Entwicklung einer hochgranularen Auslese für eine Zeitprojektionskammer und ihre Vorteile im Rahmen von Teilchenidentifikation

Dissertation
zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik und
Naturwissenschaften Fachbereich Physik
der Universität Hamburg

vorgelegt von

Ulrich Einhaus

Hamburg

2021

Gutachter der Dissertation:	Prof. Dr. Johannes Haller Dr. Ties Behnke
Zusammensetzung der Prüfungskommision:	Prof. Dr. Johannes Haller Dr. Ties Behnke Prof. Dr. Daniela Pfannkuche Prof. Dr. Gudrid Moortgat-Pick Prof. Dr. Erika Garutti
Vorsitzende der Prüfungskommission:	Prof. Dr. Daniela Pfannkuche
Datum der Disputation:	24.11.2021
Vorsitzender Fach-Promotionsausschuss PHYSIK:	Prof. Dr. Günter H. W. Sigl
Leiter des Fachbereichs PHYSIK:	Prof. Dr. Wolfgang J. Parak
Dekan der Fakultät MIN:	Prof. Dr.-Ing. Norbert Ritter

Für Oma

"I'm the Doctor!"
- The Doctor

"I am an idiot!"
- The Doctor

Abstract

This thesis adressess advanced particle identification techniques, in particular via dE/dx, in a particle detector at a future Higgs factory.

The International Linear Collider (ILC) is a proposed e^+e^- collider with a center-of-mass energy of up to 500 GeV which will provide a clean environment to perform precision measurements of the Standard Model of particle physics. This includes an unprecedented survey of the electroweak sector and in particular the Higgs boson. New detector concepts are necessary to enable the required precision measurements. One proposed concept is the International Large Detector (ILD), which has a Time Projection Chamber (TPC) as its central tracker. Compared to modern silicon-based detectors, the TPC allows continuous tracking in its gaseous sensitive volume, which provides excellent pattern recognition and tracking efficiency, as well as a measurement of the specific energy loss dE/dx. The specific energy loss of a charged particle depends on its momentum and mass and can be used for particle identification (PID). This work studies dE/dx in its different aspects in a particle detector, from a novel hardware approach to a TPC readout, through detailed simulations how such an approach would affect PID, through a full-ILD simulation of particle identification in particular via dE/dx but also in combination with time-of-flight (TOF), to the application of these PID algorithms to full-ILD physics analyses and its effect on the physics observables.

In the first part, the development of a novel approach of a TPC readout is presented. Its goal is to combine the flexibility of a pad-based readout with the possible high granularity of a pixel-based readout. It uses the direct readout of a pad plane on a printed circuit board (PCB) via a pixel ASIC. The challenging bonding process of this chip to the PCB is laid out, together with lessons learned for a future continuation of this approach. The noise in the bonded system is measured and compared to a simulated signal, which shows that a S/N ratio of more than 10 can be achieved and this approach is viable.

Next, the advantages of a high granularity TPC readout for PID are studied in an ab-initio TPC simulation. The conventional method to measure dE/dx, charge summation, improves moderately with increasing granularity. Below a few 100 µm pad size, the cluster counting approach becomes effective, which reconstructs the number of primary ionisation clusters. This number has a better correlation to the particle's momentum than its summed charge, allowing for an increase in PID, which is quantified and laid out. An extrapolation to the ILD TPC shows that with square pads of 1 mm height the effective dE/dx resolution would improve by about 20 % with charge summation, and by another 20 % with cluster counting, requiring a pad height below 300 µm.

Furthermore, tracking-based PID is studied in a full-ILD simulation, focusing on dE/dx, but also including V0-finding and TOF. A dE/dx calibration strategy for MC productions is laid out. The dE/dx performance is combined with TOF, highlighting its complementarity.

