposed by CMS, and involves applying a small forward bias on the sensor, increasing the leakage current and sensor noise from thermal fluctuations. Pixels that see no increase in noise threshold are then assumed to have broken bump bonds.

input channels [135, 136]. These chips are placed on local supports and in the optoboxes of the pixel system, and are used to monitor temperatures by measuring the resistance of negative temperature coefficient thermistors (NTCs) using voltage dividers, and voltages by direct probing.

MOPS chips can be assigned a 2-bit address that follows the controller area network (CAN) standard [137], allowing up to four MOPS chips to be combined on a single bus. In the final detector, a dedicated readout crate termed a MOPShub will collate the data from up to 36 buses (144 MOPS chips, or up to 4896 channels) using an onboard FPGA coupled to a dual-core processor, and also provide the power line for each bus via a DC-DC converter [138]. The crate is designed to host up to 18 CAN interface cards (CICs), each of which perform the necessary level conversion for two CAN buses. It forms part of the patch panel 3 (PP3) interface at the boundary of the experimental cavern.

A stripped-down version of the MOPShub crate, aptly named the MOPShub for beginners (MH4B), has been made available for loading sites and system tests for readout of up to 8 CAN buses. The FPGA has been replaced by a Raspberry Pi, and can host up to four CICs, reading these out using the native CAN interface support in the Linux operating system hosted on the Raspberry Pi.

4.4.2 Interlock

An off-detector interlock system is necessary to ensure fast reactions to operating conditions that could endanger the detector. The detector safety and interlock system for the ITk is split into three main tiers.

The Local Interlock and Safety System (LISSY) combines local temperature data in the form of analogue NTC inputs from the pixel and strip modules and the optosystem, through modular units termed Temperature Input (T2I) cards installed in a multi-slot interlock crate [139], shown in Figure 4.9. Additional digital inputs summarising the state of the local and global environments and beam are accepted from a connected Main Interlock Crate (MIC) via a digital input/output (I/O) card — the Global Safety System (GSS) card — of which there is typically one per LISSY crate. Digital output cards route output signals to the power supplies corresponding




Figure 4.9: A prototype LISSY crate used for the ITk Pixel system test. The various card types are labelled above, and example signal designations are given below.

to the monitored SP chains, with varying number of output channels depending on the power supply channel multiplicity. The interlock logic is implemented in the Interlock FPGA (ILK-FPGA) card, which receives binary inputs from the T2I and GSS cards, and routes outputs to the digital output cards.

The Main Interlock Crate (MIC) distributes summary signals coming from the Detector Safety System (DSS) and Beam Interface System (BIS) to all LISSY crates within a particular detector domain.

The Detector Safety System (DSS) assumes responsibility for the environmental stability of the detector volume, detector cavern, and services cavern [140]. DSS *alarms* are raised in response to environmental factors such as ambient temperature and humidity, rack power distribution, water leak sensors, human safety, and many others. These alarms will be reflected in the state of summary signals that are propagated by the MICs to the individual LISSY domains, which are defined such as to cause suitable responses at the local interlock level.

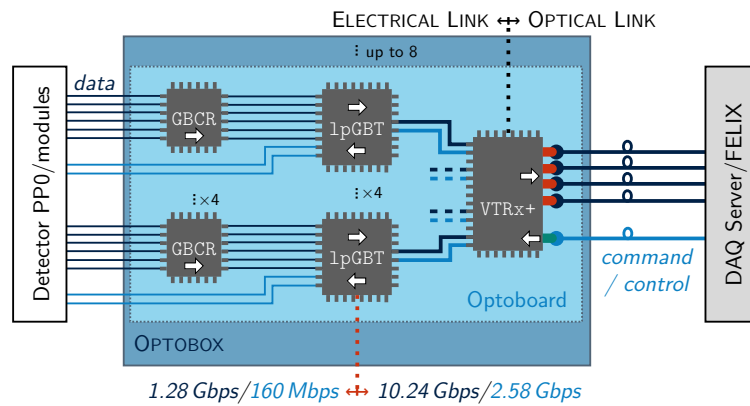
Table 4.2: Summary of total channels and data output for the ATLAS ID and ITk [103, 107, 114, 141, 142].

		channels	bandwidth / link	max. links / module
ATLAS ID	IBL	$\sim 1 \times 10^8$	160 Mbit/s	2
	Pixel	$\sim 1 \times 10^8$	40 Mbit/s	16
	SCT	$\sim 6 \times 10^6$	40 Mbit/s	2
ATLAS ITk	Pixels	$> 1 \times 10^9$	1.28 Gbit/s	16
	Strips	$\sim 6 \times 10^7$	640 Mbit/s	1

4.4.3 The optosystem

The ITk represents an order-of-magnitude change in the number of readout channels compared to the ID, as is summarised in Table 4.2. In order to handle the much greater data output without increasing cabling and corresponding detector material, both very fast electrical link speeds and early conversion to optical data are necessary. The optosystem is designed to meet these needs for the ITk Pixel detector [108], using several radiation-hard CERN- or ATLAS-designed custom ASICs mounted on a modular transceiver unit termed an *optoboard*:

- Four **Gigabit Cable Receivers (GBCRs)** [144] perform signal recovery and equalisation of up to six 1.28 Gbit/s digital signals sent from the ITkPix chips.
- Four **Low Power Gigabit Transceivers (lpGBTs)** [145] serialise data from six links fed by the GBCR output into a single 10.24 Gbit/s link. Additionally, a

**Figure 4.10:** Optosystem components as connected within the optoboard [143].

2.56 Gbit/s downlink is deserialised into two 160 Mbit/s links for command and clock that are sent to the attached ITkPix chips, with the remaining links used within the Optoboard for the internal slow control (IC) protocol.

- A single **Versatile Link Plus Transceiver (VTRx+)** [146] performs the conversion to optical for four 10.24 Gbit/s uplinks and one 2.56 Gbit/s downlink.
- A DC-DC converter translates the input 2.5 V power rail (which is used as-is for the VTRx+) to 1.2 V for the GBCR and lpGBT.
- An ERF-8 connector defines the electrical connection to the detector. Custom adapter PCBs will be used to connect the optoboards to the various flavours of Twinaxial (twinax) bundles coming from the PP0s.

In the final detector, up to eight optoboards will be combined in a single optobox, with a *power board* PCB that steps down five 9 V power domains to 2.5 V, and also hosts a MOPS chip for monitoring the optoboard voltages and temperatures. A *connector board* PCB splits the five voltage domains in a configurable way¹⁶ between the optoboards. Two NTCs, one hosted on the connector board and one on the power board, provide interlock feedback on the state of the optobox.

Sets of 28 optoboxes will be hosted in dedicated large-scale *optopanel*s situated at the end of the ECal barrel in the ATLAS endplate, and will provide cooling and services routing for the optosystem. For system tests and loaded local support quality control (QC) testing, prototype optobox enclosures have been developed that provide the same services for a single optobox.

For the ITk Strip system the same components and general connectivity scheme are used, but are hosted on an End Of Substructure (EOS) PCB mounted at the end of the local support. The decision to move the optical conversion for the pixels out of the ITk volume was driven by a combination of space constraints and the very high radiation dose in the innermost layers.




Figure 4.11: Annotated render of a partially-populated OEC half-ring, showing the modules, adhesive layer, carbon support, and services. On the right side, the outer carbon fibre facing layer and inner carbon foam layer are cut away to reveal the cooling pipe.

4.4.4 Outer endcap support structures

Quad modules destined for the endcap are loaded onto carbon fibre half-rings that consist of a carbon foam core glued to two carbon fibre facing layers, as is shown in Figure 4.11. A thin-walled titanium cooling pipe is embedded in the centre of the half-ring, which allows flow of the bi-phase CO_2 that is used to cool the modules. An adhesive chosen for its properties as a thermal interface, SE4445, is used to adhere the modules to the ring surface. The glue deposition and loading process as developed for half-ring construction at the University of Oxford will be discussed in Section 5.1.

Prior to loading, the half-rings are populated with a *bus tape* flex, which carries low/readout chip voltage (LV) and high/bias voltage (HV) for the serial power chain, and additionally routes the NTCs used for thermal monitoring of the modules to either a MOPS ASIC on the bus tape for analog-to-digital conversion, or directly out for use in the interlock system. Within the detector, rings are oriented vertically, with the CO_2 flow always from bottom to top through the cooling pipe. As a result, the NTC of the top module on each ring side will be routed out as the interlock signal for the corresponding SP chain. The input voltages to the modules are also monitored by the MOPS chip, and a CANbus comprising a power line, a ground, and a differential data pair is routed through the bus tape. A *type-1* services bundle combines the power, data, and interlock signal

¹⁶Using solder jumpers.

lines on the bus tape into a single cable bundle that interfaces to patch panel 1 (PP1) at the boundary of the ITk detector volume. These connections are propagated through PP1, PP2, and PP3 at the ITk, ATLAS detector, and experimental cavern boundaries, respectively, with the exception of the CANbus, which is consumed by the MOPShub crate positioned at PP3. The other signals reach the service cavern hosting the power supplies, interlock, and detector control system (DCS) servers.

Data from the modules is routed out via a flex termed the PP0, which combines 1, 2, or 4 data links per module depending on the expected detector occupancy in each region of the sub-detector¹⁷, and one command link per module. A custom twinax data cable bundle connects the bundled links to the optosystem, where an adapter board allows connection to an optoboard. Up to seven modules (or alternatively, up to 64 data links) will be read out per optoboard.

Half-rings are then assembled into half-cylinders before integration into the final detector. A carbon fibre shell acts as the outer boundary and retaining harness of the half-shell, and type-1 and twinax bundles are routed in the space between the half-rings and the half-cylinder shell, connecting to the PP1 and optosystem at the boundary of the ITk volume.

¹⁷Many modules in the OEC are expected to be run in either $2 \rightarrow 1$ or $4 \rightarrow 1$ data merging mode.

“But the plans were on display...”
 “On display? I eventually had to go down to the cellar to find them.”
 “That’s the display department.”
 “With a flashlight.”
 “Ah, well, the lights had probably gone.”
 “So had the stairs.”
 “But look, you found the notice, didn’t you?”
 “Yes,” said Arthur, “yes I did. It was on display in the bottom of a locked filing cabinet stuck in a disused lavatory with a sign on the door saying ‘Beware of the Leopard.’”

— Douglas Adams, *The Hitchhiker’s Guide to the Galaxy*

5

Large-scale tracking structures for the ATLAS ITk

Contents

5.1	Half-ring loading	70
5.1.1	Loading programs and the GiGo framework	72
5.1.2	Half-ring fiducial fitting algorithm	74
5.1.3	Module fiducial fitting algorithm	76
5.1.4	The thermal interface adhesive	76
5.1.5	Adhesive preparation	78
5.1.6	Adhesive deposition	80
5.1.7	Loading tests and progress	81
5.2	Half-ring QC testing	84
5.2.1	QC testing requirements and readout scheme	84
5.2.2	The test enclosure	87
5.2.3	Interlock design	89
5.2.4	Detector control system	92
5.2.5	Similarities to the SR1 RD53A demonstrator and slice test . .	96
5.2.6	Preliminary results	96
5.3	Conclusions and outlook	101

The two endcaps of the ITk Pixel system will be populated with assembled pixel modules at four total loading sites across Italy and the United Kingdom (UK). Istituto Nazionale di Fisica Nucleare (INFN) Lecce and INFN Genova will produce the ~60 half-rings required for one endcap, and Rutherford Appleton Laboratory (RAL) and the Oxford Physics Microstructure Detector (OPMD) laboratory at the University of Oxford those for the second endcap. Three separate flavours of ring will be produced, corresponding to layers 2, 3, and 4 of the ITk layout in Figure 4.2. These are populated

with 8, 11, and 13 modules per side, respectively. All sites use robotic gantries for loading, but due to differences in equipment and tooling, separate loading procedures have been developed at each. The loading process that has been developed at OPMD is described in Section 5.1, including a description of the procedure and current results on adhesive layer thickness and stability, and placement accuracy. Section 5.2 describes the quality control (QC) testing setup for half-rings at Oxford, and presents preliminary results.

5.1 Half-ring loading

The OPMD Aerotech robotic gantry [147] is equipped with a dual-vacuum rotating attachment point that can hold custom-designed tooling through an outer vacuum surface, and can pass a secondary vacuum through a central port to be used within the tooling itself. The attachment point and connected tool can be rotated about the vertical axis (denoted z in this section). The gantry head also hosts a high-resolution camera with a $5\times$ lens, a laser height measurement unit, and an attachment point for an adhesive syringe. A pressure-based adhesive dispense unit [148] coupled to a syringe and needle can be used to deposit adhesives in a controlled manner during gantry movements. A LabView program was previously developed that allows basic operations including movement, camera control, height measurement, control of a large number of vacuum lines through a manifold, and activation of the pressure line of the adhesive dispense unit. The LabView interface takes as input a numeric table of instructions and coordinates.

For the complex task of aligning and rotating half-ring components, a high-level instruction generation framework was developed, and a brief overview is given in Section 5.1.1. Placement procedures and control programs have been developed for the gantry, the details of which have been documented internally in Ref. [149].

Modules are loaded onto half-rings in a five-step process:

1. The half-ring is placed in a handling frame, which is in turn mounted on a base plate on the gantry. Assembled modules are extracted from their carriers using vacuum tooling and placed on a dedicated carrier tray on the gantry¹. The modules

¹If necessary, the module flex tabs still connecting the assembled flex to an outer ring are removed by cutting the tabs flush with the edges of the module.




Figure 5.1: Annotated render of the components and tooling used during loading on the Aerotech gantry. Initially, fiducial markers are located for the half-ring and modules on the module carrier, then the bridge tools are moved to the module carrier, and glue is deposited on the half-ring. Finally, modules are picked up from the module carrier using the pass-through vacuum on a bridge tool, and placed on the half-ring. The CAD model this render is derived from was provided by the Oxford Physics Technical Services team.

are vacuumed to the tray to prevent movement. The typical starting layout of the gantry for a loading procedure is shown in Figure 5.1.

2. Fiducial markers are located on the half-ring (4 to 6 depending on layer type) and module (4 per module), and a fit is performed to each to determine the starting position and orientation, as discussed in Sections 5.1.2 and 5.1.3.
3. Adhesive is mixed following the method in Section 5.1.5. A filled syringe is loaded into the holder on the gantry head, and the glue flow rate is measured through timed depositions in a small cup. The offset between syringe tip and camera is calibrated by depositing a small glue dot and measuring its position with the camera.
4. An adhesive pattern is placed in the module position as calculated in the reference frame of the half-ring, with the line speed adjusted based on the flow rate calibration. Details of the pattern design and necessary calibration steps are discussed in Section 5.1.6.

5. The module pick-up tooling is attached by suction to the gantry head, and is moved to contact a module. The secondary head vacuum is activated to suction the module to the tool, and the module tray vacuum is released. The module is lifted and moved to the desired position and orientation above the glue pattern. Finally, after a manual check by the operator, the module is placed in the desired position.

5.1.1 Loading programs and the GiGo framework

The Gantry Instruction Generation Object-oriented (GiGo) framework provides a high-level interface allowing loading scripts to be written in terms of complex object definitions (such as glue lines or dots, and repeated sequences of instructions) and a sophisticated system of reference frames and related transformations, that are then “compiled” to the low-level instructions supported by the LabView interface [150].

Coordinates have been implemented in a 3+1 dimensional space that allows automatic adjustment of the head rotation ϕ -coordinate (defined as being about the vertical z -axis of the gantry) to account for rotations and reference frames, which are implemented as a pseudo-4th dimension. Transformations are initially encoded as 4×4 affine transformation matrices A representing both a rotation R and translation by a vector \vec{t} , applied locally in that order,

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = A \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} R_{00} & R_{01} & R_{02} & t_x \\ R_{10} & R_{11} & R_{12} & t_y \\ R_{20} & R_{21} & R_{22} & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \Leftrightarrow \begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = R \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} + \vec{t}, \quad (5.1)$$

and reference frames \mathbb{F} define a hierarchical structure of transformations. The gantry reference frame (referred to as the *global frame* in the following) is the root node of this hierarchy, \mathbb{F}_0 . For notation, consider \mathbb{F}_{i-1}^a to be the parent of \mathbb{F}_i^a , and a to define a branch path through the hierarchy, where i denotes the distance from the root node \mathbb{F}_0 . The local transform associated with a reference frame \mathbb{F} is here denoted $A(\mathbb{F})$.

On application of a reference frame to a coordinate object, the sequence of transformations from the global frame to the given frame is resolved as

$$A^{\text{res}}(\mathbb{F}_i^a) = A(\mathbb{F}_i^a) \cdot A(\mathbb{F}_{i-1}^a) \cdot \dots \cdot A(\mathbb{F}_1^a), \quad (5.2)$$

where A^{res} defines the 4×4 matrix corresponding to the resolved sequence. A^{res} can now be extended to a 5×5 affine matrix to add the phi-rotation of the head (which must be considered only with respect to the global frame), which is denoted \tilde{A} and given as

$$\begin{bmatrix} x' \\ y' \\ z' \\ \phi' \\ 1 \end{bmatrix} = \tilde{A} \begin{bmatrix} x \\ y \\ z \\ \phi \\ 1 \end{bmatrix} = \begin{bmatrix} A_{00}^{\text{res}} & A_{01}^{\text{res}} & A_{02}^{\text{res}} & 0 & A_{03}^{\text{res}} \\ A_{10}^{\text{res}} & A_{11}^{\text{res}} & A_{12}^{\text{res}} & 0 & A_{13}^{\text{res}} \\ A_{20}^{\text{res}} & A_{21}^{\text{res}} & A_{22}^{\text{res}} & 0 & A_{23}^{\text{res}} \\ 0 & 0 & 0 & 0 & A_{\phi}^{\text{res}} + t_{\phi}^{\text{res}} \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ \phi \\ 1 \end{bmatrix}, \quad (5.3)$$

where A_{ϕ}^{res} is the extrinsic rotation about the z -axis² encoded by A^{res} , and $t_{\phi}^{\text{res}}(\mathbb{F}_i^a) = \sum_i t_{\phi}(\mathbb{F}_i^a)$ is the resolved sum of explicit gantry head rotations (ϕ -rotations) in the sequence of transformations corresponding to a hierarchy of reference frames.

Any transformation between two reference frames is applied by first transforming back to the global frame, then to the new reference frame,

$$\tilde{A}(\mathbb{F}_i^a \rightarrow \mathbb{F}_j^b) = \tilde{A}(\mathbb{F}_j^b) \cdot (\tilde{A}(\mathbb{F}_i^a))^{-1}. \quad (5.4)$$

A series of loading scripts have been developed that accept location data from the fiducials measurement in step 2 of the loading process described in Section 5.1, and then use automated fitting routines to construct the reference frames for each module and for their expected position on the half-ring. The module pick-up is then performed with respect to the start reference frame of the module, and the glue pattern deposition and




Figure 5.2: A frame of reference is generated for each module on a half-ring using the radial and angular nominal position data in the Electronic Document Management System [151, 152]. The frame of reference derived from the measured half-ring fiducials is assigned as a parent reference frame.

²This can be computed in many ways, for example from $A_{\phi}^{\text{res}} = \arctan(A_{10}^{\text{res}}/A_{00}^{\text{res}})$

module placement with respect to the end reference frame on the ring. The expected layout of these frames for an example layer 2 half-ring is shown in Figure 5.2.

As part of GiGo, a visualiser was developed that shows a real-time visualisation of the gantry head path, glue deposition, and tool orientation. An example “shadow trace” of a complete loading procedure is shown in Figure 5.3.

A complete documentation of GiGo is available in the code repository at Ref. [150].

5.1.2 Half-ring fiducial fitting algorithm

A simple Hough-transform-inspired circle fitting method has been implemented to determine the half-ring frame of reference from a set of three or more fiducial measurements at equal radius r_{fid} from the ring centre with known angular spacings. This method guarantees numerical stability and is deterministic, and unlike traditional fitting processes does not require seed parameter values for the fit. The method has been compared to a high-precision iterative circle function fit on fake data with randomised Gaussian errors, and




Figure 5.3: Complete shadow trace of a full half-ring loading procedure in the GiGo visualiser. Movement paths are shown as blue lines, with the start and end orientations of rotations indicated by the black rectangles. Glue deposition paths are shown in red.

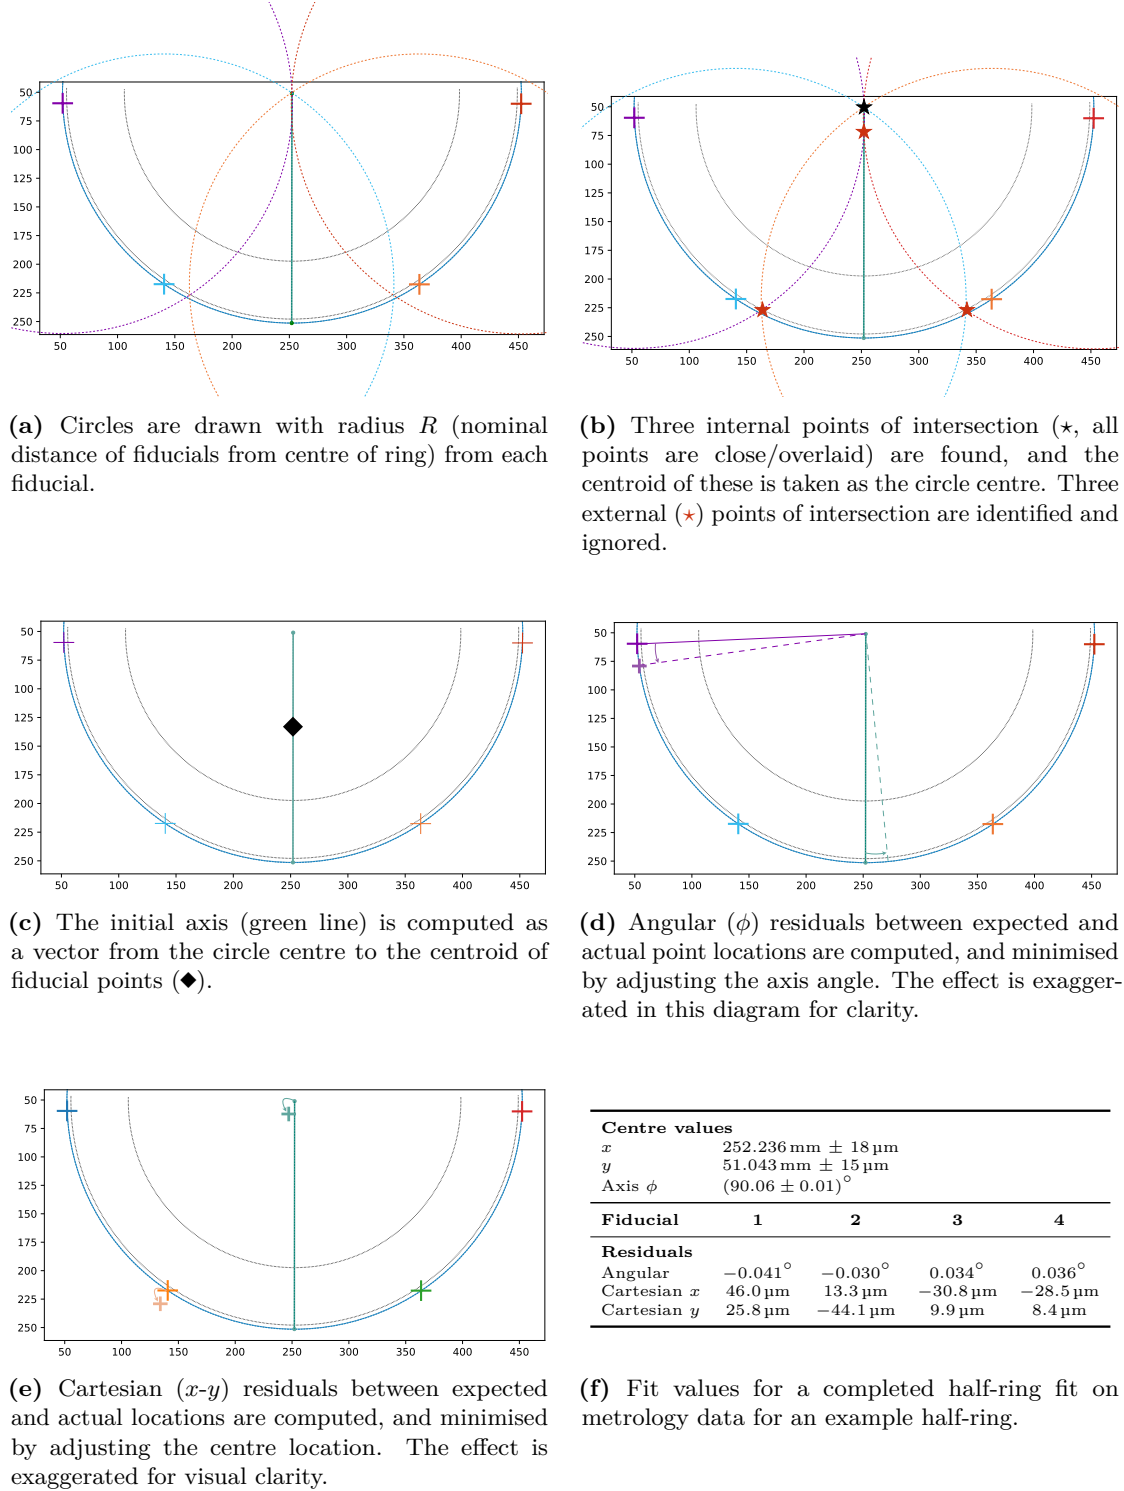



Figure 5.4: Complete procedure for a half-ring fit, as performed by GiGo on fiducials data to generate target module positions.



(a) Fit showing fiducial crosses, and axis in green.

Centre values				
x	$9.969 \text{ mm} \pm 220 \mu\text{m}$			
y	$9.877 \text{ mm} \pm 224 \mu\text{m}$			
Axis ϕ	$(22.525 \pm 0.410)^\circ$			
Fiducial	1	2	3	4
Residuals				
Angular	-0.399°	-0.007°	0.164°	0.570°
Cartesian x	$-266 \mu\text{m}$	$-39.5 \mu\text{m}$	$-39.9 \mu\text{m}$	$-266 \mu\text{m}$
Cartesian y	$116 \mu\text{m}$	$108 \mu\text{m}$	$112 \mu\text{m}$	$-337 \mu\text{m}$

(b) Fit results.

Figure 5.5: Rectangle fit as performed by GiGo on an emulated module position on the module carrier. Positions have been distorted by Gaussian noise with $\sigma = 250 \mu\text{m}$ in order to simulate the effect of the localisation uncertainty on the knowledge of the true module position. The nominal centre is at $(10 \text{ mm}, 10 \text{ mm})$ and nominal axis at 22.5° , and the observed deviations are $(31 \mu\text{m}, 123 \mu\text{m})$ and 0.025° .

both methods showed equivalent precision. The procedure used is illustrated in Figure 5.4.

The combination of a rotation encoding the axis angle and a translation to the ring centre then define the full affine transformation of the half-ring frame of reference. Uncertainties on each value can be estimated from the standard deviation of the residuals divided by \sqrt{N} , where N is the number of fiducials used, as shown in Figure 5.4f.

5.1.3 Module fiducial fitting algorithm

A similar method is used to determine initial module positions from the module fiducials, as is shown in Figure 5.5. Rather than taking the centroid of circle intercepts as the seed, the centroid of the four corners is taken as centre, and the initial axis is taken as the vector from the centre to the midpoint of the rightmost edge.

5.1.4 The thermal interface adhesive

The loading process utilises the DOW SE4445 thermal interface material [153] chosen by the ATLAS ITk project both as an adhesive and for its high thermal conductivity of $1.3 \text{ W m}^{-1} \text{ K}^{-1}$. Since thermal interfaces generally have a lower thermal conduction coefficient than the components they are connecting (silicon and carbon fibre), optimal thermal performance is achieved through a thin layer with large surface area coverage. ATLAS ITk ring-loading requirements specify a layer height of $100_{-50}^{+100} \mu\text{m}$ and at least 70%

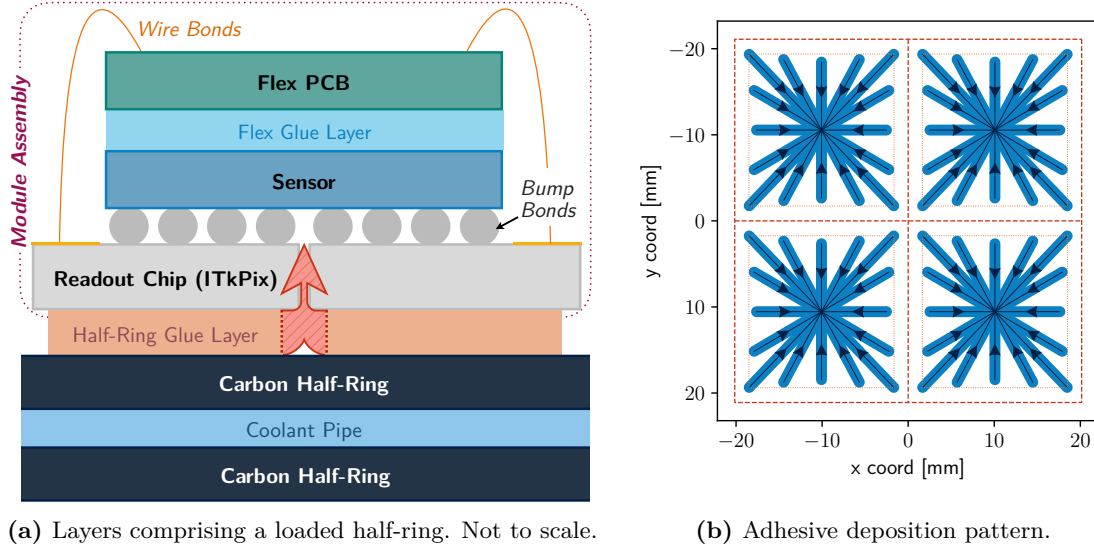


Figure 5.6: Location of adhesive layer in the loaded half-ring layer stack, and the deposition pattern used. In the layer stack, the risk region for adhesive between the modules is highlighted as a red hatched region. In the deposition pattern, deposition movements are shown as arrows, the chip outlines as dashed red lines, and the 70% coverage area as dotted orange lines.

coverage of the chip surface. Figure 5.6a shows the final module and half-ring stack and the adhesive layer, and Figure 5.6b the pattern design utilised to ensure sufficient coverage.

An additional challenge is posed by the 2×2 layout of the ITkPix chips on the module. A small gap of $\mathcal{O}(100\mu\text{m})$ is retained between the chips during bump-bonding to allow for some thermal expansion and to reduce potential chip contact and damage during hybridisation. Figure 5.6a highlights the potential issue: seepage of the thermal interface between the readout chips through capillary action. This would remove the necessary air gap, and introduces both additional risks of bump delamination during thermal cycling, and changes to the effective capacitance of the border pixels if adhesive seeps between the chip and sensor.

As a result, the pattern has been designed as a 16-spoke star centred under each chip, with carefully adjusted spoke lengths to ensure the compressed and cured pattern takes on a rectilinear form. The potential for air bubbles is minimised, since compression of the pattern will push air outward rather than allowing it to become trapped. The exact pattern definition has been optimised through a systematic study of test depositions between glass slides with varied parameters.

5.1.5 Adhesive preparation

Since SE4445 is a two-component epoxy, it is important to ensure that a consistent procedure is used for mixing. The procedure begins by separation of the SE4445 cartridge into its two components. Silica spacer spheres are mixed into the white component of the adhesive prior to the storage of both components in large syringes. Adhesive is prepared immediately before the loading process by depositing equal masses of the two components into a mixing cup, which is loaded into a planetary mixer and combined at 1500 rpm for 30 s³. After allowing the mixture to cool, it is filled into a 5 mL syringe, which is then also mixed with the planetary mixer with the same speed and duration, after which it may be loaded onto the gantry. A Nordson Pro4 adhesive dispense robot was used for test depositions to optimise the pattern and process [148], and once these were ready they were translated to the Aerotech gantry.

Spheres for layer height control

Careful control of the adhesive layer thickness is essential both to ensure thermal performance, and to constrain the final extent of the deposited adhesive pattern. The calibration methods discussed below give good control over the mass (and hence volume) of glue deposited in each pattern, and hence the uncertainty in the final pattern area is directly proportional to the uncertainty in the layer height.

A small concentration of 0.25% silica spacer spheres of nominal diameters 99 μm to 104 μm from the manufacturer Cospheric⁴ [154] are added to the adhesive prior to mixing, which in combination with the application of a strong compressive force prior to curing (tested with 500 g weights) ensure a consistent layer height. Thermal performance is not expected to be impacted for such small concentrations, and the silica spheres are expected to remain inert under high radiation doses.

To ensure no risk of oversize spheres concentrating pressure on single points of the ITkPix chip surface during compression, a sample of over 800 spheres has been surveyed under a microscope after deposition in a clear adhesive. Across the sample, 39 spheres

³The rotation speed and duration have been chosen to reduce the potential for self-heating, which could begin to cure the adhesive.

⁴Cospheric soda lime spheres, 99 μm to 104 μm nominal diameter, >95% within range, >90% spherical [154].




Figure 5.7: Surveyed glass spheres from the manufacturer Cospheric [154]. Two single-dimensional defects are highlighted. The dark cavities of varying size are air bubbles in the transparent adhesive.

were identified with defects exceeding the manufacturer upper size limit of $104\text{ }\mu\text{m}$ [154], all of which were in the form of small protrusions in a single dimension, and no spheres were oversized in two dimensions. Examples of surveyed spheres are shown in Figure 5.7. The use of a $106\text{ }\mu\text{m}$ sieve was introduced as an additional precaution, but a repetition of the study yielded 45 defects in one dimension (and no more significant defects), suggesting that sieving fails to remove spheres with these single-dimensional defects. However, such defects are not expected to exert substantial pressure on the readout chip since the spheres are likely to rotate during compression of the adhesive layer.

Over 100 test depositions have been performed during the processes of pattern and parameter optimisation, with various combinations of pressure, needle gauge, line speed, time since mixing, and using both sieved and unsieved spheres. Layer heights were




Figure 5.8: Adhesive layer height as a function of pattern mass for several adhesive samples. The target mass range and resulting thickness variation is shaded in red, and some patterns with issues during deposition (such as needle blockage) have been marked as outliers. The needle gauge is 23 AWG if not otherwise specified. *S09* and later samples used the flow rate calibration method.

measured to $\mathcal{O}(1\text{ }\mu\text{m})$ accuracy using a micrometer, and the resulting distribution is shown in Figure 5.8 as a function of pattern mass. Constraining the pattern mass to within $70\text{ }\mu\text{m}$ to $90\text{ }\mu\text{m}$ gives a layer thickness variability of $(114.0 \pm 7.4)\text{ }\mu\text{m}$.

5.1.6 Adhesive deposition

The mass (and hence volume) of adhesive deposited is directly proportional to the line speed of the robot during deposition under constant pressure, and will vary with the viscosity of the adhesive as it sets. The line speed can be related to the mass flow rate through the syringe needle using a simple linear model given as


$$v = \frac{lf}{m}, \quad \frac{\partial v}{\partial t} \approx \frac{l}{m} \frac{\partial f}{\partial t}, \quad (5.5)$$

with line speed v , total line length of pattern l , mass flow rate f , and desired pattern mass m . Test depositions for different target masses are shown in Figure 5.9. The area coverage to be expected for a given target mass m can be estimated as

$$A = \frac{m}{\rho \Delta_V h}, \quad (5.6)$$

where $\rho = 2.36\text{ g/mL}$ is the density of SE4445, h is the layer height, and Δ_V is a volumetric curing coefficient. A comparison of expected and measured areas using microscope images of 32 patterns yielded $\Delta_V = 0.969 \pm 0.002$.

The initial flow rate is measured by one or more timed depositions of adhesive into a cup under the desired dispense pressure. As it cures, the flow rate of the adhesive will decrease due to an increase in viscosity. Assuming knowledge of the evolution factor $\partial f / \partial t$, a time-dependent linear correction to the line speed may be estimated as given by



(a) Target 75 mg. (b) Target 80 mg. (c) Target 85 mg. (d) Target 90 mg. (e) Gantry, 85 mg.

Figure 5.9: Test adhesive depositions demonstrating the relationship between target mass and pattern extent. The measured mass of each pattern is overlaid. Patterns (a) to (d) were deposited by the Nordson glue robot on etched glass slides, whilst (e) was produced on the Aerotech gantry.




Figure 5.10: Flow rate evolution over time, for several adhesive mixture samples. The lines show linear fits, with parameters given in the legend. Some measurements suffering from needle blockage or other issues during deposition are included but marked as outliers. The needle gauge is 23 AWG unless specified, and the pressure is given in the legend. *S10* was an attempt to increase deposition speed using a larger diameter needle.

$\partial v / \partial t$ in Equation (5.5). Linear fits to the flow rate evolution of several adhesive samples mixed and extruded under identical conditions (23 AWG needle, 25 psi syringe pressure) converge⁵ around $(-3.0 \pm 0.6) \times 10^{-5} \text{ mg/s}^2$, and these are shown in Figure 5.10.

The model has shown high reliability for the given pressure and needle diameter, with all depositions where the model was applied within 5% of the target mass. Attempts to use a wider needle to increase the overall deposition speed (such as sample *S10* in Figure 5.10, among other attempts) were found to be subject to significant dribbling from the needle between lines of a pattern deposition, which reduced the precision of the deposited mass, and hence this idea was abandoned⁶.

5.1.7 Loading tests and progress

Several test placements on an aluminium half-ring have been performed to validate the placement accuracy and glue layer control; these placements are shown in Figure 5.11. Six positions were loaded with glass slides, and a seventh with an RD53A-based mechanical

⁵Sample *S12* has been removed from this average as it used an alternative mixing method.

⁶It was also noted that at most a $\sim 20\%$ decrease of the overall deposition time could be achieved, which was due to much of the deposition consisting of contributions (movement between lines and vertical retraction) which were already performed at a maximal reasonable speed.




Figure 5.11: Initial test placements of glass pieces (*alu2* to *alu7*) and a mechanical dummy module (*alu1*) on an aluminium half-ring prototype. The eight loading positions are denoted *alu0* through *alu7*, and the loading order is shown below.

dummy module to check the alignment of the vacuum feet with the flex.

The first placement at position *alu5* was incorrectly rotated due to operator error, and excessive adhesive was deposited for the first two placements due to early issues with the flow rate calibration on the gantry. The longer pressure lines used on the gantry compared to the Nordson Pro4 necessitated a period of adhesive extrusion (about 60 s) before the flow rate stabilised. Placements after *alu6* were performed with at least three successive calibration measurements in order to ensure the flow rate had stabilised prior to deposition.

Table 5.1: Placement accuracy in terms of Cartesian (Δx , Δy) and rotational ($\Delta\phi$) offset of module centres, for all loaded positions in chronological order. The centre and rotation of each module were determined using the rectangle fitting method described in Section 5.1.3, and compared to nominal locations determined from the half-ring fiducials. The gantry was used to locate modules after placement, and SmartScope measurements have been included where available. The rightmost columns shows differences between the gantry and SmartScope measurements.

Pos.	Gantry Offset			SmartScope Offset			Gantry - SmartScope		
	Δx [μm]	Δy [μm]	$\Delta\phi$ [$^\circ$]	Δx [μm]	Δy [μm]	$\Delta\phi$ [$^\circ$]	Δx [μm]	Δy [μm]	$\Delta\phi$ [$^\circ$]
alu5	-2085	-102	89.733	-2114	-112	89.833	28.25	9.72	-0.100
alu6	-743	-694	0.217	-747	-712	0.349	3.64	18.96	-0.132
alu7	-1163	-351	-0.030	-1190	-333	0.048	26.36	-18.24	-0.078
alu4	-126	-166	-0.194	-142	-163	-0.047	15.69	-3.04	-0.147
alu3	-46	111	-0.245	-50	107	-0.098	4.02	3.50	-0.147
alu2	-229	310	-0.315						
alu1	-130	574	0.029						

Post-placement fiducials measured with the Aerotech gantry were compared to metrology performed with a SmartScope camera system, in order to verify the accuracy of the gantry internal coordinate system. All offsets agreed to within $30\text{ }\mu\text{m}$, as shown in Table 5.1. As a result, the $\mathcal{O}(500\text{ }\mu\text{m})$ errors in placement accuracy in the first columns of Table 5.1 were attributed to slippage in the tooling or inexactness during loading. Initial attempts to address the issue involved reducing movement speed and acceleration parameters for the *alu2* and *alu1* positions, but did not substantially improve the accuracy.

During a series of dry runs (without adhesive) performed at position *alu0*, it was observed that the bridge tool may be sliding with respect to the spreader tool holding it, especially during rotations due to the large moment of inertia of the bridge tool.

Redesigned tooling

The tooling was redesigned in several steps to combat the issues with placement accuracy, and the iterations are shown in Figure 5.12. A first revision increased the surface area of the interface between spreader tool and bridge tool as shown in Figure 5.12b, but further dry placements showed only limited improvement. Tests showed that sufficient adhesion to prevent movement could be achieved within a few minutes of contact between the module and adhesive layer, and so a third iteration of the tooling was designed that forwent the bridge tool altogether in favour of a unified pickup tool, shown in Figure 5.12c.




Figure 5.12: Three iterations of module pick-up tooling used for half-ring loading.

Additionally, a careful recalibration of the offset between the gantry camera and centre of rotation has been performed using both mechanical means⁷ and by performing in-place rotations of a subject and fitting the rotation centre in the resulting data. It was also noticed that some offsets of $\mathcal{O}(10\,\mu\text{m})$ are reproducible between dry loading attempts, which has been attributed to probable local miscalibration of the gantry. A table of per-position offset values has been developed as a temporary measure until a full recalibration of the gantry coordinate system can be performed.

Dry placements performed with the most recent tooling and methodology have generally shown $\mathcal{O}(10\,\mu\text{m})$ positional accuracy, and hence achieve the local target accuracy. It should be noted that the local target of $\mathcal{O}(10\,\mu\text{m})$ is significantly more conservative than the ATLAS tolerance of $150\,\mu\text{m}$ on the position of each module corner fiducial marker [152].

5.2 Half-ring QC testing

With the exception of the IS, the components of the ITk detector will remain installed in ATLAS for the full duration of Run 4 and Run 5, for an approximately 12 yr lifespan. During this time, modules and services will not be accessible, and as a result it is critical that all installed components have undergone stringent quality control (QC) checks during construction and integration. For the OEC half-ring loaded supports, this requires both a reception test of modules and bare half-rings, and a full metrology and electrical test of the assembled half-rings. Since the reception tests generally follow the same procedure and goals as described in Section 4.3.3 and many associated references [155, 156], this section will explore the test box and infrastructure preparations for half-ring QC testing at OPMD.

5.2.1 QC testing requirements and readout scheme

OEC half-rings can be populated with up to 26 modules, and as a result the individual electrical readout of single modules is no longer feasible. Additionally, the connection to the bus-tape on the half-ring means modules must be powered in a serial powering chain, which has significant implications for the grounding and shielding scheme when compared to single-module readout.

⁷A scribe was used to trace a circle by a 360° rotation of the gantry head, then the centre was localised with the gantry camera to measure the offset.

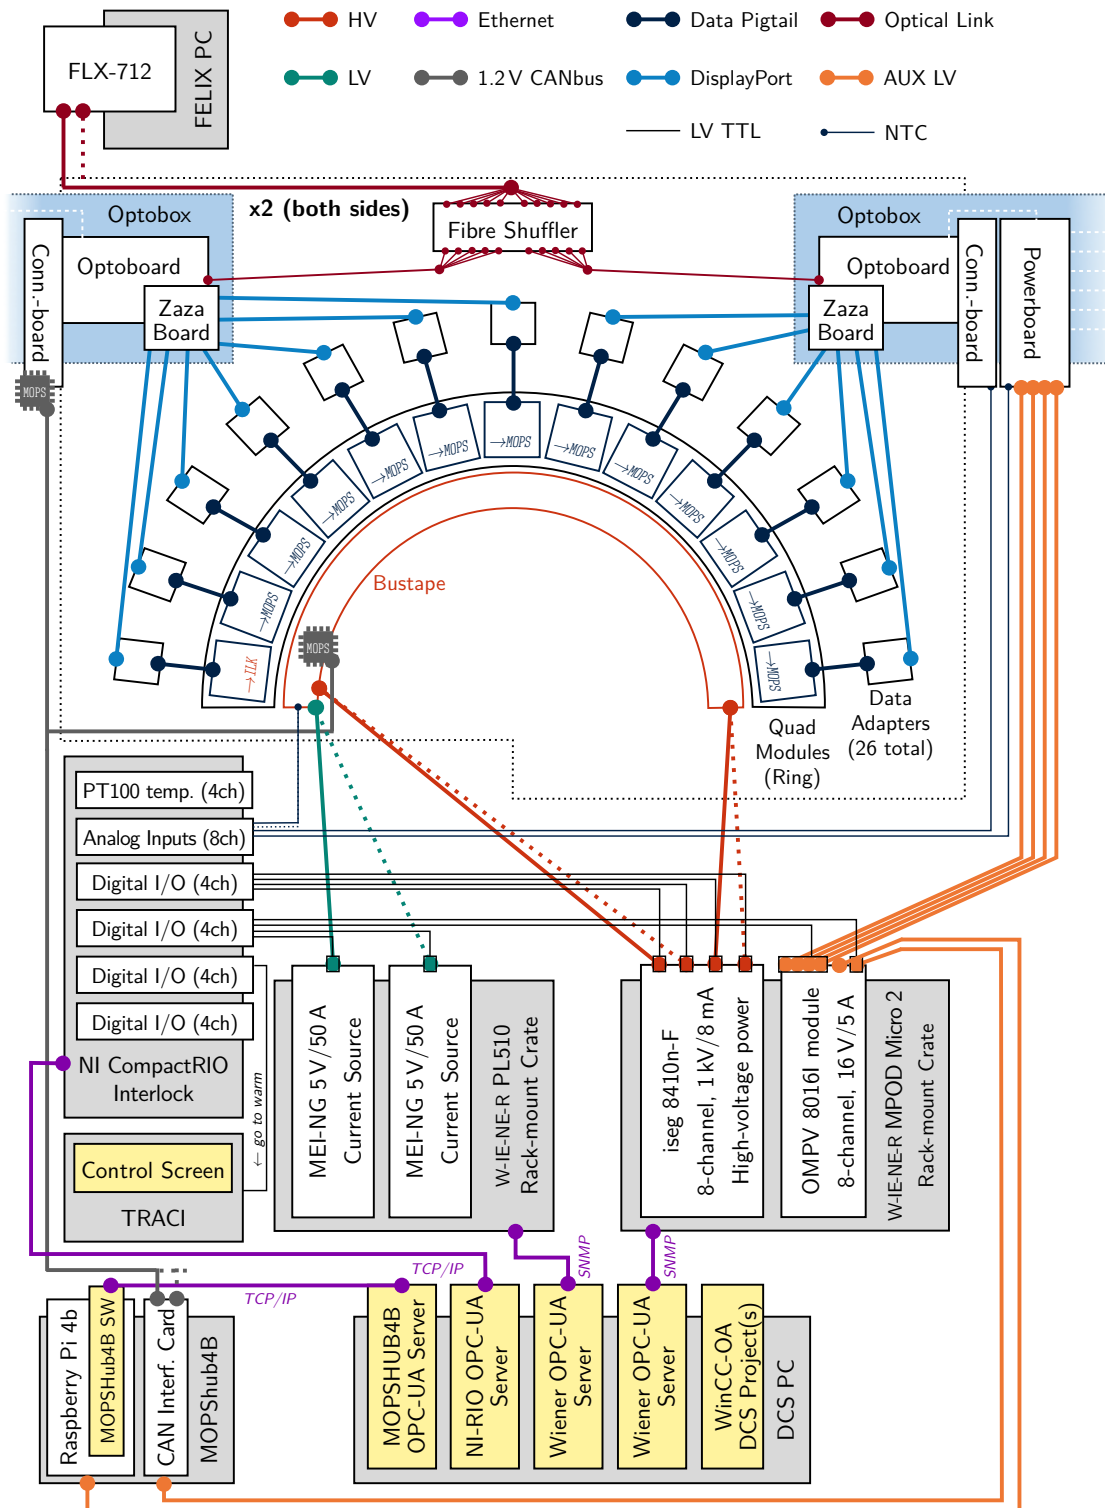


Figure 5.13: Schematic diagram of half-ring readout and services chain in the OPMD QC-testing setup.

For full QC testing, loading sites are required to be able to perform simultaneous electrical scans on all modules on a ring, at both warm ($\sim 20^\circ\text{C}$ forced cooling) and cold (around -15°C) operating points. The readout and services scheme used is very similar to that for the final detector described in Section 4.4, and is shown in Figure 5.13. An annotated photograph of the services rack is shown in Figure 5.14.

For initial readout tests, electrical signals are received from the half-ring via per-module adapter boards and DisplayPort cables, which connect to up to four electrical-to-optical conversion units (optoboards) within the optobox, as described in Section 4.4.3. Dedicated “Zaza” DisplayPort to ERF-8 adapter boards allow the connection of up to 6 modules per optoboard, but do not allow full exercising of the data-merging functionality described in Section 4.3. Toward production, this is planned to be replaced by full PP0s and attached twinax cable bundles for each half-ring side, as pre-production versions of these components become available. The optical outputs of the optobox are passed through a fibre reshuffling panel that allows the remapping of individual optical channels at a physical level.

Readout of optical output data is performed using the FELIX optical readout system. A dedicated FELIX-compatible server has been populated with an FLX-712 PCI-Express (PCIe) card connected to two 24-channel bundles of reshuffled optical fibres, and runs the YARR readout software [134]. A dedicated pixel firmware on the FELIX card converts electrical data to optical and vice versa, sending trigger and




Figure 5.14: Services rack containing readout and DCS components and power supplies.

configuration commands to the attached modules, and receiving data, in the format outlined in Section 4.3. It repackages data and streams it to a local network interface within the FELIX server, where an alternate FELIX-compatible back-end within YARR receives and decodes these packets, and allows steering of the trigger and configuration subsystems within the firmware [157, 158].

Monitoring of module and optoboard temperatures is performed using the MOPShub for beginners (MH4B) introduced in Section 4.4.1, which is populated by two CICs allowing readout of a total of three CAN buses, corresponding to one MOPS chip per bus-tape, and one within the Optobox. In each SP chain, the NTC of the module closest to the end of the cooling loop is passed to an interlock system that cuts power to this SP chain if the module temperature exceeds a set maximum of 40 °C. The interlock also accepts further signals from the Optobox and environmental sensors, and is discussed further in Section 5.2.3.

The serial power chain is supplied by a pre-production W-IE-NE-R PL510 crate [159] populated with several 50 V/10 A current sources designed for the ITk power requirements. The crate accepts per-channel interlock signals and can be remotely controlled via Simple Network Management Protocol (SNMP) over Ethernet [160]. Bias voltage is provided by an iseg 8410n-F module [161] capable of producing -10 kV and sustaining up to 8 mA of leakage current draw for each of 8 separately floating channels. This module is installed in a W-IE-NE-R MPOD Micro 2 crate [162], and provides per-channel interlock functionality. Auxiliary power for the MH4B and Optobox are provided by an 8-channel 16 V/5 A W-IE-NE-R OMPV 8016I voltage source installed in the MPOD crate [162].

To decouple the services cable routing from the half-rings to be tested, *cable saver boards* matching the design used by the OB have been produced [163], which combine all necessary power and data channels to a single 25-pin D-subminiature (D-sub) connector per SP chain. Type-1 services bundles will be produced that connect to the cable saver board at one end, and the bus tape at the other.

5.2.2 The test enclosure

To provide a controlled dry thermal environment, the insulated aluminium enclosure shown in Figure 5.15 has been built to house half-rings during electrical testing. The

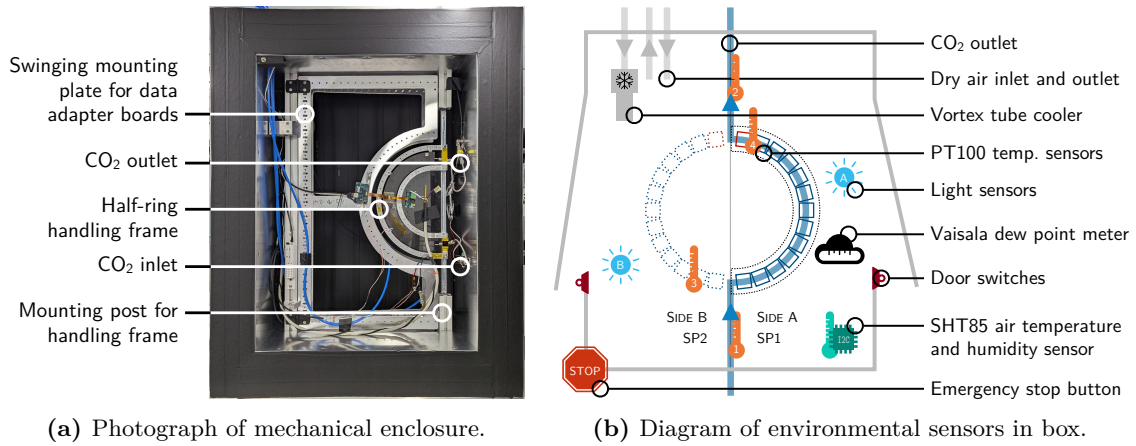


Figure 5.15: Environmental enclosure designed and built at OPMD for half-ring QC testing. The box is populated with a range of environmental sensors, and has inlets and outlets for CO₂ for cooling, and dry air to keep a low-humidity environment.

enclosure has been built from standard aluminium profile and panels, with a 10 cm-thick insulating foam layer between the outer and inner walls. A vertical attachment post provides a mounting point for a half-ring within a handling frame, and two swinging frames comprise attachment points for data adapter boards and services cables. Such frames have been designed in three different sizes, corresponding to the three sizes of half-ring that will be produced. Inlets for CO₂ cooling and dry air are provided, and cable ports have been integrated. A vortex tube cooler provides optional passive cooling of the pressurised dry air to reduce the effective heat load of the inlet.

Bi-phase CO₂ cooling is provided by a Transportable Refrigeration Apparatus for CO₂ Investigation (TRACI) cooling plant external to the enclosure [164]. The plant internally liquefies CO₂, then controls the pressure and temperature through the cooling loop so as to induce a phase transition within the loop. As a result, the heat load will be absorbed by the phase change of the coolant, allowing the temperature to remain mostly constant along the cooling loop.

The environment within the enclosure is monitored by an array of sensors outlined in Figure 5.15b, many of which are also utilised within the operational interlock system described later in Section 5.2.3. These include:

PT100 resistive temperature devices (RTDs) installed on the inlet and outlet cooling pipes, and close to the half-ring in the box environment.

Light dependent resistors (LDRs) to ensure a dark environment and prevent unexpected increases in leakage current within sensors from light leakage or operator error by interlocking the HV power supplies.

Door switches installed behind each door as an interlock mechanism to ensure a sealed environment during operation.

An SHT85 digital temperature and humidity probe to monitor air temperature to within 1 K and relative humidity to $\mathcal{O}(1\%)$, read out via an Inter-Integrated Circuit (I²C) bus [165].

A Vaisala DMT143 dew point sensor to measure the ambient dew point down to -70°C within the dry environment [166]. This additional sensor is necessary since the ambient air temperature within the enclosure ($\sim 10^\circ\text{C}$) is likely to be far above the temperature of the coldest parts of the cooling loop ($< -20^\circ\text{C}$), and so even sub-percent-level ambient humidities could cause condensation on exposed cold surfaces.

An emergency stop button as a local physical operator input for the interlock system.

5.2.3 Interlock design

Since assembled half-rings hold substantial value (both monetary and in terms of time and effort), a comprehensive interlock and safety system has been developed to ensure safe operation of the half-rings during QC testing, following similar principles to the LISSY interlock described in Section 4.4.2, but with alternative hardware. The system has been developed on the basis of the same compact reprogrammable I/O (cRIO) system from National Instruments (NI) also used for module QC and reception testing in OPMD [167]. The chosen cRIO crate can accommodate up to eight compatible I/O modules, and houses a Xilinx Zynq-series FPGA and an ARM microprocessor running NI real-time Linux. Application Programming Interfaces (APIs) in various languages are provided that allow reading and writing of the FPGA registers via the Linux host. The choice to use an FPGA-backed platform (rather than a microprocessor or software interlock solution) has been driven by the need for high reliability and low fault tolerance of this application.

FPGA-level firmware has been developed using NI LabView, which comprises three main functional blocks. The *threshold block* takes inputs from the I/O modules, provides averaging and smoothing for analog inputs, and compares these to an upper and lower threshold⁸. In the case of digital inputs, an expected “good” value is compared to instead. Thresholds are stored in registers and are hence easily programmable. Outputs from this block are passed to the *interlock matrix* block, which reads a programmable register matrix to decide which output interlock signals are triggered by which input conditions. Finally, the *output block* latches on interlock output signals once triggered, retaining the interlock condition until a reset is externally triggered via a register. This latching functionality ensures that a handover to a human operator occurs when the system experiences any interlock condition. This block is also responsible for assigning input or output roles to the bidirectional I/O channels available in the digital I/O module.

The firmware has been developed to accept the following types of I/O modules:

4-channel $0\ \Omega$ to $400\ \Omega$ RTD modules (NI-9217) in slots 1 and 2, capable of reading out PT100 temperature probes with 24-bit digital output accurate to better than 0.35 K over the expected measurement range [168].

8-channel analog input modules (NI-9201) in slots 3 and 4 with a measurement range of $-10\ \text{V}$ to $10\ \text{V}$ and 12-bit output [169].

4-channel digital I/O modules (NI-9402) in slots 5 to 8, which can accept or provide 3.3 V Low Voltage Transistor-Transistor Logic (LVTTTL) signals [170].

The interlock matrix in use is summarised in Table 5.2. When an SP chain NTC exceeds a $40\ ^\circ\text{C}$ temperature limit, the relevant serial power chain will be interlocked. Since modules should not be powered whilst connected to an optoboard without the latter also receiving power⁹, an interlock condition on an optosystem NTC will interlock both the optobox power and the power for both SP chains. Light sensors will interlock HV on both SP chains as only the sensor leakage current is affected by this condition. Inputs that

⁸For RTDs and NTCs, the lower bound encodes a maximal temperature, whilst the upper bound is used to detect a “wire break” condition at an unrealistically low temperature (generally set to $-60\ ^\circ\text{C}$).

⁹The powered ITkPix ASICs can generate capacitive charge build-up at the GBCRs on the optoboard from the synchronisation signals sent to initialise the Aurora stream, which could damage the GBCRs.

Table 5.2: Interlock matrix in use for the cRIO-based system. Inputs in the left column include both SP chain and Optobox NTCs, and environmental sensors within the enclosure. Outputs in the top row are passed to power supplies and the TRACI cooling plant. A \times symbol marks output interlock signals that are triggered by a given input exceeding its threshold.

KEY:											
\times = interlocked											
THRESHOLDS:											
SP NTCs	40 °C	SP1 LV Power	SP1A HV Power	SP1B HV Power	SP2 LV Power	SP2A HV Power	SP2B HV Power	Optobox Power	MH4B CIC Power	TRACI	Go-to-Warm
Opto NTCs	35 °C										
Dew point	−40 °C										
SP chain 1 NTC		\times	\times	\times							
SP chain 2 NTC					\times	\times	\times				
Opto conn. board NTC		\times	\times	\times	\times	\times	\times	\times			
Opto power board NTC		\times	\times	\times	\times	\times	\times	\times			
Light sensors			\times	\times		\times	\times				
Dew point sensor		\times	\times	\times	\times	\times	\times				\times
Door switches		\times	\times	\times	\times	\times	\times		\times		\times
Emergency stop		\times	\times	\times	\times	\times	\times	\times	\times		\times

signal a bad environment including the dew point sensor and door switches will interlock everything within the enclosure (but not the optobox), and will also send a *go-to-warm* signal to the TRACI cooling plant. In the case of an open enclosure, the CAN bus power for the MOPS monitoring system is also interlocked to remove all power from the half-ring circuitry. The emergency stop signal interlocks all outputs without exceptions.

Since the detector services route and expose the module NTC as bare resistive loads rather than digital or analog signals, a voltage divider and a stable reference voltage are required to allow readout as an analog voltage signal. Figure 5.16 shows a small PCB which has been designed to provide a common reference voltage generated by a linear regulator to up to eight NTCs (or LDRs), connected via two 9-pin D-sub connectors following the common ITk interlock pin mapping also used by the LISSY system. High-accuracy¹⁰ 10 k Ω resistors act as the reference resistance in the divider, and both the divided voltages and the reference voltage can be connected to analog inputs within the cRIO crate for monitoring. The reference voltages on three boards (two of which are

¹⁰A tolerance of 0.1% has been chosen as sufficient as it corresponds to a temperature error of less than 1 K across the expected operating range of the NTCs.

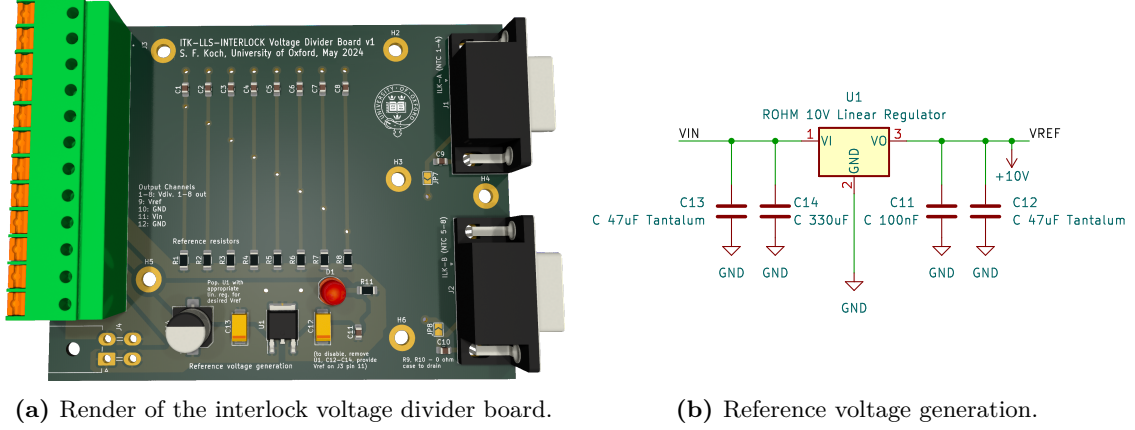


Figure 5.16: The interlock voltage divider board. Detector or Opto NTCs are attached to D-sub connectors on the right, and analog voltage signals are output on the terminal block on the left.

in active use within the setup) have been probed to determine the spread of reference voltages generated, and fall within a range of (0.999 ± 0.003) V.

5.2.4 Detector control system

Control of the power supplies, MH4B, and other components is implemented using Siemens WinCC-OA [171], which is an industry-standard Supervisory Control and Data Acquisition (SCADA) framework supported by and used at CERN for ATLAS and other detectors [172]. A cluster of existing ITk-oriented WinCC-OA projects has been developed by the system tests group at the SR1 laboratory for ITk surface integration and commissioning at CERN, and has been adjusted and extended for this use case. Three distributed projects are used, with the top-level ATLITPLLS project and the ATLITPMOPSHUB4B sub-project inherited from the SR1 setup, and the ATLITPRINGBOX project developed to interface the QC testing box with the existing structures.

Physical device interfaces are abstracted by the common Open Platform Communications – Unified Architecture (OPC-UA) protocol [173], for which servers are provided for the W-IE-NE-R power supplies and MH4B by the CERN industrial control systems (ICS) department and the MH4B developers at the University of Wuppertal, respectively. These servers provide a hierarchical address space of readable and/or writable data, which are sourced from or translated to the low-level command structure of the respective device. In the case of the W-IE-NE-R power supplies, SNMP commands and queries are used for hardware-level communication [160]. For the MH4B, a python software

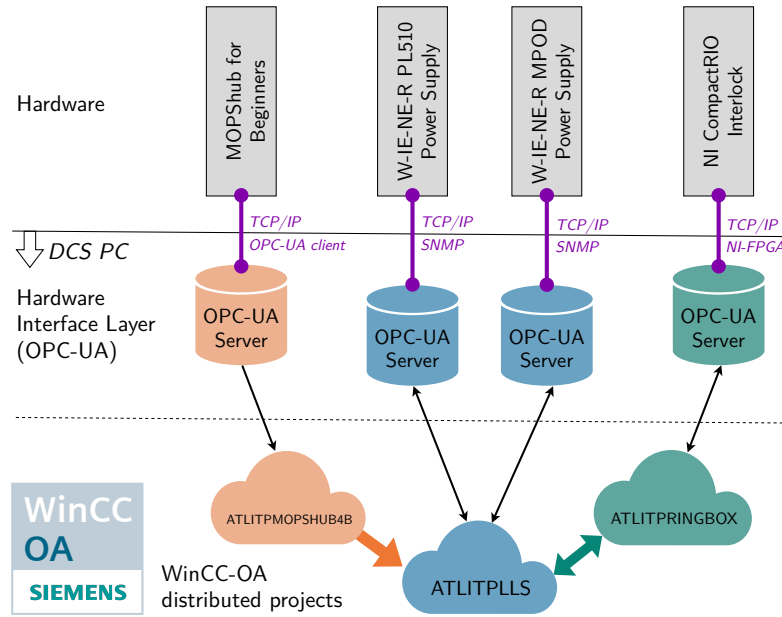


Figure 5.17: Hardware interface layers for the OPMD DCS. Software components running on the DCS server are shown below the solid line, and the protocol used for the Ethernet connection to the hardware is given. The WinCC-OA implementation is split into three distributed projects.

package running on the Raspberry Pi within the MH4B reads the MOPS chip data over the attached CAN bus (and controls the power distribution on the bus), and reads commands and writes monitoring values to the server as an OPC-UA client. An overview of these connections is shown in Figure 5.17.

A custom OPC-UA server has been developed for the cRIO interlock solution, which directly maps cRIO registers to the OPC-UA address space using the NI FPGA C API. A configuration file specifying a list of registers steers the address space structure. Each register is specified with a register name, data type and parameters, direction (read/write), and in the case of array data types, an array length. A top-level cRIO address space object also allows resetting and reprogramming of the firmware, and provides API-level communication status information.

For the interlock implementation, all interlock input and output decision registers are monitored via OPC-UA, and the interlock matrix and digital I/O configuration can be written to via a series of register arrays mapped in the OPC-UA server.

Individual address space elements can be subscribed to within an OPC-UA client, and the WinCC-OA projects generally directly map most of the available elements to internal datapoints. These can then be accessed within the project to build graphical

user interface (GUI) panels displaying datapoint values, or to influence the status and state of *device units* that form the leaves of a finite state machine (FSM) tree, shown in Figure 5.18. Power supply channels, environmental sensors, detector and Opto NTCs, and MOPS channels are all represented in this way, and each carry both a state and a status influenced by the datapoint value and metadata¹¹. State transitions can be invoked through commands available in the given device unit that write to relevant datapoints.

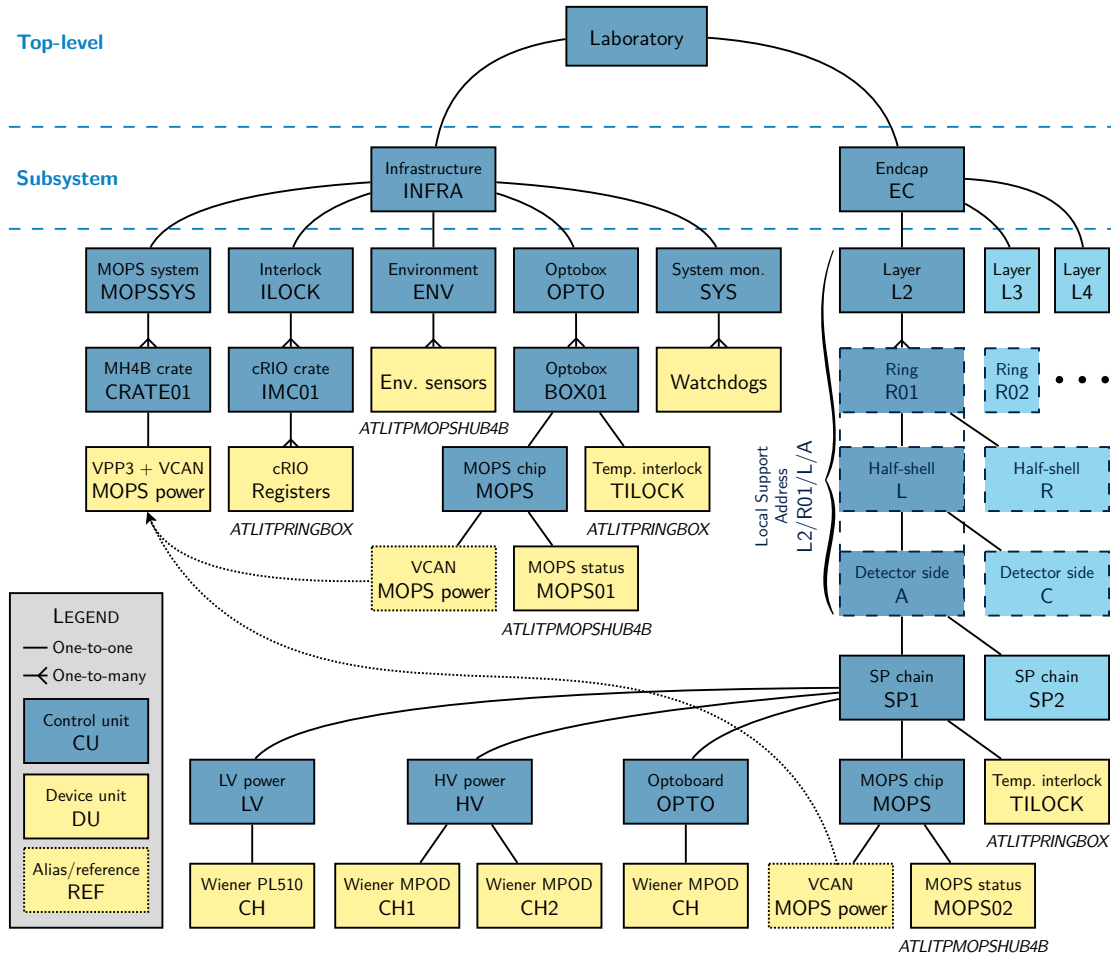


Figure 5.18: Finite state machine (FSM) structure used within the ATLITPLLS main project for the OEC subsystem. The structure is split into infrastructure and detector sub-trees, and encodes both control units (blue) and physical/device units (yellow). Where device units are sourced from other WinCC-OA projects, this has been annotated below. Each node has a defined set of states as well as a status (usually one of: *ready*, *not ready*, *warning*, *error*, *fatal*, *undefined*), and can be given commands that invoke state transitions. The status of device units is propagated upward through the tree, and commands downward. For simplicity, the local support address is concatenated into a single control unit in the final implementation, to allow cross-compatibility with the different address structures for local supports in the OB and IS subsystems.

¹¹Such as the age and provenance of the current datapoint value, the OPC-UA server status for the relevant server, or the activation of any associated alert conditions.

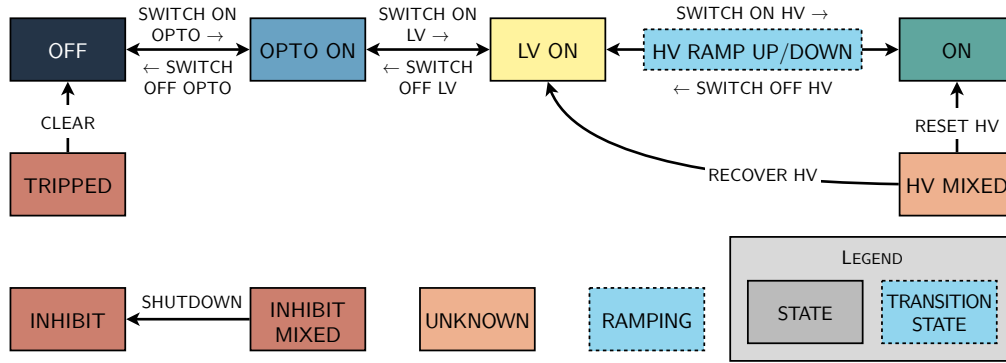


Figure 5.19: FSM states for the SP chain control unit. Typical operation involves cycling between the four states (and two transition states) in the top row. Commands invoking state transitions are shown, and some specific failure modes (HV compliance current trip, inhibit due to interlock) are represented by auxiliary states.

The status of child units is propagated upward through the tree, and commands executed at a higher level are passed down to units which support these actions.

A hierarchical definition of the OEC subsystem was added as *control units* to the tree structure in Figure 5.18 following existing local support naming conventions. The main control unit is at the level of an SP chain, with a set of states and available commands shown in Figure 5.19. The operational states of the SP chain are:

- **OFF:** all power to the detector components is off.
- **OPTO ON:** only the optoboards are powered.
- **LV ON:** in addition the LV current source for the modules is turned on.
- **ON:** in addition the HV bias is active for sensor depletion.

High-level transitions are defined which allow traversal of multiple states, such as **Go to ON**, **Go to LV ON**, and **Go to OFF**, passing through intermediate states as necessary.

The cRIO datapoints were also integrated into the **ILOCK** interlock component within the infrastructure subsystem, and appropriate panels were developed to display the state of the cRIO crate. An FSM was implemented for the cRIO variant of the interlock control unit as shown in Figure 5.20. Datapoints are used to detect whether an appropriate firmware and matrix have been set up, and the FSM defines transitions which allows these to be sent to the cRIO crate on start-up or after a reset. On an interlock condition, the crate will latch the output signals until a reset register is written to. This behaviour is

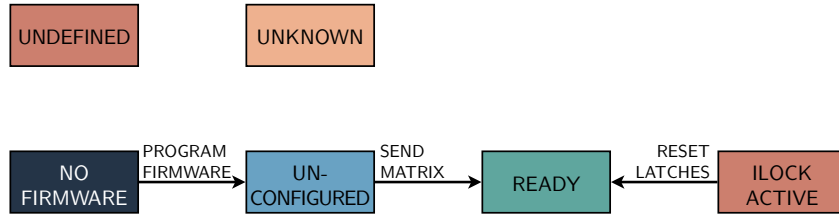


Figure 5.20: FSM states for the cRIO interlock crate control unit.

encoded by an `ILOCK ACTIVE` state and corresponding `RESET LATCHES` command, which enforce user intervention to resolve an interlock condition.

WinCC-OA provides native archiving of datapoints to a variety of database backends, and an InfluxDB database [174] stores time-series data taken during and between QC testing activities. These data can hence also be queried directly, allowing snapshots of module voltages, currents, temperatures, and other data to be retrieved and analysed or uploaded to the central production database.

5.2.5 Similarities to the SR1 RD53A demonstrator and slice test

The ring testing setup is heavily inspired by (and hence similar in many ways to) the original OB RD53A demonstrator setup in use in SR1. Other than the different local supports integrated into each setup, the main differences lie in the interlock solution design, which for SR1 was originally the same solution as used for the ATLAS IBL, but has recently been replaced by the LISSY final interlock solution, coupled to a Simatic programmable logic controller (PLC) that provides environmental monitoring functionality. The cRIO-based interlock setup developed for OPMD shows strong parallels (FPGA-backed, OPC-UA communication), but replaces the scalability of the aforementioned systems with flexibility for easy alteration and extension based on the changing needs of the half-ring QC testing programme.

5.2.6 Preliminary results

As the beginning of pre-production for half-rings is still outstanding¹², the QC testing enclosure has been commissioned with and tested on smaller representative structures including single modules and MOPS chips.

¹²The current main limitation is the availability of modules, bare half-rings, and other critical components.

MOPS readout

To qualify the MOPS readout capabilities of the system, several MOPS v1 and v2 chips have been read out and tested using the MH4B. The setup was commissioned using the MOPSV2 chip in the optobox, with voltage and temperature values of a single connected optoboard verified by direct probing and through forced cooling of the optobox in an unpowered state. The MOPS analog-to-digital converter (ADC) was calibrated by using the MH4B readout of the internal reference voltage to tune the ADC-to-voltage conversion in the MH4B OPC-UA server. Post-calibration, temperatures were found to match ambient conditions (15 °C) to within 1 °C, and voltages were accurate to within 1%.

Two MOPS chips loaded on detachable LV wings for an early prototype version of the bus tape have been successfully tested, including both v1 and v2 chips. Communication tests were performed for both variants, and the v2 chip was also connected to the bus tape with a single attached and powered RD53A module to check voltage monitoring of the module¹³, in the configuration shown in Figure 5.21. The voltage was successfully read out, but intermittent connection issues between the bus tape wing and bus tape itself were observed, which matched prior results from the University of Manchester. These




Figure 5.21: Readout test of a MOPS chip on the LV wing of a prototype bus tape, with a connected RD53A module.

¹³Unfortunately the flex on the RD53A module in question was not populated with an NTC.

issues had been previously observed, and had motivated the redesign of the final bus tape to a single integrated unit, rather than the earlier detachable wing design.

Interlock and monitoring

The interlock firmware has been extensively tested both using the LabView simulation tools, and with the real sensor inputs as these have been installed in the enclosure. Interlock output voltages for both triggered and non-triggered states have been validated by probing all channels, and the expected behaviour has been verified in the attached power supplies when manually triggering possible interlock conditions. Long-term monitoring data have been analysed for anomalous sensor values or interrupted communication, and no evidence of such behaviour has been observed.

Single loaded module readout

To ensure parity between the QC system used for module production and the ring testing setup, a full electrical QC procedure on a single module has been performed before and after loading on a prototype half-ring. In addition, these tests were designed to confirm that the loading procedure did not damage the loaded module, and to exercise the CO₂ cooling system. The loaded module is shown installed in the ring testing enclosure in Figure 5.22.

Pre-loading QC was performed in a standard electrical test stand, whilst post-loading testing used the full optical readout chain described above, and a custom Type-1 bundle replacement. Bit error rate tests (BERTs) performed on the electrical communication




Figure 5.22: Photograph of an ITkPix module on a carbon half-ring installed in the test enclosure.

Table 5.3: QC test results and pass (green) or fail (red) status before loading (QC test stand) and after (ring test enclosure) [175]. Voltages are in V and currents in A unless otherwise specified. Due to the ToT bug in ITkPix v1.1, the ToT tests trivially passed and have been omitted here.

Test	QC criteria	Pre-loading at 20 °C				Post-loading at 20 °C			
		Chip 1	Chip 2	Chip 3	Chip 4	Chip 1	Chip 2	Chip 3	Chip 4
ADC calibration		PASS				PASS			
Slope	[0.15, 0.224]	0.201	0.173	0.186	0.182	0.201	0.173	0.186	0.181
Offset	[-9, 31]	12.0	10.0	11.0	12.0	12.0	10.0	11.0	12.0
Linearity	[0.0, 4.0]	0.24	0.18	0.3	0.39	0.41	0.22	0.31	0.4
$V_{\text{ana}}^{\text{GND}}$ mean	[0.012, 0.023]	0.017	0.016	0.017	0.017	0.017	0.016	0.016	0.017
Analog readout		FAIL				PASS			
I_{ref} [μA]	[3.6, 4.2]	4.1	3.8	4	3.7	4.1	3.7	3.9	3.7
ADC gain	[0.703, 0.924]	0.874	0.749	0.806	0.788	0.866	0.74	0.8	0.781
$V_{\text{DAC}}^{\text{cal}}$	[0.668, 0.987]	0.882	0.753	0.81	0.792	0.879	0.747	0.807	0.79
$V_{\text{ref}}^{\text{core}}$	[0.435, 0.49]	0.465	0.456	0.46	0.464	0.464	0.455	0.458	0.462
$V_{\text{ref}}^{\text{OVP}}$	[1.8, 2.152]	1.974	1.976	1.999	1.995	1.965	1.966	1.99	1.985
$V_{\text{dig}}^{\text{DD}}$ trim	[2, 13]	2	10	9	12	2	11	9	13
$V_{\text{ana}}^{\text{DD}}$ trim	[2, 13]	1	10	11	10	2	11	11	10
SLDO		FAIL				FAIL			
Linearity	[0, 0.09]	0.012	0.012	0.01	0.011	0.012	0.012	0.01	0.012
$V_{\text{ana}}^{\text{in}} - V_{\text{dig}}^{\text{in}}$	[-0.05, 0.05]	-0.0	-0.0	-0.0	-0.0	0.0	-0.0	-0.0	-0.0
$V_{\text{ana}}^{\text{DD}}$	[1.18, 1.22]	1.3	1.23	1.22	1.2	1.2	1.21	1.2	1.21
$V_{\text{dig}}^{\text{DD}}$	[1.18, 1.22]	1.29	1.21	1.21	1.18	1.21	1.2	1.2	1.2
$V_{\text{ana}}^{\text{in}}$	[1.403, 1.575]	1.511	1.513	1.512	1.512	1.516	1.517	1.516	1.514
$V_{\text{dig}}^{\text{in}}$	[1.403, 1.575]	1.514	1.517	1.516	1.517	1.515	1.52	1.519	1.516
$V_{\text{offset}}^{\text{in}}$	[0.94, 1.041]	1.012	1.013	1.015	1.013	1.01	1.012	1.014	1.01
$I_{\text{ana}}^{\text{in}}$	[0.545, 0.648]	0.632	0.585	0.605	0.593	0.616	0.591	0.606	0.598
$I_{\text{dig}}^{\text{in}}$	[0.805, 0.946]	0.854	0.86	0.836	0.84	0.862	0.871	0.845	0.85
$I_{\text{ana}}^{\text{shunt}}$	≥ 0.005	0.0	0.019	0.003	0.022	0.036	0.064	0.052	0.071
$I_{\text{dig}}^{\text{shunt}}$	≥ 0.123	0.134	0.165	0.133	0.16	0.133	0.151	0.115	0.133
V_{cal} calibration		PASS				PASS			
$V_{\text{cal}}^{\text{med}}$ slope	[0.16, 0.24]	0.21	0.18	0.2	0.19	0.21	0.18	0.2	0.19
$V_{\text{cal}}^{\text{med}}$ offset	[-23, 17]	-2.0	-2.0	-3.0	-1.0	-2.0	-2.0	-3.0	-0.0
$V_{\text{cal}}^{\text{med}}$ linearity	[0.0, 4.0]	0.77	0.87	0.86	0.87	0.75	0.72	0.66	0.85
$V_{\text{cal}}^{\text{high}}$ slope	[0.16, 0.24]	0.22	0.18	0.2	0.19	0.22	0.18	0.2	0.19
$V_{\text{cal}}^{\text{high}}$ offset	[-23, 17]	-1	-2	-4	-3	-1	-2	-3	-3
$V_{\text{cal}}^{\text{high}}$ linearity	[0.0, 4.0]	0.51	0.41	0.42	0.38	0.43	0.5	0.48	0.47
$V_{\text{cal}}^{\text{med}}$ slope ratio	[0.49, 0.51]	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
$V_{\text{cal}}^{\text{high}}$ slope ratio	[0.49, 0.51]	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Inj. capacitance [6.0, 9.0]		7.7	7.65	7.98	7.69	8.29	8.23	8.58	8.25
Minimum health test		PASS				PASS			
Bad analog FE	≤ 7680	6	13	1	2	6	2	1	6
High noise	≤ 1536	1	10	5	2	2	11	8	8
Threshold tuning		PASS				PASS			
Untuned $\sigma_{\text{thr.}}$	≤ 500	48	47	48	49	48	48	50	51
Failed fits	≤ 7680	7	26	4	9	7	20	4	10
Tuned $\mu_{\text{thr.}}$	[1350, 1650]	1564	1510	1526	1529	1552	1532	1497	1529
Tuned $\sigma_{\text{thr.}}$	≤ 100	41	39	39	39	41	39	42	40
Pixel failure		PASS				PASS			
Elec. failed	≤ 7680	15	82	37	78	18	75	51	102
Disconnected	≤ 7680	3	8	0	1	3	9	0	1
Total failing	≤ 7680	18	90	37	79	21	84	51	103

channels using the psuedo-random binary sequence (PRBS) generation infrastructure in the ITkPix chip showed no errors within the measured test range, setting an upper limit on the error rate of 10^{-15} .

The chip characterisation tests outlined in Section 4.3.3 were performed in a warm environment with forced cooling at 20°C ¹⁴, following the standard ATLAS ITk Pixel QC procedure and parameters [175]. A dew point below -40°C was maintained throughout the tests. The results were analysed against the standardised QC criteria, and the results are shown in Table 5.3. Comparisons of sensor $I - V$ characteristics and threshold distributions are shown in Figure 5.23.

The module in question was chosen for this test as it had already been deemed unsuitable for operation as part of a serial powering chain due to the SLDO characteristics. One or more chip SLDOs failed the QC criteria both before and after loading, but the latter case shows greater stability in the SLDO response, likely due to the better current stability available from the WI-EN-ER current source compared to a desk-top voltage source. No significant increase in the number of disconnected pixels was observed, with the 16% increase in failing pixels mostly attributable to additional electrical failures. The $I - V$ characteristics in Figure 5.23a show very similar behaviour in leakage current in the test enclosure to the pre-loading result, with a slightly lower leakage current overall.

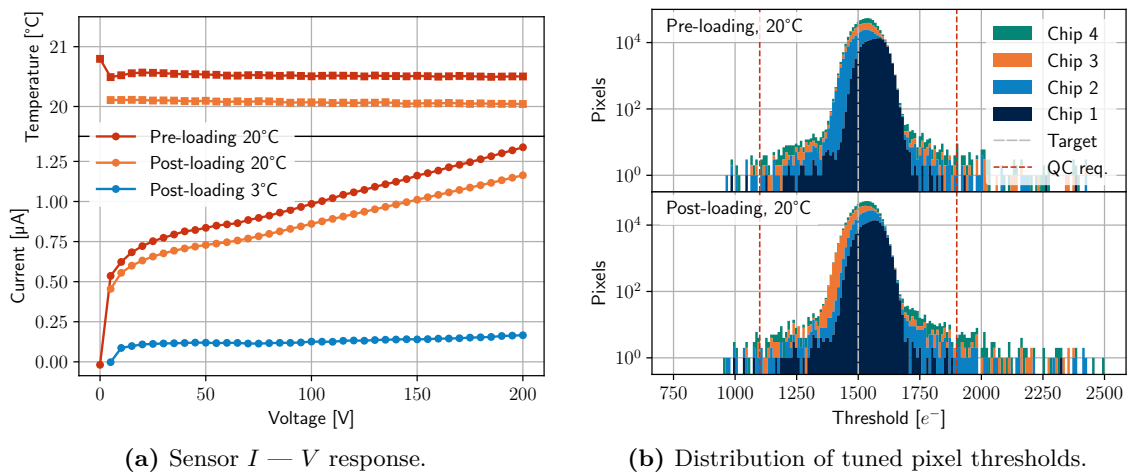


Figure 5.23: QC test results for an ITkPix v1.1 module before and after ring loading, including sensor $I - V$ response curves and pixel threshold distributions.

¹⁴Limitations on the TRACI CO_2 cooling at time of measurement prevented cold tests below 0°C . A cooling set point of 10°C corresponded to a powered module temperature of 20°C .


The decrease is likely attributable to a slight difference in sensor temperatures between the measurements. The post-loading measurement was repeated at a minimal stable temperature setpoint achievable at the time of 3 °C, and showed a six-fold decrease in the overall leakage current, matching expectations. No visible differences were observed in the tuned threshold distributions in Figure 5.23b.

5.3 Conclusions and outlook

Loading procedures and software have been developed for the local robotic gantry, and these capabilities have been exercised on a range of test objects to ensure the tolerances and precision requirements set by the ITk project are met. An electrical test setup for half-rings has been designed, developed, and commissioned. Detector readout, control, monitoring, and safety systems have been implemented that mirror the planned implementation for the final ITk detector, and these elements have been validated through targeted readout and functionality tests. A single functional module was loaded onto a carbon half-ring without bus tape, and installed into the test enclosure. It was shown that the module performance was not affected by the loading process using a modified services chain and through repetition of the standard QC test procedure in both a reference setup and the ring test enclosure. All relevant results were correctly replicated in the ring test setup, and a comprehensive range of readout functionalities have been exercised.

The progress that has been made bodes well for the OEC community as the half-ring project enters pre-production. Components for a fully loaded pre-production half-ring are expected to become available imminently, and this object will represent both the first complete test of the loading procedure, and the first operation of a full serial-power chain and readout of multiple modules in the ring test setup. The safe and effective operation of the DCS and readout software during this time will be essential to ensure timely completion of the half-ring production readiness review (PRR), and to prepare for production in the coming years.

PART II



Measurements of ATLAS

Illus. 2: *MiniMONSTAR* successor to the *MONSTAR* telescope, by Simon Koch [digital/CAD], 2025.

Multiple small-angle Coulomb scattering through silicon detectors is one of very few examples I know of for which the Central Limit Theorem is partially (but not completely) applicable. The angular distribution from 2^{20} scatters still exhibits strong non-Gaussian tails, whilst 2^{30} would already allow for a near-perfect approximation by a single Gaussian.

6

Multiple scattering and material content

Contents

6.1	Material distribution in ATLAS	103
6.1.1	Impact of material on physics results	104
6.2	Multiple scattering of charged particles	105
6.2.1	The single Coulomb scattering distribution	106
6.2.2	Molière's theory of multiple scattering	109
6.2.3	The Lynch-Dahl-Highland approximation	110
6.2.4	The Frühwirth-Regler formalism	111
6.2.5	Scattering models in GEANT4	114

6.1 Material distribution in ATLAS

In ATLAS and similar detectors at hadron colliders, vertex resolution is critical for pile-up rejection, flavour tagging, and specialised analyses such as searches for displaced vertices. In all these cases, precise knowledge of material content is also crucial to ensure the correctness of Monte-Carlo (MC) simulation datasets used for algorithm training, background estimation, and validation. As a result, alongside the minimisation of material in the design phase, the validation of material models is crucial.

The material content of a new detector has in the past generally only been estimated during the design phase, from the material composition of components¹. Such estimates

¹In the case of the original ATLAS tracker, this was cross-checked during construction by weighing the assembled sub-detectors and comparing to an estimated total weight [66].

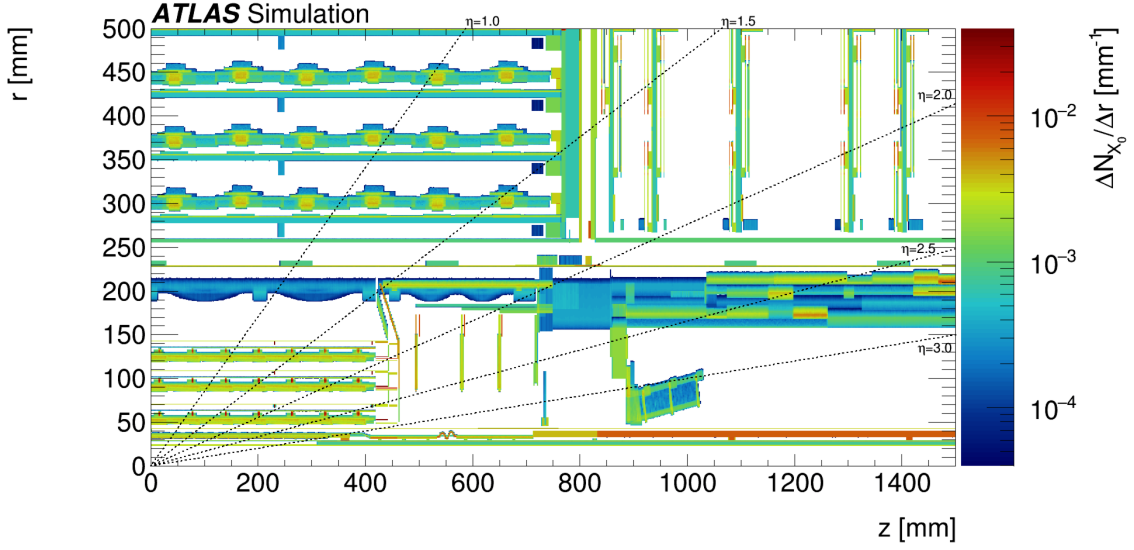


Figure 6.1: The r - z distribution of the local fractional radiation length $\Delta N_{X_0}/\Delta r$ for a recent geometry model of the current ATLAS inner detector including the IBL, pixel layers, and strip layers [176].

are highly dependent on knowledge of the composition of individual components, and are usually given as coarse-grained global averages rather than as a spatially resolved map. The material will usually be verified after detector commissioning using alignment and calibration data, using observations of photon conversions and nuclear interactions from track reconstruction to probe the radiation length X_0 and nuclear interaction length λ_0 of the detector. However, such measurements tend to compare data to simulations based on an existing material model (or modifications of a reference model) rather than directly deriving X_0 or λ_0 , and generally only provide a relatively coarse spatial resolution for material maps. Such measurements have been performed at ATLAS for the original inner detector using photon conversions, nuclear interactions, and K_S^0 decays, and have generally shown agreement with simulation models to within 10% [177–180]. Equivalent studies have been performed for the insertable B-layer, an additional pixel detector layer installed in 2014, and show similar results [176]. The material map validated in said study is shown in Figure 6.1.

6.1.1 Impact of material on physics results

Electron and photon reconstruction serves as a useful case study of the impact of detector material on physics objects, since energy measurements within the ECal are affected by

passive material in front of the calorimeter. The most recent energy calibrations performed for the ATLAS Run 2 dataset in Ref. [181] outline a commonly used uncertainty scheme for the current detector, in which a 5% uncertainty is applied to the ID material content, 10% to that of the IBL, and 25% to that of the inner detector services, in particular the patch-panel affecting the high- η region. Final uncertainties on the electron energy are generally 1‰ to 2‰, with the relative contribution of the ID material uncertainty to the total at 50% of the total for low p_T electrons and decreasing at higher p_T .

For many analyses limited by systematic uncertainties, material effects will often have a non-leading but significant impact. For some analyses that are especially susceptible to these, dedicated data-driven methods or alternative reconstruction or selection criteria are used to account for specific limitations of the current models. A recent search for displaced vertices at ATLAS defined explicit veto regions to reduce the background of secondary vertices produced in material [182]. Despite this, a substantial background of displaced vertices outside the defined veto regions was expected due to uncertainties on the material distributions used to derive the veto regions, and uncertainties on the precise vertex location of the detected jet. Another similar analysis used a fully data-driven extrapolation of charged-particle scattering events from a dedicated control region, resulting in extrapolation uncertainties associated with this estimate [183].

In general, it is conceivable that the reduction of ID material uncertainties from the current 5% to 25% range used in many analyses to the single-percent level could provide non-negligible improvements to object reconstruction and analysis results across many facets of the ATLAS physics program. A campaign of dedicated high-granularity measurements at the level of individual components could provide such improvements.

6.2 Multiple scattering of charged particles

Alongside the effects of energy loss through ionisation discussed in Section 2.1, multiple small-angle Coulomb scattering of charged particles is one of the most important material effects in tracking detectors. Most relevant for low- p_T tracks, multiple scattering in detector material has an irreducible impact on vertexing resolution, as has been shown in Section 2.2.3.

6.2.1 The single Coulomb scattering distribution

Any discussion of multiple scattering is fundamentally derived from the single-scatter case. The elastic scattering of charged particles with momentum p , velocity v , and charge q off a Coulomb potential defined by a point charge Q is well-described by the Rutherford cross-section, and both quantum and classical mechanics agree on a common form for the differential cross-section,

$$\frac{d\sigma}{d\Omega} = \left(\frac{Qq}{2pv \sin^2 \frac{\theta}{2}} \right)^2, \quad (6.1)$$

describing the scattering cross-section per infinitesimal solid angle $d\Omega = d\cos\theta \cdot d\phi$ defined by the azimuthal angle θ and polar angle ϕ .

Since scattering of charged particles within a material is generally dominated by the atomic nucleus (especially for particles heavier than the electron), the Rutherford formula may be applied by replacing Q with the nuclear charge Ze , where Z is the atomic number. For the cases of electron, muon, pion, and proton scattering, the particle charge is equal to the elementary charge e , with the cross-section invariant under the sign of q . In the following, all momenta will be considered to be given in units of GeV/c.

In the case of one or multiple small scatters, the small angle replacements $\sin \frac{\theta}{2} \rightarrow \frac{\theta}{2}$ and $\sin \theta \rightarrow \theta$ may be made to rewrite the cross-section in the form

$$\frac{d\sigma}{d\theta} = 2\pi \left(\frac{2Ze^2}{pv} \right)^2 \frac{1}{\theta^3}. \quad (6.2)$$

This gives a singularity at $\theta = 0$ due to the implications of an (idealised) infinite range Coulomb potential. In practise, the potential generated by the atomic nucleus will be screened by the electrons at larger distances, and Jackson [184] shows that this effect may be approximated by a Yukawa-type potential

$$(Ze^2/r) \cdot \exp(-r/a), \quad (6.3)$$

with $a \approx 1.4 a_0 Z^{-\frac{1}{3}}$ defined in terms of the Bohr radius $a_0 = r_e/\alpha^2$, with the fine structure constant $\alpha \approx \frac{1}{137}$, and the classical electron radius defined in terms of the electron mass m_e as $r_e = e^2/(m_e^2 c^2)$. This replacement modifies the small-angle scattering cross-section to

$$\frac{d\sigma}{d\theta} = 2\pi \left(\frac{2Ze^2}{pv} \right)^2 \frac{\theta}{(\theta^2 + \theta_{\min}^2)^2}, \quad (6.4)$$

with the quantum-mechanical cutoff angle θ_{\min} given as

$$\theta_{\min} = \frac{\hbar}{pa} \approx \frac{\alpha Z^{\frac{1}{3}}}{1.4} \cdot \frac{m_e c}{p} \approx 2.66 \times 10^6 Z^{\frac{1}{3}} p^{-1}. \quad (6.5)$$

Similarly to the small-angle singularity, the approximation of the nucleus as a point charge causes the scattering cross-section to depart from Equation (6.4) when this is integrated over the full range of θ due to deviations in predictions of large-angle scatters. The treatment of this effect given in Jackson [184] approximates the charge of the atomic nucleus as a uniform volume distribution within a sphere of radius R , giving a parabolic potential within the sphere, and a Coulomb potential outside:

$$V(r) = \begin{cases} \frac{3Ze^2}{2R} \left(1 - \frac{r^2}{3R^2}\right), & r < R, \\ \frac{Ze^2}{r}, & r > R. \end{cases} \quad (6.6)$$

The classical cross-section calculated from such a potential has a singularity at a maximum angle given approximately as $\theta_{\max}^{\text{classical}} = 2Ze^2/(pva_0)$, which is a consequence of the classical treatment. This may be rectified through introduction of a form factor $F(Q^2)$ given by the spatial Fourier transform of the charge distribution, which is defined as unity at $Q^2 = 0$ and rapidly falls off for $QR > 1$. The finite nuclear size hence sets an effective upper limit for the scattering

$$\theta_{\max} \approx \frac{\hbar}{pR} \approx \frac{274 m_e c}{A^{\frac{1}{3}} p} \quad (6.7)$$

$$\approx \frac{0.11}{Z^{\frac{1}{3}} p}, \quad (6.8)$$

with the approximation $A \approx 2Z$ in Equation (6.8) valid for all but very heavy nuclei.

This allows calculation of the integrated cross-section as

$$\sigma = 2\pi \left(\frac{2Ze^2}{pv} \right)^2 \int_0^{\theta_{\max}} \frac{\theta d\theta}{(\theta^2 + \theta_{\min}^2)^2} \approx \frac{3.68 \times 10^{-20} \cdot Z^{\frac{4}{3}}}{\beta^2}, \quad (6.9)$$

with the relativistic speed $\beta = v/c$. The single-scatter process can be described as the product of the total cross-section σ and a pdf

$$f(\theta) = \begin{cases} N \cdot \frac{\theta}{(\theta^2 + \theta_{\min}^2)^2}, & 0 \leq \theta \leq \theta_{\max}, \\ 0, & \text{otherwise.} \end{cases} \quad (6.10)$$

with a normalisation constant $N = 2\theta_{\min}^2(1 + \theta_{\min}^2/\theta_{\max}^2) \approx 2\theta_{\min}^2$. The variance of θ under $f(\theta)$ is approximately

$$\text{Var}[\theta] \approx 2\theta_{\min}^2 (\ln(\theta_{\max}/\theta_{\min}) - 0.5) \quad (6.11)$$

$$\approx \frac{2.84 \times 10^{-11} \cdot Z^{\frac{2}{3}} \ln(159 Z^{-\frac{1}{3}})}{p^2}, \quad (6.12)$$

and is especially relevant for multiple scattering, as the variance is additive for independent small-angle scatters.

For track reconstruction or simulation, it is much more convenient to work with projected scattering angles defined by a Cartesian reference system, for example $\theta_x = \theta \cos \phi$ and $\theta_y = \theta \sin \phi$, with z defined in the original direction of travel of the particle (i.e. $\vec{v} := v\hat{z}$). This allows effects on the particle x - and y -position and momenta to be decoupled, and allows each to be treated independently. The pdf may be transformed to one describing a projected scattering angle following the method in Fröhwrth and Regler [185], in which a joint density is defined as

$$g(\theta_x, \theta_y) = \begin{cases} \frac{N}{2\pi} \cdot \frac{\theta}{((\theta_x^2 + \theta_y^2) + \theta_{\min}^2)^2}, & 0 \leq (\theta_x^2 + \theta_y^2) \leq \theta_{\max}^2, \\ 0, & \text{otherwise,} \end{cases} \quad (6.13)$$

under the assumption that scattering is independent of the azimuthal angle ϕ , and hence a certain symmetry exists between θ_x and θ_y . The marginal density is then calculated through integration over one of the projected angles

$$g(\theta_x) = \frac{N}{\pi} \int_0^{\sqrt{\theta_{\max}^2 - \theta_x^2}} \frac{d\theta_y}{\theta_x^2 + \theta_y^2 + \theta_{\min}^2} \quad (6.14)$$

$$= \frac{\theta_{\min}}{\pi} \left[\frac{\sqrt{\theta_{\max}^2 - \theta_x^2}}{(\theta_{\max}^2 + \theta_{\min}^2)(\theta_x^2 + \theta_{\min}^2)} + \frac{\arccos\left(\sqrt{(\theta_x^2 + \theta_{\min}^2)/(\theta_{\max}^2 + \theta_{\min}^2)}\right)}{(\theta_x^2 + \theta_{\min}^2)^{\frac{3}{2}}} \right]. \quad (6.15)$$

Figure 6.2 shows the behaviour of the projected single scatter pdf for several common detector materials given a particle momentum of 1 GeV. The behaviour of the distribution as θ_x becomes large ($\theta_x \approx \theta_{\max}/2$) is instructive — the tails of the distribution roughly follow $1/\sin^3 \theta_x$ for single hard scatters.

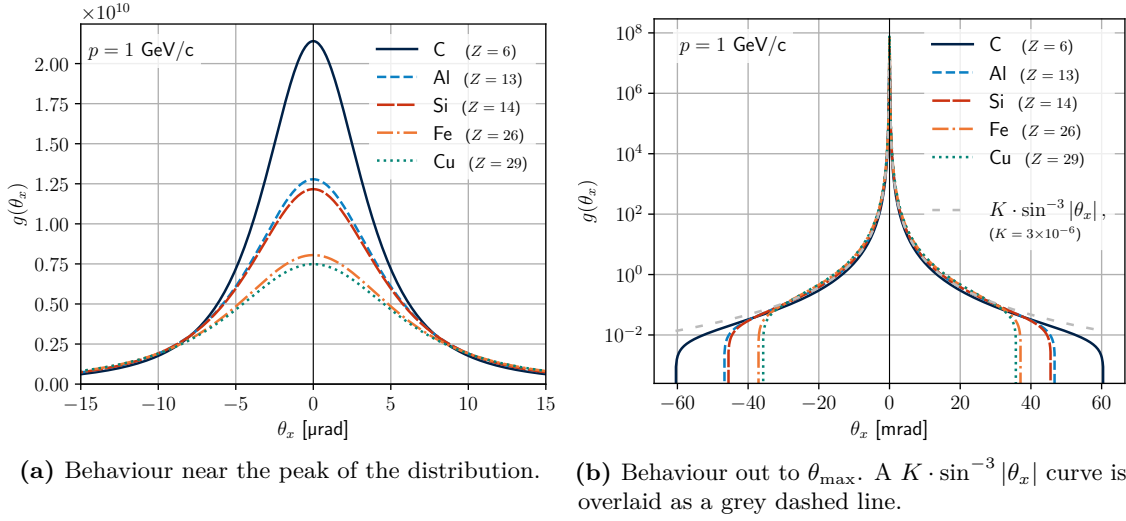


Figure 6.2: Marginal probability density of the projected single scatter angle in several materials.

The symmetry between θ_x and θ_y means the variance of the projected distribution is simply half that of the original, since $\theta_x^2 + \theta_y^2 = \theta^2$ and hence $\text{Var}[\theta_x] + \text{Var}[\theta_y] = \text{Var}[\theta]$:

$$\text{Var}[\theta_x] = \text{Var}[\theta_y] = \frac{1}{2} \text{Var}[\theta] = \theta_{\min}^2 [\ln(\theta_{\max}/\theta_{\min}) - 0.5] \quad (6.16)$$

$$= \frac{1.42 \times 10^{-11} \cdot Z^{\frac{2}{3}} \ln(159 Z^{-\frac{1}{3}})}{p^2}. \quad (6.17)$$

6.2.2 Molière's theory of multiple scattering

At least five separate extrapolations to very large numbers of small-angle scatters have been performed. The first four works by Molière [186], Snyder and Scott [187, 188], Goudsmit and Saunderson [189], and Lewis [190] all use analytic methods and are mathematically related. The first two methods apply a small-angle approximation to the single scatter cross-section, and then expand the multiple scattering cross-section in terms of Bessel functions. Goudsmit and Saunderson use a single scatter cross-section valid at all angles and extrapolate by an expansion in terms of Legendre polynomials. Lewis begins similarly but then applies the small-angle limit on the final result, hence deriving the connection between the first three methods. These methods provide analytic results but will not be explored here, though the reader is encouraged to consult Bethe [191] and Gottschalk et al. [192] for detailed derivations.

In the fifth method, Frühwirth and Regler [185] performed computational convolutions of the single scattering pdf given in Equation (6.10) for up to 2^{25} individual scatters.

They found good agreement with Molière theory for angles up to θ_{\max}^2 , and the practical usage of their methodology will be explored in Section 6.2.4.

Due to its early formulation and some slight advantages over the other theories³, Molière’s theory has become widely used as a preferred analytic formulation for comparisons to experiment and simplified calculations. Molière’s prescription predicts the scattering angle distribution projected onto a Cartesian plane as a sum of three analytic terms, comprising a Gaussian core (analogous to the *central limit* of a large number of soft single scatters), a carry-over to residual single scatter behaviour at large angles (corresponding to low-probability, large-angle single scatters), and a correction. Bethe has shown agreement of Molière theory with experiment to within 1% or better [186, 191]. Hanson, Lanz, Lyman and Scott have also shown agreement to within experimental uncertainties [193, 194].

6.2.3 The Lynch-Dahl-Highland approximation

As a result of the analytic complexity of Molière theory and dependency on empirically-determined material-dependent constants, several efforts have been made to produce an empirical fit to the theory. An original fit was performed by Highland to derive an expression for the central Gaussian core of the Molière distribution in terms of the radiation length of the material, which was later refined by Lynch and Dahl in 1990 [195, 196]. These methods have been the de-facto standard for experimental calculations by the Particle Data Group since 1982, and enjoy widespread use within the community.

The Lynch and Dahl revision⁴ of the Highland formula gives the RMS projected scattering angle of the Gaussian core⁵ in terms of the fractional radiation length x/X_0 as

$$\Theta_{\text{HL}} := \theta_{\text{proj}}^{\text{RMS}} = \frac{13.6 \text{ MeV}}{\beta c p} z \sqrt{\frac{x}{\beta^2 X_0}} \left(1 + 0.038 \ln \left(\frac{x}{\beta^2 X_0} \right) \right). \quad (6.18)$$

²They attribute the more persistent tails in the Molière theory to said formulation not properly accounting for the asymptotic behaviour at the upper limit of single scattering. It should be noted this effect is confined to the region of the pdf where the probability density is less than one thousandth of its peak value, and thus the overall effect is very small.

³Neither Molière theory nor that of Goudsmit and Saunderson require the assumption of any special form for the differential scattering cross-section.

⁴A long history of iterative improvements has converged to the constants in Equation (6.18). Alternative prescriptions are often useful in specific cases.

⁵In Lynch and Dahl’s fit, they take the RMS of the central 98% of the scattering angle distribution. Although perfectly suitable for analytic distributions, this method is shown to lack stability on experimental distributions with limited statistics in Chapter 7.

As before, βc , p , and z are the speed, momentum, and charge number of the scattering particle, respectively. To allow for the derivation of material from scattering angle distributions, Equation (6.18) can be analytically inverted by substitution of the principal branch W_0 of the Lambert W -function⁶ to give

$$\begin{aligned} \frac{x}{\beta^2 X_0} &= \exp\left(2 \cdot W_0(w) - \frac{1}{0.038}\right), \\ w &= \frac{1}{2 \cdot 0.038} \exp\left(\frac{1}{2 \cdot 0.038}\right) \frac{\beta c p}{13.6 \text{ MeV} \cdot z} \Theta. \end{aligned} \quad (6.19)$$

6.2.4 The Fröhlich-Regler formalism

Fröhlich and Regler [185] have computed projected density functions for multiple scattering in standard measure for sets of $N = 2^n$ convolutions of the single scatter pdf in Equation (6.10), for $n = 0, 1, 2, \dots, 25$. They then approximate the resulting projected distributions as Gaussian mixtures of the form

$$G(\theta_x; \sigma_1, \sigma_2, \epsilon) = (1 - \epsilon) \cdot \text{Gaus}(\theta_x; 0, \sigma_1) + \epsilon \cdot \text{Gaus}(\theta_x; 0, \sigma_2), \quad (6.20)$$

where $\text{Gaus}(\mu, \sigma)$ defines a Gaussian pdf. The conditions $\sigma_1^2 < \sigma_2^2$ and $\epsilon < 0.5$ are applied to ensure that σ_1 encodes the width of the core, and σ_2 the width of the tails. As above, θ_x is taken to represent an arbitrary projected scattering angle, and θ_y is distributed identically.

The expected variance of the mixture G is a function of the average number of scatters $N(d)$ encountered along a path of length d , given as

$$N(d) = \frac{\sigma \rho d N_A}{A}, \quad (6.21)$$

in terms of the integrated cross-section σ as given in Equation (6.9), the material density ρ , and Avogadro's number N_A . To relate $N(d)$ to a material property, a new characteristic length X_s is defined such that the total variance of the projected scattering angle for particles traversing this length is independent of the material:

$$\text{Var}[\theta_x(X_s)] = \text{Var}[\theta_x^{\text{ss}}] \cdot N(X_s) = \frac{225 \times 10^{-6}}{p^2 \beta^2}, \quad (6.22)$$

$$\implies \text{Var}[\theta_x(d)] = \frac{225 \times 10^{-6}}{p^2} \frac{d}{\beta^2 X_s} =: \mu_2(d). \quad (6.23)$$

⁶The Lambert W -function is a multi-valued function defined as the solutions to $W_k(w) \cdot \exp(W_k(w)) = w$ for any $w \in \mathbb{C}$, which form an infinite set of branches indexed by $k \in \mathbb{Z}$ [197, 198].




Figure 6.3: Behaviour of the material-dependent factor $h(Z)$ relating the radiation length X_0 to the material-dependent scaling variable X_s used in the Fröhwrth-Regler prescription.

The variance of the projected single scatter distribution $\text{Var}[\theta_x^{\text{ss}}]$ follows Equation (6.17).

Hence, X_s is defined as

$$X_s = \frac{A}{(4\alpha r_e^2 N_A) \rho Z^2 \ln(159 Z^{-\frac{1}{3}})} = \frac{716.4 A}{\rho Z^2 \ln(159 Z^{-\frac{1}{3}})} \quad (6.24)$$

where one may notice a similarity to the definition of X_0 in Equation (2.2). These are related by a factor close to 1, and the relationship may be parametrised in terms of Z by the ratio $h(Z)$, defined as

$$X_s = X_0 h(Z), \quad h(Z) = \frac{Z + 1}{Z} \frac{\ln(287 Z^{-\frac{1}{2}})}{\ln(159 Z^{-\frac{1}{3}})}. \quad (6.25)$$

The behaviour of $h(Z)$ with Z is shown in Figure 6.3.

Parametrisations of σ_1 , σ_2 , and ϵ are provided in terms of both X_s and X_0 , with an observation made that although scaling by a function of X_0 is at best an approximate encoding of material properties (whilst X_s is exact), such a parametrisation is likely sufficiently accurate for most purposes and allows for description of multiple scattering through the same parametrisation as other electromagnetic processes. The interpretations of scattering distributions using the Fröhwrth-Regler formalism in this work use the X_0 -based prescription to retain compatibility with general descriptions of material models.