Finally, the investigated PID is applied to physics analyses of ILC events in full-ILD simulation. Using TOF, kaons produced in the ILC collisions can be selected and the kaon mass can be measured, with the remaining statistical uncertainty depending on the

assumed timing resolution. Given the luminosity of the full ILC programme at 500 GeV of $4\,\mathrm{ab}^{-1}$, the statistical error on the kaon mass for a realistic timing performance would suffice to resolve the current tension of the two leading existing measurements. In the second analysis, tracking-based PID is applied to the flavour tag in hadronic W decays to separate first from second quark generation decays. It is shown to provide additional information independent of the LCFIPlus flavour tag, which is currently used as a default. This allows to select second generation decays with higher efficiency, increasing the usable statistics at high required purities by several 10 %. The quantitative gain depends strongly on the actually achievable dE/dx resolution: with the ILD design dE/dx resolution of 4.5 % this increase is about 27 % at a purity of 0.99, with a worse dE/dx resolution of 7 % it would shrink to around 10 %. If on the other hand a dE/dx resolution of 2.6 % could be achieved, i.e. with the novel readout technology pioneered in this thesis, the increase would even reach 45 %.

Zusammenfassung

Diese Doktorarbeit betrachtet fortgeschrittene Techniken zur Teilchenidentifikation, insbesondere mittels dE/dx, in einem Teilchendetektor an einer zukünftigen Higgs-Fabrik.

Der International Linear Collider (ILC) ist ein vorgeschlagener e^+e^--Collider mit einer Schwerpunktsenergie von bis zu 500 GeV, der eine saubere Umgebung für Präzisionsmessungen des Standarmodells der Teilchenphysik ermöglichen wird. Dies umschließt eine beispiellose Vermessung des elektroschwachen Sektors und insbesondere des Higgs-Boson. Neue Detektorkonzepte sind notwendig, um die dafür notwendigen Präzisionsmessungen zu ermöglichen. Ein vorgeschlagenes Konzept ist der International Large Detector (ILD), in dem als zentrale Spurkammer eine Zeitprojektionskammer (TPC) zum Einsatz kommt. Im Vergleich zu modernen Silizium-basierten Detektoren ermöglicht die TPC in ihrem gasförmigen sensitiven Volumen eine kontinuierliche Spurverfolgung, die eine exzellente Strukturerkennung und Spurrekonstruktionseffizienz sowie eine Messung des spezifischen Energieverlustes dE/dx liefert. Der spezifische Energieverlust eines geladenen Teilchens hängt von seinem Impuls und seiner Masse ab und kann zur Teilchenidentifikation (PID) genutzt werden. Diese Arbeit untersucht dE/dx in seinen verschiedenen Aspekten in einem Teilchendetektor, von einem neuartigen Ansatz einer TPC-Auslese, über detaillierte Simulationen, wie sich ein solcher Ansatz auf die PID auswirkt, über eine gesamt-ILD-Simulation von Teilchenidentifikation insbesondere mittels dE/dx, aber auch in Kombination mit Flugzeitmessung, bis hin zur Verwendung dieser PID-Algorithmen in gesamt-ILD-Physikanalysen und zu ihrer Auswirkung auf die Observablen.

Im ersten Teil wird die Entwicklung eines neuartigen Ansatzes einer TPC-Auslese präsentiert. Sein Ziel ist die Kombination der Flexibilität einer Pad-basierten Auslese mit der Möglichkeit hoher Granularität einer Pixel-basierten Auslese. Er nutzt die direkte Auslese einer Padebene auf einer Leiterplatte (PCB) mittels eines Pixel-ASIC. Der herausfordernde Bonding-Prozess dieses Chips auf das PCB wird beschrieben zusammen mit den Lehren in Hinblick auf eine zukünftige Fortsetzung dieses Ansatzes. Das Rauschen im gebondeten System wird gemessen und mit einem simulierten Signal verglichen, was zeigt, dass ein S/N-Verhältnis von mehr als 10 erreicht werden kann und dieser Ansatz daher gangbar ist.