The parametrisations of σ_1 , σ_2 , and ϵ are given in terms of a reduced thickness

$x' = \frac{x}{\beta^2 X_0}$ and a modified version encoding material dependence of the tails $x'' = \frac{Z^{\frac{2}{3}} x}{\beta^2 X_0}$ as

$$\sigma_1^2 = \mu_2(x) \cdot \left[0.8471 + 0.03347 \ln x' - 0.001843 (\ln x')^2 \right], \quad (6.26)$$

$$\sigma_2^2 = \mu_2(x) \cdot -\sigma_1^2, \quad (6.27)$$

$$\epsilon = \begin{cases} 0.04841 + 0.0063 \ln x'' - 0.00061 (\ln x'')^2, & \text{if } \ln x'' < 0.5, \\ -0.01908 + 0.1106 \ln x'' - 0.00573 (\ln x'')^2, & \text{if } \ln x'' \geq 0.5, \end{cases} \quad (6.28)$$

where the path length d has been approximated by the thickness of a thin material layer x . Focusing only on the Gaussian core, σ_1 can be rewritten into an similar form to the Highland formula in Equation (6.18):

$$\Theta_{\text{FR}} := \sigma_1 = \frac{13.81 \text{ MeV}}{\beta c p} \sqrt{\frac{x}{\beta^2 X_0} \cdot \frac{1}{h(Z)}} \cdot \left[1 + 0.03951 \ln \left(\frac{x}{\beta^2 X_0} \right) - 0.002176 \left(\ln \left(\frac{x}{\beta^2 X_0} \right) \right)^2 \right]^{\frac{1}{2}} \quad (6.29)$$

This form is referred to in the following sections as the Fröhlich-Regler model, and it can be seen by comparison to the Highland formula that the parametrisation is to higher order than Highland's, with one additional parameter in the $\ln(x/\beta^2 X_0)$ power series. The behaviour of the two models is compared in Figure 6.4.




Figure 6.4: Comparison of the Highland and Fröhlich-Regler models, for positrons at 1 GeV. The Fröhlich-Regler model is evaluated using $h(Z)$ values for carbon (light blue, $Z = 6$), silicon (red, $Z = 14$), and copper (orange, $Z = 29$). The ratio of the projected distribution width Θ_{FR} for each Fröhlich-Regler configuration to that predicted by the Highland formula Θ_{HL} for a given value of x/X_0 is shown to the right, and an equivalent ratio of the predicted x/X_0 for each model compared to that predicted by Highland is given in the lower plot. Since the models scale identically with p , the ratios are valid for any particle species at relativistic energies.

6.2.5 Scattering models in Geant4

The Geometry and Tracking toolkit (GEANT4) provides several implementations of multiple scattering models, some specialised for low momentum ($p < 100$ MeV) [199, 200]:

The *Urban* model is the default GEANT4 model for electrons and positrons below 100 MeV [201], and is based on Lewis theory [190], with separate parametrisations of the central Gaussian and the tails.

The *Goudsmit-Saunderson* model is currently only implemented for electrons and positrons, and provides better accuracy than the *Urban* model at the cost of runtime. The sampling of the pdf is done according to Goudsmit-Saunderson theory [189], and cases of zero-, one-, and many scatterings are separately parametrised.

The *CoulombScattering* model is a single scatter model based on the Wentzel single scatter cross-section (a more general form of Equation (6.4)). For very thin scatterers, it can be applied in sequence to simulate multiple scattering, but this is computationally expensive.

The *Wentzel-VI* model is used as default for all charged particles above 100 MeV, and combines the *CoulombScattering* model for large-angle scatters with a multiple scattering model for soft scatters. A dynamic cutoff angle θ_{\max} defaulting to 0.2 rad steers the choice between simulation as a single hard scatter or multiple soft scatters.

The default ATLAS detector simulation uses a physics list implementing the **Opt0** configuration, which combines the default *Urban* model for e^-/e^+ below 100 MeV with the *Wentzel-VI* model for all other cases. The simulations performed in the work described here used a physics list implementing the **Opt4** configuration, which replaces the *Urban* model with the *Goudsmit-Saunderson* model for low-energy e^-/e^+ , but retains the *Wentzel-VI* model for $\mathcal{O}(\text{GeV})$ charged particles. More importantly, **Opt4** uses stricter step limit rules that improve precision around boundaries and for very thin scattering targets.

What is this monster of a telescope?

— Valerio Dao [paraphrased], after walking into his office
to see a 2 m long telescope frame blocking his desk.

7

Measuring detector material via multiple scattering at testbeams

Contents

7.1	Experimental design considerations	116
7.2	The MONSTAR telescope	118
7.2.1	Thermal considerations	122
7.2.2	Datasets collected	123
7.3	Analysis methodology	124
7.3.1	Data preparation, alignment, and data quality checks	124
7.3.2	Tracking and angle extraction	125
7.3.3	MALTA half-chip efficiency correction	126
7.3.4	Telescope acceptance checks	127
7.3.5	Fitting angle distribution widths	128
7.3.6	Mechanics and air contribution subtractions	131
7.4	Material estimate	133
7.5	Geant4 simulations of air and telescope scattering	134
7.5.1	Comparison of scattering distribution width extraction methods	137
7.5.2	Telescope response	138
7.5.3	Air and telescope scattering validation	140
7.5.4	Variation of air and telescope contributions with plane spacing	142
7.6	Results	142
7.6.1	Resolution and beam energy uncertainties	143
7.6.2	Comparison of fit methods and scattering models	145
7.6.3	Analysis without DUT data	146
7.6.4	Comparison to the 700 MeV campaign 1 dataset	146
7.7	Implications and conclusions	147

The importance of accurate knowledge of the material content of a particle detector during its design and operation motivates the development of a direct measurement method for the material composition. Multiple scattering of low-momentum particles

provides a suitable probe for measuring the fractional radiation length of subjects in a testbeam environment. Multiple-scattering-based measurements have been performed at the University of Göttingen and the Deutsches Elektronen-Synchrotron (DESY) using a six-plane MIMOSA telescope, but have generally focussed on relatively thick calibration subjects, and did not attempt to image instrumented detectors [202–204]. A recent measurement by the CMS team at ETH Zurich of an RD53A quad module at the Paul-Scherrer Institut (PSI) PiE1 beamline used a simpler three-plane telescope with the module acting as the central plane, but a more rudimentary analysis methodology without full tracking as a result [205]. The methodology and results described in this chapter build on these previous measurements to design a four-plane telescope based on MALTA detectors, and perform the first direct measurement of the material content of an ITk Pixel module.

Following a discussion of telescope design considerations in Section 7.1, the MONSTAR telescope is introduced in Section 7.2 in relation to a testbeam campaign performed at the CERN PS in May 2023. A detailed description of the analysis methodology is given in Section 7.3, and a dedicated estimate for comparison is developed in Section 7.4. Section 7.5 compares the observed behaviour to simulations to validate the methodology. Final results are presented in Section 7.6, and a follow-up measurement which expanded this work to a more comprehensive measurement campaign including samples from ATLAS, CMS and Mu3e has been performed, but is not discussed within the scope of this thesis. The results in this chapter have been published in Ref. [2].

7.1 Experimental design considerations

Using a telescope with low- X_0 reference planes and a sufficiently long lever-arm, it is possible to perform a direct measurement of the fractional radiation length x/X_0 for a thin subject via measurement of the distribution of scattering angles due to multiple scattering. As discussed in Section 6.2, the Highland formula in Equation (6.18) and the equivalent formulation in the Frühwirth-Regler formalism in Equation (6.29) both predict the width of the central Gaussian core of the scattering distribution as a function of the thickness of a scatterer given in radiation lengths.




Figure 7.1: Relationship between beam energy and multiple scattering distribution width predicted by the Highland formula in Equation (6.18) for a positron beam, and varying spacing distances between a subject with $x/X_0 = 0.8\%$ and a downstream plane at a given distance. The upper horizontal blue line marks the limit at which the expected scattering full width at half maximum (FWHM) equals the sensor width, and the lower red line where it equals the pixel pitch. The red region shows an expectation of positron beam intensity at the CERN PS, melding to blue at the point at which the hadronic beam contributions are expected to become significant ($> 1\%$). The orange shaded region shows an expectation for the usable range of the MONSTAR telescope (with reasonable limits of 20 to 100 pixel pitches, as detailed in Section 7.2), and the purple dot denotes the configuration chosen for the measurement presented in this chapter.

Pixel detectors are particularly suitable devices under test (DUTs), as they allow for the localisation of the scattering point within the subject itself. In order to make the scattering angles measurable, a relatively low energy particle beam is required, and a suitable lever-arm length must be chosen for the telescope design. As derived in Section 2.2.3, for telescope planes with square pixel pitch d , an upper limit on the spatial resolution can be estimated as $d/\sqrt{12}$ without accounting for resolution improvements from clustering. Hence, a mean resolution uncertainty of $\mathcal{O}(1\%)$ on extracted scattering angles can be achieved by ensuring the expected RMS scattering angle Θ given by the telescope lever arm is such that the width of the projected Gaussian on the furthest downstream plane¹ corresponds to 29 pixels, since in this case the spatial resolution of $0.289d$ constitutes less than 1% of the expected deflection.

As shown in Figure 7.1, a suitable range of energies and telescope spacing distances

¹The width of the projected Gaussian distribution is related to Θ by $\sigma_{\text{plane}} = (z_{\text{plane}} - z_{\text{DUT}}) \cdot \arctan \Theta$, where z_{DUT} is the z -position of the DUT, and z_{plane} the z -position of the downstream plane.

has been determined (assuming MALTA telescope planes with square pixels of pitch $36.4\,\mu\text{m}$). The CERN PS target area 9 (T9) testbeam facility was chosen as it allows adjustable beam momenta in the range 100 MeV to 15 GeV, and can provide high-purity electron or positron beams at low momenta. Up to three 0.4 s spills of particles can be provided per 40 s supercycle, depending on the upstream infrastructure and priorities for the accelerator chain. A 1.2 GeV positron beam configuration was chosen for the main data-taking, as this was expected to give the highest rate, and the use of a non-hadronic particle species excludes an additional scattering contribution from the strong interaction that could be present with protons or pions. The beam instrumentation was removed and the beam vacuum extended up to the start of the telescope to further improve the trigger rate, but despite initially expecting a rate exceeding 10 000 particles per spill, ~ 6000 were observed within the telescope acceptance. The beam momentum band was set to 5% FWHM, chosen as a compromise between the beam rate and the uncertainty in particle momenta from the momentum band.

7.2 The MONSTAR telescope

The Multiple or Negligible Scattering Telescope, as Required (MONSTAR) is a 4-plane telescope allowing installation of a subject on a pair of linear stages in the centre providing $\mathcal{O}(\mu\text{m})$ accurate $x - y$ positioning. The telescope frame was constructed from aluminium profile with a total length of 2 m, allowing it to accommodate inter-plane spacing distances from 5 cm to 90 cm. This corresponds to accessible energies for scattering measurements ranging from 50 MeV to 1500 MeV for an expected material thickness of $0.8\% x/X_0$, as shown in Figure 7.1. This value corresponds to coarse-grained estimates for ITk and similar hybrid pixel modules.

As shown in Figure 7.2, the telescope plane mounting points are designed to be populated by MALTA monolithic pixel detectors [51] with 512×512 pixels, and the default configuration consists of two $300\,\mu\text{m}$ -thick Czochralski (Cz)² MALTA sensors as the outer planes, which provide larger cluster sizes and hence higher spatial resolution than thinner sensor types. $50\,\mu\text{m}$ and $100\,\mu\text{m}$ MALTA sensors with lower material content

²The Czochralski method is used to grow large single crystals, and is used to produce a variant of high-resistivity monocrystalline silicon used in some MALTA sensors [206].




Figure 7.2: The telescope configuration and trigger scheme used at the CERN PS, annotated with plane spacings and thicknesses. Multiple scattering within the DUT produces track kinks, the distribution of which is analysed to derive the radiation length of the DUT.

are used for the two inner telescope planes to reduce multiple scattering within the telescope, allowing the telescope to be used in multiple configurations. Placed close to the outer planes, the inner planes acts as additional track points for measuring the kink angle in the centre plane, improving track resolution and, when integrated into the trigger coincidence, reducing the rate of fake triggers from sensor noise. Placed close to the

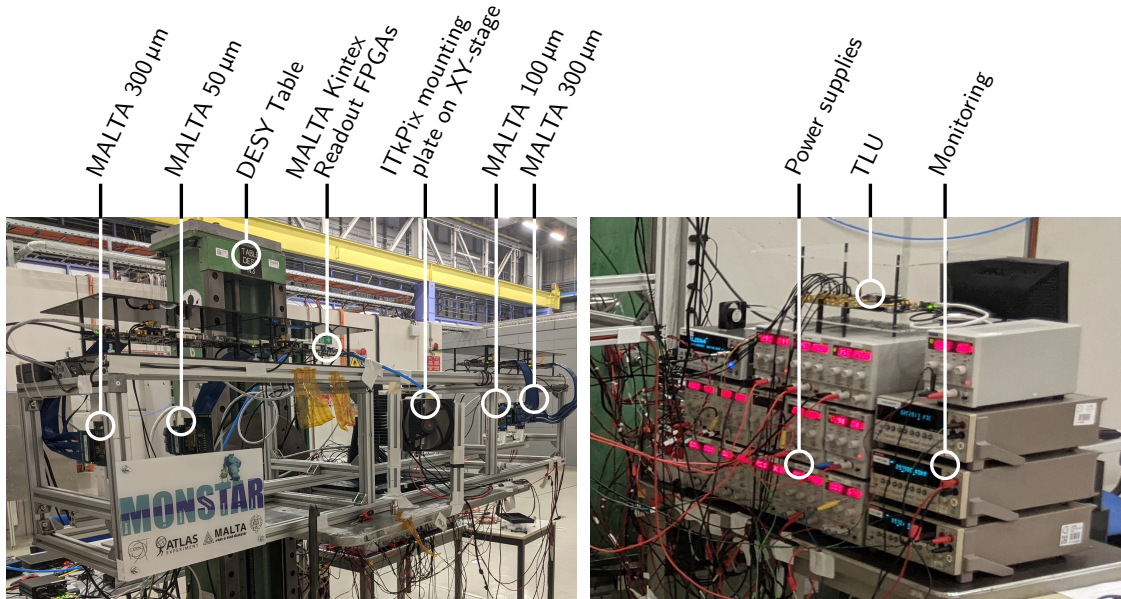



Figure 7.3: Annotated photograph of the MONSTAR telescope setup during the CERN PS testbeam measurement.


DUT, these planes allow for long-lever-arm tracking without requiring data from the DUT itself, potentially allowing the telescope to be used for radiation length measurements on non-instrumented subjects. Trigger signals are generated from the coincidence of all four MALTA planes, using the hit-or signal available on these planes. An annotated photograph of the telescope is shown in Figure 7.3.

MALTA planes require power to be provided to four separate domains (digital, analog, substrate, p -well) from LV power supplies. For the 300 μm planes, a substrate bias of 30 V was used, whilst the thinner planes were limited to a 6 V bias. Each plane was read out via a Xilinx Kintex development board using the MALTA data acquisition (DAQ) software and firmware that has been used in past MALTA-based telescopes [207], and trigger coincidence and distribution was handled by the MALTA trigger logic unit (TLU) loaded on another Xilinx Kintex development board.

The telescope planes and linear stage are connected to the outer frame by sliding mounting points for easy adjustment of the effective length of the telescope. The mounting



(a) Photograph of installation in telescope.



(b) Design render.

Figure 7.4: Photograph of the custom mounting plate for mechanical support and cooling of the ITkPix module, annotated with the relevant components. The Peltier elements hidden between the holder and heatsink are represented by the blue shaded regions, and the orange region shows the thermal contact surface between the module and plate. The size of the window beneath the module used for the measurement was 35 mm \times 25 mm.

plate shown in Figure 7.4 was designed to host an ITkPix quad module, and was based on a similar design used previously by the ATLAS testbeam group. Peltier elements and an external forced air convection fan for cooling were added to the plate, as well as mounting holes matching the linear stage. Air flow was directed by a custom-designed manifold with a dry air inlet connected to an N₂ feed to control the dew point of the air near the module. A rectangular window in the mounting plate below the module sensor allows for a region which can be measured with no interference from telescope material, with the remainder of the sensor in contact with a 500 μm -thick layer of aluminium. The adapter boards required for ITkPix readout were mounted on the holder, and were connected to LV power, HV sensor bias, and a HDMI cable providing five LVDS links (one common command uplink, and four downlinks).

Readout for the ITkPix chips was performed via optical link using the FELIX system with a single optoboard and relevant adapter boards³ [121], using a scheme analogous to that described and used in Section 5.2.1. Instead of the YARR readout software used for QC testing, a modified version of the ITk-FELIX-software framework was integrated into the MALTA MultiDAQ software [208, 209], which allowed for real time hit monitoring during data-taking. Trigger input was accepted directly into the FELIX system as an LVDS signal from the TLU, and a version of the ITkPix firmware for the FLX712 card was used that read out 16 bunch crossings with sequential IDs for each trigger [121]. The firmware was modified to reset the level-1 accept (L1A) counter back to zero after each readout sequence, allowing for easier detection of skipped or missing data frames.

The DAQ software was complemented by a simple DCS built on the Instrument Control Interface and Commands Library (ICICLE) [5], that continuously monitored currents, voltages, temperatures and humidity to ensure safe operation of the telescope, and allowed for remote control from outside the beam area. A Grafana monitoring interface was set up to display this information to shifters.

³A “Camilla” board — a predecessor of the “Zaza” board breaking out ERF-8 to DP — was used, which allowed connection of the five differential links via ten SMA cables.

7.2.1 Thermal considerations

As previously discussed in Section 4.3, ITkPix modules must be operated below a maximum temperature of 40°C and in a dry environment. A series of finite element analysis (FEA) simulations was performed, assuming a 12 W average power draw for the module, and chip to holder contact via a thin graphite thermal pad of thickness 0.127 mm with a heat transfer coefficient⁴ of $2.5\text{ W m}^{-1}\text{ K}$. The holder was designed to be clamped to the heatsink separated by a $6\text{ W m}^{-1}\text{ K}$ thermal pad of thickness 0.5 mm , and this configuration was reflected in the simulation. To represent the effect of the fan, forced air convection providing $20\text{ W m}^{-1}\text{ K}$ of cooling power was applied to the heatsink fins and backside of the module mounting plate and module window. Radiative heat transfer and still-air convection were neglected, and an ambient laboratory temperature of 20°C was assumed, which was later found to be an underestimate for the PS experimental area in the summer months.

Initial simulations were performed for a simplified geometry with no separation between readout chips, and no localisation of the heat generation in the readout chips. A check was also performed for several mesh resolutions, and showed stability for resolutions below 1.5 mm , which was then chosen as a baseline. Shown in Figures 7.5a and 7.5b, the simulations indicated feasibility of the cooling setup for window sizes up to $30\text{ mm} \times 30\text{ mm}$ whilst maintaining a headroom of 2 K . Further simulations were performed with a more

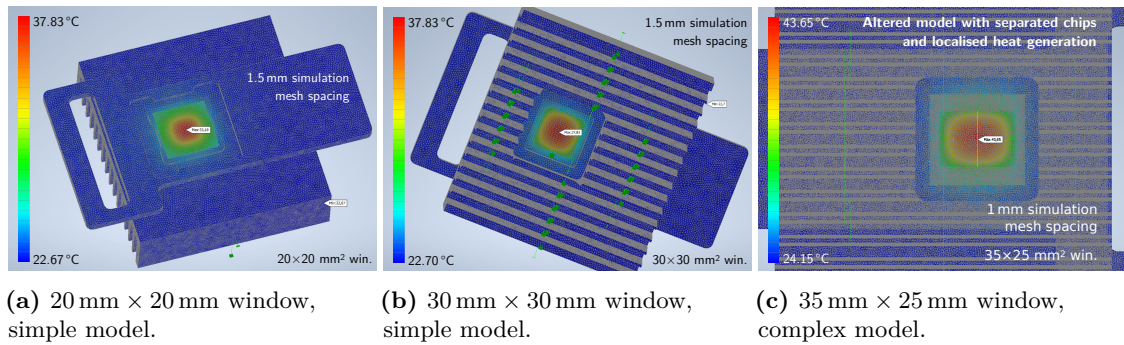


Figure 7.5: Thermal simulation results using the Autodesk Inventor® Nastran® simulation package for different window sizes [210]. The simulation in (c) used a more complex model with separated readout chips and a larger proportion of heat generation at the SLDOs.

⁴This value was underestimated compared to the relevant data sheet to account for potential lower-than-expected application pressure in this particular use case.




Figure 7.6: Aluminium holder (dark blue) and module (blue) temperature for the duration of the testbeam period.

detailed model that included the physical separation between the readout chips and a localised heat generation model, since $\sim 40\%$ of the total heat output comes from the SLDOs under the wirebond pads at the edges of the module. The refined model showed a $\sim 5\text{ K}$ higher temperature for all window sizes, but did show better results for a horizontal rectangular cutout allowing for more material under the SLDOs. A simulation of a $35\text{ }\mu\text{m} \times 25\text{ }\mu\text{m}$ cutout with the complex model is shown in Figure 7.5c.

Since these results demonstrated low thermal headroom in an uncertain environment even with small hole sizes, several variants of the module mounting plate were produced with window sizes of $20\text{ mm} \times 20\text{ mm}$, $25\text{ mm} \times 25\text{ mm}$, and $35\text{ mm} \times 25\text{ mm}$. Additionally, two Peltier element were added to provide additional cooling and operated in an underpowered but stable configuration to provide a few degrees of additional headroom. The operation of the cooling setup was tested in a laboratory environment and shown feasible with the mounting plate with the largest $35\text{ mm} \times 25\text{ mm}$ cutout, and the module was successfully operated below the 40°C temperature limit throughout the campaign, as shown in Figure 7.6.

7.2.2 Datasets collected

Two main measurement datasets were collected in the six-day testbeam period in June 2023, with differing telescope configurations summarised in Table 7.1. Campaign 1 used 700 MeV positrons with a 5% momentum band and the inner telescope planes positioned close to the outer planes, and imaged a $3\text{ cm} \times 3\text{ cm}$ region in the centre of the module. Campaign 2 was performed at 1.2 GeV to increase the particle rate, the inner planes were moved close to the DUT to allow for tracking without the DUT, and the full $4\text{ cm} \times 4\text{ cm}$

Table 7.1: Telescope configuration in Campaign 1 and 2. The z -position is measured from the first upstream plane, the centre of which defines the origin of the global coordinate system.

Telescope plane	Detector type	Sensor thickness (μm)	Camp. 1	Camp. 2
			z -position (mm)	
Plane 1	MALTA Cz	300	0	0
Plane 2	MALTA	50	194	864
DUT	ITkPix v1.1 quad	150	948	947
Plane 4	MALTA	100	1701	1027
Plane 5	MALTA Cz	300	1874	1838
Energy			700 MeV	1.2 GeV
Mom. band			5%	

area of the module was covered. In each campaign, the DUT was scanned across the beam spot using the $x - y$ stage, with runs taken in a grid of 9 positions for campaign 1, and 16 for campaign 2. A $1\text{ cm} \times 1\text{ cm}$ region was imaged at each position, and approximately 2 million triggers were collected for each. Additionally, measurements of the air in the telescope (no DUT) and of a $50\text{ }\mu\text{m}$ -thick MALTA sensor as a calibration sample were performed in the campaign 2 configuration as reference measurements. Unless otherwise specified, the methodology, results and simulations discussed here focused on the campaign 2 datasets, which formed the baseline analysis target.

7.3 Analysis methodology

The Corryvreckan testbeam analysis suite was used for data quality checks and alignment, as well as tracking [211]. A global frame was defined with z in the beam direction, and x and y chosen to align with the column and row directions of the DUT.

7.3.1 Data preparation, alignment, and data quality checks

Noisy pixels were masked per-run using the local density method, in which the occupancy of each pixel is compared to its immediate neighbourhood. If the pixel occupancy exceeded the local average by more than five standard deviations, it was masked and hits in this pixel disregarded. A correction was added to the method to account for the higher

occupancy of the central four rows and columns, which are double length to cover the inter-chip region on the quad module, and also exhibited high crosstalk.

After building clusters by grouping all pixels touching on edges or diagonally, a trigger synchronisation check was performed using the methodology described in [205]. The detection of desynchronisation and early termination of a run was implemented in a new Corryvreckan module following the statistical method described in Ref. [205]. No cases of desynchronisation were detected between MALTA planes, or between the MALTA telescope and the ITkPix DUT. To ensure validity of the method, a test analysis was performed of a run for which individual triggers had been intentionally removed from one plane.

Alignment was performed in two stages for each run, starting with a correlation-based pre-alignment followed by the application of the Millepede-II algorithm to perform residual minimisation in x , y and θ_z (the plane rotation about the z -axis) for each plane, with other degrees of freedom left constrained [212].

7.3.2 Tracking and angle extraction

The multiplet tracking method was utilised, in which straight line track segments (*tracklets*) are fit to hits *upstream* and *downstream* of the DUT, with the DUT hit (if in use) being included as both the final hit of the upstream tracklet and the first hit of the downstream tracklet. Uncertainties on the track origin and slope were computed from the linear fit, and propagated to the extracted angles between the tracklets. Two discrete analyses were performed on all datasets: one which included the DUT hit, and one which excluded it, and relied only on the data from the telescope planes. The former method shows some benefits since it allows precise localisation of the scatter point within the DUT and is hence less sensitive to scattering between the inner telescope planes and the DUT, as is discussed in the following sections. The latter method has the benefit that it could be applied on non-instrumented subjects, and hence could be used to measure detector mechanics or services, or calibration samples.

Two types of angles were extracted from the multiplet tracks. *Projected* angles θ_x and θ_y were extracted from the slope differences of the tracklets in the global $x - z$ and $y - z$

planes, respectively, which were chosen to align with the row and column directions of the ITkPix DUT. *Global* ϕ -invariant angles θ_g were also determined by computing the space angle between the upstream and downstream tracklet vectors.

7.3.3 MALTA half-chip efficiency correction

In some runs, the MALTA telescope planes suffered from light-induced efficiency differences between the left and right sides of the plane, as is visible in Figures 7.7a and 7.7b. Since the MALTA chips used were originally designed as prototypes, they use two different discriminator types for the left and right chip halves, which both suffered effective increases to thresholds from the light pollution, but to different degrees. The effect was only observed during morning hours with direct sunlight on the setup through the testbeam hall windows,

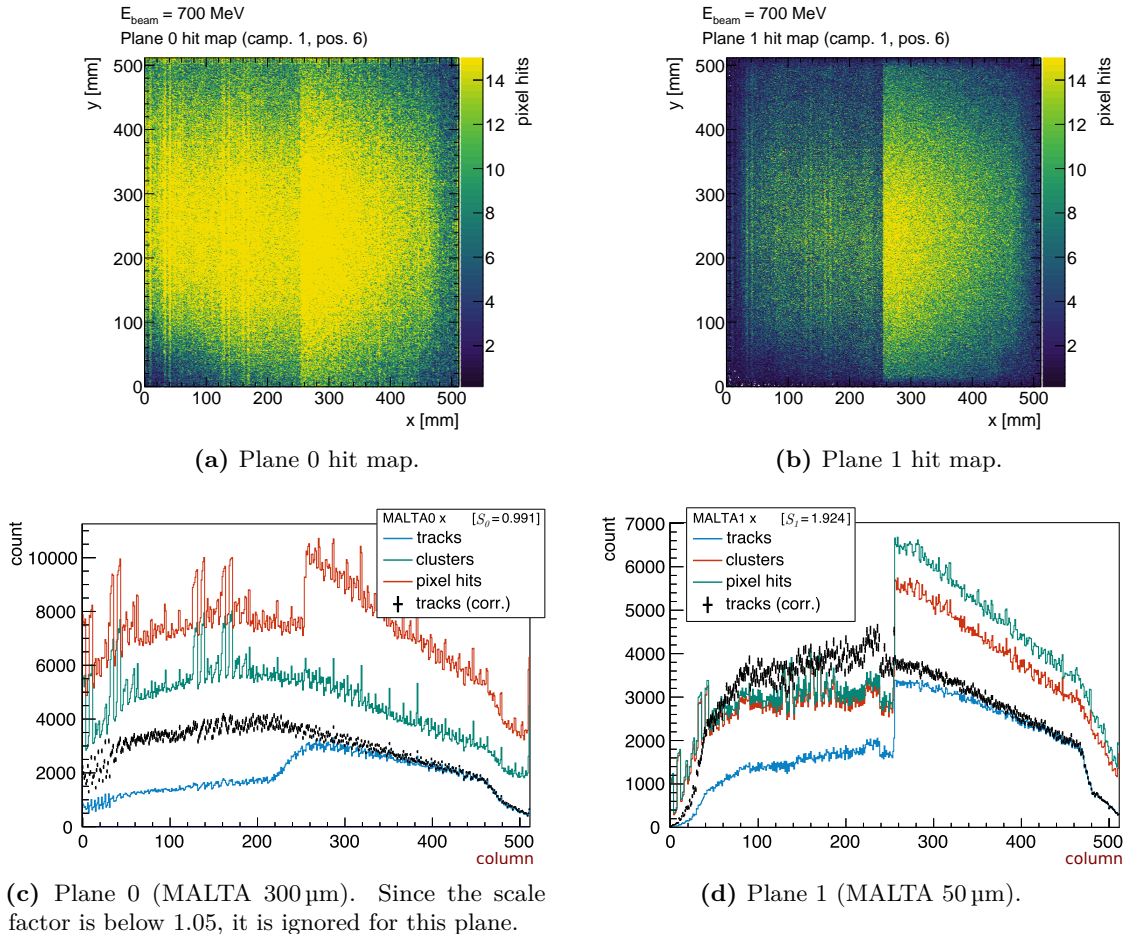


Figure 7.7: Light-influenced efficiency differences in MALTA half-planes for a heavily affected dataset from campaign 1, position 6. In (c) and (d), the scale factor derived from the plane is given in the legend, and the corrected track distribution scaled by the final corrective weight is shown in black.

and despite the addition of a shade cloth during data-taking, several datasets remained affected. It may be seen in Figure 7.7c that although the 300 μm sensors only suffered from a reduction in cluster size but did not appear to lose clusters, the 50 μm sensors in Figure 7.7d showed a large discontinuity in the number of observed clusters and tracks across this boundary. The 300 μm planes did show a soft discontinuity in the number of tracks, which was due to the correlation of track positions in these planes with the 50 μm planes when combined with the tracking requirement for hits in all four telescope planes.

The effect was corrected by taking the ratio of the number of clusters in the fifteen columns right of centre with those left of centre to derive a scale factor S_i for each MALTA plane i for a given run:

$$S_i = \left(\sum_{c=241}^{255} C_i(c) \right) / \left(\sum_{c=256}^{270} C_i(c) \right) \quad (7.1)$$

where $C_i(c)$ is the number of clusters in column c of plane i . Scale factors less than 1.05 were discarded, as a less than 5% deviation between the halves was not considered significant, and different runs were not equally affected. The scale factor for each plane was then applied as a weight to every track passing through the left (lower efficiency) half of the sensor, removing the discontinuity in the number of observed tracks, as may be seen in Figure 7.7. The final corrective weight w_{eff} applied to a track may be summarised as

$$w_{\text{eff}} = \prod_{i=0}^3 s_i(x_i^{\text{column}}), \quad \text{where } s_i(x_i^{\text{column}}) = \begin{cases} S_i, & \text{if } 0 \leq x_i^{\text{column}} \leq 255 \\ 1, & \text{if } 256 \leq x_i^{\text{column}} \leq 511 \end{cases} \quad (7.2)$$

where x_i^{column} is the track x-position at plane i , in the local column coordinates of the plane.

7.3.4 Telescope acceptance checks

In order to ensure large-angle scatters were not being undercounted toward the edges of the DUT measurement area due to the telescope acceptance, a generalisation of the method in Ref. [205] was applied to measure and correct for the acceptance. For each track, a cone anchored at the DUT position with axis equal to the upstream tracklet angle and opening angle θ_g was intersected with the downstream planes. For each downstream plane, the probability of an equivalent scatter with equal opening angle being accepted was calculated as the sum of subtended angles of all segments of the elliptical intersection

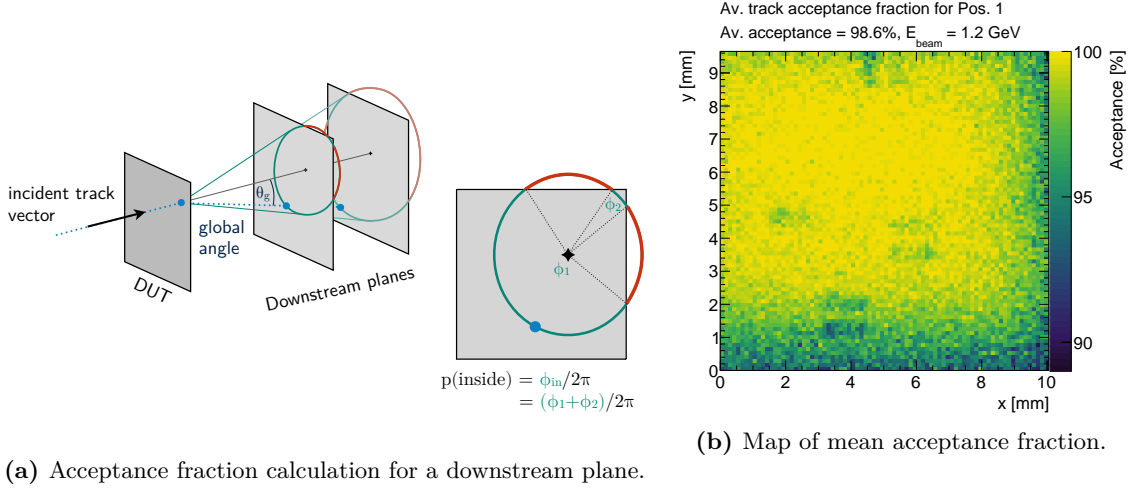


Figure 7.8: Acceptance fraction calculation, and average acceptance fraction for all measured regions in position 1 of campaign 1. Some module features are visible due to larger scattering angles caused by higher material content.

as a fraction of 2π , as shown in Figure 7.8a. The total acceptance probability for a given track is then the product of the individual probabilities for each downstream plane.

This fraction gives a probability of equivalent angle tracks being within the telescope acceptance, and is shown for an example position in Figure 7.8b. Overall, over 96% of tracks intersecting the measurement region on the DUT were estimated to be within the acceptance of the downstream planes for a measurement region of size $1\text{ cm} \times 1\text{ cm}$. However, to mitigate suppression of large angle scatters, a per-track weight equal to the inverse of the acceptance probability p_{acc} was applied to all tracks to help account for the “lost” tracks that fell outside the telescope acceptance,

$$w_{\text{acc}} = \frac{1}{p_{\text{acc}}} = \prod_{i=2}^3 \frac{1}{p_{\text{acc}}^i} = \prod_{i=2}^3 \left(\sum_j \frac{2\pi}{\phi_j^i} \right) \quad (7.3)$$

where p_{acc}^i is the acceptance probability in downstream plane i , which can be written in terms of the subtended angles ϕ_j^i of the j segments within the acceptance of plane i .

7.3.5 Fitting angle distribution widths

Extracted angles were grouped in $0.5\text{ mm} \times 0.5\text{ mm}$ geometric subregions based on the local track intercept within the DUT⁵. The choice of subregion size was driven by a

⁵In the case of the analysis without DUT data, this was extrapolated from the upstream and downstream tracklets by taking the centroid of the intercepts with the DUT z -position.

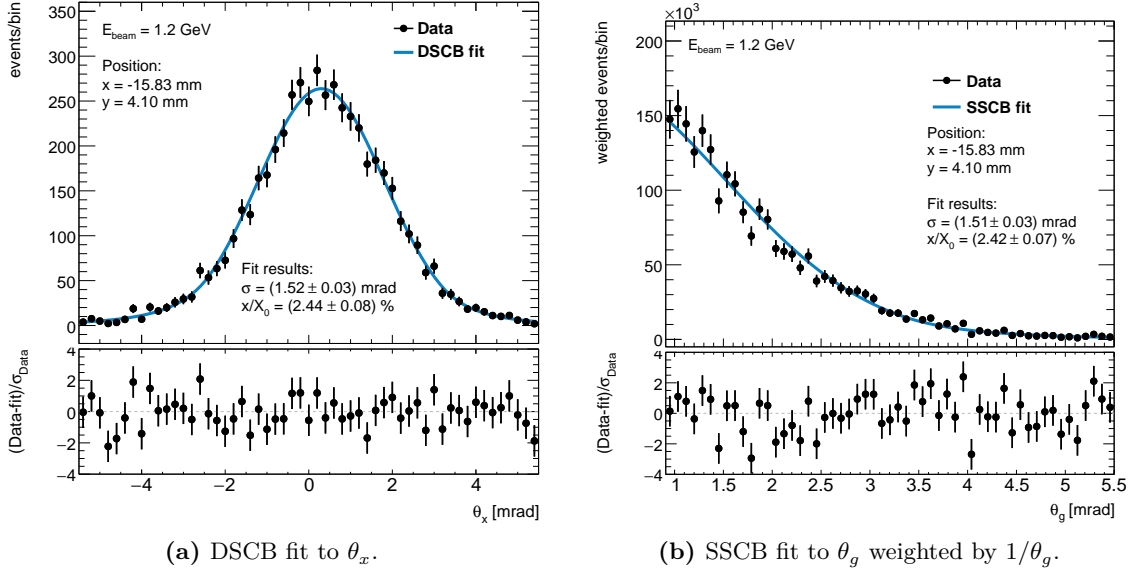


Figure 7.9: Angle distributions and Crystal Ball fits for an example region in the campaign 2 dataset. The track-weight corrections described in the previous sections have been applied.

compromise between the statistical uncertainty contribution in each subregion, and the resulting resolution of the x/X_0 map.

The projected angle distributions in each region were each fit with a double-sided Crystal Ball function (DSCB) consisting of a Gaussian core which goes over into adjustable tails of form $1/N$ past a cut-off α , and can hence model the $1/\sin^3 \theta$ tails of the Highland distribution well for small θ [213, 214]. Comparisons of binned fits on sample distributions generated from an underlying double-Gaussian distribution (following the parameterisation of the Frühwirth-Regler model described in Section 6.2.4) showed some dependence of the extracted width σ_θ on the binning parameters, which could be removed by a dynamic fit range adjustment, or by using an unbinned fit. An unbinned maximum-likelihood fitting method to the underlying angle data was implemented using the RooFit framework [215]. The fractional radiation length was then computed from σ_θ via the inverse Highland




Figure 7.10: Process of transforming the global angle θ_g to a pseudo-projected angle.




Figure 7.11: Global angle distribution prior to the reweighting process in Figure 7.10, for an example region in the campaign 2 dataset. The track-weight corrections described in the previous sections have been applied.

formula in Equation (6.19), or via a numerical inverse of the Fröhlich-Regler formalism in Equation (6.29). A fit for an example region is shown in Figure 7.9a.

In order to be able to treat the global angle distribution analogously to a projected distribution, each event was weighted by the inverse of the circumference of a circle with radius θ_g ⁶. As shown visually in Figure 7.10, this produces an effective radial projection of the data whilst maintaining the ϕ -invariant nature of θ_g . The global angle distribution prior to this reweighting process is shown in Figure 7.11. Since the global angles forming this distribution are necessarily positive, the distribution was then fit with a single-sided Crystal Ball function (SSCB), as shown in Figure 7.9b. The fit was adjusted to ignore the smallest 20% of angles due to the large weights on these tracks as θ_g becomes very small. In combination with the lower number of tracks at small global angles, these effects produce large effective uncertainties at the peak of the SSCB distribution which had a negative impact on the fit stability.

Further methods of extracting σ_θ were explored, but were not found to provide substantial benefits over the DSCB and SSCB fits. Some of these are discussed in Section 7.5 below.

Figure 7.12a shows the extracted width and Figure 7.12b a corresponding x/X_0 map for the DSCB method on the θ_x distributions reconstructed from the campaign 2 dataset. For many of the following results, since the θ_x and θ_y distributions (and extracted σ_θ for each) are independent, they have been combined as independent measurements. The

⁶The global factor of $1/2\pi$ applied to all events can be neglected.

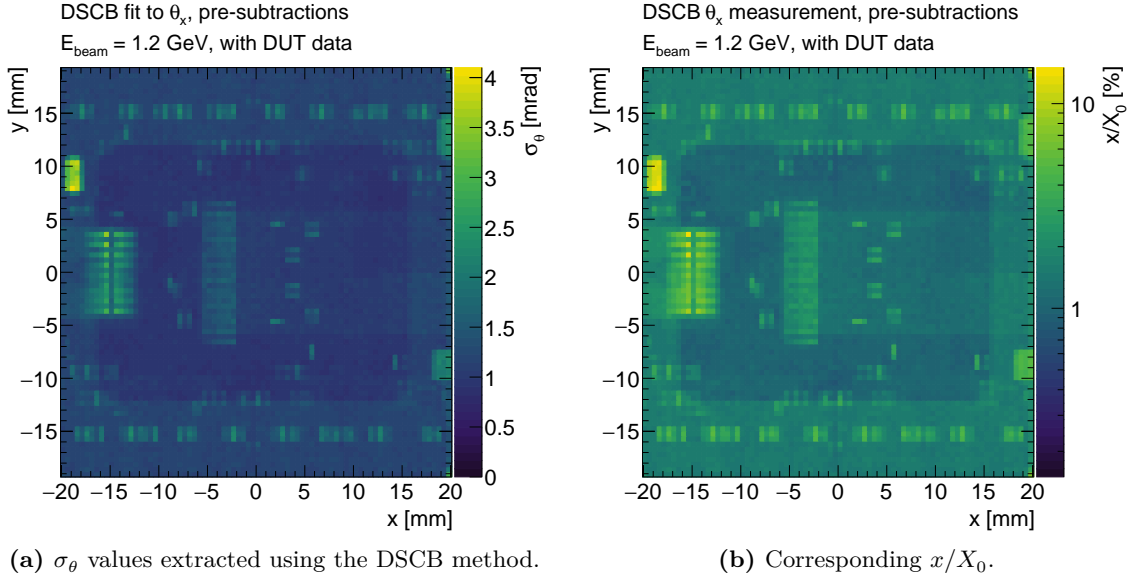


Figure 7.12: Scattering angle and radiation length maps extracted from the campaign 2 dataset, prior to the subtraction of contributions from telescope mechanics and air. The x/X_0 map in (b) was derived via the inverse Highland formula. Note that the map in (a) uses a linear color scale, whilst (b) uses the same logarithmic scale as for all result plots in the following sections.

statistical uncertainties are treated as completely uncorrelated, whilst the systematics are assumed to be fully correlated as the two distributions result from the same particles.

7.3.6 Mechanics and air contribution subtractions

In order to achieve a pure measurement of the ITkPix module, the auxiliary material contributions from the testbeam mechanics and air within the telescope must be subtracted. For the analysis utilising the hit in the DUT, the air measurement was derived from a measurement of a 50 μm reference MALTA plane shown in Figure 7.13a, subtracting the expected radiation length to determine the air contribution. The MALTA sensor thickness was estimated as consisting of 50 μm of silicon plus a 25 μm polyimide window protecting wire bonds, and a 1 μm metallisation layer. The extracted contribution from the air was computed as $(x/X_0)_{\text{air+tele}} = \langle x/X_0 \rangle_{\text{meas}} - (x/X_0)_{\text{exp}} = (0.42 \pm 0.10)\%$.

For the analysis not using the hit in the DUT, the same methodology can be applied on the same data processed with the alternative method. As may be seen in Figure 7.13b, the air contribution increases in this configuration to $(0.46 \pm 0.03)\% X_0$. The increase in both contribution and uncertainty may be attributed to the impact on track building of the DUT hit. For small-angle scatters from air close to the DUT, the inclusion of the DUT

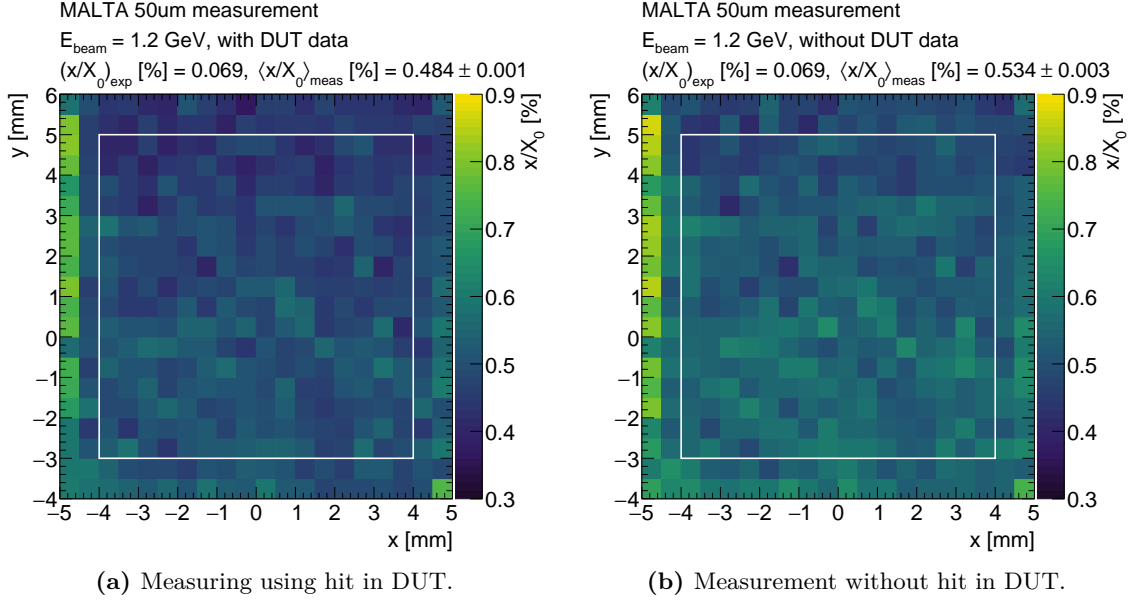


Figure 7.13: MALTA 50μm measurement, using the campaign 2 telescope configuration and hit data from the DUT in the track fit. The quoted mean and uncertainty only include data within the white rectangle to avoid including regions of overlapping PCB close to the edges of the sensor.

hit acts as a corrective factor, reducing the impact of this scatter on the upstream and/or downstream track segments. Without the information from the DUT hit, the scatter is not considered in the upstream or downstream track fits, and instead acts as a stronger additive factor to the scattering within the DUT itself. In this case, the extracted value

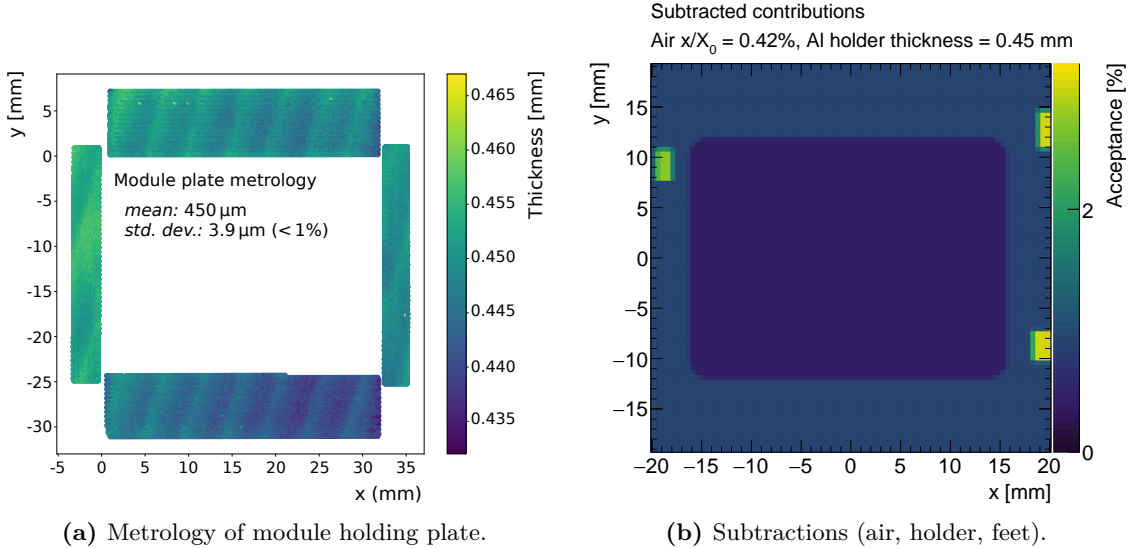


Figure 7.14: Contributions subtracted from the final measurement. The MALTA 50μm measurement shown in (a) was used to derive the air contribution, the metrology in (b) measures the module holder thickness, and the sum of all subtracted components is given in (c).

has been compared to direct measurements of scattering in air using only the four telescope planes, as both use identical telescope configurations, and showed very good agreement.

Figure 7.14a shows a metrology measurement of the thickness of the module mounting plate section surrounding the window. This was also validated by checking the size of the discontinuity in the pre-subtraction results, as the window gives a possibility to compare sections of similar material content with and without the additional contribution from the holder. The feet on the 3D-printed clamp used to hold the module were dissected and individually measured and weighed to determine exact dimensions and density, the results of which are given in Table A.3 in Appendix A. The relevant contributions were then estimated based on radiation length values for aluminium and polylactic acid (PLA)⁷, and together with the air and telescope contributions are summarised in Figure 7.14b.

7.4 Material estimate

A high-granularity estimate of the material of the quad module was produced from the design documentation by producing 2D maps of material content for 57 layers of contributions including sensor (bulk and metallisation), readout chip (bulk and metal

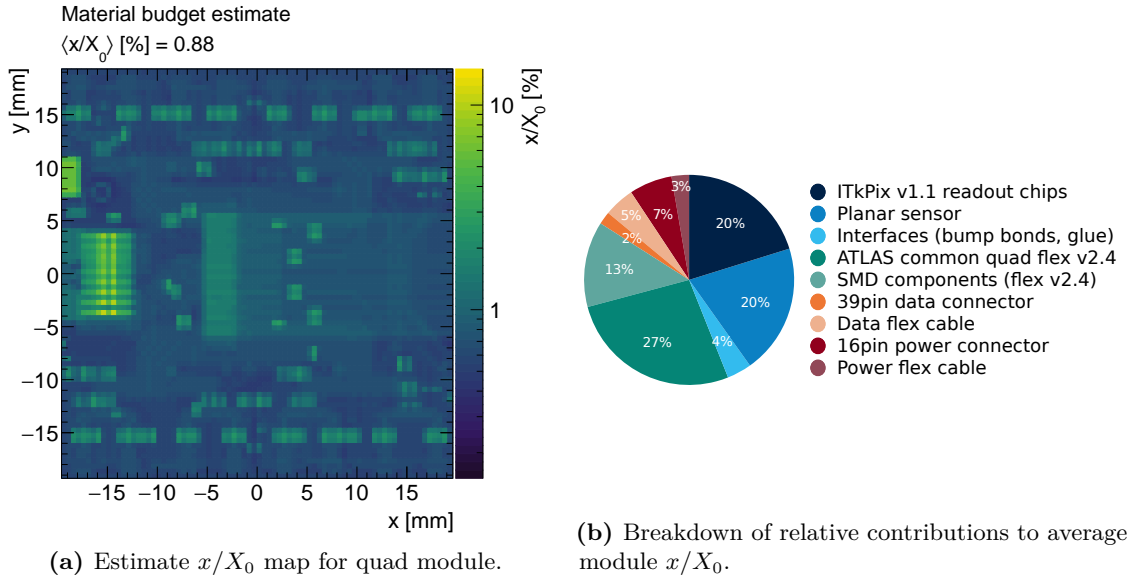


Figure 7.15: Expected x/X_0 for a prototype ITkPix pixel quad module comprising four ITkPix v1.1 readout chips, a 150 μm -thick planar sensor, and a quad flex v2.4 populated with the relevant SMD components. The preliminary data and power flex cables are included, but do not correspond to the final design or routing for these components.

⁷Using the analytic formula given in Ref. [216] on polylactide — $(\text{C}_3\text{H}_4\text{O}_2)_n$. See Table A.1.

layers), interfaces, flex, surface-mount (SMD) components, and connectors. Literature and calculated radiation length values used for the estimate are tabulated in Appendix A. Several components were identified as *uncertain elements* due to limited information on material content, or missing reference radiation length values, and these were found to contribute 21.8% of the total estimate.

The estimate and a breakdown of relative contributions of components is shown in Figure 7.15, and the overall average expected for the sensor-covered area of the module flex including the power and data flex cables is $0.88\% X_0$. The data connector and HV filter capacitor represent the largest localised contributions, and the thinnest region is expected within the HV wirebond hole.

7.5 Geant4 simulations of air and telescope scattering

Monte-Carlo simulations of scattering behaviour within the telescope have been produced using the GEANT4 toolkit [199] via the Allpix² detector simulation framework [217]. A telescope geometry matching that used for the analysis was implemented, and relevant mechanical structures within the telescope beam path⁸ were added. As the goal of the study was to characterise the telescope response across a range of localised material thicknesses, the complex structure of the ITkPix module was not simulated, and was instead replaced by a generic monolithic detector with the same pixel pitch, consisting purely of silicon. The Geant4 FTFP_BERT_EMZ physics list was used as it uses the same *Wentzel-VI* model as the standard ATLAS simulation for positrons at 1.2 GeV [200], but with step limits that are more precise for very thin subjects, as is discussed in Section 6.2.5.

To simplify the simulation and decouple the effects of interest — the impacts of auxiliary scattering, spatial resolution, and telescope acceptance — a new *TruthTransfer* module was implemented for Allpix². This module allows the usual steps of charge deposition, propagation and transfer to front-end to be skipped, instead directly determining a front-end charge from the location of a truth-level track through the sensitive volume. An average charge quantity Q of $107 e^-/\mu\text{m}$ is directly deposited into the intersecting pixel, or into the

⁸The retaining PCBs behind the edges of the MALTA sensors, and the 25 μm thick polyimide tape windows used for wirebond protection of the MALTA telescope planes.

2×2 neighbouring pixels via a charge-sharing split. An inverse of the common center-of-mass clustering method was derived [45], with allocated charge fractions $Q_{i,j}$ determined by the distance between the track passing through (x, y) and each pixel centre (i, j) ,

$$\begin{aligned} Q_{\lfloor x \rfloor, \lfloor y \rfloor} &= (1 - r_x)(1 - r_y)Q, & Q_{\lfloor x \rfloor, \lceil y \rceil} &= (1 - r_x)r_yQ, \\ Q_{\lceil x \rceil, \lfloor y \rfloor} &= r_x(1 - r_y)Q, & Q_{\lceil x \rceil, \lceil y \rceil} &= r_x r_y Q, \end{aligned} \quad (7.4)$$

where x, y, i, j are in pixel coordinates, and the 1-dimensional charge distribution ratios r_x, r_y are defined as

$$\begin{bmatrix} r_x \\ r_y \end{bmatrix} := \begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \lfloor x \rfloor \\ \lfloor y \rfloor \end{bmatrix}. \quad (7.5)$$

This guarantees recovery of the true track position (x, y) when using the usual center-of-mass charge sharing algorithm⁹,

$$\frac{1}{Q} \left(\sum_{i=\lfloor x \rfloor}^{\lceil x \rceil} \sum_{j=\lfloor y \rfloor}^{\lceil y \rceil} Q_{i,j} \cdot \begin{bmatrix} i \\ j \end{bmatrix} \right) = \begin{bmatrix} \lfloor x \rfloor + r_x \\ \lfloor y \rfloor + r_y \end{bmatrix} = \begin{bmatrix} x \\ y \end{bmatrix}. \quad (7.6)$$

This method also allowed simulations to be performed that excluded specific scattering contributions (for example, from the telescope planes) by replacing relevant detector material with unphysical ones such as vacuum. Since the readout charge depended only on the existence of a track through the sensitive volume, and not on a realistic deposition or propagation model, this allowed simulations of a “ghost telescope” that did not impact passing tracks, or of active DUTs in alternative materials such as polyimide or copper.

A realistic threshold was used for digitisation matching data-taking conditions of $300 e^-$ for the MALTA planes, and $2000 e^-$ for the ITkPix. Since both detector types only supported binary readout¹⁰, a one-bit charge-to-digital conversion (QDC) was configured for the digitisation. No attempt was made to simulate detector noise, misalignment, or threshold variances for the MALTA chip halves, as these are difficult to generalise and were not expected to have a significant impact on the extracted x/X_0 values. Since the beam momentum band correlates with spatial directions due to the effects of the momentum selection dipoles, the momentum variation was also not included in the simulation. The same analysis chain was utilised as for the main data, but skipping the noise masking, alignment, and efficiency correction, as these steps are unnecessary for the simulation data.

⁹Assuming unlimited precision on the resolution of the pixel charges $Q_{i,j}$.

¹⁰The ITkPix v1.1 chip is subject to a bug that prevents correct time over threshold readout.

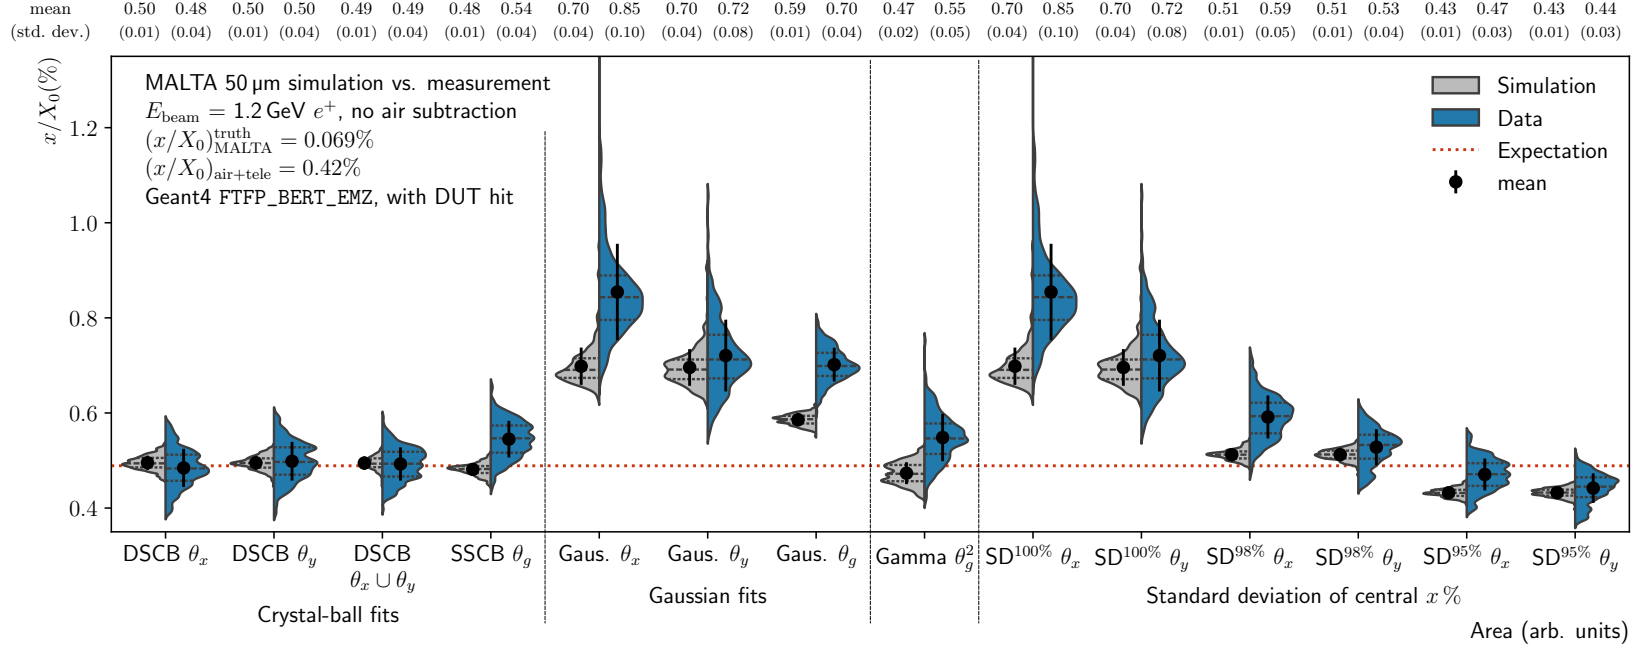


Figure 7.16: Comparison of σ_θ extraction methods, shown as x/X_0 spectra where the vertical axis represents a measured x/X_0 value, and the horizontal width of each distribution the fraction of the measured area with the given value. The dashed horizontal line gives the expected x/X_0 for the 50 μm MALTA target. The widths of extracted scattering angle distributions from simulation (grey) and measurement (blue) datasets of a 50 μm MALTA plane are estimated through several methods. The first four methods include double-sided crystal-ball fits to θ_x or θ_y , or a combination of the two measurements, and single-sided crystal-ball fits to θ_g (in which each event is weighted by $1/\theta_g$). Gaussian fits to these same distributions are given as a baseline comparison, as are the standard deviation of the central 100%, 98%, and 95%. A Gamma distribution fit to θ_g^2 used in [205] is also given as an alternative to the SSCB method. All datasets consist of 18×18 subregions of size $0.5 \text{ mm} \times 0.5 \text{ mm}$, with about 5000 tracks per subregion.

7.5.1 Comparison of scattering distribution width extraction methods

In order to validate the telescope model, equivalent simulations were compared to both the datasets containing only air, and those using a 50 μm MALTA plane as DUT. Translation back to an x/X_0 value from the Gaussian width σ_θ was performed using the inverse Highland formula in Equation (6.19). In both cases, the expected and observed scattering agreed within uncertainties for the Crystal ball based fit methods. A comprehensive comparison of several potential width extraction methods was performed, and a comparison of the simulation and measurement results for each is shown in Figure 7.16. As a baseline, Gaussian fits were performed of both the projected angles and the reweighted global angle, and in addition the standard deviations of the central 100%, 98%, and 95% were also explored as potential estimators. The latter methods mostly showed substantial disagreement between extracted values and the simulated material thickness, and generally exhibited a larger spread in values. Despite the standard deviation of the central 98% agreeing reasonably well with the simulation dataset (as may be expected from Lynch and Dahl’s definition of Equation (6.18) discussed in Section 6.2.3), it does not remain stable on a realistic dataset.

The DSCB method shows high precision as an estimator on both the simulated and real datasets, but does slightly underestimate the expected value. This deviation was small compared to the uncertainties discussed in the following section, and may have been attributable to differences between the Highland formula and the *Wentzel-VI* scattering model used in the GEANT4 EMZ physics list¹¹. The SSCB model showed some variation between the central values of the data and simulation, overestimating the true value in the former case and underestimating in the latter. Although this indicates some instability, the deviation was within uncertainties, and the method was still found to be a useful point of comparison as it does not encode projection onto any specific geometric direction, and was found to be more stable at hard borders along the x or y directions, where the perpendicular projected angle distribution in some cases showed substantial asymmetries.

Some further potential fit functions including double-Gaussians and Gaussians constrained to a predetermined central fraction of data were explored, but the former were

¹¹The GEANT4 simulation agreed better with the Fröhlich-Regler formalism, as is shown in Figure 7.17.

found to be very sensitive to statistical fluctuations in the tails of the angle distribution, and the latter were very dependent on the precise choice of the central fraction, and showed stronger variability on realistic datasets. The Gamma distribution fit to θ_g^2 used in [205] was also explored, but was found to disagree between data and simulation in a similar way to the SSCB method, but with a broader variance in extracted X_0 values.

7.5.2 Telescope response

The telescope response was simulated as a function of silicon DUT thicknesses ranging from $0.01\% X_0$ to $20\% X_0$, and the response function for the DSCB fit method is shown in Figure 7.17. The Highland and Frühwirth-Regler models are shown for comparison, and the predicted σ_θ agree to within 3% across the range studied. The combined effect of the air and telescope scattering is clearly visible in the response curve (blue), and the relative impact of this contribution increases at low target X_0 as would be expected.

The air and telescope scattering contributions were subtracted by transforming σ_θ to an X_0 value using the inverse Highland formula, then subtracting a contribution of




Figure 7.17: Simulation of σ_θ for various assumed true x/X_0 values of a homogeneous silicon DUT, estimated using DSCB fits to θ_x and θ_y . The telescope response is shown before (blue) and after (orange) the air + telescope subtraction. The latter is compared to expectations from Highland (black dashed line) and Frühwirth-Regler (purple dotted line). The reference measurement with a $50\ \mu\text{m}$ thick MALTA plane as the DUT is shown as a data comparison point for each response curve (pre- and post-subtraction). The Highland and Frühwirth-Regler lines overlap on the left plot of the figure, and their difference is only visible in the ratio plot on the right of the figure. A third-order polynomial fit to the response curve is shown (dark blue dashed line).

$(x/X_0)_{\text{air+tele}} = 0.42\%$, and then transforming back to σ_θ with the forward Highland formula. The air-subtracted response follows the Fröhwrth-Regler model closely for the central part of the investigated range, and deviates at high and low values. At high values of x/X_0 above about 15%, the limited acceptance of the telescope leads to an underestimation of σ_θ due to missing tracks in the tails of the wider scattering distribution. For low values of x/X_0 below 0.1%, the telescope resolution limits on the extracted angle lead to an overestimation of x/X_0 , and the relative contribution of uncertainties from the air and telescope scattering become dominant compared to the measurement subject.

To allow the use of the simulated telescope response curve as an alternative to the Highland or Fröhwrth-Regler models when extrapolating extracted σ_θ values back to x/X_0 , a third-order polynomial fit to the response curve was performed in lieu of an interpolation. The resulting functional forms for the DSCB and SSCB methods are shown in Figures 7.18a and 7.18b, respectively, together with fourth-order polynomial fits. No significant improvement was seen with a quartic interpolation function. The addition of this simulation-based model ensures that the obtained results could be used as inputs to a GEANT4-based simulation (such as the ATLAS detector simulation) to reproduce the same scattering profile seen in the testbeam measurement.

The telescope response curve changes slightly with exclusion of the DUT hit data from the analysis due to the larger air scattering contribution. The relevant response curves for the SSCB model and the analysis without DUT hit are available in Appendix B.

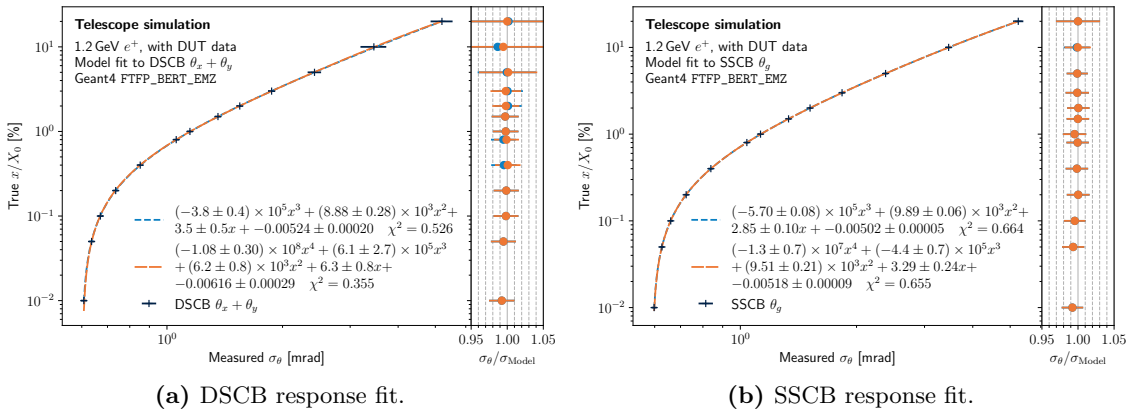


Figure 7.18: Polynomial fits to telescope response curve, for both the DSCB and SSCB fit methods for $\theta_{x/y}$ and θ_g , respectively. Both third-order and fourth-order polynomials are shown, and a ratio of the interpolation model to the simulation datapoints is shown.

7.5.3 Air and telescope scattering validation

It is interesting to note that the empirically derived value for $(x/X_0)_{\text{air+tele}}$ of 0.42% from Section 7.3.6 is significantly less than what would be expected from a naïve addition of the total air volume as a linear contribution to the scattering. Calculated from a nominal radiation length of dry air at atmospheric pressure of 3.039×10^4 cm [218], the expected contribution is 0.61% X_0 . This is already less than the observed value prior to the inclusion of the nominal thicknesses of the MALTA telescope planes, which themselves comprise an expected total of 0.859% X_0 . It was concluded that the air and telescope contributions must be some function of the geometry, and the simulation was used to validate the air scattering estimate and investigate the behaviour of these contributions.

Simulated datasets in which scattering from the telescope and air was removed by replacing both telescope planes and the surrounding environment by vacuum were compared with ones using the standard telescope definition to extract $(x/X_0)_{\text{tele+air}}$. This was repeated for a range of DUT thicknesses to test the additivity of the air scattering contribution with the DUT contribution. Figure 7.19 shows the total contribution of both air and telescope scattering, which agrees well with the value of $(0.42 \pm 0.01)\% X_0$ derived from the 50 μm MALTA dataset. No significant correlation with DUT x/X_0 is observed,




Figure 7.19: Combined air and telescope contribution to total measured x/X_0 as a function of simulated DUT x/X_0 . The shaded bands show the standard deviation of the extracted $(x/X_0)_{\text{air+tele}}$ across 20×20 subregions of size $0.5 \text{ mm} \times 0.5 \text{ mm}$, each containing a number of tracks equivalent to the average subregion statistics collected in the measurement.

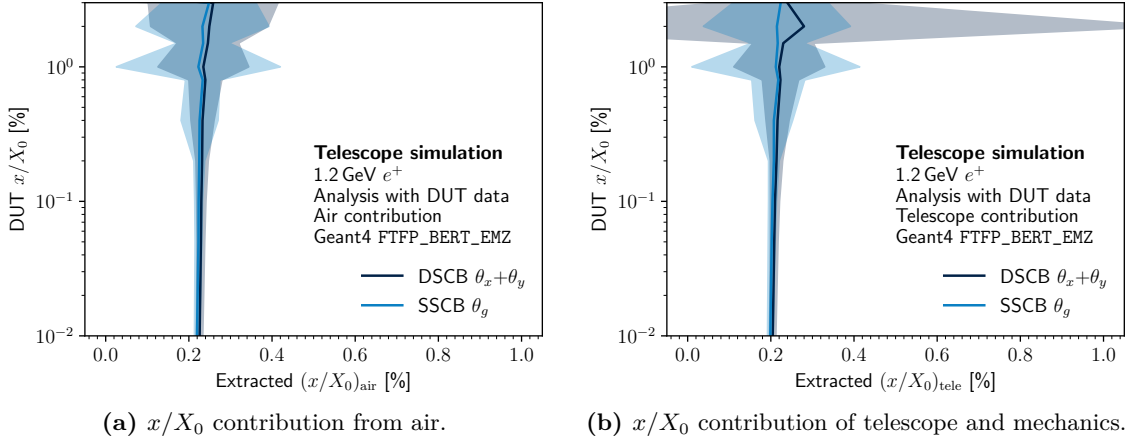


Figure 7.20: Separated air and telescope contributions to measured x/X_0 as a function of simulated DUT x/X_0 . The shaded bands show the standard deviation of the extracted value across 20×20 subregions of size $0.5 \text{ mm} \times 0.5 \text{ mm}$, each containing a number of tracks equivalent to the average subregion statistics collected in the measurement.

except that the variance on the extracted value for the air increases substantially as the relative contribution to the total uncertainty of the high- x/X_0 DUT becomes dominant, and plateaus below the point where $(x/X_0)_{\text{DUT}}$ equals $(x/X_0)_{\text{air+tele}}$.

The simulation confirmed an expected intrinsic variability within the air contribution of about $0.01\% X_0$ in the limit of a vanishingly thin DUT with the statistics and subregion sizes used in the measurement, which was in agreement with the statistical uncertainty observed on the measured dataset. The sources of systematic uncertainties on the air contribution to real datasets are necessarily completely correlated with the systematic uncertainties on the total extracted x/X_0 .

Simulations were repeated with only the telescope planes or surrounding environment replaced by vacuum, which allowed the contributions from each to be decoupled, as is shown in Figure 7.20a for $(x/X_0)_{\text{air}}$ and in Figure 7.20b for $(x/X_0)_{\text{tele}}$. The relative contributions were close to equal, with scattering from air comprising 52% of the total air and telescope contribution. The alternative analysis excluding the DUT data was also applied to all datasets, and the extracted value of $(x/X_0)_{\text{air+tele}} = 0.46\% X_0$ in this case again matched the derived value from data. The relevant results are shown in Appendix B, in Figures B.5 and B.6.




Figure 7.21: Air and telescope contribution to total measured x/X_0 as a function of the distance between the inner telescope planes and the DUT. A symmetric nominal telescope configuration was used, in which the DUT is 930 mm from the most upstream plane, and the most downstream plane is 1860 mm from the most upstream plane. The plane spacings corresponding to the campaign 1 and campaign 2 configurations are annotated. The shaded bands show the standard deviation across the 20×20 simulated subregions. A vacuum detector model with no scattering contribution was utilised as DUT for this study.

7.5.4 Variation of air and telescope contributions with plane spacing

Qualitative differences in the air and telescope scattering contributions were observed between the two different telescope configurations for which data had been collected. As a result, these contributions were studied in simulation as a function of the distance between the inner planes and the DUT. As is shown in Figure 7.21, the contribution originating from the telescope material is observed to decrease as the inner planes are moved toward the outer planes, whilst the air contribution remains relatively constant, with a slight peak when the inner planes are positioned one quarter of the way to the outer planes.

7.6 Results

The resulting fractional radiation length map post-subtractions is shown in Figure 7.22a, and a comparison to the estimate is shown in Figure 7.22b. For the comparison, the measurement and estimate are smeared by a Gaussian filter with $\sigma = 0.5$ mm to soften effects due to slight differences in component placement in reality compared to the

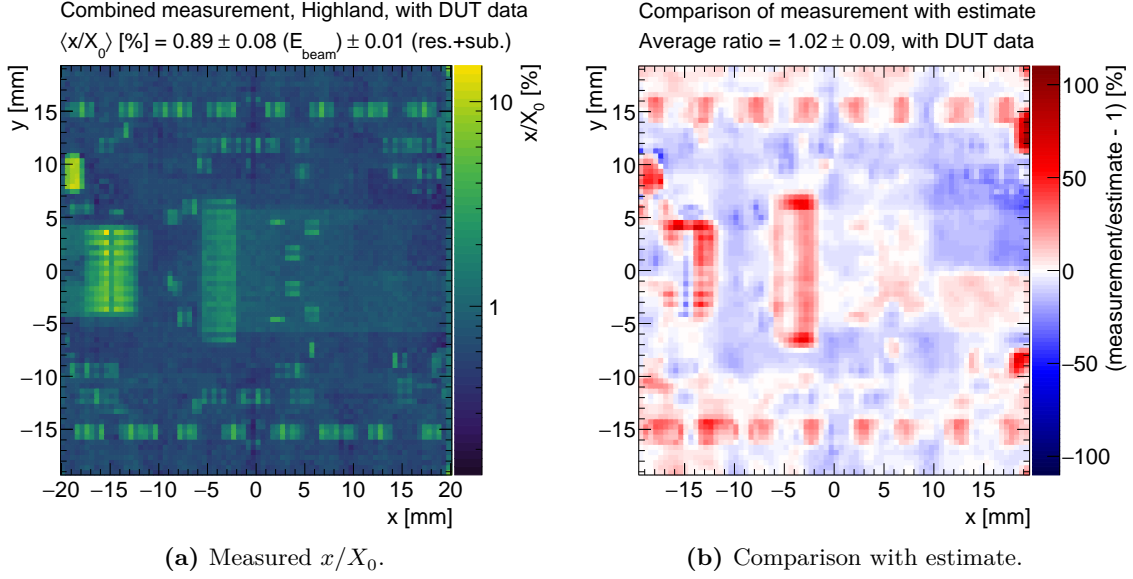


Figure 7.22: Measured fractional radiation length x/X_0 for an ITkPix v1.1 quad module from data collected in campaign 2 at 1.2 GeV. The values obtained from DSCB fits to θ_x and θ_y are combined as independent measurements. For the comparison in (b), both measurement and estimate were smeared by a Gaussian filter with $\sigma = 0.5$ mm to soften placement differences.

nominal position used for the estimate¹². The most significant differences are observed at the power and data connectors, HV decoupling capacitor, and other SMD components, the locations of which are shown in Figure 4.3. As discussed above in Section 7.4, these elements are difficult to produce accurate spatially-resolved estimates for, and can vary significantly between individual exemplars, depending on the internal tolerances used during manufacturing for each component¹³. The module-wide averages and most subregions for the measurement and estimate agree to within uncertainties, with the average fractional radiation length of the module measured as

$$\left\langle \frac{x}{X_0} \right\rangle_{\text{meas}} [\%] = 0.89 \pm 0.01 (\text{reso.}) \pm 0.01 (\text{subtraction}) \pm 0.08 (E_{\text{beam}}).$$

7.6.1 Resolution and beam energy uncertainties

The resolution uncertainty was determined by building toy angle distributions for each subregion by resampling each scattering angle from a Gaussian with width equal to the kink angle uncertainty propagated from the track fit. 500 toy distributions were built and fit per subregion, and the standard deviation of the resultant x/X_0 values was taken

¹²Such deviations are likely due to the pick-and-place accuracy during flex population.

¹³Which are often not shared publicly by the vendors in question.

as the resolution uncertainty in that subregion, which are shown in Figure 7.23. Since the dominant contribution to resolution is expected to be the pixel pitch, the relative resolution uncertainty should be uniform across the module (except the double row/columns in the centre), and the global average of the relative resolution uncertainty was found to be 0.99%.

The momentum band systematic was determined by propagating the beam momentum band of 5% (as set in the beamline configuration) through the Highland formula, resulting in a global relative uncertainty of 9.15%, which was by far the largest contribution to the systematic uncertainty in all bins. This was verified by reinterpreting the air measurement calibration dataset as an energy measurement (assuming uniform x/X_0) and computing the standard deviation, which was found to be 59 MeV, or equivalently 4.9%.

The statistical uncertainty was taken from the fit uncertainty in x/X_0 , and was negligible on the average across all subregions. The mean local statistical uncertainty across subregions was 0.05% X_0 .

The uncertainty in the subtraction was estimated from the uncertainty in the measurements or estimates of the individual components, including the air, aluminium holder, and 3D-printed clamp, as described previously in Section 7.3.6. These were combined as independent to give a final uncertainty from the subtractions of $\leq 0.01\% X_0$.




Figure 7.23: Relative resolution uncertainty for the DSCB fits to θ_x , derived from the variance of fits to 500 toys within each subregion. Toys were generated by resampling all angles in a given subregion from a Gaussian distribution with width equal to the track kink angle uncertainty.

7.6.2 Comparison of fit methods and scattering models

Figure 7.24 shows a comparison of the baseline methodology — a combination of DSCB fits to θ_x and θ_y and the inverse Highland formula — to alternative scattering models and the SSCB fit method on θ_g , in the form of x/X_0 spectra for each. The estimate is also shown, both in a pure form, and smeared by an uncertainty model matching the uncertainties derived for the measurement (rounded to 10% systematic and 2% statistical in each subregion) to allow better comparison to the results. The spectral forms are very similar in all cases, indicating no significant differences between the various estimated or measured compositions of the module. The measurements do show a stronger tail toward small values, which were determined to come from subregions close to the inter-chip region and the edge of the module. This may indicate variation of the module measured from the nominal design in terms of inter-chip spacing or dicing.

The Highland formula yielded a 6% to 7% higher average x/X_0 compared to the Frühwirth-Regler formalism or the Geant4 EMZ response curve model, but still agrees with the alternative models within uncertainties.




Figure 7.24: Comparison of the estimated x/X_0 distribution of the ITkPix quad module with three different extraction methods, based on the Highland formula, the Frühwirth-Regler formalism, and a Geant4 based mapping of the measured θ to the underlying x/X_0 values. The estimate is shown both in its pure form, and with an added Gaussian smearing emulating a 12% relative uncertainty in each subregion (10% syst. + 2% stat.) to match the expected distribution for the data. The SSCB method applied to θ_g is shown as an alternative reconstruction method for the width of the angle distribution, and the analysis including DUT data (blue) is contrasted with an alternative analysis that excludes this data from the track reconstruction (orange). The mean and uncertainty for each x/X_0 distribution are shown and quoted, and the dashed lines give quartiles.

As was expected from the simulations discussed earlier, the SSCB model generally yielded a distribution shifted to lower values, but still agreed within systematic uncertainties. The separate Geant4 fit models derived for the DSCB and SSCB cases agreed very well with each other, as may be expected from their method of derivation.

7.6.3 Analysis without DUT data

The analysis performed without DUT hits showed good agreement with the baseline for all fit methods and scattering models, as may be seen in Figure 7.24. As mentioned previously, a larger air value of $0.46\% X_0$ was subtracted for this geometry. As may be seen in Figure 7.25, the analysis without DUT data agreed with both the baseline analysis and the estimate to well within uncertainties, with a difference in average values across the module of $0.005\% x/X_0$.

7.6.4 Comparison to the 700 MeV campaign 1 dataset

As summarised in Table 7.1 above, the campaign 1 dataset was taken at a lower energy and with a different spacing between the telescope planes. Additionally, only the central $3\text{ cm} \times 3\text{ cm}$ region of the quad module was imaged in this dataset. Since no dedicated




Figure 7.25: Fractional radiation length map without inclusion of the DUT hit in tracking, and comparison to estimate. As in Figure 7.22, the values obtained from DSCB fits to θ_x and θ_y are combined as independent measurements, and for the comparison in (b), both measurement and estimate were smeared by a Gaussian filter with $\sigma = 0.5\text{ mm}$ to soften placement differences.




Figure 7.26: Comparison of spectra extracted for the campaign 1 (blue) and campaign 2 (yellow) datasets, for the DSCB fit to θ_x and the SSCB fit to θ_y . In each case, x/X_0 values derived via the Highland and Frühwirth-Regler formalisms are given. The estimate is also shown both in its pure form, and with an added Gaussian smearing in each subregion (10% syst. + 2% stat. for camp. 2, 8% syst. + 2% stat. for camp. 1) to match the expected distribution for the data. The mean and uncertainty for each x/X_0 distribution are shown and quoted, and the dashed lines give quartiles.

calibration run had been performed in this configuration, an air contribution of 0.19% X_0 was estimated from the simulations outlined in Section 7.5.4. The resulting map agrees closely with the campaign 2 result in terms of the relative differences between features, as may be seen in the spectra in Figure 7.26. The means agree within systematic uncertainties, but a slight shift in the extracted value may be due to the lack of a dedicated calibration sample for this measurement. The extracted systematic and statistical uncertainties were observed to be very close to those for the campaign 2 dataset, albeit slightly smaller as a relative contribution.

7.7 Implications and conclusions

The measurement presented in this section shows a mature methodology utilising the multiple scattering of low-energy electrons or positrons at a testbeam area, building on lessons learned in similar past measurements. A new telescope, MONSTAR, was designed and built to be configurable for a wide range of beam parameters by allowing adjustable inter-plane spacings from 5 cm to 200 cm. A measurement was performed at the CERN PS T9 beamline of an ATLAS ITk prototype quad module, which was populated with components very close to the final production specifications for the upgrade. Projected and

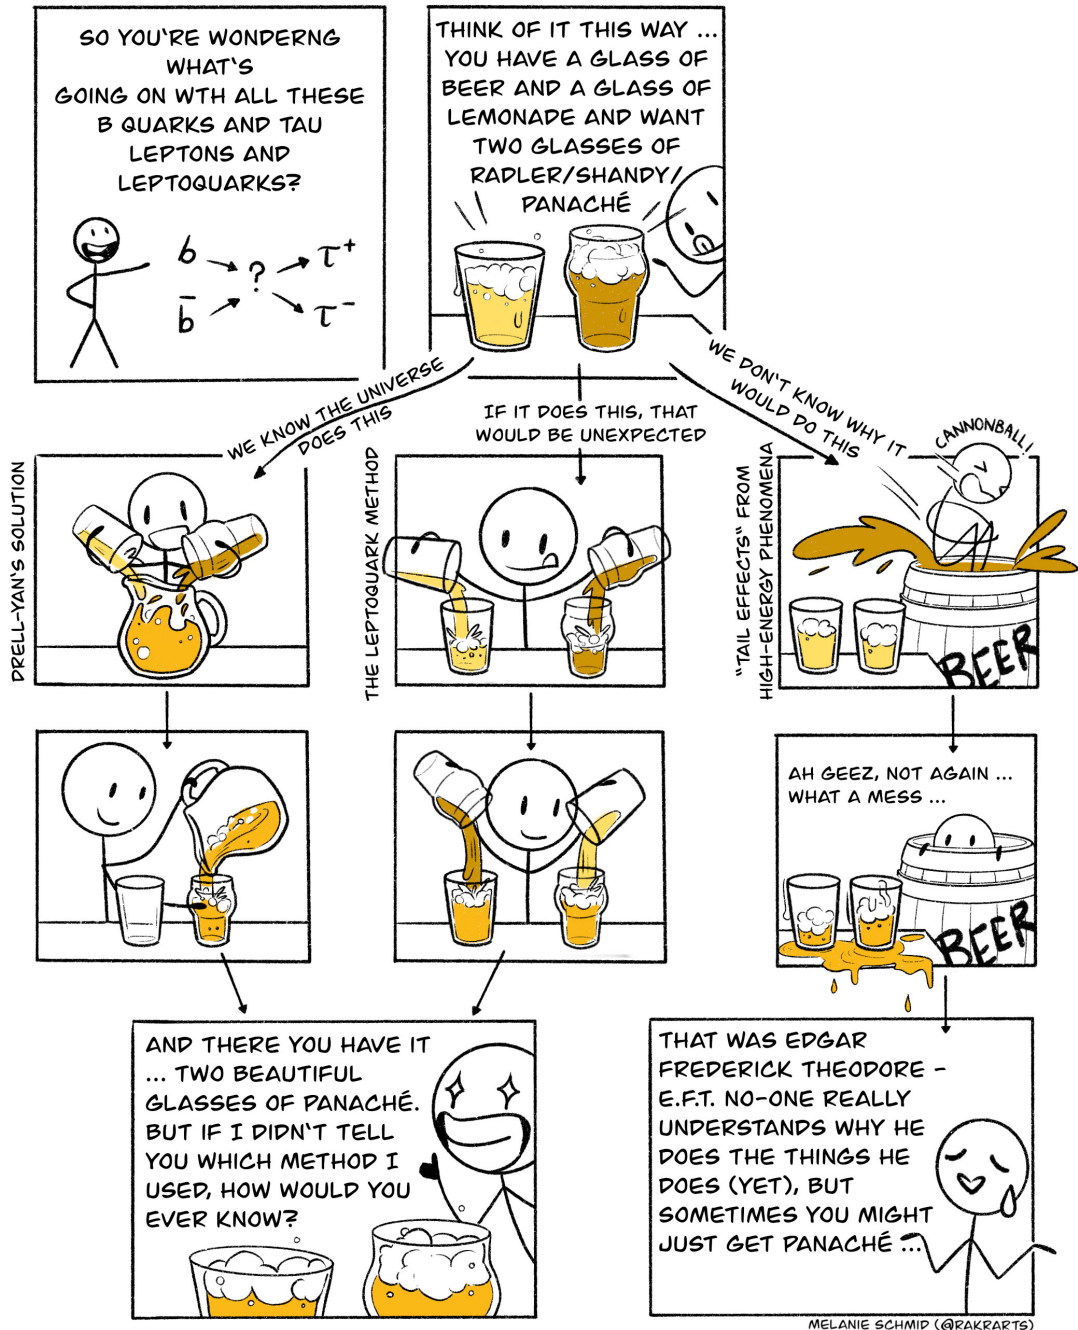
global scattering angles were extracted from multiplet tracks, and were fit using double- and single-sided Crystal-ball functions, respectively. Geant4-based simulations were performed to validate the air and telescope scattering contributions, and a response curve derived from the simulation was used as an alternative to the Highland and Frühwirth-Regler models for computing a fractional radiation length from the distribution width.

The average module thickness was measured as $(0.89 \pm 0.01 \text{ (reso.)} \pm 0.01 \text{ (subtraction)} \pm 0.08 (E_{\text{beam}})) \% X_0$, well in agreement with the estimated value of 0.88% derived as a 2D-resolved map specifically for this study. The dominant uncertainty was the beam momentum band of 5%, which was chosen to ensure a sufficient yield to achieve a granularity of $0.5 \text{ mm} \times 0.5 \text{ mm}$ on the measurement, but could be improved by choice of a different beam area that is able to provide higher rates with smaller momentum bands. A parallel analysis disregarding data from the ITkPix module demonstrated very close agreement with the main measurement across all regions, demonstrating the potential for measurements of non-instrumented subjects using the same methodology.

The measured module average exceeds the target value of $0.67\% X_0$ specified in the ITk Pixel TDR (see Table 4.1) by a significant margin [103]. Although the data and power flex cables are considered services and as such are not included the quoted target, the estimate breakdown in Figure 7.15b shows these are only expected to contribute 9% of the total material. Adjusted for this, the measured average of $0.81\% X_0$ is still 20% in excess of the original target, potentially motivating direct measurements of both module and services material contributions to the ITk detector, both for comparison to design targets and for improvements to the material model used for simulation once the detector is commissioned.


The results of this study have been made public by the ATLAS ITk collaboration [2, 4], and demonstrate the feasibility of the methodology and apparatus for measurements of both active detector components and inactive services to higher precision than is currently available. Further material measurements using the same methodology are in progress, including a measurement performed at the PSI PiM1 facility in collaboration with the CMS upgrade tracker project and Mu3e experiment.

PART III



Measurements with ATLAS

Illus. 3: Leptoquarks in layman's terms, illustrated by Melanie Schmid, 2025.



Two leptoquarks, a Z' and a b quark walk into a bar(yon).

8

Lepton flavour universality violation in BSM models and EFT

Contents

8.1	Experimental evidence	151
8.1.1	$R(D^{(*)})$	152
8.1.2	$R(K^{(*)})$	153
8.2	Relevant effective field theory couplings	155
8.2.1	$\mathcal{C}_9^{\text{NP}}$ and other $b \rightarrow s\ell^+\ell^-$ operators	155
8.2.2	$b \rightarrow c\ell\bar{\nu}$ and leptoquark-equivalent couplings	157
8.2.3	Flavour-dependent SMEFT couplings	157
8.2.4	Electric and magnetic dipole moments	157
8.3	Leptoquark models	158
8.3.1	Scalar leptoquarks	160
8.3.2	Vector leptoquarks	161
8.3.3	Leptoquark searches at collider experiments	162
8.4	Z' boson models	165

The SM predicts that the weak interaction should couple identically to all lepton generations, a phenomenon referred to as lepton flavour universality (LFU). Post-symmetry-breaking, the *charged current* electroweak coupling to leptons is of the form

$$\mathcal{L}_W = \frac{g}{\sqrt{2}} \left(\bar{\nu}_{lL} \gamma^\mu l_L W_\mu^+ + \bar{l}_L \gamma^\mu \nu_{lL} W_\mu^- \right), \quad (8.1)$$

where $g = e/\sin\theta_W$ is the coupling constant of the $\text{SU}(2)_L$ gauge group, and the lepton fields have been expanded over the $\text{SU}(2)$ doublets

$$\ell_L = \begin{pmatrix} \nu_l \\ l \end{pmatrix}_L, \quad l = e, \mu, \tau. \quad (8.2)$$

Since the coupling strength is independent of the lepton flavour l , the charged-current interaction is expected to couple identically to all generations at tree-level. Similarly, the form of the equivalent Z - and γ^* -mediated *neutral current* couplings

$$\mathcal{L}_Z = \frac{g}{\cos \theta_W} \bar{l} \gamma^\mu \left(I^3 - \sin^2 \theta_W Q \right) l Z_\mu, \quad (8.3)$$

$$\mathcal{L}_{\text{EM}} = g \sin \theta_W \bar{l} \gamma^\mu l A_\mu, \quad (8.4)$$

only shows dependence on g , the third component of weak isospin I^3 , and the weak mixing angle θ_W . Such interactions are expected to be flavour-universal at tree-level, with the caveat that many such neutral current diagrams interfere with equivalent Higgs-boson-mediated diagrams, due to the Yukawa couplings for charged leptons

$$\mathcal{L}_{\text{Yukawa}} \supset -\frac{m_l}{v} \bar{l} l H + \text{h.c.}, \quad (8.5)$$

which show a direct dependence on the lepton mass hierarchy, as expressed through the mass term m_l . For the charged leptons, the ratio of the small lepton mass of $\mathcal{O}(\text{MeV--GeV})$ to the Higgs vacuum expectation value $v = 246 \text{ GeV}$ effectively suppresses each such vertex by a factor 10^2 or greater compared to the neutral current process.

More generally, in the absence of Yukawa interactions, the SM exhibits a global flavour symmetry as a consequence of the fundamental group structure and fermion content of the SM Lagrangian. Since the latter is a sum over covariant kinetic energy terms $\sum_n \bar{\psi}_n i \not{D} \psi_n$, linear unitary transformations applied to any of the constituent fields in the SM-representation leave the Lagrangian invariant:

$$q_L \rightarrow U_q q_L, \quad u_R \rightarrow U_u u_R, \quad \dots, \quad e_R \rightarrow U_e e_R, \quad (8.6)$$

where $U_q^\dagger U_q = U_u^\dagger U_u = \dots = 1$. Given the five distinct irreducible fermion representations $(q_L, u_R, d_R, L_L, e_R)$ of the SM, this gives a flavour symmetry under $\text{U}(3)^5$, which can be reduced to $\text{SU}(3)^5$ when considering that the component $\text{U}(1)$ factors are anomalous and broken by quantum effects [219].

8.1 Experimental evidence

Direct tests of lepton universality have been performed in several channels, and universality has been verified in τ branching fractions and lifetime measurements [42], Z and W boson decays [220, 221], and several categories of meson decays [222–224].

8.1.1 $R(D^{(*)})$

Recent experimental evidence from the LHCb [227–229], Belle and Belle II [230, 231], and BaBar [232] experiments points toward a preference for heavy-lepton decays in B -meson decay fractions, with substantial deviations in the $R(D^{(*)})$ ratio of $b \rightarrow c\ell\nu$ transition rates. Defined as the ratio of the branching fractions

$$R(D^{(*)}) = \frac{\mathcal{B}(B \rightarrow D^{(*)}\tau\nu)}{\mathcal{B}(B \rightarrow D^{(*)}\ell\nu)}, \quad (8.7)$$

the average measurements provided by BaBar, Belle, and LHCb currently show a significant tension with the SM. The most recent results from the HFLAV are given in Figure 8.1, showing an $\mathcal{O}(10\%)$ preference for the τ -channel decay modes over the light-lepton channels compared to the SM expectation. The most recent world averages¹ give combined values of 0.342 ± 0.026 for $R(D)$ and 0.287 ± 0.012 for $R(D^*)$, with a correlation coefficient of $\rho = -0.39$. Several recent SM predictions based on lattice QCD calculations for $R(D)$ and $R(D^*)$ have been combined by the Heavy Flavour Averaging group (HFLAV) to give arithmetic averages of 0.298 ± 0.004 and 0.254 ± 0.005 , respectively [225, 226].





Figure 8.1: Summary of $R(D^{(*)})$ measurements at Belle, BaBar and LHCb, as collated by the Heavy Flavour Averaging group (HFLAV), from [225, 226].

¹As of writing.



(a) SM-mediated, lepton-flavour-universal transition. (b) Potential leptoquark-mediated transition. If the coupling $g_{b\ell}^{\text{LQ}}$ is not identical for all lepton generations, the transition violates LFU.

Figure 8.2: Representative leading order SM and leptoquark-mediated diagrams contributing to $\bar{B}^0 \rightarrow D^+ \ell^- \bar{\nu}$. Equivalent diagrams exist for the full family of $B \rightarrow D^{(*)} \ell \nu$ transitions, by particle-antiparticle symmetry, or by replacement of the d -quark in the original meson.

Combining both $R(D)$ and $R(D^*)$ and accounting for the correlation between the measurements, the significance of the deviation from the SM prediction² is currently 3.31σ .

The size of these discrepancies is indicative of a tree-level explanation, as theoretical calculations of higher-order SM diagrams that would preference transitions to τ leptons tend to be suppressed by factors of 10^2 or more compared to the tree-level diagrams [233]. A representative SM leading-order flavour-universal diagram contributing to $R(D^{(*)})$ is shown in Figure 8.2a, and an example of a leptoquark BSM model providing a tree-level LFUV coupling is given in Figure 8.2b.

8.1.2 $R(K^{(*)})$

Prior to 2023, substantial evidence of lepton flavour universality violation (LFUV) had also accumulated between the light lepton species in data for $b \rightarrow s \ell \ell$ transitions, with the relevant ratios $R(K^{(*)}) = \mathcal{B}(B \rightarrow K^{(*)} \mu \mu) / \mathcal{B}(B \rightarrow K^{(*)} e e)$ being measured as less than unity by the LHCb Collaboration [234, 235]. Since the SM predicts a value of $1 \pm \mathcal{O}(10^{-2})$ with high precision [236], these results had been taken as strong indications of potential tree-level or loop-level BSM behaviour in diagrams interfering with the dominant SM transition shown in Figure 8.3a.

However, a recent re-examination and combination of Run 1 and 2 data by the LHCb collaboration revised these values to $R(K) = 0.949 \pm 0.047$ and $R(K^*) = 1.027 \pm 0.075$

²Derived from a combined χ^2 of 13.95 for 2 degrees of freedom, corresponding to a p-value of 9.33×10^{-4} [225].

in the kinematic region $1.1 \text{ GeV}^2 < m_{\ell\ell}^2 < 6 \text{ GeV}^2$, yielding overall agreement at the level of 0.2σ across all studied kinematic regions³ [237, 238]. As a result, the general experimental consensus has shifted toward universality between electrons and muons for both $b \rightarrow s\ell\ell$ and $b \rightarrow c\ell\nu$ transitions. In combination with the hints of third-generation LFUV from $R(D^{(*)})$, it is reasonable to expect a potential enhancement of $b \rightarrow s\tau\tau$ transitions compared to the light lepton channels [239]. However, such transitions have not yet yielded strong experimental evidence for LFUV behaviour due to significant associated experimental challenges⁴.

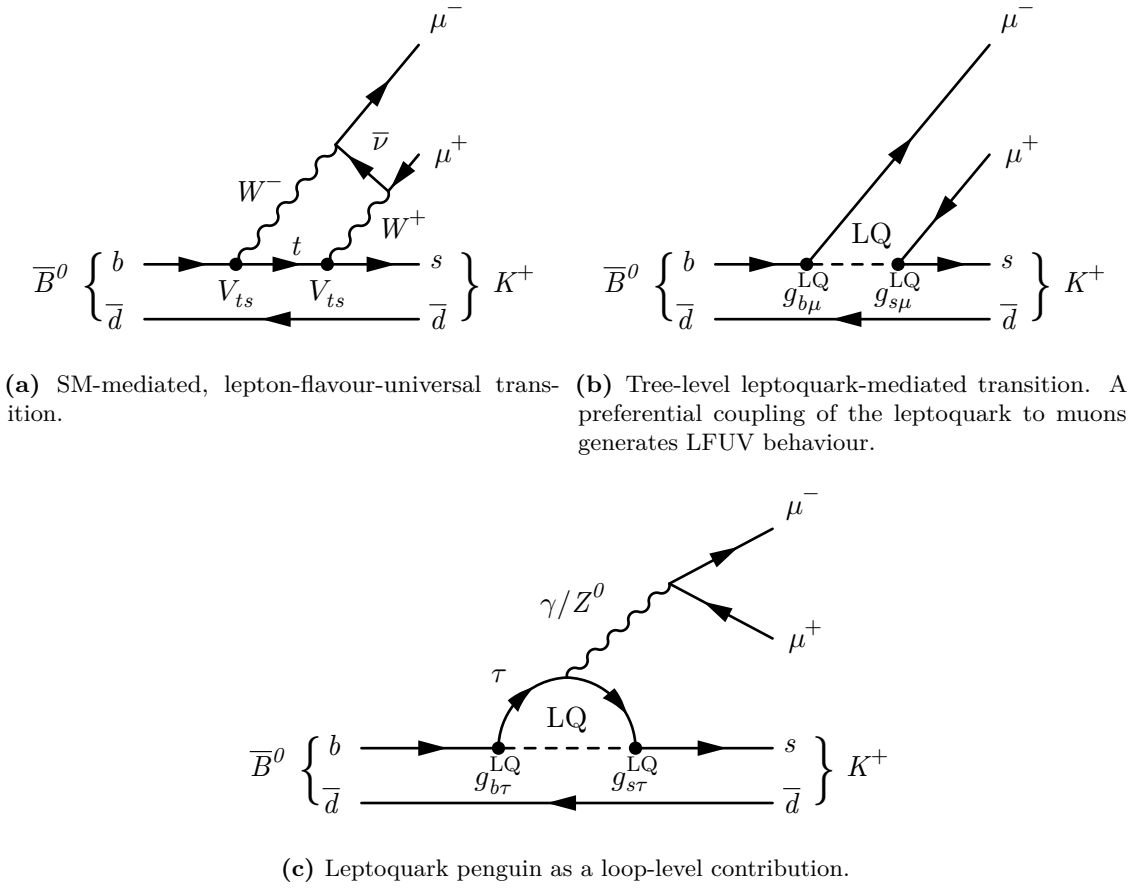


Figure 8.3: Representative diagrams for the $B^+ \rightarrow K^0 \mu^- \mu^+$ transition contributing to the $R(K^{(*)})$ ratio. The SM contribution shown in (a) is highly suppressed by the CKM matrix elements and the order of the diagram. The potential tree-level leptoquark-mediated contribution shown in (b) depends on relevant muon couplings to leptoquarks. The photon penguin diagram in (c) is a loop-level alternative that instead relies on the τ couplings of the leptoquark.

³Here $m_{\ell\ell}$ refers to the invariant mass of the final state ee or $\mu\mu$ pair.

⁴The combination of multiple neutrinos in the final state, and non-hermetic coverage in the case of the LHCb detector, make the relevant final states particularly challenging to isolate. Current limits are still several orders of magnitude from SM predictions [240, 241].

8.2 Relevant effective field theory couplings

Deviations from the SM in both lepton flavour conserving and violating modes can be studied by measuring the strengths of effective couplings of dimension 6 and higher. As has been introduced in Section 1.3, a generic EFT Lagrangian of the form

$$\mathcal{L}_{\text{eff}} = \sum_i c_i(\mu) \mathcal{Q}_i(\mu) \quad (8.8)$$

encodes a sequence of operators $\mathcal{Q}_i(\mu)$ with coupling strength given by the Wilson coefficients c_i that encode the short-distance dynamics at a given energy scale μ .

Effective descriptions of processes relevant to flavour physics are often formulated using a low-energy EFT (LEFT) [242], which is derived from the SM fields, but integrates out particles above the B meson energy scale including the W and Z bosons, the Higgs boson, and the top quark. This gives a low-energy effective field theory of the SM similar to the Fermi theory of weak interactions, including 4-fermion interactions at dimension 6. To disambiguate the SMEFT and LEFT prescriptions, SMEFT operators and coefficients are here denoted \mathcal{Q}_i and c_i , whereas LEFT operators and coefficients are labelled \mathcal{O}_i and C_i .

8.2.1 $\mathcal{C}_9^{\text{NP}}$ and other $b \rightarrow s\ell^+\ell^-$ operators

An effective Hamiltonian for $b \rightarrow s\ell\ell$ below the electroweak scale can be written following the conventions introduced in [243, 244] to give

$$\mathcal{H}_{\text{eff}} = \frac{4G_F}{\sqrt{2}} \sum_{q=u,c} V_{qb}V_{qs}^* \sum_i C_i(\mu) \mathcal{O}_i(\mu) \quad (8.9)$$

in terms of the CKM matrix elements⁵ V_{ij} , with the most relevant operators $\mathcal{O}_i(\mu)$ given in momentum space as

$$\begin{aligned} \mathcal{O}_7 &= \frac{e^2}{(4\pi)^2} (\bar{s}_L \sigma_{\mu\nu} b_R) F_{\mu\nu}, \\ \mathcal{O}_{9\ell} &= \frac{e^2}{(4\pi)^2} (\bar{s}_L \gamma_\mu b_L) (\bar{\ell} \gamma^\mu \ell), & \mathcal{O}_{S\ell} &= \frac{e^2}{(4\pi)^2} (\bar{s}_R b_R) (\bar{\ell} \ell), \\ \mathcal{O}_{10\ell} &= \frac{e^2}{(4\pi)^2} (\bar{s}_L \gamma_\mu b_L) (\bar{\ell} \gamma^\mu \gamma^5 \ell), & \mathcal{O}_{P\ell} &= \frac{e^2}{(4\pi)^2} (\bar{s}_R b_R) (\bar{\ell} \gamma^5 \ell), \end{aligned} \quad (8.10)$$

⁵It is useful to note that due to the unitarity of the CKM matrix, $\sum_{q=u,c} V_{qb}V_{qs}^* = -V_{tb}V_{ts}^*$.

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ denotes the electromagnetic field strength tensor in terms of the photon field A^μ , the tensor $\sigma^{\mu\nu} = \frac{i}{2}[\gamma^\mu, \gamma^\nu]$ is defined in terms of the commutator of gamma matrices γ^μ , s and b denote the quark fields, and ℓ is in this case summed over the both the lepton doublet and singlet states. For each of these operators, a parity-inverted equivalent \mathcal{O}'_i can be derived by applying the replacement $f_L \leftrightarrow f_R$ on all fermion fields f . The contributions to the corresponding Wilson coefficients have been calculated to next-to-next-to-leading-order (NNLO) in QCD and logarithmic electroweak corrections at the b quark scale $\mu_b = 4.8 \text{ GeV}$ [245–247], and results from EFT fits are generally quoted as deviations from these values indicative of new physics, denoted as C_i^{NP} or ΔC_i .

Global fits to these Wilson coefficients have been performed by several groups [248–250], and generally agree on a significant new physics contribution corresponding to $C_9^{\text{NP}} \approx -1$, whilst remaining in agreement with the SM for C_{10}^{NP} . Results from the ABCDMN⁶ global fits are shown in Figure 8.4. The fits also indicate a preference for a flavour-universal contribution in the light-lepton sector $C_{9\mu}^{\text{NP}} = C_{9e}^{\text{NP}}$, as may be seen in Figure 8.4b, which is largely influenced by the LFU conserving nature of recent $R(K^{(*)})$ and $B_s^0 \rightarrow \phi\mu\mu$ results [243].

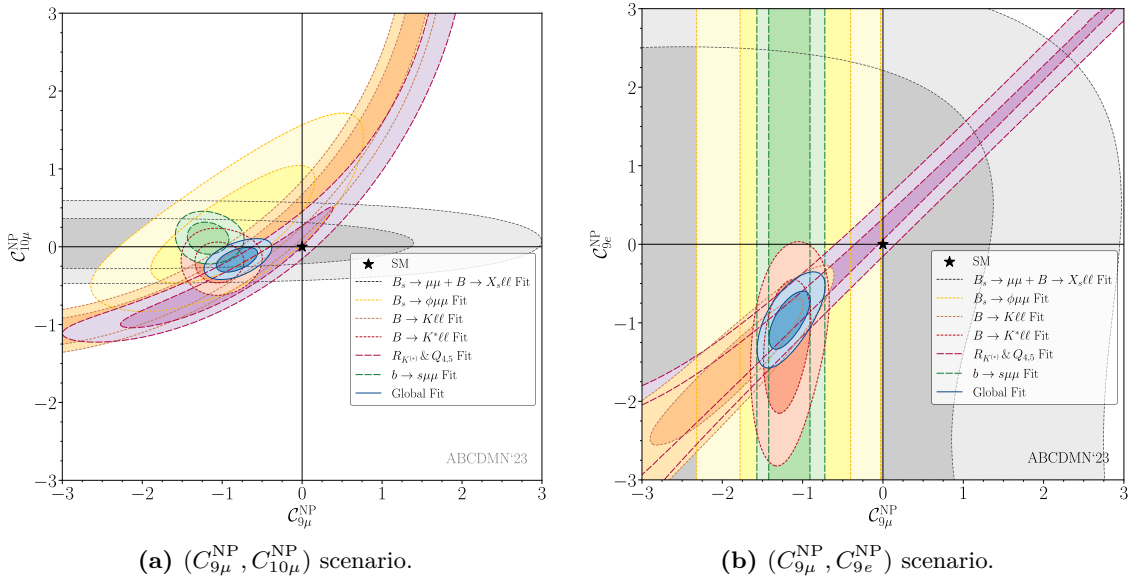


Figure 8.4: 1σ (dark shaded) and 2σ (light shaded) confidence regions for pairs of Wilson coefficients in the ABCDMN fits, from Ref. [248]. Distinct fits were performed for each $b \rightarrow s\ell^+\ell^-$ decay mode (short dashes), combined $b \rightarrow s\mu^+\mu^-$ modes or LFUV observables (long dashes), and a global fit (solid contour). Some fits shared observables, and are hence not entirely uncorrelated.

⁶M. Algueró, A. Biswas, B. Capdevila, S. Descotes-Genon, J. Matias, M. Novoa-Brunet [248].

8.2.2 $b \rightarrow c\ell\bar{\nu}$ and leptoquark-equivalent couplings

The general dimension-six LEFT Hamiltonian for $b \rightarrow c\ell\bar{\nu}$ transitions is typically written by separating it into left-chiral and right-chiral vector, scalar, and tensor currents:

$$\mathcal{H}_{\text{eff}} = -\frac{4G_F}{\sqrt{2}} V_{cb} \left[(1 + C_{V_L}(\mu)) \mathcal{O}_{V_L}(\mu) + \sum_{i=S_L, S_R, V_R, T_L} C_i(\mu) \mathcal{O}_i(\mu) \right] + \text{h.c.} \quad (8.11)$$

The operators in question are defined for each lepton field $l \in \{e, \mu, \tau\}$, as

$$\begin{aligned} \mathcal{O}_{V_L} &= (\bar{u}_L \gamma_\mu d_L) (\bar{l}_L \gamma^\mu \nu_{lL}), & \mathcal{O}_{V_R} &= (\bar{u}_R \gamma_\mu d_R) (\bar{l}_L \gamma^\mu \nu_{lL}), \\ \mathcal{O}_{S_L} &= (\bar{u}_R d_L) (\bar{l}_R \nu_{lL}), & \mathcal{O}_{S_R} &= (\bar{u}_L d_R) (\bar{l}_R \nu_L), \\ \mathcal{O}_T &= (\bar{u}_R \sigma_{\mu\nu} d_L) (\bar{l}_R \sigma^{\mu\nu} \nu_{lL}). \end{aligned} \quad (8.12)$$

At tree-level in the SM, $C_{V_L} = 1$ and all other associated Wilson coefficients are zero, and the coefficients are flavour-universal. The most recent EFT study combining measurements of several observables pertaining to the $b \rightarrow c\ell\bar{\nu}$ transition found an overall 3.3σ tension in measured values of operators coupling to third generation leptons ($l = \tau$) [251]. Several 1D and 2D hypotheses were tested: among the 1D hypotheses, $C_{V_L} = 0.07 \pm 0.02$ gave the best explanation of the data; for the 2D hypotheses, the best candidates either included a dominant contribution to C_{V_L} , or involved combinations of C_{S_R} and C_{S_L} .

8.2.3 Flavour-dependent SMEFT couplings

Several groups have computed matchings between the LEFT operators mentioned above and equivalent combinations in the full SMEFT framework to various orders of precision. These will not be discussed here, but the reader is directed to Refs. [252, 253] as useful references. Beyond operators specifically matching the b anomalies, LFUV and LFUV-generating models can also be studied in global EFT fits using operator sets that do not impose universal flavour symmetry, and comparing results on equivalent SMEFT operators for different lepton generations.

8.2.4 Electric and magnetic dipole moments

A discussion of lepton flavour universality would not be complete without a brief discussion of the potential for the electric and magnetic dipole moments of leptons to be modified by BSM physics. The electron anomalous magnetic moment $a_e = (g_e - 2)/2$ is one of the

most precisely verified predictions in the history of physics, having been determined to agree with the SM to better than one part in 10^{10} [254]. The muon anomalous magnetic moment has been measured by Fermilab National Accelerator Laboratory (FNAL) to a precision of 127 parts in 10^9 [255, 256], and has previously been in some contention with theory predictions that include corrections derived from electron-positron annihilation experiments [257]. However, since 2021, a new lattice-QCD based theoretical model shows compatibility with the measured value, but is discrepant with previous predictions [258]. The most recent update from the Muon g-2 Theory Initiative in 2025 adopted a world average of recent lattice-QCD calculations as baseline, resulting in good agreement with the measured value [259]. The leading measurements of a_τ are subject to an uncertainty of about 300% compared to the predicted value of 0.001 177 21 [260–262]. In light of the potential for new physics models with preferential second- or third-generation couplings to affect a_μ or a_τ (or, alternatively, to be excluded by stringent limits on these), it is valuable to explore the potential for LFUV in these quantities.

Within the SMEFT, a post-symmetry-breaking effective operator $\mathcal{Q}_{\tau\gamma}$ can be defined as a linear combination of $\mathcal{Q}_{\tau W}$ and $\mathcal{Q}_{\tau B}$ such that $c_{\tau\gamma} = \cos\theta_W c_{\tau B} - \sin\theta_W c_{\tau W}$, giving

$$\mathcal{Q}_{\tau\gamma} = (\bar{\ell}\sigma^{\mu\nu}\tau)\sigma^i H \left(-\sin\theta_W \sigma^i W_{\mu\nu}^i + \cos\theta_W B_{\mu\nu} \right), \quad (8.13)$$

where θ_W is the weak mixing angle. By matching to an expansion of the SM $\tau\tau\gamma$ vertex to linear order in the photon four-momentum following the method in Ref. [263], the contributions of $c_{\tau\gamma}$ can be computed as

$$\delta a_\tau = \frac{2\sqrt{2}m_\tau v}{e} \cdot \frac{\text{Re}[c_{\tau\gamma}]}{\Lambda^2}, \quad \delta d_\tau = \sqrt{2}v \frac{\text{Im}[c_{\tau\gamma}]}{\Lambda^2}, \quad (8.14)$$

where v is the Higgs vacuum expectation value, e is the electron charge, m_τ is the τ lepton mass, Λ is the suppression mass scale, and δ indicates deviations from the SM predictions for a_τ and d_τ .