Als nächstes werden die Vorteile einer hoch granularen TPC-Auslese für PID in ab-initio-Simulationen untersucht. Die konventionelle Methode dE/dx zu messen, die Ladungssummierung, wird mit höherer Granularität moderat besser. Unterhalb einer Padgröße von einigen 100 µm wird der Cluster-Zählungsalgorithmus effektiv, der die Anzahl der primären Ionisations-Cluster rekonstruiert. Diese Anzahl ist stärker mit dem Teilchenimpuls korreliert als die Ladungssumme, wodurch eine verbesserte PID ermöglicht wird, was quantifiziert und dargelegt wird. Eine Extrapolation zur ILD TPC zeigt, dass mit quadratischen Pads von 1 mm Höhe die effektive dE/dx-Auflösung um etwa 20 % verbessert würde, und um weitere 20 % mit Cluster-Zählung, was eine Padhöhe von unter 300 µm erfordern würde.

Weiterhin wird Spur-basierte PID in einer gesamt-ILD-Simulation untersucht, mit einem Fokus auf dE/dx, aber auch unter Berücksichtigung von V0-Ermittlung und

Flugzeitmessung (TOF). Eine dE/dx-Kalibrationsstrategie für MC-Produktionen wird dargelegt. Die dE/dx-Leistungsfähigkeit wird mit der von TOF kombiniert und ihre Komplementarität hervorgehoben.

Schließlich wird die untersuchte PID auf Physikanalysen von ILC-Events in gesamt-ILD-Simulationen angewandt. Mittels TOF können Kaonen, die in ILC-Kollisionen erzeugt werden, ausgewählt und die Kaonmasse gemessen werden, wobei die verbleibende statistische Unsicherheit von der genutzten Zeitauflösung abhängt. Mit der Luminosität des vollständigen ILC-Programms bei 500 GeV von $4\,\mathrm{ab}^{-1}$ würde der statistische Fehler auf die Kaonmasse bei einer realistischen Zeitauflösung ausreichen, um die derzeitige Spannung zwischen den zwei führenden existierenden Messungen aufzulösen. In der zweiten Analyse wird Spur-PID im Flavour-Tag in hadronischen W-Zerfällen verwendet, um Zerfälle der ersten Quark-Generation von solchen der zweiten zu unterscheiden. Es wird gezeigt, dass dies zusätzliche Information liefert, die unabhängig vom LCFIPlus Flavour-Tag ist, der derzeit als Standard verwendet wird. Dies ermöglicht, Zerfälle der zweiten Generation mit höherer Effizienz auszuwählen, was die nutzbare Statistik bei hohem erforderten Reinheitsgrad um mehrere 10 % erhöht. Diese Steigerung hängt dabei stark von der erreichbaren dE/dx-Auflösung ab: Mit der Design-dE/dx-Auflösung des ILD von 4.5 % ist der Anstieg etwa 27 % bei einem Reinheitgrad von 0.99, mit einer schlechteren dE/dx-Auflösung von 7 % würde er auf 10 % zurückgehen. Wenn andererseits eine verbesserte dE/dx-Auflösung von 2.6 % erzielt werden könnte, d. h. mit der neuartigen Auslesetechnologie, die in dieser Arbeit vorgestellt wird, dann würde der Anstieg sogar 45 % erreichen.

Own Contributions

I lead the collaborative effort of the development of the Ropperi system detailed in chapter 5. I developed the detailed concept of the Ropperi boards and coordinated between the collaborators the technical implementation and production. Technical assistance regarding the PCB technical drawings and the bump bonding were provided by University Bonn and the Karlruhe Insitute for Technology (KIT). University Bonn, as member of the Medipix 2 collaboration, also provided the Timepix ASICs, as well as the ASIC readout system and software. I lead the data taking campaign in the bonding lab immediately after bonding, and took additional separate data at a later point. I analysed the data and assessed the system performance. I regularly presented updates to the work in meetings of the LCTPC collaboration and communicated the development of the Ropperi system within the linear collider community at two Linear Collider Workshops [1, 2].