8.3 Leptoquark models

The SM gauge group content of $\text{SU}(3) \times \text{SU}(2)_L \times \text{U}(1)_Y$ and its associated pre-symmetry-breaking $\text{SU}(3)^5$ flavour symmetry entirely separate the lepton and quark sectors, producing separate lepton number (L) and baryon number (B) conservation laws. However,

Table 8.1: Categorisation of possible leptoquarks by their spin, couplings to fermions, and quantum numbers under the QCD representation $SU(3)$, weak isospin $SU(2)_L$, and weak hypercharge $U(1)_Y$. The carried fermion number F is equal to $3B + L$. The c superscripts on fermions in the listed couplings represent charge-conjugation. For completeness, modes coupling to a possible sixth representation ν_R representing right-chiral neutrinos that are not part of the SM field content have been added, and are highlighted in grey.

	Spin	F	$(SU(3)_c, SU(2)_W, U(1)_Y)$	Allowed couplings
Scalar	S_1	0	-2	$(\bar{q}_L^c \ell_L), (\bar{u}_R^c e_R), (\bar{d}_R^c \nu_R)$
	\tilde{S}_1	0	-2	$(\bar{d}_R^c e_R)$
	\bar{S}_1	0	-2	$(\bar{u}_R^c \nu_R)$
	R_2	0	0	$(\bar{q}_L e_R), (\bar{u}_R \ell_L)$
	\tilde{R}_2	0	0	$(\bar{d}_R \ell_L), (\bar{q}_L \nu_R)$
	S_3	0	-2	$(\bar{q}_L^c \ell_L)$
Vector	U_1	1	0	$(\bar{q}_L \gamma^\mu \ell_L), (\bar{d}_R \gamma^\mu e_R), (\bar{u}_R \gamma^\mu \nu_R)$
	\tilde{U}_1	1	0	$(\bar{u}_R \gamma^\mu e_R)$
	\bar{U}_1	1	0	$(\bar{d}_R \gamma^\mu \nu_R)$
	V_2	1	-2	$(\bar{q}_L^c \gamma^\mu e_R), (\bar{d}_R^c \gamma^\mu \ell_L)$
	\tilde{V}_2	1	-2	$(\bar{u}_R^c \gamma^\mu \ell_L), (\bar{q}_L^c \gamma^\mu \nu_R)$
	U_3	1	0	$(\bar{q}_L \gamma^\mu \ell_L)$

many higher-energy theories which reduce to the SM as an effective theory at lower energy scales will naturally couple these sectors through one or more new bosons [42, 264]. This category of hypothetical particles is referred to as *leptoquarks* (LQs), which carry both baryon number and lepton number, leading to the conservation of a more general fermion number defined as $F = 3B + L$.

Several extensions of the SM include leptoquark states, such as the Pati-Salam model containing an $SU(4)$ colour gauge group [265]. Grand unification theories based on $SU(5)$ [35], $SO(10)$ [266, 267], or larger gauge groups also typically generate leptoquark couplings, and scalar quarks in supersymmetry models with R -parity violation can exhibit a leptoquark-equivalent nature by coupling to all leptons [268]. Leptoquarks also emerge within technicolor scenarios [269] and composite-fermion models [270], where they are expected to be accessible at an $\mathcal{O}(\text{TeV})$ energy scale.

Starting from the SM fields, a possible set of leptoquark states can be derived by requiring their direct interactions with fermions to be dimensionless and invariant under

the SM gauge groups, and these are tabulated in Table 8.1. Generally, terms involving di-quark couplings to the leptoquark would be allowed for leptoquarks carrying a non-zero fermion number, but these must be forbidden by a symmetry to prevent proton decay expectations far exceeding current limits [271].

As the Dirac and/or Majorana nature of neutrinos is yet to be experimentally verified, the SM could be extended with one or more right-chiral neutrino fields. As such states could couple directly to leptoquarks, they extend the set of allowed couplings of the S_1 scalar singlet and U_1 vector singlet, and additionally generate alternate-charge \bar{S}_1 and \bar{U}_1 leptoquarks only accessible through right-handed neutrino couplings. These additional couplings and states are listed in Table 8.1 in grey for completeness, but will be omitted in the following phenomenological discussions.

A brief summary of the phenomenology of the most interesting candidates is provided in the following, focusing on compatibility with the B meson anomalies.

8.3.1 Scalar leptoquarks

The three scalar candidate classes are referred to as S_1 , R_2 , and S_3 , as categorised by their singlet, doublet, or triplet nature under weak isospin, respectively. Replacements of $\bar{u}_R \leftrightarrow \bar{d}_R$ generate distinct species that differ in weak hypercharge, which are labelled \tilde{S}_1 and \tilde{R}_2 to disambiguate. Similarly, in the case of replacement by right-handed neutrino fields, another alternative hypercharge is generated for the scalar singlet, and is denoted by \bar{S}_1 .

The S_1 weak scalar singlet with hypercharge $Y = 1/3$ has been found to be a viable explanation for the excess above the SM value in $R(D^{(*)})$ [272]. Although this flavour of leptoquark cannot contribute at tree-level to $R(K^{(*)})$ as in Figure 8.3b, the one-loop diagram given in Figure 8.3c allows contributions to this observable which are described in greater detail in Ref. [273]. Non-universal couplings to the lepton generations would allow an SM-like $R(K^{(*)})$, whilst accounting for excesses in both $R(D^{(*)})$ and the C_9^{NP} Wilson coefficient. The Lagrangian may be given as

$$\mathcal{L}_{S_1} = y_L^{ij} \bar{q}_{Li} i\sigma_2 \ell_{Lj} S_1 + y_R^{ij} \bar{u}_{Ri}^c e_{Rj} S_1 + \text{h.c.} \quad (8.15)$$

$$= S_1 \left[(V^* y_L)_{ij} \bar{u}_{Li}^c \ell_{Lj} - y_L^{ij} \bar{d}_{Li}^c \nu_{Lj} + y_R^{ij} \bar{u}_{Ri}^c e_{Rj} \right] + \text{h.c.}, \quad (8.16)$$

where y_L and y_R define matrices of Yukawa-style couplings of the S_1 leptoquark to the fermion generations, σ_2 denotes the second Pauli matrix, and V is the CKM matrix. The c superscript represents charge conjugation. Since $F = -2$, the di-quark couplings must be forbidden by an additional symmetry to preserve proton stability. The tree-level contributions to Wilson coefficients relevant to $b \rightarrow c\ell\bar{\nu}_{\ell'}$ can be computed by integrating out the leptoquark at the matching scale $\mu = m_{S_1}$ to give

$$C_{V_L} = \frac{v^2}{4V_{cb}m_{S_1}^2} y_L^{b\ell'} (V y_L^*)_{c\ell} , \quad (8.17)$$

$$C_{S_L} = -4C_T = -\frac{v^2}{4V_{cb}m_{S_1}^2} y_L^{b\ell'} (y_R^{c\ell})^* , \quad (8.18)$$

where v is the Higgs vacuum expectation value, and m_{S_1} the leptoquark mass [274].

The R_2 and \tilde{R}_2 doublets automatically avoid di-quark couplings and preserve proton stability without requiring an additional symmetry, but do generate $R(K)$ and/or $R(K^*)$ values in excess of the SM prediction at tree-level and are clearly disfavoured by data. Such problems can be avoided by restricting leptoquark corrections to only loop-level and above, by requiring $y_R = 0$ [275]. Studies of impacts on the $b \rightarrow c\tau\bar{\nu}$ transitions have shown that such scenarios can generate the observed deviations in $R(D^{(*)})$, but cannot simultaneously accommodate $R(K^{(*)})$ due to stringent limits on $\mathcal{B}(\tau \rightarrow \mu\gamma)$, which would also show an enhanced contribution [275].

The S_3 triplet can simultaneously accommodate deviations in both $R(K)$ and $R(K^*)$, but suffers from the same requirement for an additional mechanism to forbid di-quark transitions as S_1 . The dominant effect on $b \rightarrow c\tau\bar{\nu}$ is suppressive, which is clearly at odds with current measurements of $R(D^{(*)})$ [274].

8.3.2 Vector leptoquarks

Similarly to the scalar leptoquarks, three main categories U_1 , V_2 , and U_3 correspond to a weak isospin of 1, 2, and 3, respectively. As above, species with altered weak hypercharge \tilde{U}_1 , \tilde{V}_2 , and \bar{U}_1 are generated by replacements $\bar{u}_R \rightarrow \bar{d}_R$ or $e_R \rightarrow \nu_R$.

It is important to note that the vector leptoquark models are generally non-renormalisable, and hence require an ultraviolet (UV) mechanism generating the leptoquark terms to be explicitly specified to allow loop-level calculations [274].

The U_1 vector singlet has received considerable attention, as it could provide a simultaneous explanation for the historical anomalies in $R(K^{(*)})$ and $R(D^{(*)})$ [276]. The state carries no fermion number, and as a result does not permit any di-quark couplings. The most general Lagrangian (ignoring right-handed neutrino fields) is

$$\mathcal{L}_{U_1} = \frac{1}{\sqrt{2}} \left(\beta_L^{ij} \bar{q}_{Li} \gamma_\mu U_1^\mu \ell_{Lj} + \beta_R^{ij} \bar{d}_{Ri} \gamma_\mu U_1^\mu e_{Rj} \right) + \text{h.c.}, \quad (8.19)$$

where β_L and β_R are coupling matrices. The matching to LEFT operators given in Ref. [274] shows that the left-handed couplings contribute as

$$C_{9\mu}^{\text{NP}} = -C_{10\mu}^{\text{NP}} = -\frac{\pi v^2}{V_{tb} V_{ts}^* \alpha_{\text{em}}} \frac{\beta_L^{s\mu} (\beta_L^{b\mu})^*}{m_{U_1}^2}. \quad (8.20)$$

The right-handed couplings generate contributions $\Delta C'_9 = \Delta C'_{10}$, $\Delta C'_S = -\Delta C'_P$ and $\Delta C'_S = \Delta C'_P$, which are generally disfavoured by current data, and hence β_R is often set to zero to form a *minimal U_1 model*. The contribution to $b \rightarrow c \ell \bar{\nu}_{\ell'}$ arises from a contribution to the effective coefficient

$$C_{V_L} = \frac{v^2 (V \beta_L)_{c\ell'} (\beta_L^{b\ell})^*}{2 V_{cb} m_{U_1}^2} \quad (8.21)$$

$$= \frac{v^2 (\beta_L^{b\ell})^*}{2 m_{U_1}^2} \left[\beta_L^{b\ell'} + \frac{V_{cs}}{V_{cb}} \beta_L^{s\ell'} + \frac{V_{cd}}{V_{cb}} \beta_L^{d\ell'} \right], \quad (8.22)$$

where the leading term ensures $C_{V_L} > 0$, hence agreeing with $R(D^{(*)})$ measurements.

The U_3 vector triplet generates negative contributions to $R(D^{(*)})$ in a similar way to S_3 , and hence cannot explain the observed deviations [274].

8.3.3 Leptoquark searches at collider experiments

In general, leptoquark searches at high p_T in general-purpose collider experiments can be categorised in three types. *Pair-produced leptoquarks* from diagrams such as that shown in Figure 8.5a will generally form final states of two leptons and at least two associated jets. Leptoquarks preferentially coupling to third generation fermions will show stronger signals in final states containing $\tau^+ \tau^-$ and at least two jets with b quark content. An effective resonant threshold of $m_{q\ell\ell} > 2m_{\text{LQ}}$ suppresses this process at lower energy scales. Single-leptoquark production (Figure 8.5b) produces very similar final states⁷, but with a

⁷Especially since b quarks in the initial state are generally produced in association with a quasi-free \bar{b} quark which will form an additional jet not associated with the leptoquark interaction.




Figure 8.5: Example diagrams for pair, single, and non-resonant production of U_1 leptoquarks preferentially coupling to third-generation fermions.

lower resonant threshold of m_{LQ} . Non-resonant contributions such as that in Figure 8.5c can be probed off-shell at energies substantially lower than m_{LQ} , and at these energies they will dominantly contribute to interference with SM Drell-Yan production of di-lepton final states. The cross-section can be separated into SM, interference, and BSM contributions:

$$\sigma \propto |\mathcal{M}|^2 = \left| \mathcal{M}_{\text{SM}} + \left(\frac{g_{LQ}}{m_{LQ}} \right)^2 \mathcal{M}_{LQ} \right|^2 \quad (8.23)$$

$$= \underbrace{|\mathcal{M}_{\text{SM}}|^2}_{\sigma_{\text{SM}}} + 2 \underbrace{\left(\frac{g_{LQ}}{m_{LQ}} \right)^2 \text{Re}(\mathcal{M}_{\text{SM}}^* \mathcal{M}_{LQ})}_{\sigma_{\text{Int}}} + \underbrace{\left(\frac{g_{LQ}}{m_{LQ}} \right)^4 |\mathcal{M}_{LQ}|^2}_{\sigma_{\text{BSM}}}, \quad (8.24)$$

where g_{LQ} is the leptoquark coupling for the given diagram with matrix element \mathcal{M} , and m_{LQ} the leptoquark mass. The matrix element has been separated into SM components \mathcal{M}_{SM} and leptoquark interaction terms \mathcal{M}_{LQ} . For the S_1 and U_1 leptoquark models and a $\tau^+ \tau^-$ final state, the interference term σ_{Int} is negative and dominant for low di-lepton invariant mass $m_{\ell\ell}$. At $m_{\ell\ell}$ closer to the on-shell mass m_{LQ} , the necessarily positive definite σ_{BSM} will become dominant. The level of cancellation between the interference and pure-BSM terms is dependent on the mass of the leptoquark and transferred energy Q^2 in the interaction.

Previous searches at ATLAS and CMS have studied leptoquark models across all three production modes [277–283]. Pair-production limits from $b\tau b\tau$ final states have generally set lower limits on leptoquark masses of 1.2 TeV to 2 TeV for S_1 and U_1 leptoquark models with various coupling parameters. In general, analyses targeting non-resonant production in the $\tau\tau$ channel have neglected the impact of the interference term σ_{Int} . Final states including $t\bar{t}$ plus light leptons or $t\tau$ are also common targets for pair-production searches [284–286]. In some cases, searches for specific leptoquark species


(a) Left-handed-only scenario: $|\beta_L| = 1$, $|\beta_R| = 0$.(b) Pati-Salam-like scenario: $\beta_L = -\beta_R = 1$.

Figure 8.6: Preferred region for the U_1 vector leptoquark hypotheses in terms of mass M_U and coupling strength g_U , at 90% confidence level (CL) from low-energy charged-current data. The grey regions and solid lines give current limits from high-energy searches at 95% CL. The dotted line shows an extrapolation to the HL-LHC, with a total integrated luminosity of 3 ab^{-1} . The parameter g_U represents a universal scale factor for all active couplings. Reproduced from Ref. [291].

have been performed across multiple final states, generally by reinterpreting existing results. Examples include those for the S_1 and S_3 models [287, 288], or for scalar and vector leptoquarks produced by pair production [289, 290].

Isidori et. al. [291] have computed preferred parameter space regions⁸ for two U_1 vector leptoquark hypotheses given current data from the b anomalies. A preferred region contour for a U_1 vector leptoquark coupling only to left-handed b and c quarks and τ leptons is shown in Figure 8.6a, and recent summary limits from high-energy searches for pair production and Drell-Yan results are shown on the plot. An alternative full coupling scenario including both left- and right-handed couplings (similar to the Pati-Salam case) is shown in Figure 8.6b, with the preferred region at couplings approximately a factor of 2 lower than the left-handed-only case. Consequently, limits on the latter scenario are significantly stronger, with part of the preferred region already excluded.

⁸As of November 2022.

8.4 Z' boson models

Many BSM models⁹ introduce a heavier equivalent of the Z boson, acting as an electrically neutral colour singlet of spin 1 [294]. The interaction terms for a generic Z' coupling to SM fermions may be written as

$$\mathcal{L}_{Z'} = \frac{1}{2\sqrt{6}} \left(\zeta_{qL}^{ij} \bar{q}_{Li} \gamma^\mu Z'_\mu q_{Lj} + \zeta_{uR}^{ij} \bar{u}_{Ri} \gamma^\mu Z'_\mu u_{Rj} + \zeta_{dR}^{ij} \bar{d}_{Ri} \gamma^\mu Z'_\mu d_{Rj} \right. \\ \left. - 3 \zeta_{\ell L}^{ij} \bar{\ell}_{Li} \gamma^\mu Z'_\mu \ell_{Lj} - 3 \zeta_{eR}^{ij} \bar{e}_{Ri} \gamma^\mu Z'_\mu e_{Rj} \right), \quad (8.25)$$

where $\zeta_{q,L}$, $\zeta_{u,R}$, $\zeta_{d,R}$, $\zeta_{\ell,L}$ and $\zeta_{e,R}$ are 3×3 coupling matrices. In general a quantum field theory will only remain well-behaved if the additional heavy spin-1 particle is a gauge boson associated with a spontaneously broken gauge symmetry. This new symmetry could take many forms¹⁰, but in general it will imply the existence of a new conserved quantity, which could be chosen in such a way as to distinguish between leptons and quarks, or between fermion generations.

In many models the Z' is assumed to couple identically to all fermion generations (as does the Z boson), but a Z' boson coupling preferentially to third-generation fermions would introduce LFUV phenomena in Drell-Yan di-lepton production through interference with the SM Z boson, or on-shell production at $Q^2 \sim m_{Z'}^2$. Current constraints on such models are generally driven by limits on flavour-changing neutral currents, due to the implications of a Z' -mediated exchange at tree-level¹¹ [297]. Such models can affect Drell-Yan scattering cross-sections as a function of invariant mass, and can hence be probed through precision studies of this process.

The following chapter will introduce a Run 2 ATLAS measurement of Drell-Yan $\tau^+ \tau^-$ production at high invariant mass, a region of phase space that lends itself well to searches for leptoquarks and Z' bosons with preferential couplings to third-generation leptons.

⁹Including grand unified theories, superstring theories, and theories with extra dimensions, among others. For relevant reviews see Refs. [292, 293].

¹⁰This could manifest as a new $U(1)'$ gauge symmetry, or a higher-order symmetry. For example, if the SM electroweak gauge bosons propagate in extra dimensions, they produce Kaluza-Klein excitations which include a series of Z' boson pairs, each associated with a different $SU(2) \times U(1)$ symmetry [295].

¹¹Such constraints can be somewhat relaxed if the first two fermion generations share the same charge under the new gauge group [296].

“See first, think later, then test. But always see first.
Otherwise you will only see what you were expecting.
Most scientists forget that...”

The wisdom of Wonko the Sane

— Douglas Adams, *So Long, and Thanks for All the Fish*

9

Measurement of the $pp \rightarrow \tau\tau$ cross-section at high $m_{\tau\tau}$ and search for leptoquarks

Contents

9.1	Analysis design and observable	167
9.2	Data and simulation samples	169
9.2.1	MC-simulated background samples	169
9.2.2	Signal models	172
9.3	Object reconstruction and event selection	173
9.3.1	Trigger conditions	173
9.3.2	Tau candidate reconstruction	174
9.3.3	Jets	175
9.3.4	b -tagging	175
9.3.5	Overlap removal	175
9.4	Background estimation for jets faking taus	176
9.4.1	The Universal Fake Factor method	178
9.4.2	Application to the $e\tau_{\text{had}}$ and $\mu\tau_{\text{had}}$ channels	182
9.4.3	Application to the $\tau_{\text{had}}\tau_{\text{had}}$ channels	184
9.5	Systematic uncertainties	186
9.5.1	Experimental uncertainties	187
9.5.2	Modelling uncertainties on background and Drell-Yan samples	188
9.5.3	Modelling uncertainties on BSM samples	189
9.6	Modelling validation and detector-level cross-section	189
9.7	Measurement of the inclusive $\tau_{\text{had}}\tau_{\text{had}}$ cross-section	192
9.7.1	Accounting for detector effects	192
9.8	Statistical model	195
9.9	Search for non-resonant leptoquark and Z' interactions	198
9.9.1	Signal distribution correction for the ≥ 2 b -jet category	198
9.9.2	Results	201
9.10	SMEFT interpretation and limits on Wilson coefficients	205
9.10.1	Parametrisation of the cross-section in terms of couplings	207
9.10.2	Signal samples	209
9.10.3	Results	209
9.11	Conclusions and outlook	212




Figure 9.1: Non-resonant leptoquark contributions to the 0 b -jet, 1 b -jet and 2 b -jet categories. The initial state partons are sourced from the sea quarks and gluons of the colliding protons.

region is shown in blue (dark and light) in Table 9.1¹. The unfolding was performed as a function of the visible di-tau invariant mass $m_{\tau\tau}^{\text{vis}}$, which was defined using the two lepton momenta $\vec{p}_\ell^{\text{vis}}$. For leptonically-decaying τ candidates (τ_{lep}), this was taken as the lepton momentum, whereas for hadronically-decaying candidates (τ_{had}) this was defined as the sum of the momenta of all visible decay products. Particle masses were neglected as they were insignificant with respect to the $m_{\tau\tau}^{\text{vis}}$ range under study.

Alongside the unfolded measurement, dedicated BSM searches for non-resonant leptoquark and Z' interactions were performed using the detector-level reconstructed $m_{\tau\tau}^{\text{vis}}$. The signal regions were split in bins of b -jet multiplicity to improve the sensitivity to third-generation couplings, and relevant diagrams for non-resonant leptoquark interactions in association with zero, one, or two b -jets are shown in Figures 9.1a to 9.1c, respectively.

The main backgrounds in categories with no associated b -jets are di-boson production and jets misidentified as τ leptons. Single-top and $t\bar{t}$ production become significant in the 1 b -jet category, and dominant when ≥ 2 b -jets are present in the final state. To study the consistency of the background estimates and data, several validation regions were defined. Despite the small Drell-Yan and BSM model yields in the ≥ 2 b -jet category compared to the $t\bar{t}$ background, this region was included to act as a constraint for the $t\bar{t}$ cross-section normalisation in the profile likelihood fits described in Section 9.8.

¹Initially the analysis additionally aimed to produce unfolded cross-sections in the $e\tau_{\text{had}}$ and $\mu\tau_{\text{had}}$ regions, but a lack of sensitivity to BSM physics due to large $t\bar{t}$ background contributions in these regions motivated a stronger emphasis on the $\tau_{\text{had}}\tau_{\text{had}}$ channel.

9.2 Data and simulation samples

The analysed dataset was collected by the ATLAS detector during the Run 2 operating period of the LHC in 2015 to 2018, corresponding to an integrated luminosity of $(140.0 \pm 1.2) \text{ fb}^{-1}$ of pp collision data with centre-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. Only runs with fully-operational tracking detectors and calorimetry were accepted, as catalogued in the relevant Good Run Lists (GRLs) [62].

9.2.1 MC-simulated background samples

With the exception of the contribution from jets misidentified as τ_{had} candidates in multi-jet events, all background processes are estimated using MC generation. In general, all background samples were simulated to next-to-leading-order (NLO) or better precision for the matrix element (ME). These were then normalised to

$$N = \frac{\sigma \cdot \epsilon_{\text{filter}} \cdot \mathcal{B}}{\sum_i w_i} \quad (9.1)$$

where N is the normalisation factor applied to all event weights, σ is the integrated cross-section², \mathcal{B} is the branching ratio for a given decay mode, ϵ_{filter} is the product of the efficiencies of all applied event filters for the sample, and the generator-defined event weights are denoted w_i for each event i .

After simulating the central inelastic collision, minimum-bias events are superimposed over the generated events to emulate the effects of pile-up interactions within the same or neighbouring bunch crossings. These events are sourced from a soft QCD sample produced with PYTHIA 8.186 [307] using the A3 underlying event tune [308] and the NNPDF2.3_{LO} PDF set [309]. Events are reweighted to reproduce the pile-up conditions of the relevant data runs. All background samples are processed by a high-fidelity “full” simulation of the ATLAS detector based on GEANT4 [199, 310].

The MC generators used to produce particular samples are summarised in Table 9.2, with some additional details given in the following.

²Unless otherwise noted, σ is taken as the cross-section computed by the generator itself.

Table 9.2: MC generators used for the simulation of background samples. Nominal samples are shown in bold, and secondary generators used for two-point systematic uncertainties are italicised. For the parton distribution functions (PDFs), HS refers to the hard-scattering process generator, whereas PS refers to the parton shower simulation. For single-top production, DS refers to the diagram-subtraction method used to remove interference and overlap with $t\bar{t}$ production [315, 316]. The alternative diagram-removal (DR) scheme is used to estimate an uncertainty for this process. N_p refers to the number of additional parton emissions at matrix-element-level.

Process	Matrix element generator	PDFs for matrix element	Parton shower and matching	Underlying event tune	Hadron decay	QED FSR	Cross-section precision
Drell-Yan $Z/\gamma^* \rightarrow \ell\ell + \text{jets}$ $W \rightarrow \ell\nu + \text{jets}$	SHERPA 2.2.11 [311] + COMIX [312] + OPENLOOPS [313]	NNPDF3.0 _{NNLO} [309]	SHERPA 2.2.11 [317] + MEPS@NLO [318–321]	Default	—	—	NLO ^($N_p \leq 2$) LO ^($N_p \leq 5$)
<i>Modelling uncertainty</i>	POWHEG BOX v1 [314]	CTEQ6L1 ^(PS) [322] + CT10 _{NNLO} ^(HS) [323]	PYTHIA 8.186 [307]	AZNLO [324]	EVTGEN 1.2.0 [325]	PHOTOS++ 3.52 [326, 327]	NLO
$t\bar{t}$	POWHEG BOX v2 [315] $h_{\text{damp}} = 1.5m_{\text{top}}$	NNPDF3.0 _{NLO} [328]	PYTHIA 8.230 [307] + NNPDF2.3 _{LO} [309]	A14 [329]	EVTGEN 1.6.0 [325]	—	NLO
<i>Parton shower/hadronisation</i>	POWHEG BOX v2 $h_{\text{damp}} = 1.5m_{\text{top}}$	NNPDF3.0 _{NLO}	HERWIG 7.13 [330, 331] + MMHT2014 _{LO} [332]	Default	EVTGEN 1.6.0	—	NLO
<i>ME matching</i> ($p_T^{\text{hard}} = 1$)	POWHEG BOX v2 $h_{\text{damp}} = 1.5m_{\text{top}}$	NNPDF3.0 _{NLO}	PYTHIA 8.230 + NNPDF2.3 _{LO}	A14	EVTGEN 1.6.0	—	NLO
Single top-quark (tW using DS)	POWHEG BOX v2	NNPDF3.0 _{NLO}	PYTHIA 8.230 + NNPDF2.3 _{LO}	A14	EVTGEN 1.6.0	—	NLO
<i>tW (DR scheme)</i>	POWHEG BOX v2	NNPDF3.0 _{NLO}	PYTHIA 8.230 + NNPDF2.3 _{LO}	A14	EVTGEN 1.6.0	—	NLO
<i>Parton shower/hadronisation</i>	POWHEG BOX v2	NNPDF3.0 _{NLO}	HERWIG 7.04 [333] + MMHT2014 _{LO}	H7UE [333]	EVTGEN 1.6.0	—	NLO
<i>ME matching</i>	MADGRAPH5 AMC@NLO 2.6.2 [334]	NNPDF2.3 _{NLO} [309]	PYTHIA 8.230 + NNPDF2.3 _{LO}	A14	EVTGEN 1.6.0	—	NLO
Diboson VV	SHERPA 2.2.11 + OPENLOOPS	NNPDF3.0 _{NNLO}	SHERPA 2.2.11 + MEPS@NLO	Default	—	—	NLO ^($N_p \leq 1$) LO ^($N_p \leq 3$)
Higgs boson (gluon fusion)	POWHEG BOX v2 [335] + MINLO[336]	NNPDF3.0 _{NLO}	PYTHIA 8.212 + PDF4LHC15 _{NNLO} [337]	AZNLO	EVTGEN 1.6.0	—	NNLO

Drell-Yan lepton production with associated jets is modelled using SHERPA 2.2.11 with MEs calculated to NLO in QCD for events containing up to two additional partons [311], and to leading-order (LO) for up to five additional partons using the COMIX and OPENLOOP libraries [312, 313]. The $\tau\tau$ + jets sample serves a dual purpose as both the desired dominant contribution for the unfolded measurement, and the main background for BSM interpretations. Alongside the main sample, generator-related uncertainties were estimated using an alternative sample generated using POWHEG BOX at NLO in QCD, coupled to PYTHIA [307, 314].

To extend the statistical power of the $Z \rightarrow \tau\tau$ samples in the relevant high- $m_{\tau\tau}$ region, a high-mass SHERPA sample was produced with $m_{\tau\tau} > 120$ GeV, with identical settings to the existing sample. An equivalent sample was produced using POWHEG BOX + PYTHIA in $m_{\tau\tau}$ bins starting at 120 GeV. The inclusive samples were used for the $m_{\tau\tau}^{\text{vis,truth}} < 120$ GeV region, and the dedicated samples were used above that threshold, making the latter the primary contributions to the analysis signal regions.

Top quark pair production ($t\bar{t}$) is simulated to NLO accuracy in QCD, with the h_{damp} resummation damping factor³ set to $1.5 m_t$, where m_t is the top quark mass. Parton shower uncertainties are estimated through a two-point uncertainty in which the default PYTHIA parton shower is replaced with that of HERWIG [330, 331]. The uncertainty in the ME matching is estimated from an alternative sample with the PYTHIA p_T^{hard} option enabled⁴.

Single top quark production is simulated to NLO in QCD, using the four-flavour scheme for t -channel production, and the five flavour scheme for s -channel production and production in association with a W -boson (tW). The diagram-subtraction (DS) scheme for tW production is used as the nominal method to remove interference and overlap with $t\bar{t}$, and the diagram-removal (DR) scheme is used to derive an

³The parameter h_{damp} controls the matching of the POWHEG ME to the parton shower, and hence determines the p_T cutoff of the parton radiation against which the $t\bar{t}$ system recoils.

⁴This option changes the ME matching scale in PYTHIA from the scale at which POWHEG claims to have emitted the parton to smallest scale at which POWHEG *could* have emitted the parton (as calculated from the minimum p_T of final state partons in the POWHEG calculation), thus reducing the total parton radiation and acting as an estimate of a lower bound systematic variation on this quantity.

uncertainty for this process [315, 316]. Parton shower uncertainties are estimated by replacement of the default PYTHIA shower model by HERWIG, and uncertainties on the ME matching are evaluated by replacing the POWHEG-BOX matrix-element generator by MADGRAPH5 AMC@NLO.

Diboson production (VV) is generated for both fully leptonic and semileptonic final states using SHERPA, with matrix elements calculated up to NLO for up to one additional parton, and LO for up to three additional partons. Loop-induced $gg \rightarrow VV$ processes are modelled to LO for up to one additional parton emission.

Higgs boson production via gluon-gluon fusion is simulated using POWHEG-BOX to NNLO accuracy in QCD [335, 336]. The MC sample is normalised to the next-to-NNLO cross-section in QCD, plus electroweak corrections at NLO, through a global reweighting of the sample.

9.2.2 Signal models

The signal samples for the leptoquark and Z' models and EFT operators explored in this analysis were generated using MADGRAPH5 2.9.9 with the NNPDF2.3_{LO} set of PDFs [309, 334], to LO accuracy on MEs with up to two additional final-state partons. The parton shower and hadronisation modelling were performed using PYTHIA 8.306 with the NNPDF2.3_{LO} PDF set and the A14 tune of PYTHIA [329, 338]. Overlap between the ME calculation and parton shower emissions was removed using the CKKW-L merging procedure outlined in Refs. [339, 340]. Decays of bottom and charm hadrons were simulated using EVTGEN 1.2.0. The detector simulation for the signal samples was performed using the less resource-intensive ATLFast2 simulation [341], rather than the full GEANT4 simulation used for the above samples.

Leptoquark and Z' samples were generated using the `vector_LQ_UF0` model to calculate the ME at LO [342]. Events were initially generated from a pure-SM amplitude with up to two additional jets using CKKW-L merging. These were then reweighted to match the expected ME calculation for either the pure-BSM or interference term, since the SM and BSM processes were expected to be kinematically similar in the off-shell regime. The merging process was validated by comparing the SM predictions to a recent ATLAS

measurement of $Z + b$ -jet production outlined in Ref. [343], and the estimate agreed to within the 5% uncertainty quoted on the measurement. To verify the compatibility of the reweighted SM sample to the leptoquark and Z' calculations, the total cross-section of $\tau\tau b$ via leptoquark production was calculated separately for the interference and pure-BSM contributions. These were then compared to the corresponding contributions produced via the reweighting procedure, and agreed to better than 10%. Uncertainties on the cross-section were derived from variations of the renormalisation and factorisation scales and PDF set used, and were of this order or larger.

EFT couplings involving two τ -leptons were modelled using the SMEFTSIM 3.0 `topU31` model to calculate the matrix element [39], and otherwise followed the same process outlined above. Events were initially generated with the $c_{\tau W}$ operator set to $1/\text{TeV}^2$ and all others set to zero, and then reweighted to the per-event MEs calculated for each of the other operators to generate a weight-set for each operator under study.

9.3 Object reconstruction and event selection

In general, the reconstruction and selection requirements applied to detector objects followed the conventions and recommendations associated with the *Release 21* version of the ATLAS software framework (Athena) [86], with some specific choices of identification requirements and lower limits on reconstructed p_T for particular objects.

9.3.1 Trigger conditions

Although the final analysis results only targeted the $\tau_{\text{had}}\tau_{\text{had}}$ channel, the $\tau_{\text{lep}}\tau_{\text{had}}$ and $\tau_{\text{lep}}\tau_{\text{lep}}$ channels were also utilised as validation regions. In all cases, unrescaled triggers with the loosest available selection criteria were used. In the $\tau_{\text{had}}\tau_{\text{had}}$ channel, a logical OR of available single- τ and di- τ triggers was used, which generally corresponded to a threshold of 80 GeV on the leading⁵ τ_{had} candidate, and 60 GeV on the subleading candidate, with some deviations between data-taking periods. Channels containing a lepton utilised single-lepton triggers, which generally imposed a lower p_T threshold of 20 GeV to 26 GeV on the leading- p_T lepton for the tightest offline reconstruction requirement. A universal threshold

⁵Meaning the τ_{had} candidate with the greatest p_T for a given event.

on the offline reconstructed p_T of 27 GeV was applied to all leptons matching these triggers prior to any other selections, in order to standardise the effective trigger thresholds.

Trigger matching was applied for data events: trigger objects were required to have close overlap with corresponding offline-reconstructed particles. The nature of the required matched offline objects for a given event was determined by the type of trigger used in the selection of this event. In simulated events, trigger conditions were applied as they would be on data, and trigger efficiency scale factors extracted from prior measurements were applied as event weights [344–346].

9.3.2 Tau candidate reconstruction

τ_{had} candidates were identified by matching jets reconstructed by the anti- k_T algorithm with a radius of $R = 0.4$ to reconstructed tracks with a charge sum of ± 1 . The candidate was required to be in either the region $|\eta| < 1.37$ or in $1.52 < |\eta| < 2.47$, hence excluding the transition region between the barrel and endcap of the calorimeters⁶. The selected candidates were then processed by the tau identification RNN described previously in Section 3.3.3, yielding an output score that was used to categorise candidates. Candidates used in this analysis were required to pass a *tight* working point. In order to suppress a contribution from electron-induced jets misclassified as τ_{had} , candidates were also required to pass an electron identification algorithm as a secondary discriminant (eBDT). The TES corrections described in Section 3.3.3 were applied on a per-candidate basis, and several corrective factors were derived and applied to the simulation, as described in more detail in Ref. [101]. To suppress potential trigger-turn-on effects and ensure a high trigger efficiency, lower limits of $p_T > 90$ GeV and $p_T > 60$ GeV were applied to the leading and subleading τ_{had} candidates, respectively.

τ_{lep} decays cannot easily be distinguished from prompt lepton production⁷. As a result, the $\tau_{\text{lep}}\tau_{\text{had}}$ and $\tau_{\text{lep}}\tau_{\text{lep}}$ channels were defined without attempting to separate prompt light leptons from τ_{lep} decays. Since light leptons from a τ_{lep} decay tend

⁶The “calorimeter gap” region was excluded due to the large amount of material in front of the first active calorimeter layers for candidates with $1.37 < |\eta| < 1.52$.

⁷Here, *prompt* refers to particles originating directly from the Drell-Yan resonance or other weak boson decays.

to receive a small fraction of the original τ lepton p_T , candidate light leptons were required to have $p_T > 7 \text{ GeV}$.

Electron candidates were filtered with the same $|\eta|$ restrictions as τ_{had} candidates. The *medium* identification working point and the `FCHighPtCaloOnly` isolation criteria were utilised [347, 348].

Muons were reconstructed either by matching Muon Spectrometer (MS) objects to ID tracks or by extrapolating back to the original interaction point, following the requirements of the *medium* working point [349]. The `FCLoose` isolation requirement was additionally applied.

9.3.3 Jets

Jets were reconstructed from Particle Flow (PFlow) objects using the anti- k_T algorithm with a radius parameter of $R = 0.4$, as is further described in Section 3.3.1. A multivariate JVT algorithm was applied to jets with $p_T < 60 \text{ GeV}$ to reject those associated with pile-up vertices [350–352]. JES and JER calibrations were applied, and jets with $p_T < 20 \text{ GeV}$ or $|\eta| > 4.4$ were removed post-calibration.

9.3.4 b -tagging

The DL1r b -tagging algorithm was applied to all reconstructed jets to identify those with a high likelihood of originating from b -quarks [95]. The 77% operating point was utilised, with the b -tagging efficiency for the working point measured using an inclusive sample of b -jets in $t\bar{t}$ -enriched events as described in Section 3.3.2. Correction factors were applied to simulated samples to account for the differences between data and simulation in b -tagging efficiencies on b , c and light-quark jets⁸.

9.3.5 Overlap removal

Since tracks or calorimeter clusters could be associated with multiple reconstructed objects, an overlap removal algorithm was applied to disambiguate and remove objects

⁸These show negligible deviation from unity across a majority of the studied p_T range, but do introduce uncertainties at the level of a few percent.

Table 9.3: Overlap removal criteria applied to each possible pair of objects in an event, listed in order of application.

Retained object	Removed object	Criterion for removal
e	τ_{had}	$\Delta R_y < 0.2$
μ	τ_{had}	$\Delta R_\eta < 0.2$
e/μ	μ/e	If the muon is associated with a signal in the MS, the electron is removed. Else, the muon is removed.
e	jet	$\Delta R_y < 0.2$
μ	jet	$\Delta R_y < 0.2$
jet	e	$\Delta R_y < 0.4$
jet	μ	$\Delta R_y < 0.4$
τ_{had}	jet	$\Delta R_y < 0.2$

reliant on the same detector signals. The algorithm generally relies on the angular separation ΔR , defined as

$$\Delta R_y = \sqrt{(\Delta y)^2 + (\Delta \phi)^2}, \quad (9.2)$$

where the objects' rapidity y is replaced by the pseudorapidity η for the $\mu - \tau_{\text{had}}$ overlap condition. The overlap removal criteria are summarised in Table 9.3.

9.4 Background estimation for jets faking taus

The substantial rates of multi-jet and $W(\rightarrow \tau\nu)$ +jets production in ATLAS contribute a significant background across the $\tau_{\text{had}}\tau_{\text{had}}$, $e\tau_{\text{had}}$, and $\mu\tau_{\text{had}}$ final states, due to the misidentification of jets in such events as τ_{had} candidates (also referred to as τ_{had} *fakes*). Modelling such backgrounds via MC simulation poses substantial challenges, as correct modelling requires both high-precision simulation of soft jet properties in the non-perturbative regime, and accurate simulation of the detector response to hadrons [353]. Additionally, as misidentification rates of τ_{had} candidates improve with better identification and discrimination algorithms, the size of the simulated multi-jet and W +jets samples required to produce a given number of events containing τ_{had} fakes can be orders of magnitude larger than the number of simulated events with true τ leptons.

Identification of the source of a given jet generally depends on both the jet energy and associated shape variables. In particular, a summary measure of the jet spread can be defined as the jet width j , given as

$$j = \frac{\sum_i \Delta R^i p_{\text{T}}^i}{\sum_i p_{\text{T}}^i}, \quad (9.3)$$

where i iterates over the jet constituents, and the ΔR^i here defines the angular distance from the jet axis for a given constituent,

$$\Delta R^i = \sqrt{\left(y_{\text{jet}} - y^i\right)^2 + \left(\phi_{\text{jet}} - \phi^i\right)^2}, \quad (9.4)$$

where y denotes the rapidity and ϕ the polar angle of the constituent particle or jet axis, as labelled. Jets originating from τ leptons vary significantly in their structure depending on the kinematics of the parent particle and decay. The jet width is strongly influenced by the momentum of the τ parent, and the 1-prong or 3-prong nature of the τ decay. Misidentified jets are generally separable into three main categories. Quark-origin jets tend toward smaller j , whilst jets originating from gluons tend slightly toward larger j due to a higher parton multiplicity and softer fragmentation function [354]. Jets originating from pile-up within the same bunch-crossing will generally show a broad distribution of widths kinematically distinct from the other signals, but the rate of production of such jets is highly dependent on the jet energy and rapidity. In general, the relative fractions of these contributions will vary significantly with small changes in energy, rapidity, or substructure (1- or 3-prong) of the τ candidate.

Challenging backgrounds such as these are generally approached through *data-driven* methods, in which the frequency and kinematic distributions of the background processes are measured in dedicated *control regions* orthogonal to the signal region, in which sources of the true particles are either kinematically suppressed or forbidden by conservation laws. Several such methods for estimating backgrounds have been developed for similarly challenging final states in the past, including the matrix method [355], the ABCD method [356], and the fake factor method [357–359]. These methods all rely on the provision of one or several control regions with a high degree of kinematic and compositional similarity to the signal region, which can be difficult to achieve in the case of unfolded

measurements intending to cover a wide kinematic region. In the case of this analysis, one could consider using the same-charge $\tau\tau$ regions for such purposes. However, since the same-sign requirement successfully suppresses τ_{had} fakes from W +jets production as well as real τ leptons from Drell-Yan production, these regions are not a good match for the composition of the opposite-charge regions. In addition, available data yields in these regions are limited, leading to substantial statistical uncertainties.

9.4.1 The Universal Fake Factor method

Instead, this analysis implemented the *Universal Fake Factor* methodology published in Ref. [360], which was developed by the ATLAS Fake Tau Task Force (FTTF) specifically for τ_{had} final states. In order to pass the **tight** working point, candidate τ_{had} objects were required to achieve an RNN score greater than 0.8 and pass the **loose** eBDT veto condition. As a result, a background-enriched *anti-ID* ($\overline{\text{ID}}$) region could be defined for each signal and validation region as a complementary region in which one or two τ_{had} candidates failed this cut, but all other selections were maintained. As is the case in other fake factor methods, a transfer factor f was applied to extrapolate from yields in the anti-ID region to obtain an estimate for the ID region,

$$N_{\text{ID}}^{\text{SR}} = f \cdot N_{\overline{\text{ID}}}^{\text{SR}}. \quad (9.5)$$

These transfer factors were parametrised as a function⁹ of several factors, including the transverse momentum, rapidity, τ_{had} decay type (1 or 3-prong), and whether or not the candidate τ_{had} was matched to a trigger object. Additionally, separate fake factors were derived for the opposite-sign (OS) and same-sign (SS) regions, due to non-negligible differences in the kinematic distributions of W +jets events in these regions.

Instead of choosing a single control region, three separate *primitive* transfer factors were derived for quark-like (f_q), gluon-like (f_g) and pile-up-like (f_p) jets in dedicated control regions heavily enriched in each type. The transfer factors were extracted as

$$f = \frac{N_{\text{ID}}^{\text{CR}}}{N_{\overline{\text{ID}}}^{\text{CR}}} \quad (9.6)$$

⁹In practise, the transfer factors f were defined with a four-dimensional binning, and on application an interpolation between bins was used. In the implementation for this analysis, no attempt to smooth the interpolation was made, but rather the value for the closest bin was used directly.

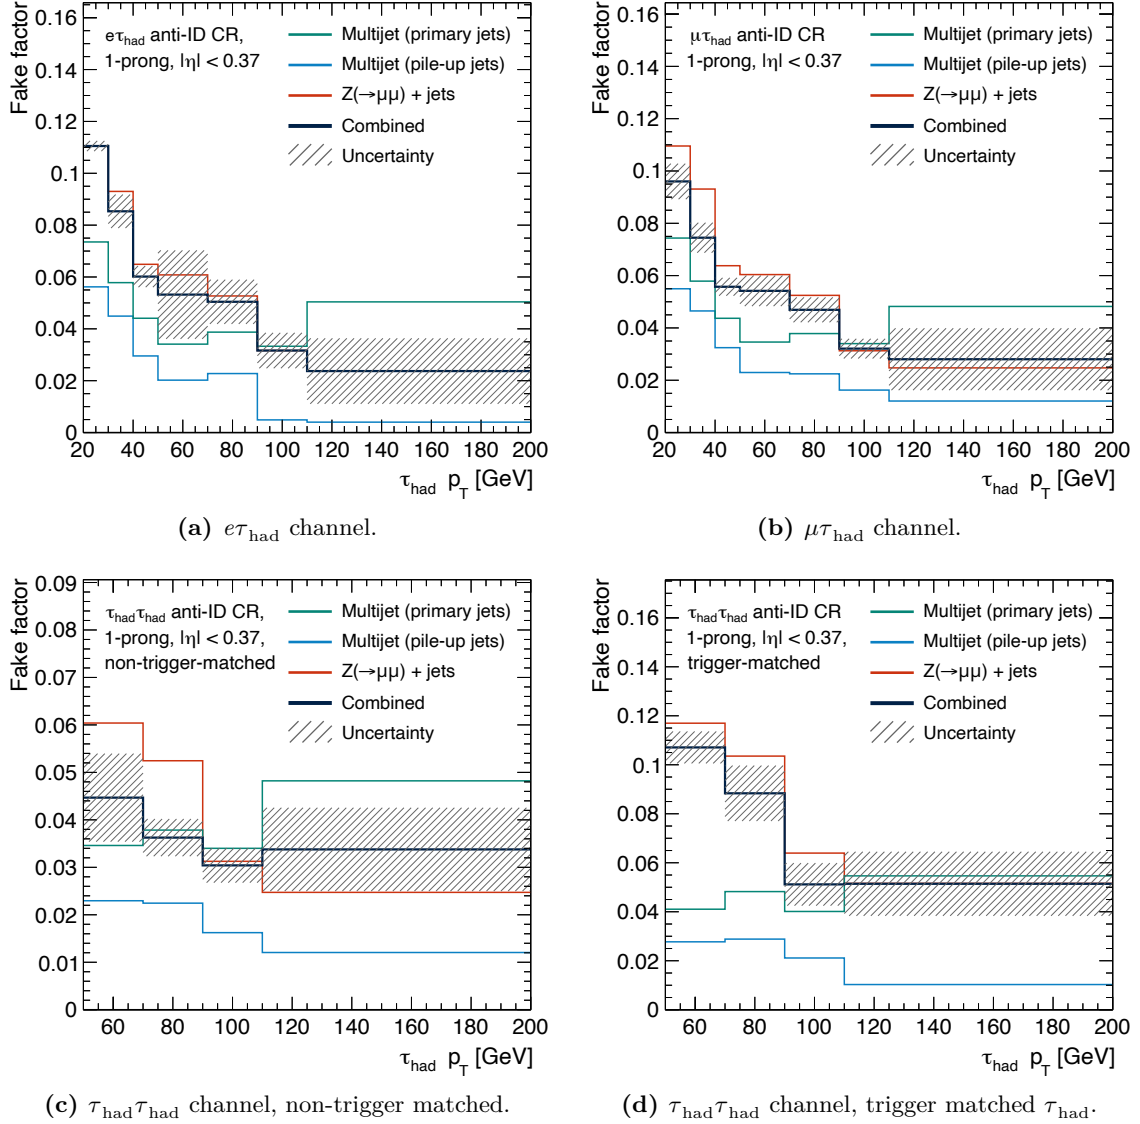


Figure 9.2: Fake factors derived from candidate OS di-tau events in samples enriched with quark-like (red), gluon-like (green), and pile-up (blue) jets, for 1-prong candidates in the barrel region ($|\eta| < 1.37$). A combined fake factor (black) has been derived from a fit of jet width templates for each primitive transfer factor to the jet width profile in the anti-ID complement of the given signal or validation region.

in each region. The control regions used were alternate event topologies enriched in each category of background jet. Samples of quark-like jets were obtained from $Z \rightarrow \mu^+ \mu^-$ events, where the additional final-state jets frequently originate from quarks produced in the hard interaction. For the other types, multi-jet events were selected with two jets that were back-to-back in the transverse plane. Events with additional electrons, muons, photons, or b -tagged jets were vetoed, and a requirement was placed on events with three jets that the p_T of the third jet must be less than that of the τ_{had} candidate. Candidates

were then separated by JVT score: jets with a JVT score above 0.8 were classed as gluon-like, whilst jets with a JVT score below 0.8 were considered likely to originate from pile-up. Fake factors derived from each of these samples for 1-prong candidates in the barrel region ($\eta < 1.37$) are shown in Figure 9.2.

Prior to application on τ_{had} candidates in the anti-ID complement of each signal and validation region, the primitive transfer factors were combined to single fake factor f_{comb} , shown as the black lines in Figure 9.2. The individual transfer factors were weighted by the proportions of quark-like α_q , gluon-like α_g , and pile-up jets α_p in a given anti-ID region¹⁰:

$$f_{\text{comb}} = \alpha_q f_q + \alpha_g f_g + (1 - \alpha_q - \alpha_g) f_p. \quad (9.7)$$

To derive the relative fractions of each jet species, template distributions of jet width j were extracted from each control sample and fit to anti-ID data with floating normalisation factors. The α -coefficients were extracted using the functional form

$$N_{\text{data}}^{\overline{\text{ID}}} = \alpha_q N_{q,i}^{\overline{\text{ID}}} + \alpha_g N_{g,i}^{\overline{\text{ID}}} + (1 - \alpha_q - \alpha_g) N_{p,i}^{\overline{\text{ID}}}, \quad (9.8)$$

where $N_{\text{data}}^{\overline{\text{ID}}}$ is the data yield of τ_{had} candidates in the anti-ID complement of the given signal or validation region, and $N_{q,i}^{\overline{\text{ID}}}$, $N_{g,i}^{\overline{\text{ID}}}$, and $N_{p,i}^{\overline{\text{ID}}}$ are the yield of anti-ID τ_{had} candidates in the quark-like, gluon-like, and pile-up jet enriched samples in a given jet width bin i . Since the transfer factors f were parametrised as described above, template fits were performed independently for each bin of this parametrisation, yielding a separate set of α coefficients for each.

The uncertainty in the evaluated fake factors was estimated from three independent components. First, the statistical uncertainty in the templates was estimated by generating varied templates in which the value in each bin of the template was sampled from a Gaussian with mean equal to the bin content, and width equal to the statistical uncertainty of the bin. The same fit procedure as above was performed on varied templates for 30 total fits, and the standard deviation of the resultant f_{comb} was taken as the uncertainty contribution. The second component was the statistical uncertainty in the primitive fake factors themselves. These were obtained from the statistical uncertainties on the derivation

¹⁰The anti-ID sample was assumed to consist solely of these jet species, implying $\alpha_p = 1 - \alpha_q - \alpha_g$.

regions, and propagated as 1σ up- and down-variations through the fit procedure, after which the two bands were symmetrised to give a single uncertainty contribution. The third component related to the uncertainty in the coefficients α in the fit. Since the α parameters are in general correlated, this component was estimated by varying the data yields in the anti-ID region through the same method used for the first component. A total of 30 fits were performed to varied datasets, and the standard deviation of the fit results was used as the uncertainty contribution. The three contributions were combined in quadrature.

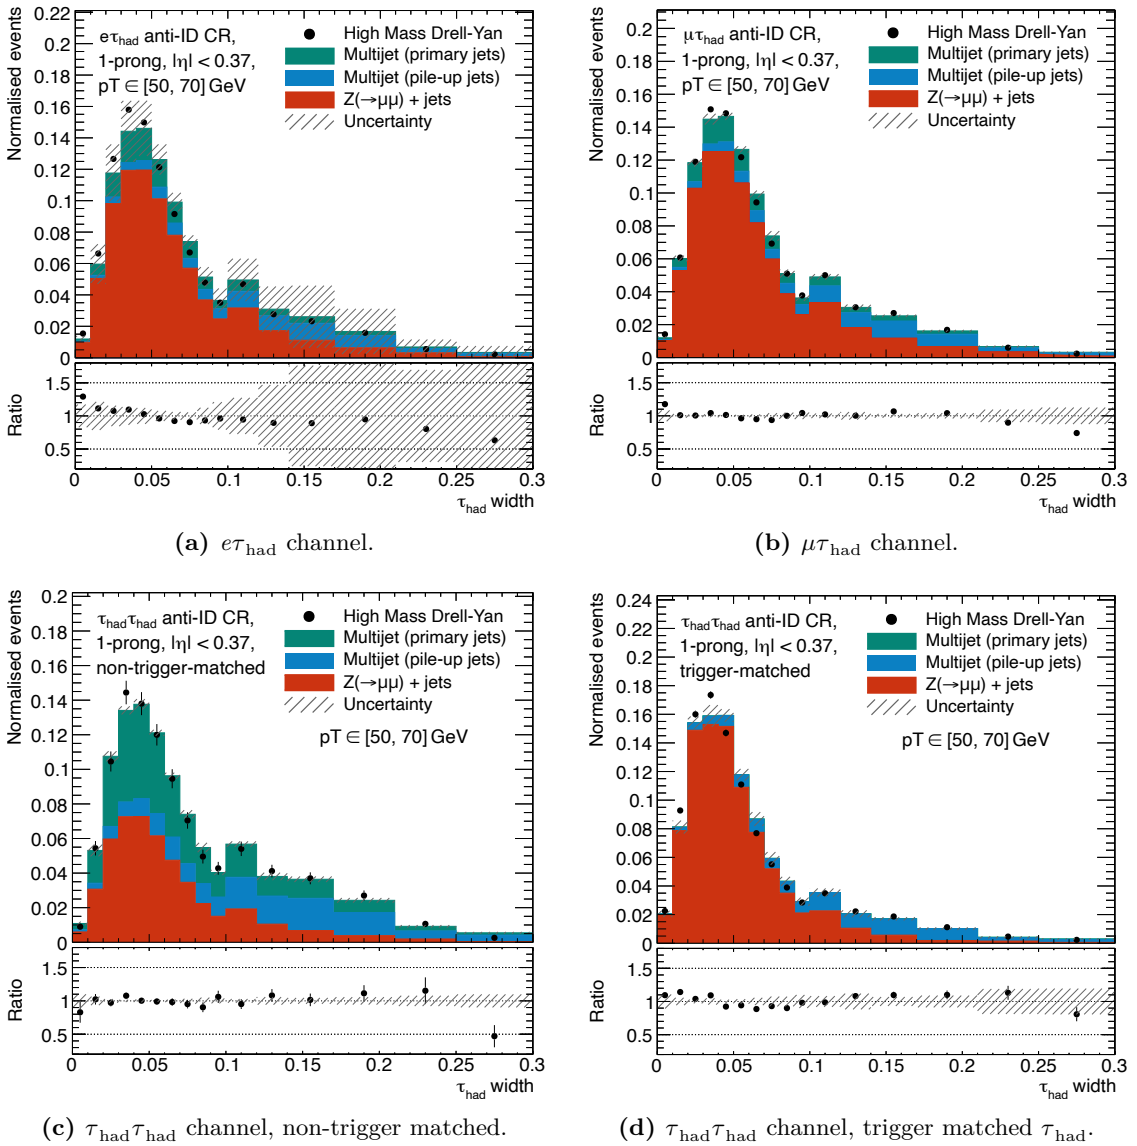


Figure 9.3: Template fits of anti-ID τ_{had} candidate yields in the $p_T \in [50, 70]$ GeV bin, for 1-prong candidates in the barrel region ($|\eta| < 1.37$), in OS events. Quark-like (red), gluon-like (green), and pile-up (blue) contributions were fit to data yields in the anti-ID complement (black points) of the given signal or validation region.

Example fits for the $p_T \in [50, 70]$ GeV bin are shown in Figure 9.3. For one-prong decays, the $\mu\tau_{\text{had}}$ channel is largely dominated by quark-like jets at low p_T and low η , with an increasing fraction of gluon-like jets at higher values. For three-prong decays, the candidates are almost exclusively quark-like. The $e\tau_{\text{had}}$ channel is dominated by quark-like jets across all p_T and η in both 1 and 3-prong decays. Since the τ_{had} candidates could be matched to trigger-level objects in the $\tau_{\text{had}}\tau_{\text{had}}$ channel, an additional distinction was made between matched and non-matched candidates. Non-trigger-matched candidates generally follow the behaviour of the $\mu\tau_{\text{had}}$ candidates, being quark-dominated at low η and p_T , and becoming gluon dominated at higher values. The three-prong candidates are again predominantly quark-like, except at large η . Trigger-matched candidates are dominantly quark-like across the kinematic phase space and both decay types.

9.4.2 Application to the $e\tau_{\text{had}}$ and $\mu\tau_{\text{had}}$ channels

An estimate for the total yield of τ_{had} fakes in each signal and validation region was obtained by applying the transfer function in Equation (9.5) on a per- τ_{had} -candidate basis in the anti-ID regions.

The anti-ID region also included a proportion of real τ leptons decaying hadronically that failed the **tight** τ identification RNN requirements. To avoid applying the transfer factor to this component¹¹, an estimate of the number of true τ leptons failing the τ RNN was derived from the MC datasets for all MC-modelled backgrounds in the analysis, as well as the Drell-Yan MC samples. In all cases, these yields are significantly smaller than the data-driven estimate.

The final estimated fakes yield in the $e\tau_{\text{had}}$ and $\mu\tau_{\text{had}}$ channels can be summarised as

$$N_{\text{bkg}}^{\ell\tau} = \sum_{\text{data}} \delta\left(\ell\tau^{\overline{\text{ID}}}\right) \cdot f_{\text{comb}}(\tau) - \sum_{\text{MC}} \delta\left(\ell\tau^{\overline{\text{ID}}}\right) \cdot w_{\text{MC}} \cdot f_{\text{comb}}(\tau), \quad (9.9)$$

where $\delta\left(\ell\tau^{\overline{\text{ID}}}\right)$ denotes a selection function which is equal to unity for events containing a lepton and an anti-ID τ_{had} candidate, and zero otherwise, hence isolating the anti-ID region. The fake factor $f_{\text{comb}}(\tau)$ is determined from the properties of the τ_{had} candidate, and for each MC event, w_{MC} denotes the product of all applicable event weights.

¹¹The source samples for the primitive transfer factors were carefully chosen to exclude true τ leptons.

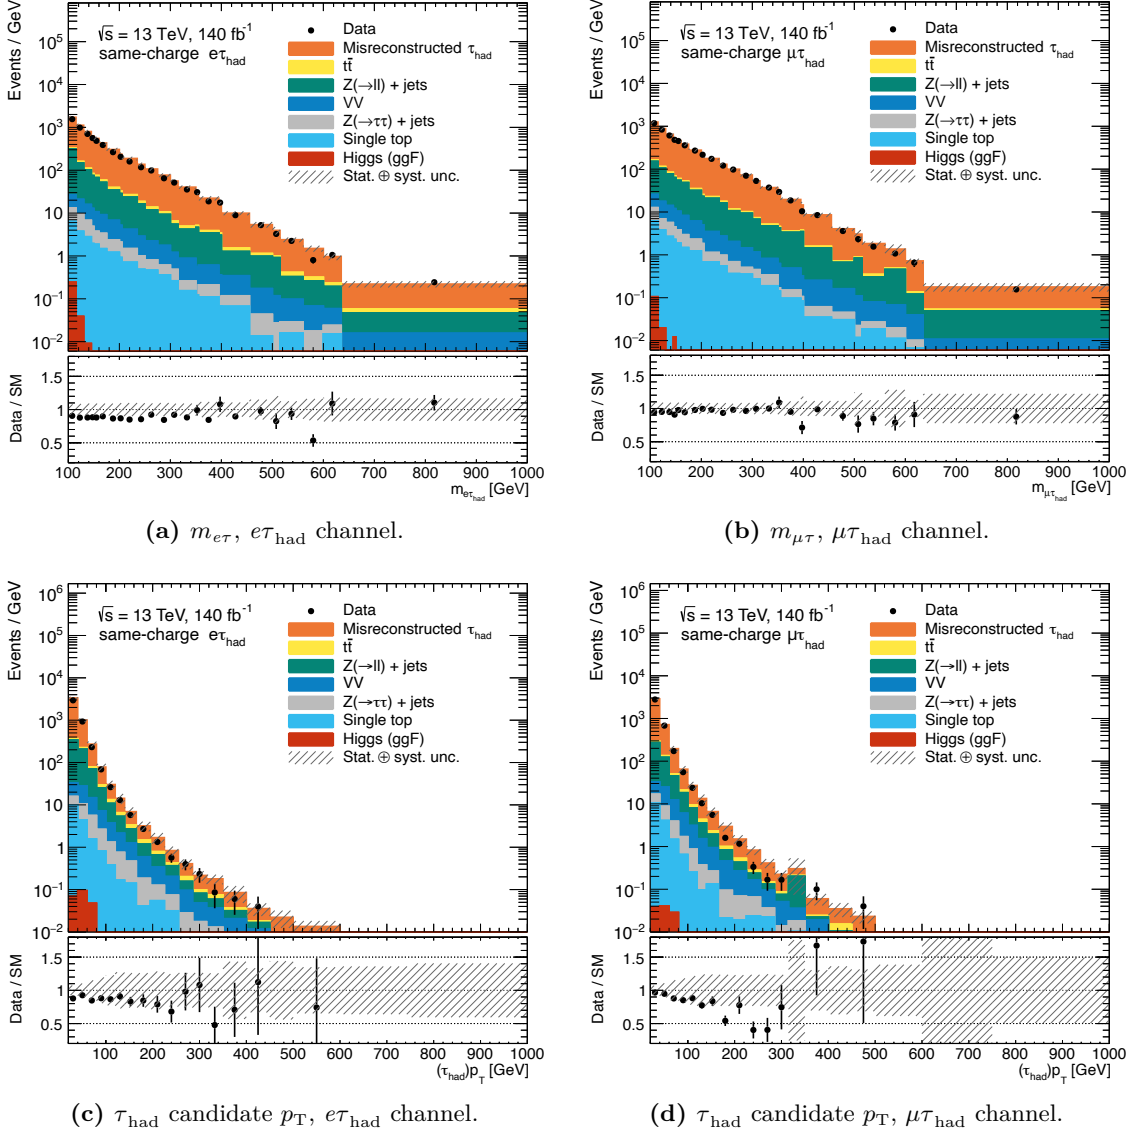


Figure 9.4: Comparison between data (black dots) and MC (coloured regions) in the SS validation region for the $e\tau_{\text{had}}$ (a, c) and $\mu\tau_{\text{had}}$ (b, d) channels, as a function of $m_{e\tau}$ (a), $m_{\mu\tau}$ (b), and the p_T of the τ_{had} candidate (c, d). The combined statistical and systematic uncertainties on the simulation are shown as a black hatched region.

Data-MC comparisons for the SS validation region are shown in Figure 9.4 for both the di-lepton invariant mass and the p_T of the τ_{had} candidate, and demonstrate good agreement across both channels. A small systematic overestimate of the total τ_{had} fakes yield is observed in both channels, but the prediction agrees with the data within the uncertainties on the total background estimate¹².

¹²A hypothesised explanation for the $e\tau_{\text{had}}$ discrepancy was non-trivial correlations in the validation region between the eBDT veto for τ_{had} identification, the electron identification, and the overlap removal rules removing potential τ_{had} fakes close to electrons. Attempts to verify this behaviour proved inconclusive due to the eBDT veto being crucial for correct reconstruction in this channel.




Figure 9.5: Representation of the three anti-ID regions used in the fake factor application in the $\tau_{\text{had}}\tau_{\text{had}}$ channel. The application and resulting yield fractions are shown in blue for positive contributions, and red for the double-counting subtraction. Region A is the analysis signal region.

9.4.3 Application to the $\tau_{\text{had}}\tau_{\text{had}}$ channels

The application of fake factors in the $\tau_{\text{had}}\tau_{\text{had}}$ channel was complicated by the 2-multiplicity in τ_{had} candidates. Events containing τ_{had} fakes fell into one of three categories: either the leading τ_{had} candidate τ_1 was fake, or the subleading τ_{had} candidate τ_2 was fake, or both were fake. The parameterised fake factors $f_{\text{comb}}(\tau)$ must be applied to a given event for each τ_{had} candidate individually. Hence, an event with two fakes was included in the fake-factor calculation twice: once in the application for events where the leading τ_{had} fails the τ ID requirements but the subleading τ_{had} does not (region B), and once in the complementary case where the subleading τ_{had} fails the τ ID but the leading passes (region C). This effect is demonstrated visually in Figure 9.5. To rectify this, a third region (D) was defined consisting only of events where both τ_{had} candidates fail the τ ID. A subtractive contribution to the total estimate was defined by applying the fake factors of both τ_{had} candidates to the events in region D, giving

$$\begin{aligned}
 N_{\text{bkg}}^{\tau\tau} = & \underbrace{\sum_{\text{data}} \delta(\tau_1^{\text{ID-bar}} \tau_2^{\text{ID}})}_{\text{region B data}} \cdot f_{\text{comb}}(\tau_1) - \underbrace{\sum_{\text{MC}} \delta(\tau_1^{\text{ID-bar}} \tau_2^{\text{ID}})}_{\text{region B MC}} \cdot w_{\text{MC}} \cdot f_{\text{comb}}(\tau_1) \\
 & + \underbrace{\sum_{\text{data}} \delta(\tau_1^{\text{ID}} \tau_2^{\text{ID-bar}})}_{\text{region C data}} \cdot f_{\text{comb}}(\tau_2) - \underbrace{\sum_{\text{MC}} \delta(\tau_1^{\text{ID}} \tau_2^{\text{ID-bar}})}_{\text{region C MC}} \cdot w_{\text{MC}} \cdot f_{\text{comb}}(\tau_2) \\
 & - \left[\underbrace{\sum_{\text{data}} \delta(\tau_1^{\text{ID-bar}} \tau_2^{\text{ID-bar}})}_{\text{region D data}} \cdot f_{\text{comb}}(\tau_1) \cdot f_{\text{comb}}(\tau_2) \right. \\
 & \quad \left. - \underbrace{\sum_{\text{MC}} \delta(\tau_1^{\text{ID-bar}} \tau_2^{\text{ID-bar}})}_{\text{region D MC}} \cdot w_{\text{MC}} \cdot f_{\text{comb}}(\tau_1) \cdot f_{\text{comb}}(\tau_2) \right].
 \end{aligned} \tag{9.10}$$

Comparisons of modelled background to data in the SS region are shown in Figure 9.6 as a function of $m_{\tau\tau}$, both inclusively and in b -jet bins. Good closure was observed for the inclusive and 0 b -jet categories, with a slight overestimate of τ_{had} fakes in the first bin. This bin contains relatively few events compared to the central regions of the distribution, and suffers from selection acceptance effects due to the lower edge of 100 GeV already being very close to the minimum p_T requirement for the leading τ_{had} of 90 GeV. Both the 1 b -jet and the ≥ 2 b -jet categories suffer from low yields in both data and prediction

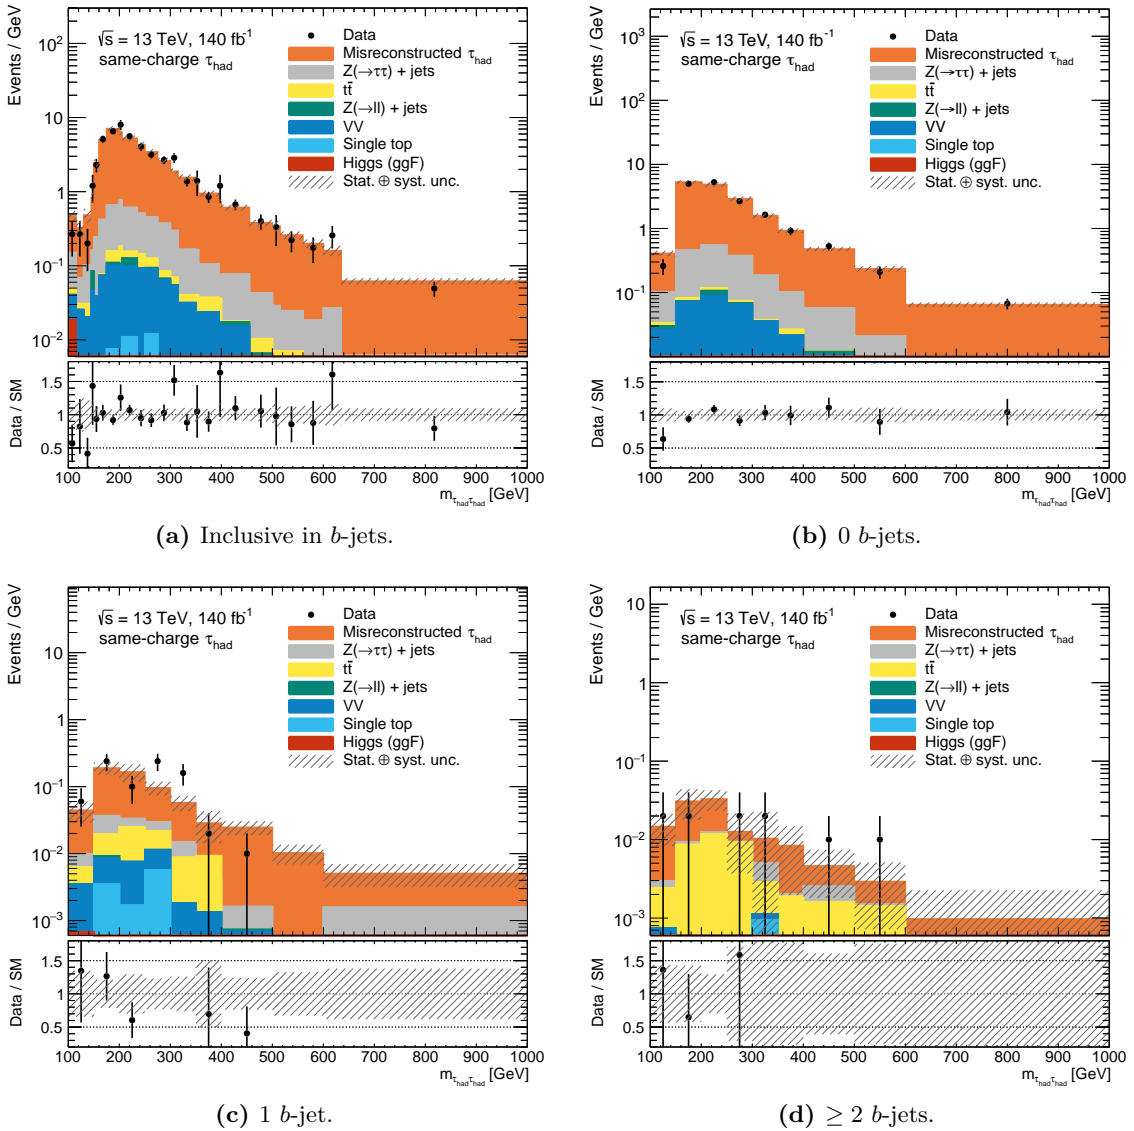


Figure 9.6: Comparison between data (black dots) and prediction (coloured regions) in the SS region for the $\tau_{\text{had}}\tau_{\text{had}}$ channel as a function of $\tau^+\tau^-$, both inclusive in b -jets (a), and split into 0, 1, and ≥ 2 b -jet categories (b, c, d, respectively). The combined statistical and systematic uncertainties on the simulation are shown as a black hatched region. Bins with no visible data had zero yield.

due to the heavy suppression of W +jets and multi-jet backgrounds when requiring one or more b -jets. Two bins in the 1 b -jet category showed deviations of size $\sim 2\sigma$, but are not considered a significant issue due to the large data uncertainties in this region.

9.5 Systematic uncertainties

Systematic uncertainties arise in two main categories. *Experimental uncertainties* pertain to limitations on the detector resolution for measured quantities used in reconstruction, the performance of the event reconstruction itself, and the calibration and validation of algorithmic or multivariate methods used for classification or discrimination on reconstructed objects. By contrast, *modelling uncertainties* on MC simulations used for signal and background modelling parametrise the level of analytical accuracy of the event generators and their underlying physics models. A breakdown of uncertainty contributions is shown in Figure 9.7 for the unfolded b -jet-inclusive cross-section, and in Figure 9.8 split by b -jet multiplicity as used in the detector-level BSM search and EFT interpretation. The individual components are discussed in the following sections.




Figure 9.7: Sources of systematic and statistical uncertainties on the unfolded inclusive cross-section, as a function of $m_{\tau\tau}^{\text{vis}}$.

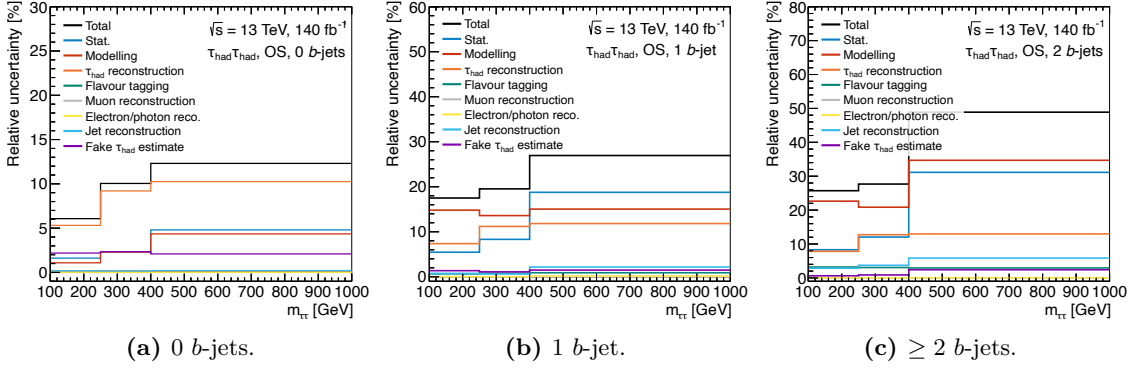


Figure 9.8: Uncertainty breakdown split by b -jet multiplicity in the OS $\tau_{\text{had}}\tau_{\text{had}}$ signal region.

9.5.1 Experimental uncertainties

Applicable experimental systematic uncertainties were considered for each of the object categories used within the analysis, including hadronically-decaying τ leptons, electrons, muons, jets, and b -jet identification.

The total set of recommended systematic uncertainties included over 260 individual classes, but many of these had negligible effects on the $m_{\tau\tau}$ distributions targeted by this analysis. A pruning method was applied to reduce this set to 69 non-negligible contributions, by neglecting systematic uncertainties for which the relative variation¹³ was smaller than 10% of the relative statistical uncertainty in the simulation prediction for the SM yield¹⁴ in all $m_{\tau\tau}$ bins. This method was applied to the full set of experimental uncertainties both before and after unfolding, and only uncertainties which were negligible in both cases were removed.

The most important experimental uncertainties pertain to the τ_{had} candidates themselves, and were propagated from uncertainties on the tau energy scale (TES) and τ_{had} identification algorithms. The latter varied substantially with the kinematics of the τ_{had} candidate, but were in general about 2% to 6% for the RNN and 1% to 2% for the electron-veto eBDT. The TES has been previously derived by ATLAS in on-shell $Z \rightarrow \mu\tau_{\text{had}}$ decays, using a *tag-and-probe* method to determine the energy calibration required to match the invariant mass of the $\tau_{\text{had-vis}}$ decay and tagged muon to the data [361].

¹³Defined here as the distance between the symmetrised $\pm 1\sigma$ bound and the nominal value, divided by the nominal value.

¹⁴Approximated using the Poissonian $1/\sqrt{N}$ in each $m_{\tau\tau}$ bin, where N is the prediction from MC samples and the data-driven τ_{had} fakes estimate.

Toward high $m_{\tau\tau}$, the statistical uncertainty in the MC background estimate is a significant factor. This is related to the difficulty of simulating events at very high $m_{\tau\tau}$ within inclusive samples, and has been addressed to some extent by the use of dedicated MC samples for Drell-Yan production with $m_{\tau\tau} > 120$ GeV, as described in Section 9.2.1.

Other significant uncertainties include those on the fake factors applied in the τ_{had} fakes estimate, which have been described in detail in Section 9.4.1, and are at the level of 2% to 3% across the $m_{\tau\tau}$ range. In categories with high b -jet multiplicity, uncertainties related to b -, c -, and light-jet tagging efficiencies become relevant [95], and are in general quantified by decomposing into independent contributions. Three of these components relevant to b -tagging passed the pruning method described in the next section, as did two for light-jet tagging and one for c -tagging. A similar decomposition is used for some of the uncertainties pertaining to the JES and JER [91], which were derived from a combination of simulation and in-situ measurements. Uncertainties arising from the measurement of the integrated luminosity and the variability in pile-up during data-taking were also taken into consideration [362, 363].

9.5.2 Modelling uncertainties on background and Drell-Yan samples

Uncertainties on the modelling of background and Drell-Yan samples were either estimated from comparisons at truth-level, or through two-point comparisons of different event generator combinations. The uncertainties considered in the latter case have been outlined in Table 9.2 above, and have been described in Section 9.2.1.

QCD scale uncertainties were estimated for Drell-Yan and $t\bar{t}$ production at truth-level from an envelope of scale variations around the nominal prediction, yielding a relatively flat 8% variation band across much of the $m_{\tau\tau}$ range. To remain on the conservative side, a universal 10% QCD scale uncertainty was assumed. Similarly, PDF uncertainties and variations of the strong coupling α_s were estimated from uncertainty eigenvectors provided for the NNPDF3.0_{NNLO} PDF set used for Drell-Yan and $t\bar{t}$ production, which were at the level of 2% to 4% across the central $m_{\tau\tau}$ range. These uncertainties were combined in quadrature to give a relatively flat 10% total uncertainty, which was then applied to the normalisation of the combined Drell-Yan and $t\bar{t}$ predictions.

9.5.3 Modelling uncertainties on BSM samples

Since the SM Drell-Yan predictions of the MADGRAPH_AMC@NLO MC generation used for the BSM signal samples were validated against the SHERPA 2.2.11 prediction, the QCD and PDF scale uncertainties were assumed to be of the same order. The PDF uncertainty was also checked by comparing the predictions of the nominal NNPDF3.0_{NLO} PDF set to the MSHT20_{NLO} [364] and CTEQ18_{NLO} [365] alternatives at truth-level. Since both methods yielded uncertainties $\lesssim 10\%$, a flat 10% uncertainty has been applied to all BSM signal samples to account for these.

An additional modelling uncertainty is associated with the renormalisation scale used to determine $\alpha_{\text{QCD}}(\mu)$, μ_R , and factorisation scale used to evaluate the PDFs, μ_F . For both, the nominal value was set to the order of the hard scattering scale in the event [366]. Up- and down-variations were generated for the varied scales $2\mu_{R/F}$ and $0.5\mu_{R/F}$, respectively. At truth-level, a linear fit to the signal scale variation as a function of $m_{\tau\tau}$ yielded

$$\begin{aligned} V_{\text{Int}}^{\uparrow} &= 0.983 - 0.000171 m_{\tau\tau} & V_{\text{Int}}^{\downarrow} &= 0.973 - 0.000396 m_{\tau\tau} \\ V_{\text{BSM}}^{\uparrow} &= 1.018 - 0.000103 m_{\tau\tau} & V_{\text{BSM}}^{\downarrow} &= 0.953 - 0.000211 m_{\tau\tau} \end{aligned} \quad (9.11)$$

where $m_{\tau\tau}$ is in GeV, V^{\uparrow} represents the relative 1σ up-variation and V_{\downarrow} the equivalent down-variation, given for both the BSM-only and interference terms in Equation (8.23). Due to the underlying signal models only being computed to LO, the scale variations were expected to be universally applicable across all signal models reweighted from the same event set, as is described in Section 9.2.2. The scale variations were applied as a weight-based uncertainty on an event-by-event basis to give up- and down-variation distributions post-reconstruction, and the weights were calculated with reference to the truth-level invariant mass of the di-tau system.

9.6 Modelling validation and detector-level cross-section

Agreement between data and modelled Drell-Yan and background contributions was initially verified in the validation regions outlined in Table 9.1 prior to unblinding¹⁵ the measurement regions. Validation was performed through investigation of data-MC

¹⁵As is standard in this style of analysis, the signal regions were *blinded* during the design and optimisation stages, and the data were not considered until the analysis methodology had been finalised.

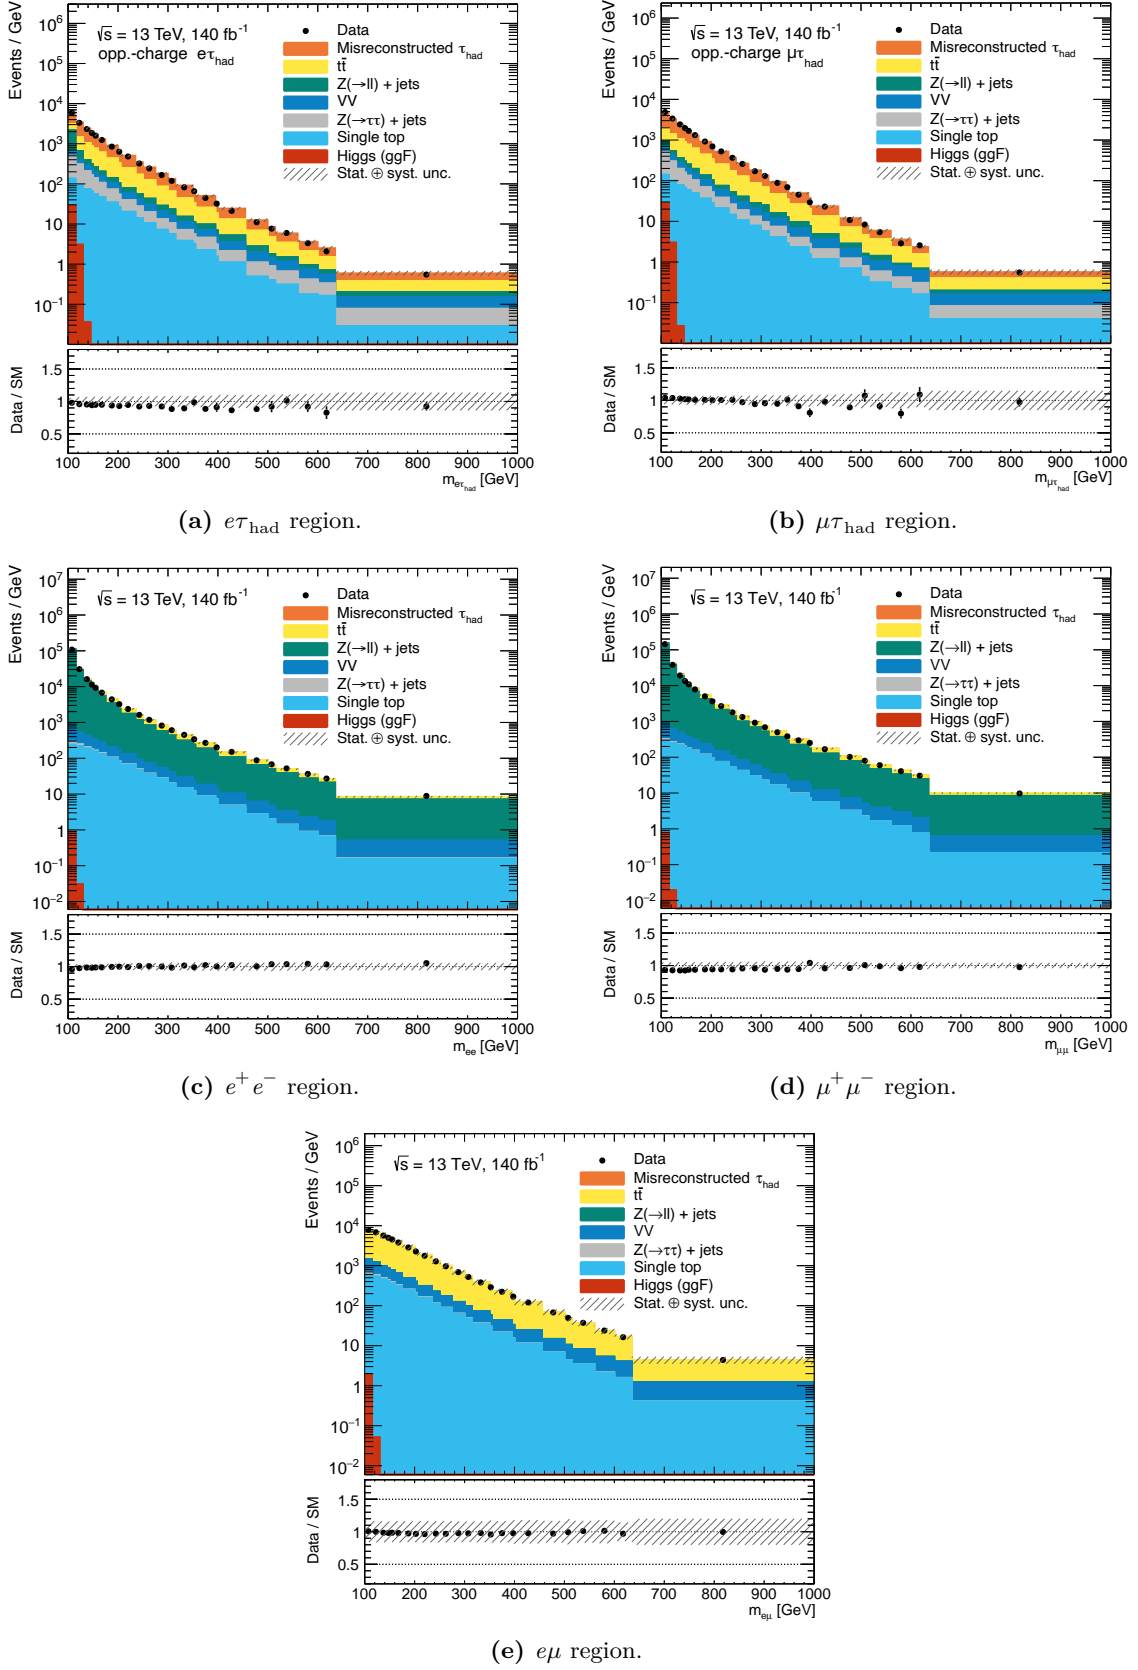


Figure 9.9: Comparison of data yields to Drell-Yan and background modelling in the validation regions, as a function of reconstructed $m_{\ell\ell}^{\text{vis}}$.

agreement of the detector-level cross-section differential in several variables, including the leading and sub-leading τ_{had} p_T , the p_T of the di-lepton system, the jet and b -jet multiplicities, and of course $m_{\ell\ell}^{\text{vis}}$. In general, all distributions agreed to within reasonable fluctuations defined by the statistical and systematic uncertainty bands. For brevity, only the $m_{\ell\ell}^{\text{vis}}$ distributions are discussed here.

The same-sign regions have already been discussed above in Section 9.4 as they are dominated by τ_{had} fakes. Figure 9.9 shows $m_{\tau\tau}$ distributions for the alternative channels

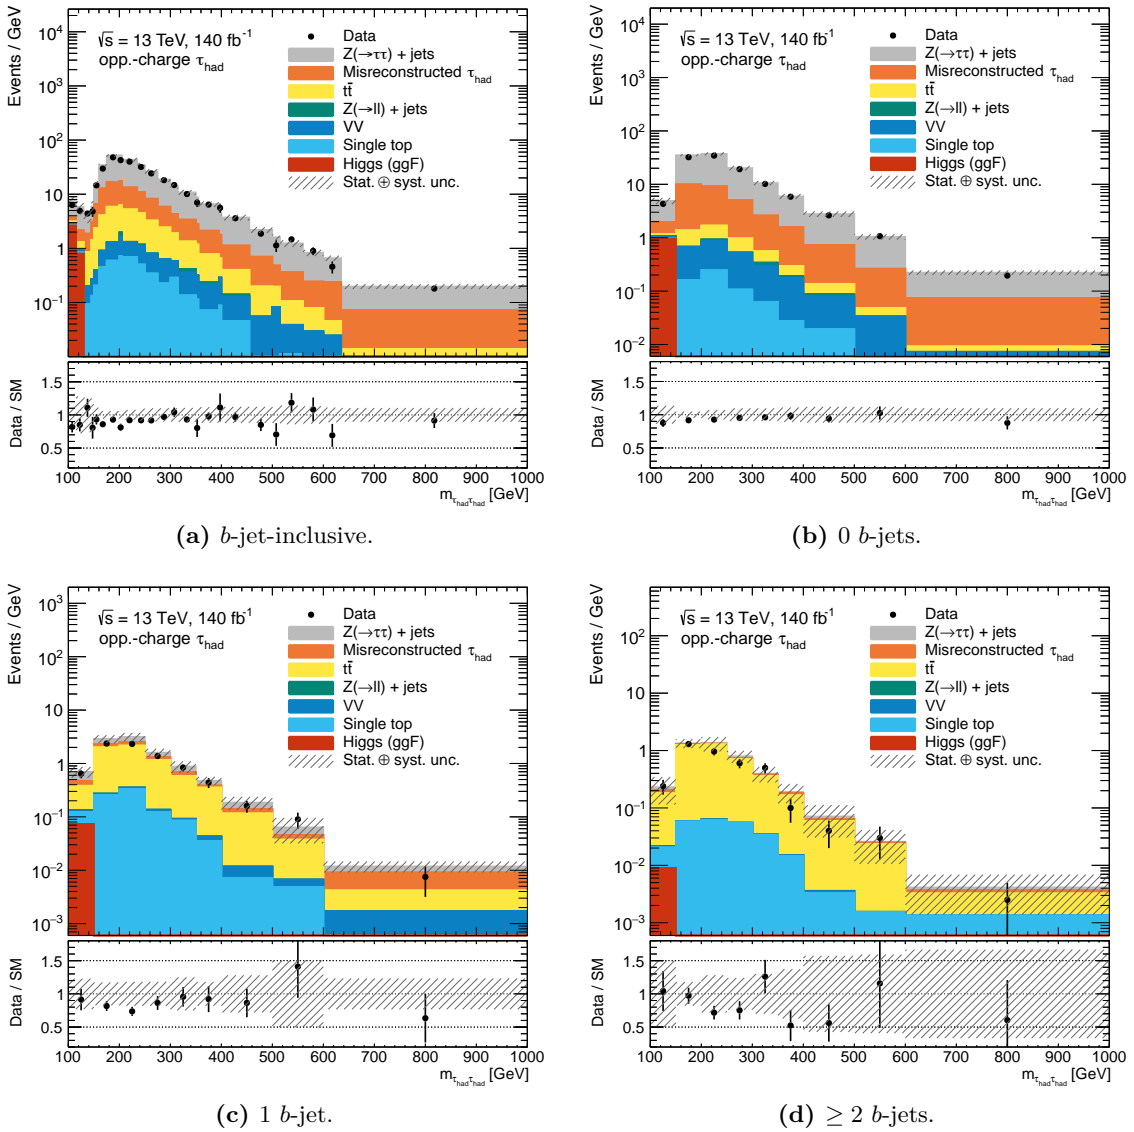


Figure 9.10: Comparison of data yields to Drell-Yan and background modelling in the signal region, as a function of reconstructed $m_{\ell\ell}^{\text{vis}}$. The b -jet-inclusive distribution used for the unfolded measurement is shown in (a), whilst the regions in (b), (c), and (d) split by b -jet-multiplicity were used for the BSM and EFT interpretations.

containing one lepton and one τ_{had} , or two leptons. In the former case, these are taken to support the accuracy of the $t\bar{t}$ and single-top modelling, whilst the latter regions show good agreement of the Drell-Yan modelling across the full $m_{\ell\ell}$ range. The $e\mu$ region validates the non-Drell-Yan components, as direct Drell-Yan production is forbidden in this channel, and tau decays to leptons are kinematically disfavoured at high $m_{\ell\ell}$.

The measurement and BSM-search signal regions are shown in Figure 9.10. Excellent agreement is observed in all variables both inclusively and split into b -jet multiplicity categories.

9.7 Measurement of the inclusive $\tau_{\text{had}}\tau_{\text{had}}$ cross-section

The single-differential cross-section $\partial\sigma/\partial m_{\tau\tau}^{\text{vis}}$ was measured in a fiducial phase space chosen to closely match the acceptance of the detector, minimising the necessary extrapolation between the two phase-space regions. *Particle-level* τ_{had} objects were constructed from all visible decay products of simulated τ_{had} decays. To construct the fiducial volume, events were required to contain exactly two τ_{had} objects, which were required to meet a minimum p_{T} threshold of 90 GeV for the leading τ_{had} and 60 GeV for the sub-leading τ_{had} . Only τ_{had} objects originating from the barrel or endcaps were accepted, excluding the gap between these using the same η requirement as for the reconstructed τ_{had} objects, as is described in Section 9.3.2. The particle-level objects were required to be opposite-sign. No requirement was placed that fiducial events must originate from any particular source process, meaning that although the unfolded cross-section was predominantly from Drell-Yan production, a non-negligible contribution of $t\bar{t}$, single-top, and di-boson events were included in the fiducial volume. The contribution from jets misidentified as τ_{had} objects is the only background subtracted from the data yield prior to the unfolding.

9.7.1 Accounting for detector effects

Detector effects are corrected for through an iterative Bayesian unfolding procedure. This procedure uses a Bayesian technique to estimate the true¹⁶ cross-section distri-

¹⁶The *truth* here refers to the underlying physical behaviour leading to the observed detector behaviour. The definition of this quantity should optimally be entirely model-independent, but with methods such as iterative Bayesian unfolding, model-dependent simulations or assumptions are invariably necessary in the implementation.

bution T , given a measurement suffering from detector effects D . Bayes' theorem is invoked in the form

$$P(T_i|D_j) = \frac{P(D_j|T_i) \cdot P_0(T_i)}{\sum_{l=1}^{N_T} P(D_j|T_l) \cdot P_0(T_l)}, \quad (9.12)$$

to compute the probability of a particular observed true distribution given a measured distribution, $P(T_i|D_j)$. The indices i and j here refer to the N_T and N_D discretised bins of the truth and measured differential distributions, respectively. The ingredients required are the migration probability $P(D_j|T_i)$, which gives the likelihood of a given event belonging to the true bin i migrating to an observed bin j through detector effects, and a prior expectation for the true distributions $P_0(T_i)$. The former is generally estimated from MC-based detector simulation in the form of a response matrix, whereas the latter may also be estimated from SM predictions or simulations, but can be updated through several iterations in which the unfolded measurement result is used as a prior for a subsequent iteration of the unfolding. The unfolded distribution is determined from the $P(T_i|D_j)$ as

$$\hat{N}(T_i) = \frac{1}{\varepsilon_i} \sum_{j=1}^{N_D} f_j N(D_j) \cdot P(T_i|D_j), \quad (9.13)$$

where the expected distribution of truth-level events $\hat{N}(T_i)$ also depends on a measurement $N(D_j)$, and the efficiencies $\varepsilon_i \neq 0$ for the observation of a given event belonging to the truth bin T_i in any bin of a measurement. The fiducial fraction in each bin is given as f_j . More comprehensive references on Bayesian unfolding (including the correct treatment of uncertainties) can be found in Refs. [367, 368].

Events containing jets misreconstructed as τ_{had} were subtracted from the detector-level distribution prior to unfolding. The response matrix used during this measurement is shown in Figure 9.11, which was populated using the truth-level and reconstructed $m_{\tau\tau}$ for fiducial simulated events. Cross-checks were performed in which the relative proportions of the Drell-Yan and background contributions were varied within uncertainties, and these were used to estimate an additional uncertainty for the response matrix from the composition of the fiducial region. Two iterations were utilised, as this represented an optimal compromise between the reduction of detector bias, and increased statistical uncertainties built up in successive iterations. Systematic uncertainties were propagated by




Figure 9.11: The transfer matrix encoding the probabilities $P(D_j|T_i)$ that an event from a given truth-level $m_{\tau\tau}$ bin T_i is observed in each reconstructed $m_{\tau\tau}$ bin D_j at detector-level, as estimated from MC simulation.

repeating the unfolding process with varied response matrices corresponding to individual up- and down-varied observables on events. Statistical uncertainties were propagated by varying the expected yield in each bin according to a Gaussian distribution.

Reconstruction efficiencies and fractions of non-fiducial events in detector-level distributions were estimated from simulation. In general, the reconstruction efficiency was 10% to 15% for $\tau_{\text{had}}\tau_{\text{had}}$ events across the measured $m_{\tau\tau}$ range, and non-fiducial events constituted 4% of the total, except in the first $m_{\tau\tau}$ bin, in which they contributed 27%.

Additional *signal injection* tests were performed with pseudodata to assess potential bias from the SM modelling used in the procedure. Events generated using one of the leptoquark models discussed in the following sections were added to an SM pseudodata sample and unfolded using the nominal response matrix. The signal significance was generally preserved and biases were found to be small.

The final unfolded distribution is shown in Figure 9.12, alongside SM predictions with the Drell-Yan process simulated using SHERPA 2.2.11 (MEPS@NLO) or POWHEG BOX + PYTHIA 8 (NLOPS). The models generally agree with the unfolded distribution within uncertainties, with a slight (less than 1σ) tendency for the models to predict greater yields than the data at low $m_{\tau\tau}$. The p-value for the model to yield the data distribution was estimated from the covariance matrix of the uncertainties, and was found to be 0.78.




Figure 9.12: Unfolded fiducial differential cross-section $\partial\sigma/\partial m_{\tau\tau}^{\text{vis}}$, alongside SM predictions with Drell-Yan simulated by MEPS@NLO (SHERPA 2.2.11, red) or NLOPS (POWHEG BOX + PYTHIA 8, blue). All contributing SM processes are included, except in the case of the green points, which show MEPS@NLO predictions for Drell-Yan production exclusively.

9.8 Statistical model

Alongside the differential cross-section measurement, a search for BSM behaviour was performed that utilised the models described in Chapter 8, with a particular emphasis on those showing a preference for third-generation leptons to violate lepton-flavour universality. For each model, a profile likelihood fit was performed on the detector-level $m_{\tau\tau}^{\text{vis}}$ distribution split by b -jet multiplicity, corresponding to the signal regions outlined in Table 9.1. The binning was optimised to reduce the impact of the MC statistical uncertainty and hence increase sensitivity, with bin edges at 100 GeV, 250 GeV, 400 GeV, and 1000 GeV. The `cabinetry` toolkit was used to steer the likelihood fit [369], utilising the `pyhf` framework as a back-end to compute the likelihood function itself and find the maximum [370].

A likelihood function was built from Poissonian and Gaussian probability terms as

$$\mathcal{L} = \prod_{i \in \text{bins}} \text{Pois} \left(N_i^{\text{data}}; \lambda = g^4 N_i^{\text{BSM}}(\vec{\theta}) + g^2 N_i^{\text{Int}}(\vec{\theta}) + N_i^{\text{SM}}(\vec{\theta}) + N_i^{\text{bkg}}(\vec{\theta}) \right) \cdot \prod_{k \in \text{nuis. params.}} \text{Gaus}(\theta_k; \mu = 0, \sigma = 1) \cdot \prod_{i \in \text{bins}} \text{Gaus} \left(\frac{\gamma_i}{(\sigma_i^{\text{rel}})^2}; \mu = 0, \sigma = \sqrt{\frac{\gamma_i}{(\sigma_i^{\text{rel}})^2}} \right), \quad (9.14)$$

where N_i^{data} is the observed number of data events for a given bin i , and N_i^{BSM} , N_i^{Int} , N_i^{SM} , and N_i^{bkg} are the pure-BSM term, interference term, Drell-Yan, and background yields, respectively. The latter were predicted by the nominal simulation samples, or in the case of the misreconstructed τ_{had} estimate described in Section 9.4.1, from data. The pure-BSM and interference terms were scaled by g^4 and g^2 , respectively, to match the coupling strength scaling behaviour described in Equation (8.23) for models with non-negligible interference terms. The coupling strength g forms the parameter of interest (POI).

The second term relates to the vector of nuisance parameters $\vec{\theta}$ which encode the effects of the systematic uncertainties, and were implemented as relative scaling parameters such that $\theta_i = 1$ corresponds to the upper 1σ bound predicted by the systematic uncertainty, and similarly $\theta_i = -1$ corresponds to the lower -1σ bound. The nuisance parameter values affect the relevant contributions to the predicted simulation yields, and approximating the behaviour of these uncertainties as Gaussian led to the Gaussian penalty term shown here. Some nuisance parameters were left unconstrained¹⁷, for which penalty terms were not added to the likelihood.

The third term relates to the impact of statistical uncertainties on the simulation predictions in each bin, which were parameterised by a set of nuisance parameters referred to as γ -parameters following the Barlow-Beeston method [371]. These modify the expected yield in each bin as $N_i(\vec{\theta}) \rightarrow \gamma_i N_i(\vec{\theta})$, constrained by the relative statistical uncertainty σ_i^{rel} through the Gaussian penalty in the third term of Equation (9.14).

Since the likelihood can generally become numerically large, minimisation¹⁸ is typically performed on the logarithm of \mathcal{L} , yielding estimators \hat{g} , $\hat{\vec{\theta}}$, and $\hat{\vec{\gamma}}$. A profile-likelihood test statistic was constructed following the prescription in Ref. [373]:

$$t(g) = -2 \log \frac{\mathcal{L}\left(g, \hat{\vec{\theta}}(g), \hat{\vec{\gamma}}(g)\right)}{\mathcal{L}\left(\hat{g}, \hat{\vec{\theta}}, \hat{\vec{\gamma}}\right)}, \quad (9.15)$$

where $\hat{\vec{\theta}}$ and $\hat{\vec{\gamma}}$ are the maximum likelihood estimate of the nuisance parameters for a fixed value of g . As is described in Ref. [373], the test statistic is asymptotically

¹⁷Specifically, the normalisations of the Drell-Yan and $t\bar{t}$ contributions, which were initially assumed to be 1, but were left to float freely so they can be determined from the data.

¹⁸In the case of `pyhf`, using the `Minuit` optimiser [372].

distributed¹⁹ according to a χ^2 distribution with number of degrees of freedom equal to the number of parameters of interest (in this case, one), and this approximation was used during the determination of confidence intervals. To this end, p -values were computed from the test statistic as

$$p(g) = \int_{t_{g,\text{obs}}}^{\infty} f(t_g|g) dt_g \approx \int_{t_{g,\text{obs}}}^{\infty} \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{t_g}} e^{-t_g/2} dt_g, \quad (9.16)$$

where f is the pdf of the test statistic $t_g \equiv t(g)$.

In the absence of substantial deviations from the SM, exclusion limits were computed using the modified frequentist CLs technique described in Ref. [374]. For each BSM or EFT model, the test statistic was computed as

$$t_{\text{CLs}}(g) = -2 \log \frac{\mathcal{L}\left(g, \hat{\hat{\theta}}(g), \hat{\hat{\gamma}}(g)\right)}{\mathcal{L}\left(0, \hat{\hat{\theta}}(0), \hat{\hat{\gamma}}(0)\right)}, \quad (9.17)$$

where the denominator encodes the likelihood of the observed data distribution with respect to the SM and background contributions, whereas the numerator also includes signal contributions with a non-zero coupling strength g . The CLs p -value $p_{\text{CLs}}(g)$ corresponding to a given value of g was computed from the asymptotic approximation, and is related to $p(g)$ in Equation (9.16) by

$$p_{\text{CLs}}(g) = \frac{p(g)}{1 - p_{\text{SM}}(0)}, \quad \text{where } p_{\text{SM}}(g) = \int_{-\infty}^{t_{g,\text{obs}}} f(t_g|g) dt_g, \quad (9.18)$$

in terms the probability of the SM-only hypothesis $p_{\text{SM}}(g)$. The test statistic was scanned to determine a value of g corresponding to a confidence level of 95% ($p = 0.05$).

In addition to fits on real datasets, an *Asimov* dataset was constructed from the SM prediction as the sum of all SM and background contributions [373]. The Asimov dataset was used with nuisance parameter values determined from the fit to data (referred to as *post-fit*) to determine the *expected* sensitivity of the dataset, which can be compared to the *observed* limits from the data itself²⁰.

¹⁹In the large sample limit.

²⁰A significantly higher *observed* upper limit than *expected* can be indicative of BSM behaviour, and generally merits computation of a dedicated p -value for the excess.

Table 9.4: Leptoquark and Z' models tested, including the coupling constants and mass points generated. Where multiple coupling constants are given, these were parametrised as a single POI, and all other coupling constants were set to zero. Note that for the β and y coupling matrices, these are also referred to by numeric indices in the following, i.e. $\beta_L^{b\tau} = \beta_L^{33}$, $\beta_L^{s\tau} = \beta_L^{23}$.

Model	Coupling constants / POI	$m_{LQ}/m_{Z'}$ [TeV]	Lagrangian
U_1	$\beta_L^{b\tau}$	1, 1.5, 2, 2.5, and 3	Eq. 8.19
	$\beta_L^{s\tau}$	3	
	$\beta_L^{b\tau} = -\beta_R^{b\tau}$	3	
	$\beta_R^{b\tau}$	3	
S_1	$y^{b\tau}$	1.4	Eq. 8.15
Z'	$\zeta_{qL}^{bb} = \zeta_{qL}^{\tau\tau}$	3	Eq. 8.25
	$\zeta_{dR}^{bb} = \zeta_{eR}^{\tau\tau}$	3	

9.9 Search for non-resonant leptoquark and Z' interactions

Five leptoquark hypotheses and two Z' models with varying coupling configurations were studied, as outlined in Table 9.4. Simulated signal samples were produced as described in Section 9.2.2, and the pure-BSM and interference terms were produced separately and scaled as described in Equations (8.23) and (9.14) to correctly treat the interference between the SM and BSM cross-section contributions. The scaling of the predicted cross-section with the leptoquark mass is given in Equation (8.23), and was applied to the interference term yields ($\sim m_{LQ}^{-2}$) and pure-BSM term yields ($\sim m_{LQ}^{-4}$) for all models to extrapolate from the $m_{LQ} = 3$ TeV mass point to lower masses. Several mass points were generated for the $\beta_L^{b\tau}$ model to confirm the expected scaling behaviour.

9.9.1 Signal distribution correction for the ≥ 2 b -jet category

The MC event generation technique for these signal samples used a sampling of the SM Drell-Yan kinematic phase space. The low likelihood of very high $m_{\tau\tau}$ events in the SM and the low reconstruction efficiency for $\tau_{\text{had}}\tau_{\text{had}}$ events, combined with the low b -tagging efficiencies in the ≥ 2 b -jet bin²¹ led to a strong suppression of detector-level events in the high- $m_{\tau\tau}$ region in this bin. In addition, the reweighting procedure from the SM

²¹Since the 77% working point was used, the 2 b -jet identification efficiency is expected to be on the order of 50% to 60%.

prediction to a given signal model will reject many events based on allowable couplings of the model²², effectively reweighting these to zero and further increasing statistical uncertainty. The latter effect was found to be most prevalent for the pure-BSM terms, and as a result a large but statistically insignificant upward fluctuation was observed in the 400 GeV to 1000 GeV bin of the ≥ 2 b -jet category for all leptoquark signal models, caused by a very small number of very high-weight events. Figure 9.13a shows this behaviour in the ≥ 2 b -jet category alongside the 1 b -jet category, which is expected to have a similar line-shape. Additionally, the distributions at truth-level (prior to the effect of reconstruction efficiencies) shown in Figure 9.13b follow a more reasonable line shape.

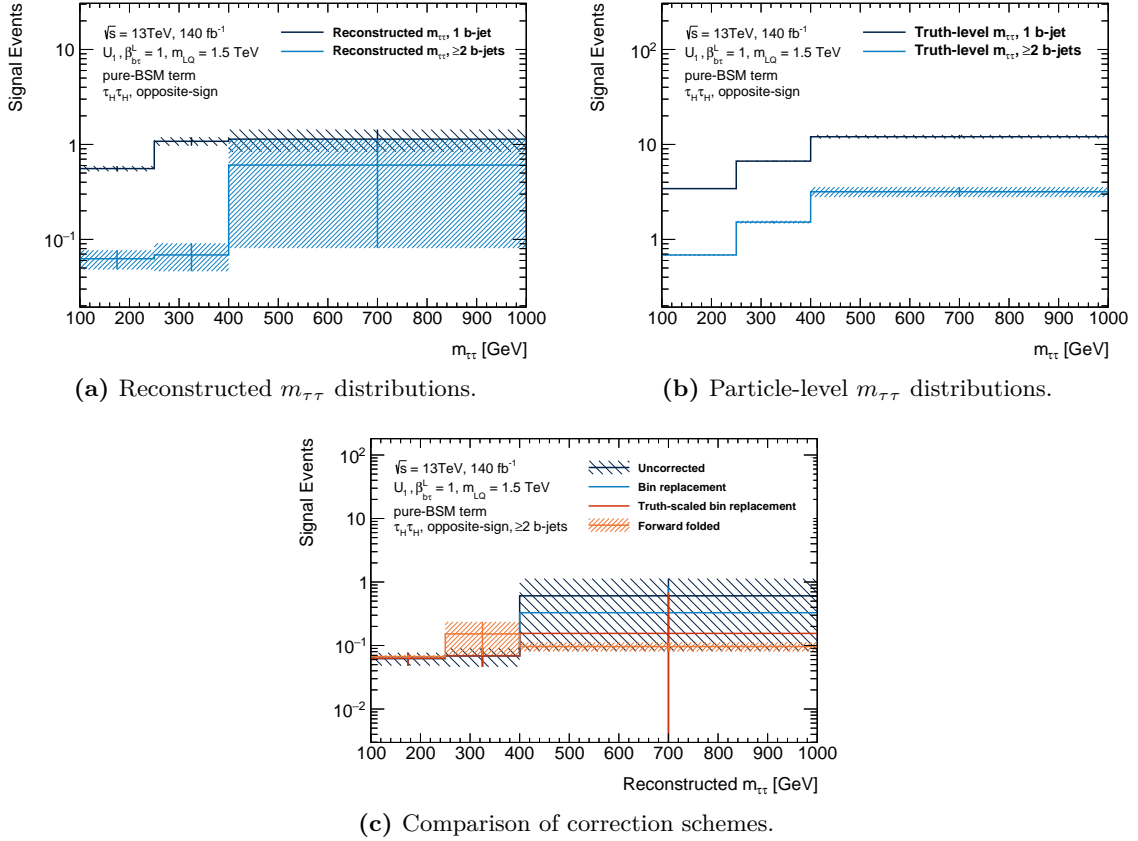


Figure 9.13: Opposite-sign, $\tau_{\text{had}}\tau_{\text{had}}$, ≥ 2 b -jet category of the BSM-term signal model with parameters $\beta_{b\tau}^L = 1$, $m_{LQ} = 1.5$ TeV. The reconstructed $m_{\tau\tau}$ distributions for the 1 b -jet and ≥ 2 b -jet categories are shown in (a), highlighting the upward fluctuation seen in the last bin, whilst (b) shows equivalent data at truth-level, in which the same discrepancy is not seen. The yields resulting from the three correction methods discussed in this section are shown in (c). For readability, only the uncertainties on the original distribution (dark blue) and the forward-folded method (orange) are shown. The bin replacement and truth-scaled bin replacement methods retain the dark blue uncertainty band of the uncorrected distribution.

²²In terms of chirality or initial-/final-state kinematics, for example.

Due to the very low yields in this category for all three of the signal, SM and background predictions, it was observed that the relative impact of the fluctuation of the central value in this bin was very large on the fit result. In order to improve the leptoquark signal estimate in the ≥ 2 b -jet category, three alternative methods for estimating the events in this bin were developed and studied, and the resulting distributions in the ≥ 2 b -jet category are shown in Figure 9.13c.