I am the developer and analyser of the high granularity TPC simulation laid out in chapter 6. I put together the simulation chain, which is based on the MarlinTPC framework [3–5]. I adapted the individual components of this chain for the purposes of this thesis, in particular to enable a variable granularity. I regularly presented the analysis results in meetings of the LCTPC collaboration and they are currently being published in an LCTPC collaboration paper currently in preparation.

I analysed the PID capabilities given in the full ILD simulation detailed in chapter 7. I adapted the `ComputedEdxProcessor` and performed the dE/dx calibration for the 2020 MC production. This adaption and calibration have become part of iLCSoft and its default configuration and have thus been actively used in a number of analyses which utilise dE/dx. The time-of-flight assessment is based on code provided by Sukeerthi Dharani in the context of her bachelor thesis [6]. I adapted this code and created the combination with dE/dx. I regularly presented this PID assessment in meetings of the ILD collaboration. I coordinated the dE/dx calibration with the ILD and LCTPC collaborations. My PID assessment has become part of the standard ILD performance presentation and as such the corresponding plots created by me are used or referenced in a number of publications, in particular the ILD Intermediate Design Report [7].

I developed and performed the analyses of the kaon mass and the hadronic W-decay, chapter 8. I regularly presented the results in meetings of the ILD collaboration. The kaon mass analysis was highlighted in the DESY Particle Physics Annual Report 2020 [8]. I presented the hadronic W-decay analysis at the EPS-HEP 2021 (European Physical Society conference of high energy physics), including proceedings [9].

Contents

1. Introduction

No generation of humanity has ever known as much about the nature of the universe as the present one. We can describe molecules and massive machines, gluons and galaxies. We extrapolated the expansion of the universe to the Big Bang 13.8 billion years ago and to infinite expansion as its ultimate fate. The Standard Model of particle physics (SM), together with general relativity, is at the core of this description and provides the most fundamental theory to describe the existence and interaction of all known matter. At the same time, it is considered the most precise physics model ever created, allowing to calculate particle properties to a precision of many orders of magnitude down to the level of quantum fluctuations. In 2012, a candidate largely seen as the last puzzle piece of the Standard Model, the Higgs boson, was found at the Large Hadron Collider (LHC) [10, 11] at CERN. This discovery at the largest machine ever built in a massive collaborative effort of scientists and technicians from all around the world was celebrated by many as the triumphant completion of the Standard Model.

And yet, we know there's more.
Dark Matter dominates gravity at large scales, galaxy rotation cannot be explained without it, neither can structure formation in the early universe. Yet, the Standard Model does not offer a candidate particle which could make up Dark Matter. Many extended theories include new particles which could account for the invisible matter, but none has been found so far. The Standard Model can also not account for the overabundance of matter in the universe. While most particles have annihilated in the early universe, the known CP violation in the SM is insufficient to explain the remaining matter-antimatter asymmetry. An adaptation of the theory is inevitable. Another hint to new physics is the Hierarchy Problem. In the Standard Model, the mass of the Higgs boson receives contributions from loop corrections from all massive particles, including those at the Planck scale, if we believe that there must be a Grand Unifying Theory making up all known forces and connecting them with gravity. Why would those contributions exactly cancel to give a Higgs mass at the electroweak scale? Naturalness suggests a mechanism close to this scale.

To address this variety of open questions, to find new particles, new laws of nature, in short Physics Beyond the Standard Model (BSM), we propose a new particle accelerator, tackling the energy frontier. The International Linear Collider (ILC) will provide electron-positron collisions at unprecedented center-of-mass energies. It will create a large number of Higgs bosons, which are at the root of many of the open questions and will work as a window into BSM. Colliding elementary leptons, as opposed to composite protons, provides a clean environment with low background, allowing detectors to be optimised for precision rather than readout speed or radiation hardness.

If the ILC – or any other Higgs factory – is built, billions will be spent to provide these particle collisions. They should be used as well as possible and all available information should be gathered by the detector. The two proposed detector concepts for the ILC follow this calling. The International Large Detector (ILD) and the Silicon Detector (SiD) are typical multi-purpose detectors, with µm-vertex tracking and high-efficiency tracking for momentum measurement. Unprecedented calorimeter granularity

will enable the first full implementation of particle flow at the energy frontier - all inside the strong solenoid. Muon chambers in the yoke and a tiered high-granularity forward system providing nearly 4π-solid angle coverage complement the concept.

But ILD goes one step further: its central tracker, the time projection chamber (TPC), measures not only the position and curvature of charged particles to determine their momentum, but also their specific energy loss en passant. This allows to determine the mass and thus flavour of the particles, and to reconstruct to much better precision what exactly happened in those precious interactions in the heart of the machine. It remains to be shown what these particle identification capabilities mean in terms of physics observables, and how the specific energy loss performance in ILD compares to other PID measurements. Even further, would an improved performance of the ILD TPC beyond its current aim lead to a better exploitation of the collisions, and if yes, how could this be achieved? This work will shed light on these questions.

In chapter 2, the theoretical background of the Standard Model, ILC and ILD are described, with a focus on PID-sensitive physics channels. The technical details of a TPC and the technology for further readout improvements are laid out in chapter 3. In chapter 4 the software framework used throughout this work is introduced.

A key factor in the performance of the specific energy loss measurement in a TPC, which can still be significantly improved upon, is the granularity, i.e. the segmentation of the anode planes into readout channels, called pads or pixels. Decreasing the pad size, which currently lies in the order of several mm, and thus sampling each track with higher statistics allows for an improved dE/dx measurement with the conventional method of charge summation. The charge, however, is governed by a – relatively broad – Landau distribution. An alternative approach, the so-called cluster counting, aims to count the number of primary ionisation events of an incident particle, governed by a much more narrow Gaussian. This is only possible, if the pad size can be reduced below the typical ionisation step length in the order of a few 100 µm.

Within the Linear Collider TPC (LCTPC) collaboration, several readout systems for the ILD TPC are being developed. The conventional pad-based readout provides a granularity of about 6 mm. On the other hand, the pixelised TPC readout has a pitch of only 55 µm, but it comes with the disadvantages of vastly more readout channels and less flexibility, since it uses ASICs as anode instead of printed circuit boards (PCBs), which are much cheaper to develop and produce. In this work, a novel hybrid readout system was developed, built and tested. It focuses on intermediate pad sizes of a few 100 µm in order to resolve primary ionisation clusters, but to provide at the same time the flexibility of a dedicated PCB pad plane and to limit the number of active channels to what is necessary for a high performance. This system and the results of its characterisation are laid out in chapter 5.

Given such a hybrid system with a flexible pad size, detailed simulation studies were performed in order to show the dependence of the PID performance of a TPC on the readout granularity. The simulation shows the improvement high granularity can provide, both with the conventional charge summation, as well as with cluster counting, revealing the transition region between the two and connecting the test beam results of

the existing pad-based and pixel-based readout systems. In chapter 6, the simulation and reconstruction chain are laid out in detail and the results on PID performance are presented.

Finally, the impact of PID on physics observables was studied. This required first a characterisation of the dE/dx performance of the ILD TPC based on full-detector simulation, provided in chapter 7. A calibration of the dE/dx parameters in this simulation was conducted in order for the ILD simulation to provide a performance comparable to test beam results by the LCTPC collaboration. In addition, the full-detector simulation allows to relate the PID performance based on dE/dx to the one based on time-of-flight (TOF), which comprises a more recent development in the ILD context.