- **Bin replacement:** the yield in the ≥ 2 b -jet 400 GeV to 1000 GeV bin was replaced by the yield in the equivalent 1 b -jet category bin.
- **Truth-scaled bin replacement:** the yield in the affected bin was replaced by the yield in the equivalent 1 b -jet category bin, then scaled by the ratio between the yields of the corresponding bins in the ≥ 2 b -jet and 1 b -jet categories at truth-level.
- **Forward-folding:** all bins of the ≥ 2 b -jet category in the signal model were derived by forward-folding the truth distribution as

$$N_i^{\text{pred}} = \varepsilon_{2b/\text{inc}} \sum_{j \in \text{bins}} M_{i,j} \cdot N_j^{\text{truth}} \quad (9.19)$$

where N_i^{pred} , the i^{th} bin of the forward-folded prediction, is calculated from a migration matrix M and the truth-level distribution N^{truth} . Since the detector response should depend only on event kinematics and not the underlying production mechanism, a response matrix populated with the corresponding SM weight for each event in the MADGRAPH signal sample was utilised to achieve lower uncertainties on the migration matrix due to the greater statistical power in the SM prediction. Similarly, the $m_{\tau\tau}$ distributions between b -jet multiplicity categories were considered similar enough that the b -jet-inclusive response matrix could be used to reduce statistical uncertainties. The migration matrix was computed from the row-wise normalisation of the response matrix R ,

$$M_{i,j} = \frac{R_{i,j}}{\sum_k R_{i,k}}. \quad (9.20)$$

The binomial approximation was used to propagate the statistical uncertainties to the matrix entries (which are efficiency-like ratios). To account for efficiency

differences, the b -tagging efficiency ratio $\varepsilon_{2b/\text{inc}}$ was estimated from the ratio of the total number of reconstructed events in the ≥ 2 b -jet and inclusive distributions for the SM prediction. This efficiency was found to be between 47% and 50% for the various samples and channels, matching expectations for a double selection of b -tagged jets at the $\sim 77\%$ efficiency working point.

In the first two cases, the original statistical uncertainty was retained, and an equivalent replacement was performed in the case of each systematic uncertainty distribution. In the third, the statistical uncertainties on the truth distribution and the migration matrix were propagated to the result as independent uncertainty contributions²³, as

$$\left(\sigma_i^{\text{pred}}\right)^2 = \sum_j \left((\varepsilon_{2b/\text{inc}})^2 \cdot (\sigma_{i,j}^M)^2 \cdot (N_j^{\text{truth}})^2 + (\varepsilon_{2b/\text{inc}})^2 \cdot (\sigma_j^{N_{\text{truth}}})^2 \cdot (M_{i,j})^2 + (\sigma_{2b/\text{inc}}^\varepsilon)^2 \cdot (M_{i,j})^2 \cdot (N_j^{\text{truth}})^2 \right), \quad (9.21)$$

where $\sigma_{i,j}^M$, $\sigma_j^{N_{\text{truth}}}$, $\sigma_{2b/\text{inc}}^\varepsilon$, and σ_i^{pred} are the uncertainties on the migration matrix element, truth-level yield, b -tagging efficiency ratio, and the final prediction, respectively. In the case of non-agreement between the forward folded distribution and the direct prediction within uncertainties, the uncertainty was instead taken as the difference between the predictions. The dominant uncertainty is the propagated statistical uncertainty in the truth-level prediction.

Due to the factor of 10 to 20 higher event yields at truth level in the ≥ 2 b -jet category, the forward-folded distribution shows substantially reduced statistical uncertainties. Since all leptoquark and Z' signal models were derived as reweighted samples from the same underlying SM simulation dataset, the same procedure was applied to all signal models.

9.9.2 Results

Figure 9.14 shows post-fit distributions in the three b -jet multiplicity categories, using the binning outlined above. Normalisation factors of 0.88 ± 0.07 and 0.8 ± 0.1 were extracted for the Drell-Yan and $t\bar{t}$ background contributions, respectively. Leading uncertainties

²³This is not an entirely correct treatment, as the underlying dataset used for derivation was used in the derivation of both the transfer matrix and the truth distribution. However, the impact of correlations was expected to be negligible, as the relative uncertainties on the transfer matrix and efficiency were small compared to the those on the truth distribution.

post-fit included the Drell-Yan and $t\bar{t}$ modelling uncertainties, followed by the statistical uncertainty in the 400 GeV to 1000 GeV bin in the 1 b -jet category. A U_1 leptoquark signal with $\beta_L^{b\tau}/\sqrt{2} = 2.5$ is overlaid, as well as the pure-BSM and interference terms both separately. The importance of the correct treatment of the interference term can be seen in this breakdown: the total yield is significantly less than the pure-BSM yield in all three categories, and hence limits set using only the pure-BSM yields would be expected

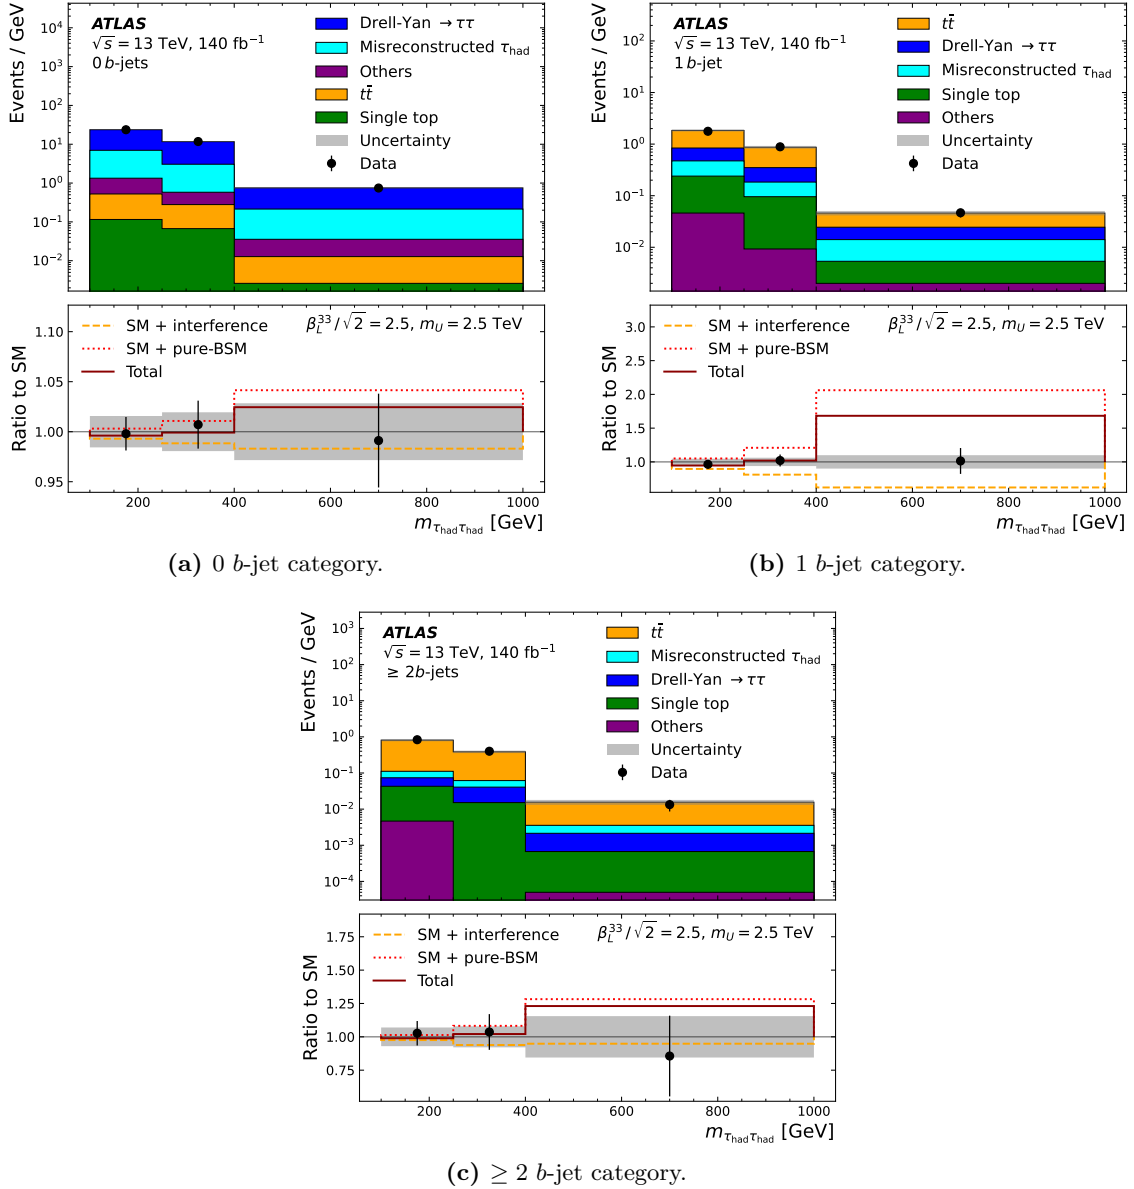


Figure 9.14: The post-fit SM $m_{\tau\tau}^{\text{vis}}$ distribution, split into b -jet multiplicity categories. The individual contributions to the SM prediction are shown alongside the data, and a U_1 vector leptoquark model is overlaid with $\beta_L^{b\tau}/\sqrt{2} = 2.5$ and all other couplings set to zero. Signal contributions are shown in the ratio plot in terms of the pure-BSM term (red dots), interference term (yellow dashes), and combination of both (black line).

to be stronger than those from a correct treatment of the interference.

Exclusion limits on couplings were set for all leptoquark and Z' models for masses between 1.5 TeV and 3 TeV²⁴, and are shown in Table 9.5 for all studied models.

Limits on the U_1 vector leptoquark models preferred by the b -anomalies and discussed in Ref. [291] (see Section 8.3.3 and Figure 8.6) are shown in Figure 9.15. In the case of the left-handed-only coupling with $\beta_R^{b\tau} = 0$, the results begin to exclude a small part of the preferred region outlined in Ref. [291], whereas for the $\beta_L^{b\tau} = -\beta_R^{b\tau}$ scheme the expected and observed limits are not yet sensitive to the preferred region.

Breakdowns of the uncertainty contributions in each fit showed the statistical uncertainty in the data to be dominant, with statistical sources²⁵ comprising 80% of the total uncertainty at the 1.5 TeV mass point for the vector leptoquark model with nonzero $\beta_L^{b\tau}$, and increasing for greater leptoquark masses. Dominant sources of systematic uncertainty post-fit include the $t\bar{t}$ modelling uncertainties (36% of the total uncertainty for the 1.5 TeV vector leptoquark model with nonzero $\beta_L^{b\tau}$), the Drell-Yan modelling uncertainties (30%), the signal and background MC statistical uncertainties (20% and 22%, respectively), and the uncertainties related to object reconstruction (31%)²⁶.

Table 9.5: Expected and observed 95% confidence exclusion limits of the POI g for the leptoquark and Z' signal models. Expected limits were derived from an Asimov data set consisting of the nominal Drell-Yan prediction plus all nominal background predictions.

Mass		1.5 TeV		2 TeV		2.5 TeV		3 TeV	
Model/parameters		exp.	obs.	exp.	obs.	exp.	obs.	exp.	obs.
U_1	$\beta_L^{b\tau}$	2.18	2.14	2.86	2.82	3.55	3.50	4.24	4.18
	$\beta_L^{s\tau}$	2.56	2.55	3.41	3.40	4.26	4.24	5.11	5.09
	$\beta_R^{b\tau}$	1.98	2.00	2.64	2.67	3.30	3.34	3.97	4.01
	$\beta_L^{b\tau} = -\beta_R^{b\tau}$	1.74	1.73	2.32	2.31	2.90	2.88	3.48	3.46
S_1	$y^{b\tau}$	1.65	1.72	2.36	2.45	2.95	3.06	3.55	3.68
Z'	$\zeta_{qL}^{bb} = \zeta_{qL}^{\tau\tau}$	3.04	3.04	4.05	4.05	5.07	5.06	6.08	6.07
	$\zeta_{dR}^{bb} = \zeta_{eR}^{\tau\tau}$	3.44	3.49	4.59	4.66	5.74	5.82	6.89	6.98

²⁴The region below 1.5 TeV has generally been excluded for leptoquarks in pair-production searches [375].

²⁵Including the $t\bar{t}$ and Z +jets floating normalisation.

²⁶The uncertainty contributions are quoted here as fractions of the total, but the underlying contributions are added in quadrature (as independent uncertainties) to obtain the total.

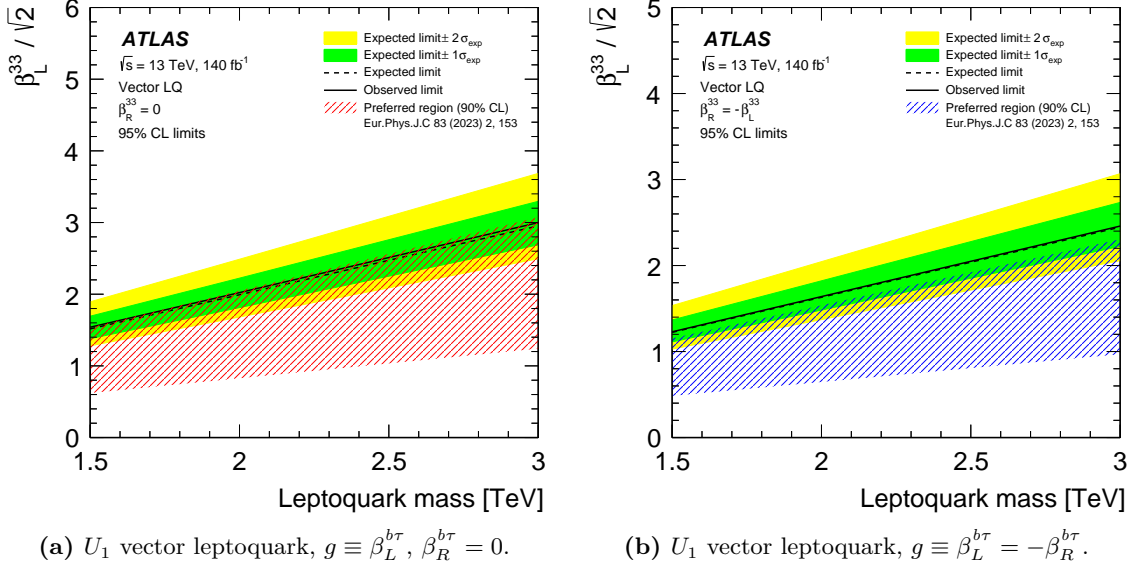


Figure 9.15: Limits on coupling constants for U_1 vector leptoquark models coupling b quarks and τ leptons, as a function of the leptoquark mass. The expected (dashed) and observed (solid) limits are shown alongside the $\pm 68\%$ (green band) and $\pm 95\%$ (yellow band) confidence intervals on the expected limit. The preferred regions (hatched) for b -anomalies discussed in Ref. [291] are shown.

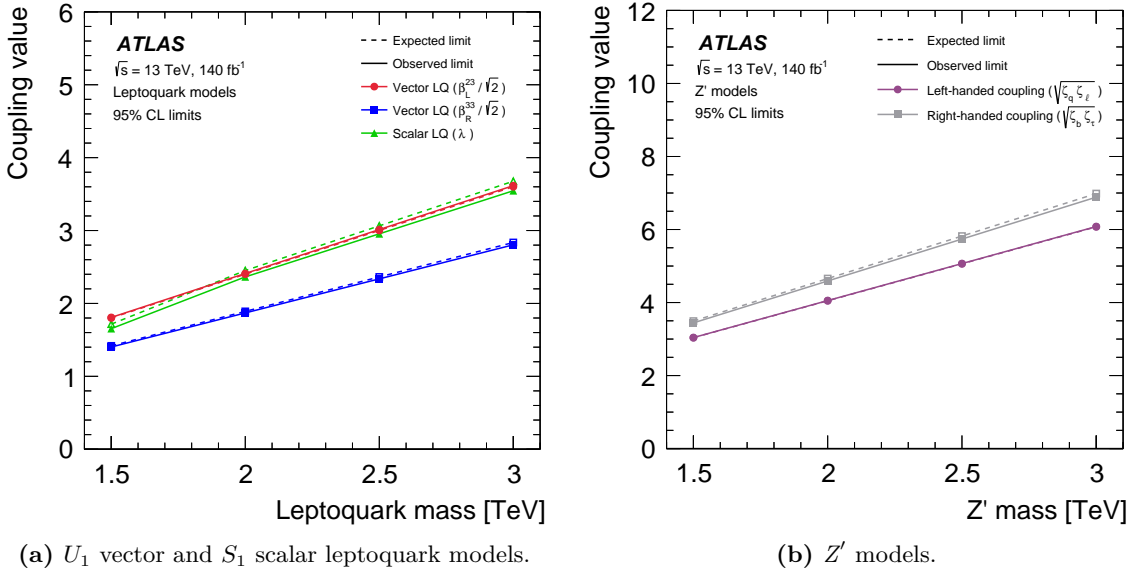


Figure 9.16: Coupling limits as a function of mass for the U_1 vector leptoquark models coupling left-handed s quarks and τ leptons (red), the U_1 model coupling right-handed b quarks and τ leptons (blue), the S_1 scalar leptoquark coupling b quarks and τ leptons (green), the Z' model coupling left-handed third-generation fermion doublets (grey), and the Z' model coupling right-handed third-generation fermion singlets (purple). Expected (dashed) and observed (solid) limits are shown for each.

Constraints on the other leptoquark and Z' couplings are shown in Figure 9.16, and improve on previous results from ATLAS [278] and CMS [277]. A local excess of 2.8σ was observed in the CMS result for a benchmark leptoquark mass of 2.5 TeV and a coupling strength of $\beta_L^{b\tau} = 2.5$. These parameter values have been excluded in this analysis. It is also of note that previous results did not include a full treatment of the interference term for non-resonant production.

9.10 SMEFT interpretation and limits on Wilson coefficients

As has been discussed in Sections 1.3 and 8.2, effective field theory formulations can be used to probe for the effects of potential BSM particles or behaviours at higher mass scales than the SM. Wilson coefficients can be measured and compared to SM predictions to search for or exclude new physics contributing to the diagrams corresponding to the studied operators.

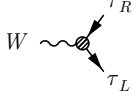
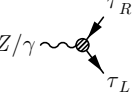
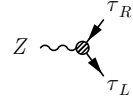
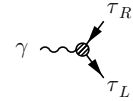
Fifteen dimension-6 Standard Model Effective Field Theory (SMEFT) operators relevant to the $\tau^+\tau^-$ final state were chosen from the `top` formulation, the coefficients of which have been studied previously by ATLAS in global EFT fits [376]. The chosen operators are tabulated in Tables 9.6 and 9.7. As discussed in Section 1.3, the `top` formulation reduces the number of operators by assuming equal couplings to the first two quark generations. Given the $\tau^+\tau^-$ final state, only couplings to τ leptons were considered relevant, and since τ - b couplings can be mapped to the scalar and vector leptoquark and Z' boson results shown in the last section, these were omitted.

The operator classes studied here include four-fermion interactions, Higgs-lepton current operators, and leptonic dipole operators. The four-fermion operators ($\psi\psi\psi\psi$) include the lepton-quark currents, which can be categorised by their current chirality combinations as $(\bar{L}L)(\bar{L}L)$, $(\bar{R}R)(\bar{R}R)$, and $(\bar{L}L)(\bar{R}R)$. The four-lepton vertex corresponding to $\mu \rightarrow e$ decays was also considered, since new physics contributions to muon decay have strong implications for the Fermi coupling G_F , and hence also impact $Z \rightarrow \tau^+\tau^-$ vertices. As a result, modifications to $c_{\ell\ell}$ can be constrained in final states with τ leptons. The Higgs-lepton current operators ($\psi^2 H^2 D$) separate into left-handed isospin singlet or triplet operators, and a right-handed singlet operator. The leptonic dipole operators

Table 9.6: Dimension-6 four-fermion and Higgs-lepton current SMEFT operators $\mathcal{Q}_n^{(6)}$ studied in this analysis. The operator and coefficient notation follows the conventions of the SMEFTSim 3.0 *top* model [376], with e replaced by τ to explicitly denote the lepton flavour. Only the first two quark generations are considered, with u_R representing a right-handed up or charm quark, and d_R representing a right-handed down or strange quark. Left-handed lepton doublets are labelled q_L , representing $(u_L \ d_L)$ or $(c_L \ s_L)$. Only operators coupling to τ leptons were considered, with τ_R representing a right-handed τ lepton, and ℓ_L representing the τ doublet state $(\tau_L \ \nu_\tau)$.

		Coefficient	Operator	Relevant vertex
$\psi\psi\psi\psi$	$(\bar{L}L)(\bar{L}L)$	$c_{\ell q}^{(3)}$	$(\bar{\ell}_L \sigma^i \gamma_\mu \ell_L)(\bar{q}_L \sigma^i \gamma^\mu q_L)$	
		$c_{\ell q}^{(1)}$	$(\bar{\ell}_L \gamma_\mu \ell_L)(\bar{q}_L \gamma^\mu q_L)$	
	$(\bar{R}R)(\bar{R}R)$	$c_{\tau u}$	$(\bar{\tau}_R \gamma_\mu \tau_R)(\bar{u}_R \gamma^\mu u_R)$	
		$c_{\tau d}$	$(\bar{\tau}_R \gamma_\mu \tau_R)(\bar{d}_R \gamma^\mu d_R)$	
	$(\bar{L}L)(\bar{R}R)$	$c_{\ell u}$	$(\bar{\ell}_L \gamma_\mu \ell_L)(\bar{u}_R \gamma^\mu u_R)$	
		$c_{\ell d}$	$(\bar{\ell}_L \gamma_\mu \ell_L)(\bar{d}_R \gamma^\mu d_R)$	
		$c_{q\tau}$	$(\bar{\tau}_R \gamma_\mu \tau_R)(\bar{q}_L \gamma^\mu q_L)$	
$\psi^2 H^2 D$	$(\bar{L}L)(\bar{L}L)$	$c_{\ell\ell}$	$(\bar{\ell}_L \gamma_\mu \ell_L)(\bar{\ell}_L \gamma^\mu \ell_L)$	
		$c_{H\ell}^{(3)}$	$(H^\dagger i \overleftrightarrow{D}_\mu^i H)(\bar{\ell}_L \sigma^i \gamma^\mu \ell_L)$	
		$c_{H\ell}^{(1)}$	$(H^\dagger i \overleftrightarrow{D}_\mu H)(\bar{\ell}_L \gamma^\mu \ell_L)$	
		$c_{H\tau}$	$(H^\dagger i \overleftrightarrow{D}_\mu H)(\bar{\tau}_R \gamma^\mu \tau_R)$	

Table 9.7: Leptonic dipole operators in the dimension-6 SMEFT affecting the τ anomalous electric and magnetic moments at the probed energy scales. The same nomenclature is used as in Table 9.6, and a discussion of $c_{\tau Z}$ and $c_{\tau\gamma}$ is given in Section 8.2.4. The cosine and sine of the weak mixing angle are denoted c_W and s_W , respectively, and only the real parts of the coupling coefficients are considered.

	Coefficient	Operator	Relevant vertex
$\psi^2 XH$	$c_{\tau W}$	$(\bar{\ell}_L \sigma^{\mu\nu} \tau_R) \sigma^I H W_{\mu\nu}^I + \text{h.c.}$	
	$c_{\tau B}$	$(\bar{\ell}_L \sigma^{\mu\nu} \tau_R) H B_{\mu\nu} + \text{h.c.}$	
	$c_{\tau Z}$	$(\bar{\ell}_L \sigma^{\mu\nu} \tau_R) \sigma^i H (c_W \sigma^i W_{\mu\nu}^i + s_W B_{\mu\nu}) + \text{h.c.}$	
	$c_{\tau\gamma}$	$(\bar{\ell}_L \sigma^{\mu\nu} \tau_R) \sigma^i H (-s_W \sigma^i W_{\mu\nu}^i + c_W B_{\mu\nu}) + \text{h.c.}$	

$(\psi^2 XH)$ produce modifications to the SM three-point couplings of vector bosons to fermion currents after symmetry breaking. Table 9.7 gives these couplings in both pre- and post-symmetry-breaking forms. As described in Section 8.2.4, these operators can be related to the τ anomalous magnetic and electric moments.

9.10.1 Parametrisation of the cross-section in terms of couplings

Just as the leptoquark and Z' operators generate both interference and pure-BSM terms when calculating cross-sections, so do EFT operators. The lowest order of additional contributions to the cross-section is associated with a suppression scale Λ^{-2} in $\sigma \propto |\mathcal{M}|^2$, corresponding to interference terms between dimension-6 operators and the SM. Pure-EFT terms appear only at Λ^{-4} , and at this order there are also pairs of dimension-6 operators, and the interference between dimension-8 operators and the SM. These effects can be studied more rigorously by augmenting the SM matrix element \mathcal{M}_{SM} with all allowable EFT couplings up to $\mathcal{O}(\lambda^{-4})$, yielding

$$\mathcal{M} = \mathcal{M}_{\text{SM}} + \sum_i \frac{c_i^{(6)}}{\Lambda^2} \cdot \mathcal{M}_i^{(6)} + \sum_{i,j} \frac{c_i^{(6)} c_j^{(6)}}{\Lambda^4} \cdot \mathcal{M}_{ij}^{(6)} + \sum_i \frac{c_i^{(8)}}{\Lambda^4} \cdot \mathcal{M}_i^{(8)} + \dots, \quad (9.22)$$

where $\mathcal{M}_k^{(d)} = \langle f | \mathcal{Q}_k^{(d)} | i \rangle$ and $\mathcal{M}_{kl}^{(d)} = \langle f | \mathcal{Q}_k^{(d)} \mathcal{Q}_l^{(d)} | i \rangle$ for an initial state $|i\rangle$ and a final state $\langle f|$. Computing $|\mathcal{M}|^2$ yields

$$\begin{aligned}
 |\mathcal{M}| = & \underbrace{|\mathcal{M}_{\text{SM}}|^2}_{\sigma_{\text{SM}}} + \underbrace{\frac{2}{\Lambda^2} \sum_i \text{Re} \left(c_i^{(6)} \mathcal{M}_{\text{SM}}^* \mathcal{M}_i^{(6)} \right)}_{\sigma_{\text{Int}}} \\
 & + \underbrace{\frac{1}{\Lambda^4} \left| \sum_i c_i^{(6)} \mathcal{M}_i^{(6)} \right|^2 + \frac{2}{\Lambda^4} \sum_{i,j} \text{Re} \left(c_i^{(6)} c_j^{(6)} \mathcal{M}_{\text{SM}}^* \mathcal{M}_{ij}^{(6)} \right) + \frac{2}{\Lambda^4} \sum_i \text{Re} \left(c_i^{(8)} \mathcal{M}_{\text{SM}}^* \mathcal{M}_i^{(8)} \right)}_{\sigma_{\text{BSM}}}
 \end{aligned} \tag{9.23}$$

The first two terms relate to SM and interference term contributions to the total cross-section, and rely only on the effects of dimension-6 operators, scaling linearly with the Wilson coefficients $c_i^{(6)}$. The $\mathcal{O}(\Lambda^{-4})$ term is split into three contributions, the first two of which rely on combinations of dimension-6 operators, and the latter on dimension-8 operators. For the expected suppression scales of $\Lambda = 1 \text{ TeV}$ or greater, the interference term is typically expected to be dominant, and depends linearly on the Wilson coefficients under study. However, for specific operators and coupling scenarios, the pure-BSM terms can become dominant despite suppression at $\mathcal{O}(\Lambda^{-4})$, especially if the SM cross-section \mathcal{M}_{SM} is highly suppressed²⁷.

A complete calculation of the pure-BSM term in Equation (9.23) would rely on the inclusion of both dimension-6 and dimension-8 operators²⁸ during calculation of the matrix element in the signal sample MC generation. The latter are currently beyond the reach of the SMEFTSIM 3.0 model used, but an effort has been made nonetheless to estimate the relative impact of the pure-BSM contribution from the term proportional to $|\mathcal{M}_i^{(6)}|^2$, following the formulation suggested in the SMEFTSIM guide [377], and successfully used in previous studies such as Refs. [378, 379]. In these cases, confidence intervals are computed for the interference term exhibiting *linear* scaling with $c_i^{(6)}$, and separately for a *linear + quadratic* combination of this with the first two terms of the pure-BSM cross-section, which both scale quadratically with $c_i^{(6)}$ in the case of only a

²⁷Some examples include operators inducing LFUV couplings, or SM-allowed operators studied in particularly improbable kinematic regions.

²⁸Including cross-interference terms between different dimension-6 operators $\mathcal{Q}_i^{(6)} \mathcal{Q}_j^{(6)}$, which relies on the simulation/calculation of several non-zero operators at once. In addition, these terms are only renormalisable by inclusion of dimension-8 interactions at loop-level.

single active coupling. The difference between these scenarios is taken as a qualitative measure of the relative importance of the $\mathcal{O}(\Lambda^{-4})$ terms.

For the purposes of this study, only one non-zero Wilson coefficient was considered at a time. The final parametrisation is given as

$$\sigma_{\tau^+\tau^-}^{\text{tot},i} = \sigma_{\text{SM}} \left(1 + c_i^{(6)} \cdot \frac{\sigma_{\text{Int}}(c_i^{(6)} = 1/\text{TeV}^2)}{\sigma_{\text{SM}}} + \left(c_i^{(6)} \right)^2 \cdot \frac{\sigma_{\text{BSM}'}(c_i^{(6)} = 1/\text{TeV}^2)}{\sigma_{\text{SM}}} \right). \quad (9.24)$$

9.10.2 Signal samples

Tree-level signal samples corresponding to the interference term $\sigma_{\text{Int}}(c_i^{(6)} = 1/\text{TeV}^2)$ and partial pure-BSM term $\sigma_{\text{BSM}'}(c_i^{(6)} = 1/\text{TeV}^2)$ were simulated using MADGRAPH, with the EFT matrix elements computed at leading order using the SMEFTSIM 3.0 `topU31` model via the setup described in Section 9.2.2. Both terms were generated for all fifteen Wilson coefficients described in Tables 9.6 and 9.7, with the relevant operator in each case set to $1/\text{TeV}^2$.

Early studies of the signal samples for these operators showed similarly large statistical uncertainties for high- $m_{\tau\tau}$ events in the ≥ 2 b -jet category for the leptoquark and Z' samples, but did not suffer from localised fluctuations of the size seen for these samples. A comparison of the expected sensitivity to Wilson coefficients to that achieved using the forwarded folding method developed for the leptoquark and Z' boson samples in Section 9.9.1 showed substantial improvements in sensitivity for the EFT samples, and hence this method was also applied to these samples.

9.10.3 Results

The same statistical framework and profile likelihood fit methodology described in Section 9.8 was applied to extract 68% confidence intervals for the Wilson coefficients under study, which are tabulated in Table 9.8. A summary of the constraints and best-fit values is given in Figure 9.17. The four-fermion operators show little sensitivity to the quadratic term, with the interference term dominating the sensitivity. This was also true of the Higgs-lepton current operators, with the exception of $c_{H\tau}$ which shows some suppression of the interference term by σ_{SM} . The interference terms of the leptonic

dipole operators are also heavily suppressed by σ_{SM} , and the partial pure-BSM term dominates the sensitivity to these couplings.

Figure 9.18 shows the excluded scale at 95% confidence level for coupling values of 0.01, 1, and $4\pi^2$, corresponding to weakly coupled²⁹, strongly coupled, or maximally strongly coupled³⁰ scenarios. Only the *linear* scenarios are included, as the *linear+quadratic* scenarios scale non-trivially with the coupling strength. The four-fermion coupling scenarios are generally constrained to one or several TeV given unit couplings, whilst constraints for the Higgs-lepton current operators and leptonic dipole operators are significantly weaker, at the order of several hundred GeV for unit couplings.

Table 9.8: Best-fit values and 68% confidence intervals for the Wilson coefficients studied, parameterised both in terms of the interference term only (*linear*), and the interference plus partial pure-BSM contributions (*linear + quadratic*). The values are given for a suppression scale of $\Lambda = 1$ TeV.

Coefficient	<i>linear</i>	<i>linear + quadratic</i>
$c_{\ell q}^{(3)}$	$-0.002^{+0.018}_{-0.016}$	$-0.002^{+0.018}_{-0.016}$
$c_{\ell q}^{(1)}$	$0.01^{+0.06}_{-0.07}$	$0.01^{+0.18}_{-0.05}$
$c_{\tau u}$	$0.01^{+0.04}_{-0.04}$	$0.01^{+0.04}_{-0.04}$
$c_{\tau d}$	$0.07^{+0.32}_{-0.25}$	$0.07^{+0.32}_{-0.25}$
$c_{q\tau}$	$0.01^{+0.07}_{-0.08}$	$0.01^{+0.07}_{-0.16}$
$c_{\ell u}$	$-0.05^{+0.15}_{-0.20}$	$-0.05^{+0.15}_{-0.20}$
$c_{\ell d}$	$-0.02^{+0.21}_{-0.21}$	$-0.02^{+0.21}_{-0.21}$
$c_{\ell\ell}$	$-1.1^{+1.1}_{-1.1}$	$-1.0^{+1.1}_{-1.1}$
$c_{H\ell}^{(3)}$	$0.7^{+0.8}_{-0.8}$	$0.7^{+0.8}_{-0.8}$
$c_{H\ell}^{(1)}$	$-1.7^{+1.6}_{-1.6}$	$-1.5^{+1.6}_{-1.6}$
$c_{H\tau}$	6^{+6}_{-6}	$0.7^{+2.1}_{-1.9}$
$c_{\tau W}$	15^{+15}_{-15}	$0.0^{+0.4}_{-0.5}$
$c_{\tau B}$	-30^{+31}_{-30}	$-0.2^{+1.8}_{-1.5}$
$c_{\tau Z}$	20^{+19}_{-20}	$0.0^{+0.4}_{-0.5}$
$c_{\tau\gamma}$	-17^{+17}_{-16}	$-0.1^{+1.0}_{-0.8}$

²⁹For example, couplings arising from loop-suppressed new physics scenarios.

³⁰Using Naïve Dimensional Analysis, a maximal natural size can be estimated at which the perturbative nature of the coupling breaks down [380]. Larger couplings pose problems for potential UV-completions of the theory, and additionally invalidate the dominance of leading-order couplings assumed here.




Figure 9.17: Summary of the best-fit values and confidence intervals obtained for the Wilson coefficients studied, in terms of the first-order *linear* or first-order plus partial second-order *linear+quadratic* contributions to the cross-section. The upper panel shows the symmetrised 1σ uncertainty at a scale of $\Lambda = 1$ TeV (left y -axis), with the probed scale given a unit coefficient for the linear results also shown (right y -axis). In the lower panel, the best-fit coefficient and 68% and 95% confidence intervals are given relative to the symmetrised 1σ uncertainty.

Implications for the τ lepton anomalous magnetic moment

As discussed in Section 8.2.4, the operator $\mathcal{Q}_{\tau\gamma}$ affects the anomalous magnetic moment of the τ lepton. Since the definition of the anomalous magnetic moment a_τ is only valid for on-shell photons with $q^2 = 0$, the matching given in Equation (8.14) is only valid at low energies, but does allow measurements of a_τ to be interpreted as constraints on the SMEFT operator $\mathcal{Q}_{\tau\gamma}$, under the assumption of no new physics below the assumed mass scale for the comparison, Λ .




Figure 9.18: The new physics scale Λ excluded at 95% confidence level by the *linear* confidence intervals derived from the interference terms, given coupling strengths of 0.01 (lower bars), 1 (middle bars), or $4\pi^2$ (upper bars).




Figure 9.19: Comparison of best-fit values and 95% confidence intervals of $c_{\tau\gamma}$ for this measurement and LHC measurements probing the anomalous magnetic moment of the τ lepton via $\gamma\gamma \rightarrow \tau\tau$ processes in heavy-ion or pp collisions [260, 261, 384].

A more precise determination of the limit on $c_{\tau\gamma}$ was obtained from a χ^2 fit performed at truth-level to the unfolded cross-section using CONTUR 3.1³¹ [381–383], yielding $c_{\tau\gamma} = -0.40^{+0.63}_{-0.24}$, and a 95% confidence interval of $[-0.88, 0.86]$. This value is compared to previous LHC constraints on a_τ in Figure 9.19, with the caveat that since the energy scale this analysis is performed at implies an off-shell photon exchange, the interpretation of this comparison should be restricted to the EFT parameter $c_{\tau\gamma}$ at higher energy scales.

9.11 Conclusions and outlook

Measuring the $m_{\ell\ell}^{\text{vis}}$ distribution in di-lepton Drell-Yan production channels provides an important check of the predicted SM continuum, and can act as a probe for interference effects from off-shell BSM particles well above the mass scales currently accessible for resonant production. This analysis is the first unfolded differential cross-section measurement of $\tau_{\text{had}}\tau_{\text{had}}$ production at high $m_{\tau\tau}$ that has been performed at the LHC, and complements a history of measurements at lower $m_{\tau\tau}$ and in the light-lepton channels. The measurement is consistent with the SM, showing agreement with several SM predictions from modern MC-based modelling schemes.

³¹Similar fits were performed for all Wilson coefficients studied. Due to current limitations in CONTUR, only either the interference term or the partial pure-BSM term could be treated. These fits were consistent with the detector-level fits, and in some cases showed greater constraining power for the couplings that do not enhance b -jet production due to a finer binning in $m_{\tau\tau}^{\text{vis}}$.

In addition, a detector-level search for non-resonant leptoquark and Z' boson interactions has been performed using a profile likelihood fit method. This analysis introduced the first full treatment of both interference and pure-BSM contributions for these models, with the former observed to significantly impact the true sensitivity when included. Several leptoquark models with preferential couplings to τ leptons and b quarks have been studied, with a particular focus on those preferred by the lepton flavour universality violating b anomalies discussed in Chapter 8. To increase sensitivity to event-topology dependent effects on $m_{\tau\tau}$, fits were performed in three categories split by b -jet multiplicity. No significant excesses were observed, and 95% confidence limits were set on leptoquark and Z' boson coupling parameters as a function of their respective masses. The results have been compared to the regions preferred by the b anomalies, and begin to exclude parts of these regions.

The same method has also been applied to derive constraints on relevant SMEFT operators involving $\tau^+\tau^-$ pairs. Interactions involving the first two quark generations have been constrained, with a symmetry applied between these two generations. In addition, interactions involving coupling to a weak boson or the combination of a weak boson and the Higgs boson have also been considered. In all cases, extracted 68% confidence intervals agree with SM predictions. The extracted limits on the leptonic dipole operators are interpreted as constraints on the τ magnetic dipole moment for off-shell photons ($q \gg 0$), and are competitive with extrapolations of a_τ measurements to the $q^2 \gg 0$ case.

The detector-level interpretations are fundamentally limited by the statistical power of the 13 TeV, 140 fb^{-1} dataset used. In general, statistical sources comprise 80% or more of the total uncertainty, indicating that searches for BSM behaviour and limits on EFT couplings in this channel are positioned to profit greatly from the HL-LHC program. The dominant systematic uncertainties related to the reconstruction of τ_{had} candidates, and recent improvements in algorithms and techniques promise a bright future for final states with τ leptons in Run 3 and beyond.

Conclusions and Outlook

For the past decade, the LHC experiments explored phenomena at the highest energies and shortest length scales that have ever been accessible at particle colliders. The ATLAS experiment has collected valuable datasets both verifying the behaviour of the SM, and constraining the possibilities for physics beyond. However, despite the substantial successes of the LHC program, a vast potential still lies ahead with the high-luminosity upgrade. Alongside an order-of-magnitude increase in statistical power for all experiments, the upgrade represents an opportunity to renew and optimise each of the detectors. For ATLAS, this will involve the replacement of the entire tracking subsystem with the largest silicon-based detector ever planned, the ITk. The work presented in this thesis gives a glimpse into the lifecycle of a detector, following various aspects from design and construction through to data analysis and physics reach.

PART I has given a comprehensive overview of the current ATLAS detector, and the ambitious ITk upgrade programme, with a particular focus on the construction of larger detector structures populated with pixel sensors — *outer endcap half-rings*. In this context, a variety of challenges related to mechanical stability, thermal performance, detector safety and data readout have been encountered and addressed. An automated method for detector construction utilising a robotic gantry provides repeatable and reliable results, and consistent quality control testing ensures these components will survive the harsh environment of the HL-LHC over their twelve-year operating lifespan.

The process of converting a detector design into a commissioned and operating machine involves many practical challenges and pragmatic compromises, and the final product often deviates to some extent from the initial plan. The discussions in Chapters 2 and 6 have shown that the physics output of a tracking detector depends strongly on the material composition and structural layout of the detector, and on accurate models of these which can be used in simulations for physics analyses. The material measurement technique

in PART II introduces a method to build and validate highly granular material models during the design phases of detectors through testbeam measurements, and applies this technique to detector components for the ITk upgrade. The quad module measurement presented achieves a 9% uncertainty limited by the momentum band of the positron beam used, but exceeds the original TDR design specification for this component by about 20%.

As has been discussed in the context of the ITk, particle detectors work by converting new high-energy phenomena to a well-understood low-energy regime. In order to relate experimental datasets back to the initial predictions of a particular model, several methods of extrapolation in both directions are required, including MC simulations to estimate the detector response, and statistical methods to unfold the original model from these detector effects. The measurement in PART III combines new unfolded cross-sections of the Drell-Yan spectrum in a final state with two τ leptons with a search for new physics motivated by recent hints of lepton flavour universality violation. The analysis has built robust estimates for the SM and new physics models from a combination of MC simulated datasets and data-driven methodologies, and has applied the first search for non-resonantly produced leptoquarks preferred by the b -anomalies which treats interference effects with the SM cross-section as well as pure-BSM contributions. The results once again reaffirm the strong standing of the SM, with predictions well matched by the unfolded cross-section, and competitive new constraints set on several leptoquark couplings. In addition, a range of SMEFT operators have been studied, and these constraints will go on to become part of a global effort to isolate the form of potential new physics in this model-independent framework.

The final conclusions of the $pp \rightarrow \tau^+\tau^-$ measurement and search harken back to the beginnings of this thesis — the statistical power of current datasets forms the fundamental limitation for many current endeavours toward better understanding rare processes in the SM or finding new physics at ATLAS and the other LHC experiments. The HL-LHC upgrade represents a natural next step toward a more complete understanding of particle physics, and the ITk upgrade will play a critical role in enabling this programme, thus taking a large step forward toward models yet to be tested and surprises yet to be discovered.

Appendices



Literature and calculated material X_0 values

This appendix gives further details of the literature or calculated radiation length values used for the ITkPix v1.1 quad module estimate introduced in Section 7.4. Table A.1 gives the values used for the materials considered in the estimate, including both pure materials and mixtures.

Table A.2 specifies values for the estimated radiation length and layer thickness for individual surface mount device (SMD) components, including resistors, capacitors, NTCs, and connectors. The connectors are broken down into housing and contacts, and

Table A.1: Literature and calculated radiation length values used for layers within the ITkPix v1.1 quad module structure. In the layer designations, *flex* refers to both the module flex, and the data and power flex cables. Phosphor bronze has been estimated assuming a typical value of 5% Sn and 95% Cu, with the $< 0.5\%$ P neglected. Electroless nickel/immersion gold (ENIG) has been estimated as 92% Ni, 7% P, and a conservative upper limit of 1% Ag [385]. Polylactide has been estimated from its chemical composition $(C_3H_4O_2)_n$. The accuracy of these estimates is limited by manufacturer variations in materials and processes.

Material	X_0 [cm]	Source	Layers
Al	8.897	PDG [218]	module holder
Ag	0.8543	PDG [218]	<i>ENIG</i>
Cu	1.436	PDG [218]	flex internal layers
Ni	1.424	PDG [218]	sensor metallisation
P	9.693	PDG [218]	<i>Phosphor bronze, ENIG</i>
Si	9.370	PDG [218]	sensor, readout chip
Sn	1.206	PDG [218]	bump bonds
Acrylic glue	42.6	PDG [218]	internal flex glues, flex to readout chip adhesive
Polyimide	28.57	PDG [218]	flex core, coverlay
Phosphor Bronze	1.41	mixture [216]	connector contacts
ENIG	1.54	mixture [216]	flex finishing and pads
Polylactide	39.11	mixture [216]	PLA module clamp

Table A.2: Radiation lengths and thicknesses estimated for different classes of SMD components used in the ATLAS quad module flex v2.4 bill of materials. The radiation lengths for capacitors, resistors, and NTCs are estimated from typical relative fractions of metallic and ceramic composition, but are likely to be quite inaccurate, and to vary significantly between different manufacturers and production processes. Footprint maps of the 39-pin and 16-pin connectors were used to define regions corresponding to acrylic housing and metallic contacts with a uniform thickness given in the table. The contacts of the 16-pin connector were split into vertical rods and horizontal feet. Due to a lack of literature values, the radiation length of the acrylic housing was estimated to be the same as for polyimide since both are composed of similar long polymers.

SMD component	Size/part	X_0 [cm]	Thickness [mm]
Capacitor	0201	2.857	0.3
	0402	2.857	0.5
	1206 (HV filter)	2.98	1.6
Resistor	0201	4.003	0.23
	0402	4.003	0.3 to 0.35
NTC	0201	3.4775	0.3
	0402	2.857	0.6
39-pin data connector	housing	29	0.1
	contacts	1.41	0.35
	sides	29	0.5
16-pin power connector	housing	29	0.03
	rods	1.41	298
	feet	1.41	3

2-dimensional footprints have been derived that allow the different contributions to be combined whilst retaining the resolution achieved by the measurement.

Table A.3 shows the mass and size measurements of the 3D-printed module clamp, and the calculated fractional radiation length used for the subtraction described in Section 7.3.6.

Table A.3: Module clamp subcomponent measurements.

Component	mass [g]	density [g/cm ³]	x/X_0
Top left HV capacitor leg	0.045	1.096	2.24%
Bottom left rectangular leg	0.069	1.046	2.63%
Bottom right square leg	0.097	1.052	2.64%
Top left HV capacitor arm	0.037	1.038	0.47%
Bottom left rectangular arm	0.038	0.958	0.45%
Bottom right square arm*	0.021	0.963	0.43%

B

Further results from telescope simulations

This appendix complements the results presented in Section 7.5, showing the telescope response for the SSCB model in Figure B.1, and for the case where no data from the DUT is used in the analysis of the simulated data in Figures B.2 and B.3. Response curve fits excluding the DUT data are shown in Figure B.4. The air and telescope scattering contributions in simulations excluding the DUT data are shown both combined and separately in Figure B.5 and Figure B.6, respectively.




Figure B.1: Simulation of σ_θ for various assumed true x/X_0 values of a homogeneous silicon DUT, estimated using an SSCB fit to θ_x and θ_y . The telescope response is shown before (blue) and after the air + telescope subtraction (orange). The latter is compared to an expectation from Highland (black dashed line) and Frühwirth-Regler (purple dotted line). The reference measurement with a 50 μm thick MALTA plane as the DUT is shown as data comparison point for each response curve (pre- and post-subtraction). The Highland and Frühwirth-Regler lines overlap on the left plot of the figure, and their difference is only visible in the ratio plot on the right of the figure. A third-order polynomial fit to the response curve is shown (dark blue dashed line).




Figure B.2: An equivalent simulation of telescope response to Figure 7.17, but excluding the DUT hit from the tracking.




Figure B.3: An equivalent simulation of telescope response to Figure B.1, but excluding the DUT hit from the tracking.

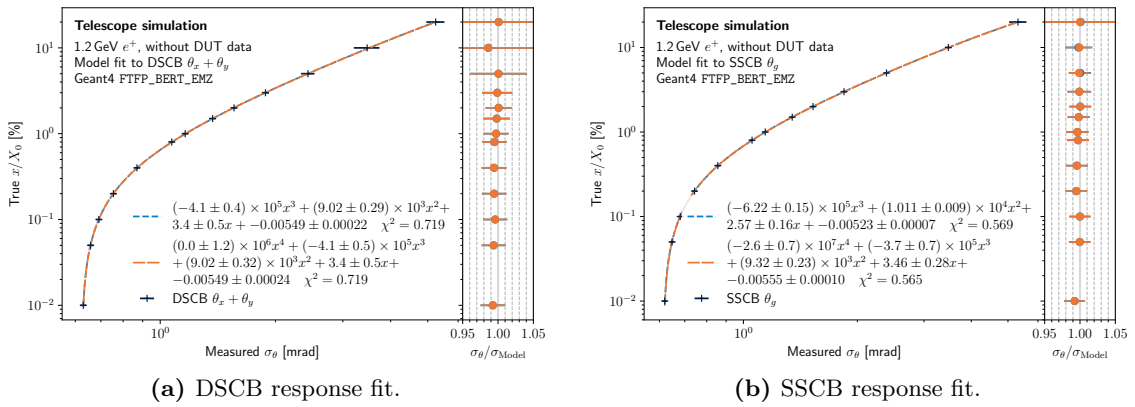


Figure B.4: Polynomial fits to telescope response curve for the analysis excluding the DUT hit, presented equivalently to Figure 7.18. The curves were used for the GEANT4 model applied to extracted widths derived without DUT data in Figure 7.24, to account for the different magnitude of the air and telescope scattering in these cases.




Figure B.5: Air and telescope contribution to total measured x/X_0 as a function of simulated DUT x/X_0 , for the case of no DUT data used in the track reconstruction. The shaded bands show the standard deviation of the extracted $(x/X_0)_{\text{air+tele}}$ across 20×20 subregions of size $0.5 \text{ mm} \times 0.5 \text{ mm}$, each containing a number of tracks equivalent to the average subregion statistics achieved for the measurement.

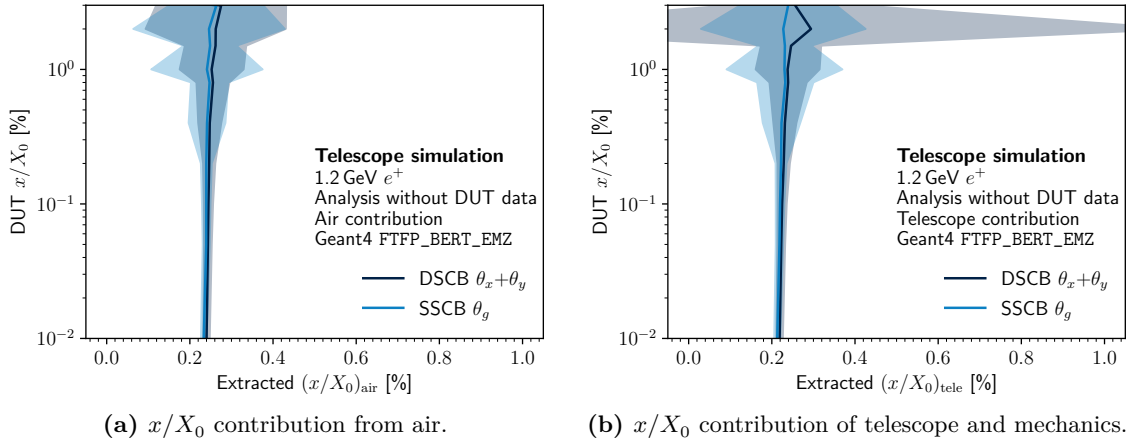


Figure B.6: Separate air and telescope contributions to measured x/X_0 as a function of simulated DUT x/X_0 , for the case of no DUT data used in the track reconstruction. The shaded bands show the standard deviation of the extracted value across 20×20 subregions of size $0.5 \text{ mm} \times 0.5 \text{ mm}$, each containing a number of tracks equivalent to the average subregion statistics achieved for the measurement.

References

- [1] R. Munroe, *xkcd: Elementary Physics Paths*, reproduced under CC-BY-NC 2.5, (2024) <https://xkcd.com/2933/> (visited on 05/06/2025).
- [2] S. F. Koch et al., ‘Measuring the ATLAS ITk pixel detector material via multiple scattering of positrons at the CERN PS’, *EPJ C* **85**, 381 (2025).
- [3] ATLAS Collaboration, ‘A measurement of the high-mass $\tau\bar{\tau}$ production cross-section at $\sqrt{s} = 13$ TeV with the ATLAS detector and constraints on new particles and couplings’, Submitted to JHEP Mar. 2025, arXiv:2503.19836 (2025).
- [4] ATLAS ITk Collaboration, *Radiation length measurement of an ATLAS ITkPix v1.1 module using 1.2 GeV positrons at the CERN PS*, ATLAS Public plots: ITK-2023-002, 2023.
- [5] D. Bacher et al., *ICICLE - Instrument Control Interface and Commands Library*, CERN GitLab code repository, <https://gitlab.cern.ch/icicle/icicle>, 2023.
- [6] L. Perivolaropoulos and F. Skara, ‘Challenges for Λ CDM: an update’, *New Astronomy Reviews* **95**, 101659 (2022).
- [7] M. K. Gaillard, P. D. Grannis and F. J. Sciulli, ‘The standard model of particle physics’, *Rev. Mod. Phys.* **71**, S96–S111 (1999).
- [8] L. d. Broglie, ‘XXXV. A tentative theory of light quanta’, *Phil. Mag.* **47**, 446–458 (1924).
- [9] A. H. Compton, ‘A quantum theory of the scattering of X-rays by light elements’, *Phys. Rev.* **21**, 483–502 (1923).
- [10] E. Hecht, *Optics* (Pearson, 2014).
- [11] W. N. Cottingham and D. A. Greenwood, *An Introduction to the Standard Model of Particle Physics*, 2nd ed. (Cambridge University Press, 2007).
- [12] D. Griffiths, *Introduction to elementary particles*, 2nd ed. (Wiley, 2008).
- [13] A. Rubbia, *Phenomenology of Particle Physics* (Cambridge University Press, 2022).
- [14] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory* (Addison-Wesley, Reading, USA, 1995).
- [15] M. D. Schwartz, *Quantum field theory and the standard model* (Cambridge University Press, 2013).
- [16] P. A. M. Dirac, ‘Quantum theory of emission and absorption of radiation’, *Proc. Roy. Soc. Lond. A* **114**, 243 (1927).
- [17] E. Noether, ‘Invariante Variationsprobleme’, *ger*, *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* **1918**, 235–257 (1918).
- [18] E. P. Wigner, ‘On Unitary Representations of the Inhomogeneous Lorentz Group’, *Annals Math.* **40**, edited by Y. S. Kim and W. W. Zachary, 149–204 (1939).
- [19] O. Klein, ‘Quantentheorie und fünfdimensionale relativitätstheorie’, *Zeitschrift für Physik* **37**, 895–906 (1926).
- [20] W. Gordon, ‘Der Comptoneffekt nach der Schrödingerschen Theorie’, *Zeitschrift für Physik* **40**, 117–133 (1926).

- [21] S. Coleman and J. Mandula, ‘All Possible Symmetries of the S Matrix’, *Phys. Rev.* **159**, 1251–1256 (1967).
- [22] W. Rarita and J. Schwinger, ‘On a theory of particles with half-integral spin’, *Phys. Rev.* **60**, 61–61 (1941).
- [23] H. Weyl, ‘Gravitation and the electron¹’, *Proc. of the National Academy of Sciences* **15**, 323–334 (1929).
- [24] J. D. Jackson and L. B. Okun, ‘Historical roots of gauge invariance’, *Rev. Mod. Phys.* **73**, 663–680 (2001).
- [25] J. D. Jackson, ‘From Lorenz to Coulomb and other explicit gauge transformations’, *American Journal of Physics* **70**, 917–928 (2002).
- [26] J. C. Romão and J. P. Silva, ‘A resource for signs and Feynman diagrams of the Standard Model’, *International Journal of Modern Physics A* **27**, 1230025 (2012).
- [27] N. Arkani-Hamed, T.-C. Huang and Y.-t. Huang, ‘Scattering amplitudes for all masses and spins’, *J. High Energy Phys.* **2021**, 70 (2021).
- [28] R. P. Feynman, ‘Space-time approach to non-relativistic quantum mechanics’, *Rev. Mod. Phys.* **20**, 367–387 (1948).
- [29] S. Mandelstam, ‘Determination of the pion-nucleon scattering amplitude from dispersion relations and unitarity. General theory’, *Phys. Rev.* **112**, 1344–1360 (1958).
- [30] L. M. Brown and A. S. Wightman, ‘Renormalization: from Lorentz to Landau (and beyond)’, *American Journal of Physics* **62**, 765–766 (1994).
- [31] F. Englert and R. Brout, ‘Broken symmetry and the mass of gauge vector mesons’, *Phys. Rev. Lett.* **13**, 321–323 (1964).
- [32] P. W. Higgs, ‘Broken Symmetries and the Masses of Gauge Bosons’, *Phys. Rev. Lett.* **13**, 508–509 (1964).
- [33] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, ‘Global Conservation Laws and Massless Particles’, *Phys. Rev. Lett.* **13**, 585–587 (1964).
- [34] N. Cabibbo, ‘Unitary Symmetry and Leptonic Decays’, *Phys. Rev. Lett.* **10**, 531–533 (1963).
- [35] H. Georgi and S. L. Glashow, ‘Unity of all elementary-particle forces’, *Phys. Rev. Lett.* **32**, 438–441 (1974).
- [36] H. E. Haber and L. S. Haskins, ‘Supersymmetric theory and models’, in *Anticipating the next discoveries in particle physics* (2018) Chap. 6, pp. 355–499.
- [37] C. N. Yang, ‘Fermi’s β -decay Theory’, *Asia Pacific Physics Newsletter* **01**, 27–30 (2012).
- [38] I. Brivio and M. Trott, ‘The standard model as an effective field theory’, *Physics Reports* **793**, 1–98 (2019).
- [39] I. Brivio, Y. Jiang and M. Trott, ‘The SMEFTsim package, theory and tools’, *JHEP* **12**, 070 (2017).
- [40] S. D. Drell and T.-M. Yan, ‘Massive lepton-pair production in hadron-hadron collisions at high energies’, *Phys. Rev. Lett.* **25**, 316–320 (1970); Erratum: *Phys. Rev. Lett.* **25**, 902 (1970).
- [41] CMS Collaboration, ‘Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV’, *JINST* **7**, P10002 (2012).
- [42] S. Navas et al. (Particle Data Group), ‘Review of particle physics’, *Phys. Rev. D* **110**, 030001 (2024).
- [43] H. Bethe, ‘Zur Theorie des Durchgangs schneller Korpuskularstrahlen durch Materie’, *Annalen der Physik* **397**, 325–400 (1930).