Two physics benchmarks were selected to show where PID leads to improvements of measurements, or even entirely new possibilities, which are laid out in chapter 8. The most precise measurements of the mass of the charged kaon were done in the early 1990s via spectroscopy of kaonic atoms. The two leading measurements disagree significantly and limit the overall uncertainty on the mass to 13 keV. TOF and dE/dx PID allow to select kaons from arbitrary events at the ILC, and their mass can be directly measured via TOF. This work provides an estimate, what statistical uncertainty on the kaon mass the ILC could achieve, depending on the TOF time resolution. This would be the first major contribution to the kaon mass in decades and done in a completely independent way.

Another novel use-case for kaon identification is flavour tagging. Conventional flavour tagging uses a number of jet properties, of which the existence and position of secondary and tertiary vertices is the most powerful. Jets that originate from s and c quarks produce additional high-momentum kaons, which can be identified via dE/dx. As benchmark, this was studied in hadronic decays of W bosons, separating d/u decays from c/s decays. Adding PID information to the conventional flavour tag improves the separation curves and increases the usable statistics significantly. This outcome depends strongly on the dE/dx resolution used in the simulation.

This work draws a direct line from the TPC hardware capabilities in terms of granularity shown in the hardware studies, via the improved dE/dx performance in a detailed TPC simulation, to the resulting increase in physics potential by the flavour-tag shown in full-detector simulation. Similarly, for TOF, the impact of the available time resolution on a physics benchmark is laid out, albeit in view of pending work on TOF algorithms. These examples provide an assessment of what hardware developments would be required in the not-too-distant future in order to allow ILD to utilise the valuable particle collisions and rare events even better and to make even more precise measurements.

2. The International Linear Collider

2.1. The Standard Model of Particle Physics

The Standard Model of Particle Physics is generally considered one of the most successful models to describe our universe in history. It consists of matter particles and fundamental interactions, which are a product of the underlying Quantum Field Theory (QFT). The associated particles are shown in Figure 1. Many properties of these particles including their interaction cross sections are predicted by theory and were confirmed with precision over the last decades. The discovery of the Higgs boson in 2012 can be seen as one of the most significant successes of the SM after its prediction in 1964. The theory covers electro-magnetic, weak and strong interactions, but is yet missing a grand unification of the three, and does not include gravity at all. Furthermore, there is strong evidence for new physics, beyond the Standard Model, be it Dark Matter, the hierarchy problem or matter-antimatter asymmetry.