- [44] H. Bethe, ‘Bremsformel für Elektronen relativistischer Geschwindigkeit’, *Zeitschrift für Physik* **76**, 293–299 (1932).
- [45] H. Kolanoski and N. Wermes, *Particle Detectors: Fundamentals and Applications* (Oxford University Press, 2020).
- [46] L. D. Landau, ‘On the energy loss of fast particles by ionisation’, *J. Phys. (USSR)* **8**, edited by D. ter Haar, 201–209 (1944).
- [47] P. V. Vavilov, ‘Ionization losses of high-energy heavy particles’, *Sov. Phys. JETP* **5**, 749–751 (1957).
- [48] U. Fano, ‘Ionization yield of radiations. II. The fluctuations of the number of ions’, *Phys. Rev.* **72**, 26–29 (1947).
- [49] R. Bellwied et al., ‘The STAR silicon vertex tracker: a large area silicon drift detector’, *NIM A* **499**, 640–651 (2003).
- [50] F. Reidt, ‘Upgrade of the ALICE ITS detector’, *NIM A* **1032**, 166632 (2022).
- [51] R. Cardella et al., ‘MALTA: an asynchronous readout CMOS monolithic pixel detector for the ATLAS High-Luminosity upgrade’, *JINST* **14**, C06019–C06019 (2019).
- [52] W. Snoeys et al., ‘A process modification for CMOS monolithic active pixel sensors for enhanced depletion, timing performance and radiation tolerance’, *NIM A* **871**, 90–96 (2017).
- [53] T. Flick, ‘The phase II ATLAS Pixel upgrade: the Inner Tracker (ITk)’, *JINST* **12**, C01098 (2017).
- [54] M. Garcia-Sciveres (RD53), *The RD53A Integrated Circuit*, CERN-RD53-PUB-17-001 (2017).
- [55] V. A. Van Lint, ‘The physics of radiation damage in particle detectors’, *NIM A* **253**, 453–459 (1987).
- [56] C. Grupen and B. Shwartz, *Particle detectors*, 2nd ed. (Cambridge University Press, 2008).
- [57] S. Watts, ‘Overview of radiation damage in silicon detectors — models and defect engineering’, *NIM A* **386**, VERTEX1996, 149–155 (1997).
- [58] M. Weber, ‘Introduction to Silicon Detectors’, part of RAL (Rutherford Appleton Laboratories) Graduate Lecture Series, slides available at https://www.ppd.stfc.ac.uk/Pages/lecture02_weber.pdf, 2006.
- [59] J. Große-Knetter, ‘Vertex measurement at a hadron collider: the ATLAS pixel detector’, Habilitation (Uni Bonn, 2007), INIS:39088841.
- [60] L. Evans and P. Bryant, ‘LHC Machine’, *JINST* **3**, S08001 (2008).
- [61] E. Lopienska, ‘The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022’, image, CERN CDS:2800984, 2022.
- [62] ATLAS Collaboration, ‘ATLAS data quality operations and performance for 2015–2018 data-taking’, *JINST* **15**, P04003 (2020).
- [63] ATLAS Collaboration, ‘Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC’, *EPJ C* **77**, 332 (2017).
- [64] ATLAS Collaboration, *ATLAS luminosity public results*, available online at <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>, accessed: 2025-06-05, 2025.
- [65] ATLAS Collaboration, ‘Operation of the ATLAS trigger system in Run 2’, *JINST* **15**, P10004 (2020).
- [66] ATLAS Collaboration, ‘The ATLAS Experiment at the CERN Large Hadron Collider’, *JINST* **3**, S08003 (2008).

- [67] ATLAS Collaboration, ‘Experiment Briefing: Keeping the ATLAS Inner Detector in perfect alignment’, image, CERN CDS:2723878, 2020.
- [68] ATLAS Collaboration, *Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-018, 2015.
- [69] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, Volume 1*, ATLAS-TDR-4; CERN-LHCC-97-016 (1997).
- [70] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, Volume 2*, ATLAS-TDR-5, CERN-LHCC-97-017 (1997).
- [71] P. Fischer (ATLAS), ‘The ATLAS pixel front end chip FEI in 0.25 μ m technology’, in 8th Workshop on Electronics for LHC Experiments (2002).
- [72] ATLAS Collaboration, *ATLAS Insertable B-Layer: Technical Design Report*, ATLAS-TDR-19; CERN-LHCC-2010-013 (2010); Addendum: ATLAS-TDR-19-ADD-1; CERN-LHCC-2012-009 (2012).
- [73] B. Mindur, ‘ATLAS Transition Radiation Tracker (TRT): straw tubes for tracking and particle identification at the LHC’, NIM A **845**, VCI2016, 257–261 (2017).
- [74] R. E. Kalman, ‘A new approach to linear filtering and prediction problems’, Transactions of the ASME–Journal of Basic Engineering **82**, 35–45 (1960).
- [75] ATLAS Collaboration, *Early Inner Detector tracking performance in the 2015 data at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-051, 2015.
- [76] ATLAS Collaboration, *ATLAS Calorimeter Performance: Technical Design Report*, ATLAS-TDR-1; CERN-LHCC-96-040 (1996).
- [77] F. Sefkow and F. Simon, ‘Calorimeters’, in *Handbook of particle detection and imaging*, edited by I. Fleck et al. (Springer, Cham, 2020), pp. 1–33.
- [78] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter: Technical Design Report*, ATLAS-TDR-2; CERN-LHCC-96-041 (1996).
- [79] N. Ilic, ‘Performance of the ATLAS Liquid Argon calorimeter after three years of LHC operation and plans for a future upgrade’, JINST **9**, C03049 (2014).
- [80] J. P. Archambault et al., ‘Energy calibration of the ATLAS Liquid Argon Forward Calorimeter’, JINST **3**, P02002 (2008).
- [81] ATLAS Collaboration, *ATLAS Tile Calorimeter: Technical Design Report*, ATLAS-TDR-3; CERN-LHCC-96-042 (1996).
- [82] ATLAS Collaboration, *ATLAS Barrel Toroid: Magnet Project Technical Design Report, Volume 2*, ATLAS-TDR-7; CERN-LHCC-97-019 (1997).
- [83] ATLAS Collaboration, *ATLAS End-Cap Toroids: Magnet Project Technical Design Report, Volume 3*, ATLAS-TDR-8; CERN-LHCC-97-020 (CERN, 1997).
- [84] ATLAS Collaboration, *ATLAS Muon Spectrometer: Technical Design Report*, ATLAS-TDR-10; CERN-LHCC-97-022 (CERN, 1997).
- [85] E. Diehl, ‘Calibration and performance of the ATLAS Muon Spectrometer’, arXiv:1109.6933 (2011).
- [86] ATLAS Collaboration, *Athena*, version 21.2.129, 2019.
- [87] M. Cacciari, G. P. Salam and G. Soyez, ‘The anti- k_t jet clustering algorithm’, JHEP **04**, 063 (2008).
- [88] M. Cacciari, G. P. Salam and G. Soyez, ‘FastJet user manual’, EPJ C **72**, 1896 (2012).
- [89] ATLAS Collaboration, ‘Jet reconstruction and performance using particle flow with the ATLAS Detector’, EPJ C **77**, 466 (2017).

- [90] ATLAS Collaboration, ‘A precise measurement of the jet energy scale derived from single-particle measurements and in situ techniques in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, arXiv:2407.15627 (2024).
- [91] ATLAS Collaboration, ‘Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, EPJ C **81**, 689 (2021).
- [92] ATLAS Collaboration, *Forward Jet Vertex Tagging: a new technique for the identification and rejection of forward pileup jets*, ATL-PHYS-PUB-2015-034, 2015.
- [93] ATLAS Collaboration, *Forward jet vertex tagging using the particle flow algorithm*, ATL-PHYS-PUB-2019-026, 2019.
- [94] S. Mondal and L. Mastrolorenzo, ‘Machine learning in high energy physics: a review of heavy-flavor jet tagging at the LHC’, EPJ Special Topics **233**, 2657–2686 (2024).
- [95] ATLAS Collaboration, ‘ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset’, EPJ C **83**, 681 (2023).
- [96] ATLAS Collaboration, ‘ATLAS b -jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV’, EPJ C **79**, 970 (2019).
- [97] ATLAS Collaboration, *Secondary vertex finding for jet flavour identification with the ATLAS detector*, ATL-PHYS-PUB-2017-011, 2017.
- [98] ATLAS Collaboration, *Topological b -hadron decay reconstruction and identification of b -jets with the JetFitter package in the ATLAS experiment at the LHC*, ATL-PHYS-PUB-2018-025, 2018.
- [99] L. Pereira Sanchez, ‘Calibration of flavour tagging algorithms in ATLAS on $t\bar{t}$ and Z+jets final states’, PoS **EPS-HEP2021**, 10.22323/1.398.0434 (2022).
- [100] ATLAS Collaboration, *Energy scale calibration of b -tagged jets with ATLAS Run 2 data using $t\bar{t}$ lepton+jets events*, ATLAS-CONF-2022-004, 2022.
- [101] ATLAS Collaboration, *Reconstruction, identification, and calibration of hadronically decaying tau leptons with the ATLAS detector for the LHC Run 3 and reprocessed Run 2 data*, ATL-PHYS-PUB-2022-044, 2022.
- [102] ATLAS Collaboration, *Reconstruction, energy calibration, and identification of hadronically decaying tau leptons in the ATLAS experiment for Run 2 of the LHC*, ATL-PHYS-PUB-2015-045, 2015.
- [103] ATLAS Collaboration, *ATLAS Inner Tracker Pixel Detector: Technical Design Report*, ATLAS-TDR-030; CERN-LHCC-2017-021 (2017).
- [104] ATLAS and CMS Collaborations, *Highlights of the HL-LHC physics projections by ATLAS and CMS*, ATL-PHYS-PUB-2025-018, CMS-HIG-25-002, arXiv:2504.00672, 2025.
- [105] ATLAS Collaboration, ‘A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery’, Nature **607**, 52–59 (2022); Erratum: Nature **612**, E24 (2022).
- [106] ATLAS Collaboration, ‘Sensor response and radiation damage effects for 3D pixels in the ATLAS IBL Detector’, JINST, 10008 (2024).
- [107] L. Gonella, ‘The ATLAS ITk detector system for the Phase-II LHC upgrade’, NIM A **1045**, 167597 (2023).
- [108] D. Santo et al., ‘Test of the Optosystem for the ATLAS ITk data transmission chain’, JINST **18**, C03021 (2023).
- [109] A. Garcia Alonso (ATLAS), *ATLAS ITk Strip Detector for the Phase-II LHC Upgrade* (CERN, Geneva, 2024).
- [110] ATLAS Collaboration, *ATLAS Inner Tracker Strip Detector: Technical Design Report*, ATLAS-TDR-025; CERN-LHCC-2017-005 (2017).

- [111] ATLAS ITk Collaboration, *Fluence distributions for ITk extended@4 layout*, ATLAS Public plots: ITK-2016-002, 2016.
- [112] ATLAS Collaboration, ‘Expected tracking performance of the ATLAS Inner Tracker at the High-Luminosity LHC’, JINST **20**, P02018 (2025).
- [113] ATLAS Collaboration, *Expected tracking and related performance with the updated ATLAS Inner Tracker layout at the High-Luminosity LHC*, ATL-PHYS-PUB-2021-024, 2021.
- [114] S. Grinstein, ‘Overview of the ATLAS insertable B-layer (IBL) project’, NIM A **699**, **HSTD-8 2011**, 61–66 (2013).
- [115] L. P. Rossi, ‘The phase-2 ATLAS ITk pixel upgrade’, NIM A **924**, 11th International Hiroshima Symposium on Development and Application of Semiconductor Tracking Detectors, 270–274 (2019).
- [116] I. Asensi Tortajada et al., ‘Development of the BCM system and readout for ATLAS’, PoS **VERTEX2023**, 023 (2024).
- [117] P. Fernandez Martinez, ‘Overview of the ATLAS High-Granularity Timing Detector: project status and results’, PoS **EPS-HEP2023**, 525 (2024).
- [118] ATLAS Collaboration, *A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-031; CERN-LHCC-2020-007 (2020).
- [119] F. Pastore (ATLAS), *ATLAS TDAQ upgrades for Phase-2*, ATL-DAQ-PROC-2024-008, 2024.
- [120] ATLAS Collaboration, *ATLAS TDAQ Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-029; CERN-LHCC-2017-020 (2017).
- [121] A. Paramonov (ATLAS TDAQ), ‘FELIX: the detector interface for the ATLAS experiment at CERN’, EPJ Web Conf. **251**, 04006 (2021).
- [122] E. Mazzeo, ‘ATLAS Liquid Argon Calorimeter Frontend electronics Phase-2 upgrade’, PoS **LHCP2023**, 249 (2024).
- [123] ATLAS Collaboration, *ATLAS LAr Calorimeter Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-027; CERN-LHCC-2017-018 (2017).
- [124] F. Carrió (ATLAS Tile Calorimeter System), ‘The Data Acquisition System for the ATLAS Tile Calorimeter Phase-II Upgrade Demonstrator’, IEEE Trans. Nucl. Sci. **69**, 687–695 (2022).
- [125] ATLAS Collaboration, *ATLAS Tile Calorimeter Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-028; CERN-LHCC-2017-019 (2017).
- [126] A. Policicchio, ‘The Phase-II upgrade of the ATLAS Muon Spectrometer’, PoS **EPS-HEP2019**, 149 (2020).
- [127] ATLAS Collaboration, *ATLAS Muon Spectrometer Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-026; CERN-LHCC-2017-017 (2017).
- [128] M. Garcia-Sciveres, F. Loddo and J. Christiansen (RD53), *RD53B Manual*, CERN-RD53-PUB-19-002 (2019).
- [129] M. Garcia-Sciveres (RD53), *RD53C Chip Manual*, CERN-RD53-PUB-24-001 (2024).
- [130] C. Solans Sánchez et al., ‘MALTA monolithic pixel sensors in TowerJazz 180 nm technology’, NIM A **1057**, 168787 (2023).
- [131] A. Schöning et al., ‘MuPix and ATLASPix – Architectures and Results’, PoS **VERTEX2019**, 024 (2020).
- [132] *Aurora 64B/66B Protocol Specification*, v1.3, Spec. no. SP011, Xilinx (2014).
- [133] M. Karagounis et al., ‘An integrated Shunt-LDO regulator for serial powered systems’, in 2009 proceedings of ESSCIRC (2009), pp. 276–279.

- [134] T. Heim, ‘YARR - A PCIe based Readout Concept for Current and Future ATLAS Pixel Modules’, *Journal of Physics: Conference Series* **898**, 032053 (2017).
- [135] R. Ahmad, ‘The Monitoring of Pixel System chip for the detector control system of the ATLAS ITk Pixel Detector’, *Journal of Physics: Conference Series* **2374**, 012094 (2022).
- [136] R. Ahmad et al., ‘Second generation Monitoring of Pixel System chip for the detector control system of the ATLAS ITk Pixel detector’, *JINST* **18**, C04015 (2023).
- [137] *Road vehicles — Controller area network (CAN)*, tech. rep., Standard ISO 11898-1:2024 (International Organization for Standardization, 2024).
- [138] A. Qamesh et al., ‘An FPGA-based data aggregator for the new ATLAS ITK Pixel DCS’, arXiv:2410.24057 **TWEPP-24** (2024).
- [139] S. Kersten et al., ‘The ITk interlock hardware protection system’, *NIM A* **1045**, 167613 (2023).
- [140] O. Beltramello et al., ‘The detector safety system of the ATLAS experiment’, *JINST* **4**, P09012 (2009).
- [141] T. Heim, ‘Status and performance of the ATLAS Pixel Detector after 3 years of operation’, *NIM A* **765**, **HSTD-9 2013**, 227–231 (2014).
- [142] F. Campabadal et al., ‘Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker’, *NIM A* **552**, 292–328 (2005).
- [143] S. Möbius and on behalf of the ATLAS ITk Group, ‘The Optosystem: validation and testing of the high-speed electro-optical conversion system for the readout of the ATLAS ITk Pixel upgrade’, *JINST* **19**, C04015 (2024).
- [144] L. Zhang et al., ‘The design and test results of A Giga-Bit Cable Receiver (GBCR) for the ATLAS Inner Tracker Pixel Detector’, *JINST* **18**, C03005 (2023).
- [145] D. Hernandez Montesinos et al., ‘Overview of the production and qualification tests of the lpGBT’, *JINST* **19**, C04048 (2024).
- [146] J. Troska et al., ‘The VTRx+, an Optical Link Module for Data Transmission at HL-LHC’, *PoS TWEPP-17*, 048 (2018).
- [147] *AGS10000 Series Gantries*, Data Sheet (Aerotech, 2021).
- [148] *PROPlus / PRO Series Automated Dispensing Systems*, Data Sheet v040122 (Nordson EFD, 2022).
- [149] S. Koch and M. Mironova (ATLAS ITk), *Oxford ITk Pixel Ring Loading Procedure Note*, version 0.1, Internal documentation, accessible at <https://gitlab.cern.ch/opmd/ring-loading/opmd-ring-loading-note>, 2022.
- [150] S. Koch and M. Miranova, *Gantry Instruction Generation, Object-Oriented*, CERN GitLab code repository, <https://gitlab.cern.ch/itkcommoninterlock/opmd/ring-loading/gigo>, 2022.
- [151] P. Sutcliffe, *OEC Bare Half Ring Assembly Drawings*, EDMS Document 2377394 v.1 | AT2-IP-ED-0013 v.1 (ATLAS internal), 2022.
- [152] D. Giugni, *OEC Loaded Half Ring drawings set*, EDMS Document 3231683 v.1 | AT2-IP-ED-0112 v.1 (ATLAS internal), 2025.
- [153] *Information about Dow Corning® brand thermally conductive materials*, Data Sheet 10-900N-01 (Dow Corning, 2008).
- [154] Cospheric, *Soda Lime Solid Glass Microspheres 2.5g/cc - 1µm to 4400µm (4.4mm)*, https://www.cospheric.com/SLGMSc_solid_glass_spheres_beads_microns.htm, Online Catalogue, accessed: 2022-07-04, 2022.
- [155] A. Petrukhin (ATLAS ITk), ‘RD53A/B pixel module assembly and testing experience’, *PoS VERTEX2023*, 065 (2024).

- [156] S. Möbius (ATLAS ITk), ‘Module development for the ATLAS ITk pixel detector’, JINST **17**, C03042 (2022).
- [157] J. Schumacher, C. Plessl and W. Vandelli, ‘High-Throughput and Low-Latency Network Communication with NetIO’, Journal of Physics: Conference Series **898**, 082003 (2017).
- [158] W. Alkakhli and on behalf of the ATLAS ITk collaboration, ‘ATLAS ITk-Pixel DAQ system’, JINST **19**, C11013 (2024).
- [159] *PL512/PL506 modular power supply system*, Manual 00679.A4 (W-IE-NE-R, 2013).
- [160] J. Case et al., *Introduction to community-based SNMPv2*, RFC 1901, Prop. Standard, 1996.
- [161] *EHS series: versatile high precision high voltage module with multiple floating options*, Manual ver. 4.1 (iseg, 2024).
- [162] *MPOD HV & LV power supply system*, Manual ver. 3.2 (W-IE-NE-R, 2020).
- [163] B. Vormwald (ATLAS ITk Pixel), ‘Performance of the ATLAS ITK Pixel Detector Prototype’, JPS Conf. Proc. **42**, 011015 (2024).
- [164] B. Verlaet, L. Zwalinski, R. Dumps et al., ‘TRACI, a multipurpose CO2 cooling system for R&D’, in Proceedings of the 10th IIR-Gustav Lorentzen Conference on Natural Refrigerants (GL2012) (2012).
- [165] *Datasheet SHT85: humidity and temperature sensor*, Data Sheet ver. 4 (Sensirion, 2021).
- [166] *DMT143L dew point transmitter for low dew points*, Data Sheet B211207EN-K (Vaisala, 2021).
- [167] *cRIO-9066 Specifications: Embedded Real-Time Controller with Reconfigurable FPGA for C Series Modules*, Data Sheet (National Instruments, 2022).
- [168] *NI-9217 Specifications*, Data Sheet (National Instruments, 2024).
- [169] *NI-9201 and sbRIO-9201 Specifications*, Data Sheet (National Instruments, 2024).
- [170] *NI-9402 and sbRIO-9402 Specifications*, Data Sheet (National Instruments, 2024).
- [171] Siemens, *WinCC-OA 3.19 documentation*, available at https://www.winccoa.com/documentation/WinCCOA/3.19/en_US/GettingStarted/GettingStarted-02.html (2024), accessed 14 Jan. 2025.
- [172] S. Schlenker et al., ‘The ATLAS Detector Control System’, Conf. Proc. **C111010**, MOBAUST02 (2011).
- [173] OPC Foundation, *OPC 10000-1: UA Part 1 – Overview and Concepts*, tech. rep. 1.04, OPC UA Specification Version 1.04 (2017).
- [174] InfluxData, *InfluxDB OSS 1.11.8*, https://docs.influxdata.com/influxdb/v1/about_the_project/release-notes/, 2025.
- [175] M. Marjanović et al., *Electrical specification and QC procedures for ITkPixV1.1 modules*, EDMS Document 2786321 v.1 | AT2-IP-QA-0025 v.2.2.2 (ATLAS internal), 2025.
- [176] ATLAS collaboration, ‘Study of the material of the ATLAS inner detector for Run 2 of the LHC’, JINST **12**, P12009 (2017).
- [177] K. Tackmann (ATLAS), ‘ATLAS inner detector material studies’, in Physics at the LHC 2010 (2010), pp. 63–66.
- [178] ATLAS Collaboration, *Study of the Material Budget in the ATLAS Inner Detector with K_S^0 decays in collision data at $\sqrt{s} = 900$ GeV*, ATLAS-CONF-2010-019, 2010.
- [179] ATLAS Collaboration, *Probing the response of the ATLAS electromagnetic calorimeter and material upstream with energy flow from $\sqrt{s} = 7$ TeV minimum bias events*, ATLAS-CONF-2010-037, 2010.
- [180] ATLAS collaboration, ‘Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data’, The European Physical Journal C **74**, 3071 (2014).

- [181] ATLAS Collaboration, ‘Electron and photon energy calibration with the ATLAS detector using LHC Run 2 data’, JINST **19**, P02009 (2023).
- [182] ATLAS Collaboration, ‘Search for long-lived, massive particles in events with displaced vertices and multiple jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, JHEP **06**, 200 (2023).
- [183] ATLAS Collaboration, ‘Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, EPJ C **82**, 606 (2022).
- [184] J. D. Jackson, *Classical electrodynamics*, 3rd ed. (Wiley, New York, NY, 1999).
- [185] R. Frühwirth and M. Regler, ‘On the quantitative modelling of core and tails of multiple scattering by gaussian mixtures’, NIM A **456**, 369–389 (2001).
- [186] G. Molière, ‘Theorie der Streuung schneller geladener Teilchen II. Mehrfach- und Vielfachstreuung’, Zeitschrift Naturforschung Teil A **3**, 78–97 (1948).
- [187] H. S. Snyder and W. T. Scott, ‘Multiple Scattering of Fast Charged Particles’, Phys. Rev. **76**, 220–225 (1949).
- [188] W. T. Scott, ‘Mean-Value Calculations for Projected Multiple Scattering’, Phys. Rev. **85**, 245–248 (1952).
- [189] S. Goudsmit and J. L. Saunderson, ‘Multiple scattering of electrons’, Phys. Rev. **57**, 24–29 (1940).
- [190] H. W. Lewis, ‘Multiple Scattering in an Infinite Medium’, Phys. Rev. **78**, 526–529 (1950).
- [191] H. A. Bethe, ‘Molière’s Theory of Multiple Scattering’, Phys. Rev. **89**, 1256–1266 (1953).
- [192] B. Gottschalk et al., ‘Multiple Coulomb scattering of 160 MeV protons’, NIM B **74**, 467–490 (1993).
- [193] A. O. Hanson et al., ‘Measurement of Multiple Scattering of 15.7-Mev Electrons’, Phys. Rev. **84**, 634–637 (1951).
- [194] W. T. Scott, ‘The Theory of Small-Angle Multiple Scattering of Fast Charged Particles’, Rev. Mod. Phys. **35**, 231–313 (1963).
- [195] V. L. Highland, ‘Some practical remarks on multiple scattering’, NIM **129**, 497–499 (1975).
- [196] G. R. Lynch and O. I. Dahl, ‘Approximations to multiple Coulomb scattering’, NIM B **58**, 6–10 (1991).
- [197] R. M. Corless et al., ‘On the LambertW function’, Advances in Computational Mathematics **5**, 329–359 (1996).
- [198] D. Veberič, ‘Lambert W function for applications in physics’, Computer Physics Communications **183**, 2622–2628 (2012).
- [199] S. Agostinelli et al., ‘GEANT4 – a simulation toolkit’, NIM A **506**, 250 (2003).
- [200] V. N. Ivanchenko et al., ‘Geant4 models for simulation of multiple scattering’, Journal of Physics: Conference Series **219**, 032045 (2010).
- [201] L. Urbán (GEANT4), *A model for multiple scattering in GEANT4*, CERN-OPEN-2006-077 (CERN, Geneva, 2006).
- [202] U. Stolzenberg et al., ‘Radiation length imaging with high-resolution telescopes’, NIM A **845**, Proceedings of the Vienna Conference on Instrumentation 2016, 173–176 (2017).
- [203] L. Poley et al., ‘Mapping the material distribution of a complex structure in an electron beam’, JINST **16**, P01010 (2021).
- [204] C. Y. Qu et al., ‘Analysis of pixel detector material budget based on test beam’, JINST **16**, T06004 (2021).

- [205] Tracker Group of the CMS Collaboration, ‘Measurement of the fractional radiation length of a pixel module for the CMS Phase-2 upgrade via the multiple scattering of positrons’, JINST **19**, P10023 (2024).
- [206] J. A. Czocharlski, ‘Ein neues Verfahren zur Messung der Kristallisationsgeschwindigkeit der Metalle’, Zeitschrift für Physikalische Chemie **92U**, 219–221 (1918).
- [207] M. van Rijnbach et al., ‘Performance of the MALTA telescope’, EPJ C **83**, 581 (2023).
- [208] ITk-FELIX-SW developers, *ITk-FELIX-SW readout software*, CERN GitLab code repository, <https://gitlab.cern.ch/itk-felix-sw/itk-felix-sw>, 2024.
- [209] ATLAS Collaboration, *MALTA DAQ Software Repository*, CERN GitLab code repository, <https://gitlab.cern.ch/malta/MaltaDAQ>, 2020.
- [210] Autodesk, *Autodesk Inventor Nastran*, Version 23.0, San Rafael, CA, 2023.
- [211] D. Dannheim et al., ‘Corryvreckan: a modular 4d track reconstruction and analysis software for test beam data’, JINST **16**, P03008 (2021).
- [212] V. Blobel and C. Kleinwort, ‘A new method for the high precision alignment of track detectors’, in Conference on Advanced Statistical Techniques in Particle Physics (2002).
- [213] M. Oreglia, ‘A study of the reactions $\psi' \rightarrow \gamma\gamma\psi$ ’, SLAC-0236, PhD thesis (Stanford University, 1980).
- [214] J. E. Gaiser, ‘Charmonium spectroscopy from radiative decays of the J/ψ and ψ' ’, SLAC-0255, UMI-83-14449-MC, PhD thesis (Stanford University, 1982), pp. 178–180.
- [215] W. Verkerke and D. P. Kirkby, ‘The RooFit toolkit for data modeling’, eConf **C0303241**, MOLT007 (2003).
- [216] M. Gupta, *Calculation of radiation length in materials*, PH-EP-Tech-Note-2010-013 (CERN, Geneva, 2010).
- [217] S. Spannagel et al., ‘Allpix²: a modular simulation framework for silicon detectors’, NIM A **901**, 164–172 (2018).
- [218] S. Navas et al. (Particle Data Group), ‘Review of particle physics’, Phys. Rev. D **110**, Ch. 6: Atomic and Nuclear Properties of Materials, 030001 (2024).
- [219] Y. Nir, ‘Flavour physics and CP violation’, in 2013 CERN - Latin-American School of High-Energy Physics (2015), pp. 123–156.
- [220] ATLAS Collaboration, ‘Search for charged-lepton-flavour violation in Z -boson decays with the ATLAS detector’, Nature Phys. **17**, 819 (2021).
- [221] ATLAS Collaboration, ‘Precise test of lepton flavour universality in W -boson decays into muons and electrons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, The European Physical Journal C **84**, 993 (2024).
- [222] M. Ablikim et al. (BESIII Collaboration), ‘Precision measurements of $B[\psi(3686) \rightarrow \pi^+\pi^- J/\psi]$ and $B[J/\psi \rightarrow l^+l^-]$ ’, Phys. Rev. D **88**, 032007 (2013).
- [223] A. Aguilar-Arevalo et al. (PIENU Collaboration), ‘Improved Measurement of the $\pi \rightarrow e\nu$ Branching Ratio’, Phys. Rev. Lett. **115**, 071801 (2015).
- [224] C. Lazzeroni et al., ‘Precision measurement of the ratio of the charged kaon leptonic decay rates’, Physics Letters B **719**, 326–336 (2013).
- [225] S. Banerjee et al. (HFLAV), ‘Averages of b -hadron, c -hadron, and τ -lepton properties as of 2023’, arXiv:2411.18639 (2024).
- [226] B. Clerbaux, ‘Experimental Summary of the Moriond 2024 conference – Electroweak Interaction & Unified Theories’, in 58th Rencontres de Moriond on Electroweak Interactions and Unified Theories (2024).
- [227] LHCb Collaboration, ‘Measurement of the Branching Fraction Ratios $R(D^+)$ and $R(D^{*+})$ Using Muonic τ Decays’, Phys. Rev. Lett. **134**, 061801 (2025).

- [228] LHCb Collaboration, ‘Measurement of the Ratios of Branching Fractions $\mathcal{R}(D^*)$ and $\mathcal{R}(D^0)$ ’, Phys. Rev. Lett. **131**, 111802 (2023).
- [229] LHCb Collaboration, ‘Test of lepton flavor universality using $B^0 \rightarrow D^{*-} \tau^+ \nu_\tau$ decays with hadronic τ channels’, Phys. Rev. D **108**, 012018 (2023).
- [230] Belle Collaboration, ‘Measurement of $\mathcal{R}(D)$ and $\mathcal{R}(D^*)$ with a Semileptonic Tagging Method’, Phys. Rev. Lett. **124**, 161803 (2020).
- [231] Belle II Collaboration, ‘Test of lepton flavor universality with a measurement of $R(D^*)$ using hadronic B tagging at the Belle II experiment’, Phys. Rev. D **110**, 072020 (2024).
- [232] BABAR Collaboration, ‘Measurement of an excess of $\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau$ decays and implications for charged Higgs bosons’, Phys. Rev. D **88**, 072012 (2013).
- [233] C. Cornella et al., ‘Reading the footprints of the b-meson flavor anomalies’, Journal of High Energy Physics **2021**, 50 (2021).
- [234] LHCb Collaboration, ‘Tests of lepton universality using $B^0 \rightarrow K_S^0 \ell^+ \ell^-$ and $B^+ \rightarrow K^{*+} \ell^+ \ell^-$ decays’, Phys. Rev. Lett. **128**, 191802 (2022).
- [235] S. Celani (LHCb), ‘Lepton flavour universality tests and lepton flavour violation searches at LHCb’, in 16th International Workshop on Tau Lepton Physics (2021).
- [236] M. Bordone, G. Isidori and A. Pattori, ‘On the standard model predictions for R_K and R_{K^*} ’, EPJ C **76**, 440 (2016).
- [237] LHCb Collaboration, ‘Measurement of lepton universality parameters in $B^+ \rightarrow K^+ \ell^+ \ell^-$ and $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays’, Phys. Rev. D **108**, 032002 (2023).
- [238] LHCb Collaboration, ‘Test of lepton universality in $b \rightarrow s \ell^+ \ell^-$ decays’, Phys. Rev. Lett. **131**, 051803 (2023).
- [239] B. Capdevila et al., ‘Searching for new physics with $b \rightarrow s \tau^+ \tau^-$ processes’, Phys. Rev. Lett. **120**, 181802 (2018).
- [240] LHCb Collaboration, ‘Search for the decays $B_s^0 \rightarrow \tau^+ \tau^-$ and $B^0 \rightarrow \tau^+ \tau^-$ ’, Phys. Rev. Lett. **118**, 251802 (2017).
- [241] J. P. Lees et al. (BaBar Collaboration), ‘Search for $B^+ \rightarrow K^+ \tau^+ \tau^-$ at the BaBar experiment’, Phys. Rev. Lett. **118**, 031802 (2017).
- [242] E. E. Jenkins, A. V. Manohar and P. Stoffer, ‘Low-energy effective field theory below the electroweak scale: operators and matching’, JHEP **2018**, 16 (2018).
- [243] B. Capdevila, A. Crivellin and J. Matias, ‘Review of semileptonic B anomalies’, EPJ Special Topics **233**, 409–428 (2024).
- [244] G. Buchalla, A. J. Buras and M. E. Lautenbacher, ‘Weak decays beyond leading logarithms’, Rev. Mod. Phys. **68**, 1125–1244 (1996).
- [245] C. Bobeth et al., ‘Complete NNLO QCD analysis of $\bar{B} \rightarrow X_s \ell^+ \ell^-$ and higher order electroweak effects’, JHEP **2004**, 071 (2004).
- [246] T. Huber et al., ‘Electromagnetic logarithms in $\bar{B} \rightarrow X_s \ell^+ \ell^-$ ’, Nuclear Physics B **740**, 105–137 (2006).
- [247] T. Huber, T. Hurth and E. Lunghi, ‘Logarithmically enhanced corrections to the decay rate and forward–backward asymmetry in $\bar{B} \rightarrow X_s \ell^+ \ell^-$ ’, Nuclear Physics B **802**, 40–62 (2008).
- [248] M. Algueró et al., ‘To (b)e or not to (b)e: no electrons at LHCb’, EPJ C **83**, 648 (2023).
- [249] M. Ciuchini et al., ‘Constraints on lepton universality violation from rare B decays’, Phys. Rev. D **107**, 055036 (2023).
- [250] T. Hurth et al., ‘More indications for lepton nonuniversality in $b \rightarrow s \ell^+ \ell^-$ ’, Phys. Lett. B **824**, 136838 (2022).

- [251] M. Blanke et al., ‘Addendum to “Impact of polarization observables and $B_c \rightarrow \tau \nu$ on new physics explanations of the $b \rightarrow c \tau \nu$ anomaly”’, Phys. Rev. D **100**, 035035 (2019).
- [252] W. Dekens and P. Stoffer, ‘Low-energy effective field theory below the electroweak scale: matching at one loop’, JHEP **2019**, 197 (2019).
- [253] S. Hamoudou, J. Kumar and D. London, ‘Dimension-8 smeft matching conditions for the low-energy effective field theory’, JHEP **2023**, 157 (2023).
- [254] X. Fan et al., ‘Measurement of the electron magnetic moment’, Phys. Rev. Lett. **130**, 071801 (2023).
- [255] Muon $g - 2$ Collaboration, *Measurement of the positive muon anomalous magnetic moment to 127 ppb*, Submitted to PRL, arXiv:2506.03069, 2025.
- [256] D. P. Aguillard et al. (Muon $g - 2$ Collaboration), ‘Detailed report on the measurement of the positive muon anomalous magnetic moment to 0.20 ppm’, Phys. Rev. D **110**, 032009 (2024).
- [257] D. Giusti et al., ‘Strange and charm HVP contributions to the muon $g - 2$ including QED corrections with twisted-mass fermions’, JHEP **2017**, 157 (2017).
- [258] S. Borsanyi et al., ‘Leading hadronic contribution to the muon magnetic moment from lattice QCD’, Nature **593**, 51–55 (2021).
- [259] R. Aliberti et al., *The anomalous magnetic moment of the muon in the standard model: an update*, Submitted to Physics Reports, arXiv:2505.21476, 2025.
- [260] CMS Collaboration, ‘Observation of $\gamma\gamma \rightarrow \tau\tau$ in proton–proton collisions and limits on the anomalous electromagnetic moments of the τ lepton’, Reports on Progress in Physics **87**, 107801 (2024).
- [261] ATLAS Collaboration, ‘Observation of the $\gamma\gamma \rightarrow \tau\tau$ process in Pb+Pb collisions and constraints on the τ -lepton anomalous magnetic moment with the ATLAS detector’, Phys. Rev. Lett. **131**, 151802 (2023).
- [262] S. Eidelman and M. Passera, ‘Theory of the τ lepton anomalous magnetic moment’, Modern Physics Letters A **22**, 159–179 (2007).
- [263] U. Haisch, L. Schnell and J. Weiss, ‘LHC tau-pair production constraints on a_τ and d_τ ’, SciPost Phys. **16**, 048 (2024).
- [264] I. Doršner et al., ‘Physics of leptoquarks in precision experiments and at particle colliders’, Physics Reports **641**, 1–68 (2016).
- [265] J. C. Pati and A. Salam, ‘Lepton number as the fourth “color”’, Phys. Rev. D **10**, 275–289 (1974).
- [266] H. Georgi, ‘The state of the art—gauge theories’, AIP Conf. Proc. **23**, 575–582 (1975).
- [267] H. Fritzsch and P. Minkowski, ‘Unified interactions of leptons and hadrons’, Annals of Physics **93**, 193–266 (1975).
- [268] R. Barbier et al., ‘ R -parity-violating supersymmetry’, Physics Reports **420**, 1–195 (2005).
- [269] E. Farhi and L. Susskind, ‘Technicolour’, Physics Reports **74**, 277–321 (1981).
- [270] B. Schrempp and F. Schrempp, ‘Light leptoquarks’, Physics Letters B **153**, 101–107 (1985).
- [271] N. Assad, B. Fornal and B. Grinstein, ‘Baryon number and lepton universality violation in leptoquark and diquark models’, Physics Letters B **777**, 324–331 (2018).
- [272] Y. Sakaki et al., ‘Testing leptoquark models in $\bar{B} \rightarrow D^{(*)} \tau \bar{\nu}$ ’, Phys. Rev. D **88**, 094012 (2013).
- [273] M. Bauer and M. Neubert, ‘Minimal leptoquark explanation for the $R_{D^{(*)}}$, R_K , and $(g - 2)_\mu$ anomalies’, Phys. Rev. Lett. **116**, 141802 (2016).

- [274] A. Angelescu et al., ‘Closing the window on single leptoquark solutions to the B-physics anomalies’, *JHEP* **2018**, 183 (2018).
- [275] D. Bečirević and O. Sumensari, ‘A leptoquark model to accommodate $R_K^{\text{exp}} < R_K^{\text{SM}}$ and $R_{K^*}^{\text{exp}} < R_{K^*}^{\text{SM}}$ ’, *JHEP* **2017**, 104 (2017).
- [276] D. Buttazzo et al., ‘B-physics anomalies: a guide to combined explanations’, *JHEP* **2017**, 44 (2017).
- [277] CMS Collaboration, ‘Search for a third-generation leptoquark coupled to a τ lepton and a b quark through single, pair, and nonresonant production in proton–proton collisions at $\sqrt{s} = 13$ TeV’, *JHEP* **05**, 311 (2023).
- [278] ATLAS Collaboration, ‘Search for leptoquarks decaying into the $b\tau$ final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, *JHEP* **10**, 001 (2023).
- [279] ATLAS Collaboration, ‘Search for pair production of third-generation leptoquarks decaying into a bottom quark and a τ -lepton with the ATLAS detector’, *EPJ C* **83**, 1075 (2023).
- [280] CMS Collaboration, ‘Searches for additional Higgs bosons and for vector leptoquarks in $\tau\tau$ final states in proton–proton collisions at $\sqrt{s} = 13$ TeV’, *JHEP* **07**, 073 (2023).
- [281] ATLAS Collaboration, ‘Search for excited τ -leptons and leptoquarks in the final state with τ -leptons and jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, *JHEP* **06**, 199 (2023).
- [282] CMS Collaboration, ‘Search for heavy resonances decaying to tau lepton pairs in proton–proton collisions at $\sqrt{s} = 13$ TeV’, *JHEP* **02**, 048 (2017).
- [283] ATLAS Collaboration, ‘Search for additional heavy neutral Higgs and gauge bosons in the ditau final state produced in 36 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, *JHEP* **01**, 055 (2018).
- [284] ATLAS Collaboration, ‘Search for pair production of scalar leptoquarks decaying into first- or second-generation leptons and top quarks in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, *EPJ C* **81**, 313 (2021).
- [285] ATLAS Collaboration, ‘Search for pair production of third-generation scalar leptoquarks decaying into a top quark and a τ -lepton in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, *JHEP* **06**, 179 (2021).
- [286] CMS Collaboration, ‘Search for third-generation scalar leptoquarks decaying to a top quark and a τ lepton at $\sqrt{s} = 13$ TeV’, *EPJ C* **78**, 707 (2018).
- [287] ATLAS Collaboration, ‘Searches for third-generation scalar leptoquarks in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector’, *JHEP* **06**, 144 (2019).
- [288] ATLAS Collaboration, ‘Combination of searches for pair-produced leptoquarks at $\sqrt{s} = 13$ TeV with the ATLAS detector’, *Phys. Lett. B* **854**, 138736 (2024).
- [289] ATLAS Collaboration, ‘Search for pair-produced scalar and vector leptoquarks decaying into third-generation quarks and first- or second-generation leptons in pp collisions with the ATLAS detector’, *JHEP* **06**, 188 (2023).
- [290] ATLAS Collaboration, ‘Search for pairs of scalar leptoquarks decaying into quarks and electrons or muons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector’, *JHEP* **10**, 112 (2020).
- [291] J. Aebischer et al., ‘Confronting the vector leptoquark hypothesis with new low- and high-energy data’, *EPJ C* **83**, 153 (2023).
- [292] M. Cvetič and P. Langacker, ‘ Z' physics and Supersymmetry’, in *Perspectives on supersymmetry* (1998), pp. 312–331.
- [293] P. Langacker, ‘ Z' physics from strings’, in 6th International Symposium on Particles, Strings and Cosmology (1998), pp. 587–596.

- [294] P. Langacker, ‘The physics of heavy Z' gauge bosons’, *Rev. Mod. Phys.* **81**, 1199–1228 (2009).
- [295] M. Masip and A. Pomarol, ‘Effects of standard model Kaluza-Klein excitations on electroweak observables’, *Phys. Rev. D* **60**, 096005 (1999).
- [296] A. J. Buras, F. De Fazio and J. Girrbach, ‘The anatomy of Z' and Z with flavour changing neutral currents in the flavour precision era’, *JHEP* **2013**, 116 (2013).
- [297] P. Langacker and M. Plümacher, ‘Flavor changing effects in theories with a heavy Z' boson with family nonuniversal couplings’, *Phys. Rev. D* **62**, 013006 (2000).
- [298] J. J. Aubert et al., ‘Experimental Observation of a Heavy Particle J' ’, *Phys. Rev. Lett.* **33**, 1404–1406 (1974).
- [299] J. E. Augustin et al., ‘Discovery of a Narrow Resonance in e^+e^- Annihilation’, *Phys. Rev. Lett.* **33**, 1406–1408 (1974).
- [300] S. W. Herb et al., ‘Observation of a Dimuon Resonance at 9.5 GeV in 400-GeV Proton-Nucleus Collisions’, *Phys. Rev. Lett.* **39**, 252–255 (1977).
- [301] G. Arnison et al., ‘Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider’, *Physics Letters B* **126**, 398–410 (1983).
- [302] P. Bagnaia et al., ‘Evidence for $Z_0 \rightarrow e^+e^-$ at the CERN pp collider’, *Physics Letters B* **129**, 130–140 (1983).
- [303] ATLAS Collaboration, ‘Measurement of the forward-backward asymmetry of electron and muon pair-production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector’, *JHEP* **09**, 049 (2015).
- [304] T. Aaltonen et al. (CDF Collaboration and D0 Collaboration), ‘Tevatron Run II combination of the effective leptonic electroweak mixing angle’, *Phys. Rev. D* **97**, 112007 (2018).
- [305] CMS Collaboration, ‘Measurement of the differential and double-differential Drell-Yan cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV’, *JHEP* **12**, 030 (2013).
- [306] ATLAS Collaboration, ‘Determination of the parton distribution functions of the proton from ATLAS measurements of differential W^\pm and Z boson production in association with jets’, *JHEP* **07**, 223 (2021).
- [307] T. Sjöstrand, S. Mrenna and P. Skands, ‘A brief introduction to PYTHIA 8.1’, *Comput. Phys. Commun.* **178**, 852–867 (2008).
- [308] ATLAS Collaboration, *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie-Landshoff diffractive model*, ATL-PHYS-PUB-2016-017, 2016.
- [309] NNPDF Collaboration, R. D. Ball et al., ‘Parton distributions with LHC data’, *Nucl. Phys. B* **867**, 244 (2013).
- [310] ATLAS Collaboration, ‘The ATLAS simulation infrastructure’, *EPJ C* **70**, 823 (2010).
- [311] E. Bothmann et al., ‘Event generation with Sherpa 2.2’, *SciPost Phys.* **7**, 034 (2019).
- [312] T. Gleisberg and S. Höche, ‘Comix, a new matrix element generator’, *JHEP* **12**, 039 (2008).
- [313] F. Bucci et al., ‘OpenLoops 2’, *EPJ C* **79**, 866 (2019).
- [314] S. Alioli et al., ‘A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX’, *JHEP* **06**, 043 (2010).
- [315] ATLAS Collaboration, *Studies on top-quark Monte Carlo modelling for Top2016*, ATL-PHYS-PUB-2016-020, 2016.
- [316] S. Frixione et al., ‘Single-top hadroproduction in association with a W boson’, *JHEP* **07**, 029 (2008).

- [317] S. Catani and M. Grazzini, ‘Next-to-next-to-leading-order subtraction formalism in hadron collisions and its application to Higgs-boson production at the Large Hadron Collider’, *Phys. Rev. Lett.* **98**, 222002 (2007).
- [318] S. Höche et al., ‘QCD matrix elements + parton showers. The NLO case’, *JHEP* **04**, 027 (2013).
- [319] S. Höche et al., ‘A critical appraisal of NLO+PS matching methods’, *JHEP* **09**, 049 (2012).
- [320] S. Höche et al., ‘QCD matrix elements and truncated showers’, *JHEP* **05**, 053 (2009).
- [321] S. Catani et al., ‘QCD matrix elements + parton showers’, *JHEP* **11**, 063 (2001).
- [322] J. Pumplin et al., ‘New Generation of Parton Distributions with Uncertainties from Global QCD Analysis’, *JHEP* **07**, 012 (2002).
- [323] J. Gao et al., ‘CT10 next-to-next-to-leading order global analysis of QCD’, *Phys. Rev. D* **89**, 033009 (2014).
- [324] ATLAS Collaboration, ‘Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector’, *JHEP* **09**, 145 (2014).
- [325] D. J. Lange, ‘The EvtGen particle decay simulation package’, *NIM A* **462**, 152 (2001).
- [326] P. Golonka and Z. Was, ‘PHOTOS Monte Carlo: a precision tool for QED corrections in Z and W decays’, *EPJ C* **45**, 97–107 (2006).
- [327] N. Davidson, T. Przedzinski and Z. Was, ‘PHOTOS interface in C++: Technical and physics documentation’, *Comput. Phys. Commun.* **199**, 86–101 (2016).
- [328] NNPDF Collaboration, R. D. Ball et al., ‘Parton distributions for the LHC run II’, *JHEP* **04**, 040 (2015).
- [329] ATLAS Collaboration, *ATLAS Pythia 8 tunes to 7 TeV data*, ATL-PHYS-PUB-2014-021, 2014.
- [330] M. Bähr et al., ‘Herwig++ physics and manual’, *EPJ C* **58**, 639 (2008).
- [331] J. Bellm et al., ‘Herwig 7.1 release note’, arXiv:1705.06919 (2017).
- [332] L. A. Harland-Lang et al., ‘Parton distributions in the LHC era: MMHT 2014 PDFs’, *EPJ C* **75**, 204 (2015).
- [333] J. Bellm et al., ‘Herwig 7.0/Herwig++ 3.0 release note’, *EPJ C* **76**, 196 (2016).
- [334] J. Alwall et al., ‘The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations’, *JHEP* **07**, 079 (2014).
- [335] K. Hamilton et al., ‘NNLOPS simulation of Higgs boson production’, *JHEP* **10**, 222 (2013).
- [336] K. Hamilton, P. Nason and G. Zanderighi, ‘Finite quark-mass effects in the NNLOPS POWHEG+MiNLO Higgs generator’, *JHEP* **05**, 140 (2015).
- [337] J. Butterworth et al., ‘PDF4LHC recommendations for LHC Run II’, *J. Phys. G* **43**, 023001 (2016).
- [338] C. Bierlich et al., ‘A comprehensive guide to the physics and usage of PYTHIA 8.3’, *SciPost Phys. Codeb.*, 8 (2022).
- [339] L. Lönnblad, ‘Correcting the colour-dipole cascade model with fixed order matrix elements’, *JHEP* **05**, 046 (2002).
- [340] L. Lönnblad and S. Prestel, ‘Matching tree-level matrix elements with interleaved showers’, *JHEP* **03**, 019 (2012).
- [341] J. Schaarschmidt and on behalf of the ATLAS Collaboration, ‘The new ATLAS fast calorimeter simulation’, *Journal of Physics: Conference Series* **898**, 042006 (2017).
- [342] M. J. Baker et al., ‘High- p_T signatures in vector-leptoquark models’, *EPJ C* **79**, 334 (2019).

- [343] ATLAS Collaboration, ‘Measurements of the production cross-section for a Z boson in association with b -jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector’, JHEP **07**, 044 (2020).
- [344] ATLAS Collaboration, *The ATLAS Tau Trigger in Run 2*, ATLAS-CONF-2017-061, 2017.
- [345] ATLAS Collaboration, ‘Performance of the ATLAS muon triggers in Run 2’, JINST **15**, P09015 (2020).
- [346] ATLAS Collaboration, ‘Performance of electron and photon triggers in ATLAS during LHC Run 2’, EPJ C **80**, 47 (2020).
- [347] ATLAS Collaboration, ‘Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment’, JHEP **05**, 162 (2024).
- [348] ATLAS Collaboration, ‘Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data’, JINST **14**, P12006 (2019).
- [349] ATLAS Collaboration, ‘Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV’, EPJ C **76**, 292 (2016).
- [350] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, ATLAS-CONF-2014-018, 2014.
- [351] ATLAS Collaboration, ‘Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector’, EPJ C **76**, 581 (2016).
- [352] ATLAS Collaboration, ‘Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector’, EPJ C **77**, 580 (2017); Erratum: EPJ C **77**, 712 (2017).
- [353] ATLAS Collaboration, ‘Tools for estimating fake/non-prompt lepton backgrounds with the ATLAS detector at the LHC’, JINST **18**, T11004 (2023).
- [354] G. Rauco (ATLAS, CMS), ‘Distinguishing quark and gluon jets at the LHC’, in Parton radiation and fragmentation from LHC to FCC-ee (2017), pp. 73–78.
- [355] ATLAS Collaboration, *Estimation of non-prompt and fake lepton backgrounds in final states with top quarks produced in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, ATLAS-CONF-2014-058, 2014.
- [356] M. R. Rahman, ‘Background estimation in the search for single production of vector-like quarks decaying into Wb in pp collisions using a data-driven method at $\sqrt{s} = 13$ TeV with the ATLAS detector’, arXiv:2106.03961, MA thesis (Bonn U., 2020).
- [357] K. Lehmann and B. Stelzer, ‘The Fake Factor Method and its relation to the Matrix Method’, NIM A **1054**, 168376 (2023).
- [358] M. Bahmani, ‘Data-driven estimation of fake τ background in Higgs searches in ATLAS’, PoS **CHARGED2018**, 018 (2019).
- [359] J. Alison, ‘The Road to Discovery: Detector Alignment, Electron Identification, Particle Misidentification, WW Physics, and the Discovery of the Higgs Boson’, CERN-THESIS-2012-295, PhD thesis (Pennsylvania U., 2012).
- [360] ATLAS Collaboration, ‘Estimation of backgrounds from jets misidentified as τ -leptons using the Universal Fake Factor method with the ATLAS detector’, Submitted to EPJ C Feb. 2025, arXiv:2502.04156. (2025).
- [361] ATLAS Collaboration, *Measurement of the tau lepton reconstruction and identification performance in the ATLAS experiment using pp collisions at $\sqrt{s} = 13$ TeV*, ATLAS-CONF-2017-029, 2017.
- [362] ATLAS Collaboration, ‘Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC’, EPJ C **83**, 982 (2023).
- [363] G. Avoni et al., ‘The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS’, JINST **13**, P07017 (2018).

- [364] S. Bailey et al., ‘Parton distributions from LHC, HERA, Tevatron and fixed target data: MSHT20 PDFs’, EPJ C **81**, 341 (2021).
- [365] T.-J. Hou et al., ‘New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC’, Phys. Rev. D **103**, 014013 (2021).
- [366] MadGraph Collaboration, *What are the default dynamic factorization and renormalization scales in MadEvent?*, MadGraph wiki, last modified: 2012-05-16, accessed: 2025-06-27, <https://cp3.irmp.ucl.ac.be/projects/madgraph/wiki/FAQ-General-13>, 2012.
- [367] K. Bierwagen, U. Blumenschein and A. Quadt, ‘Bayesian Unfolding’, in PHYSTAT 2011 (2011), pp. 260–263.
- [368] G. D’Agostini, ‘A multidimensional unfolding method based on Bayes’ theorem’, NIM A **362**, 487–498 (1995).
- [369] K. Cranmer and A. Held, *Building and steering template fits with cabinetry*, Zenodo:4627038, 2021.
- [370] L. Heinrich et al., ‘Pyhf: pure-python implementation of histfactory statistical models’, Journal of Open Source Software **6**, 2823 (2021).
- [371] R. Barlow and C. Beeston, ‘Fitting using finite Monte Carlo samples’, Comput. Phys. Commun. **77**, 219 (1993).
- [372] F. James and M. Roos, ‘Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations’, Comput. Phys. Commun. **10**, 343–367 (1975).
- [373] G. Cowan et al., ‘Asymptotic formulae for likelihood-based tests of new physics’, EPJ C **71**, 1554 (2011); Erratum: EPJ C **73**, 2501 (2013).
- [374] A. L. Read, ‘Presentation of search results: the CL_S technique’, J. Phys. G **28**, 2693 (2002).
- [375] ATLAS Collaboration, *Summary Plots from ATLAS Searches for Pair-Produced Leptoquarks*, ATL-PHYS-PUB-2022-012, 2022.
- [376] ATLAS Collaboration, ‘Interpretations of the ATLAS measurements of Higgs boson production and decay rates and differential cross-sections in pp collisions at $\sqrt{s} = 13$ TeV’, Journal of High Energy Physics **2024**, 97 (2024).
- [377] I. Brivio, ‘SMEFTsim 3.0 —a practical guide’, JHEP **2021**, 73 (2021).
- [378] C. Hays et al., ‘On the impact of dimension-eight SMEFT operators on Higgs measurements’, JHEP **2019**, 123 (2019).
- [379] ATLAS Collaboration, *Combined effective field theory interpretation of Higgs boson and weak boson production and decay with ATLAS data and electroweak precision observables*, ATL-PHYS-PUB-2022-037, 2022.
- [380] H. Georgi, ‘Generalized dimensional analysis’, Physics Letters B **298**, 187–189 (1993).
- [381] A. Buckley et al., ‘Testing new physics models with global comparisons to collider measurements: the Contur toolkit’, SciPost Phys. Core **4**, 013 (2021).
- [382] J. M. Butterworth et al., ‘Constraining new physics with collider measurements of Standard Model signatures’, JHEP **03**, 078 (2017).
- [383] J. Y. Araz, ‘Spey: smooth inference for reinterpretation studies’, SciPost Phys. **16**, 032 (2024).
- [384] CMS Collaboration, ‘Observation of τ Lepton Pair Production in Ultraperipheral Pb–Pb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV’, Phys. Rev. Lett. **131**, 151803 (2023).
- [385] K. Johal et al., *Electroless Nickel / Immersion Gold process technology for improved ductility of flex and rigid-flex applications*, Report (National Center for Manufacturing Sciences, 2005).