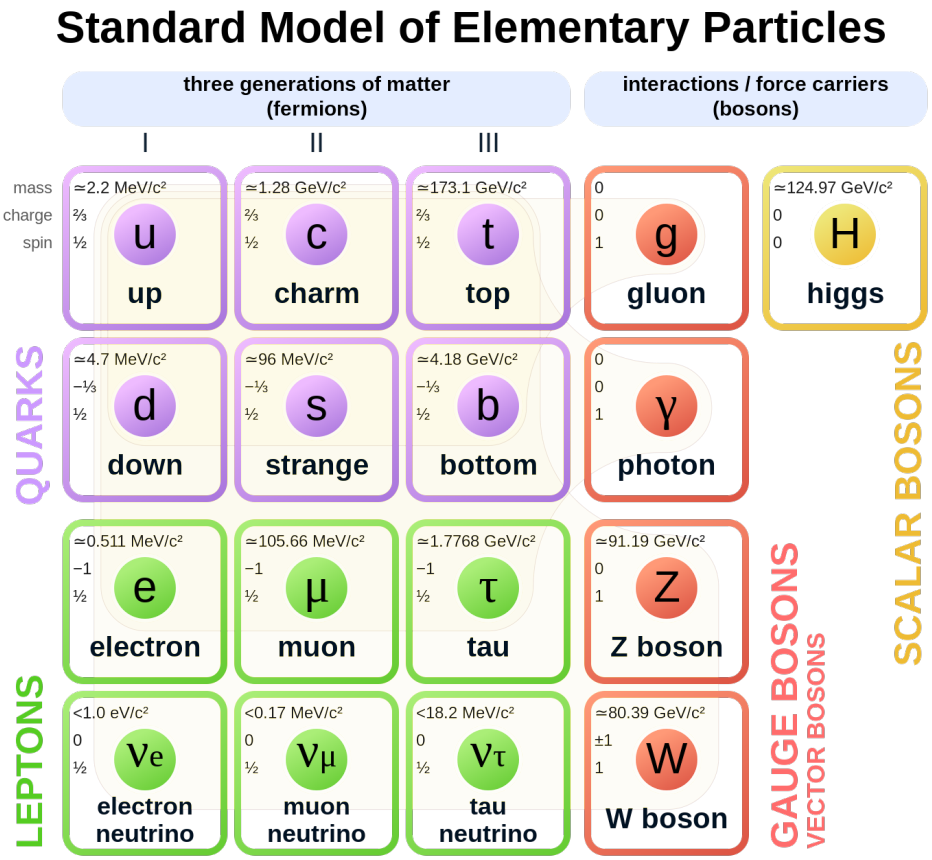

Figure 1: Particle content of the Standard Model, from [12].

Cosmological observations of stars and galaxies, in particular their rotational curves and gravitational lensing effects, suggest that there should be a lot more mass in and around galaxies than is visible. Therefore, this material is called Dark Matter (DM), and there should be about 5 times as much DM compared to our regular baryonic matter. Most theories assume DM to consist of new elementary particles, which could be found with particle colliders, given their mass and interaction cross section suffice.

The hierarchy problem deals with the mass of the Higgs boson: since it is a scalar, the Higgs receives loop corrections to its bare mass value from all physics scales, including up to the Planck scale at $10^{19}\,\mathrm{GeV}$, if the theory includes gravity. While it is possible to renormalise such contributions, this still requires considerable fine-tuning of positive and negative contributions to achieve a total value of the observed 125 GeV, more than 16 orders of magnitude smaller, which is often considered unnatural. To avoid this, different mechanisms have been suggested. One could add compact extra dimensions, which are only accessible for gravity, and would effectively scale it to the same scale of the other fundamental interactions. Another possibility is to counter each loop contribution to the Higgs mass by an inverse from a supersymmetrical particle with opposing spin statistic. This means, Supersymmetry (SUSY) adds a counterpart boson to each fermion and vice versa, keeping most particle properties the same. These new particles would have been found by now if they had the same mass, so the symmetry must be broken in this aspect, allowing for masses which are larger than currently accessible, but still small enough to consider their contributions to the Higgs mass natural. This mass scale is considered to be at the few-TeV range. In addition, in most versions of SUSY there is a lightest, stable supersymmetric particle, which would be a natural candidate for the Dark Matter.

As far as we can observe, our universe consists of only matter, but (nearly) no anti-matter. Given their fundamental symmetry in the SM, this begs the question, why most, but not all matter and anti-matter annihilated into photons after the Big Bang. Given a ratio of the baryon-to-photon number in the universe of 10^{-7}, this means that the asymmetry is small, but still significantly larger than can be explained through known CP-violation in the electroweak sector. The nature of this asymmetry thus remains a mystery to be solved.

One way of addressing these questions is to provide ever higher energy collisions and attempt to produce the particles associated with BSM physics - or exclude them. Another way is to measure quantities very precisely at lower energies, in particular lower than the mass of new particles. New physics influences particle scattering events via virtual corrections to the scattering amplitudes. With a sufficiently accurate detector and a clean environment with a low background, the impact of these corrections, e.g. cross sections or angular distributions, can be observed. Individual measurements can provide evidence for new physics, but the most general approach is to perform a global fit over all SM observables and compare the result to theoretical predictions. This way, effects that impact several SM observables and might not be considered significant in only a single measurement may become significant by the overall combination. This requires, however, a high precision for all observables, including rather well-known ones.

One example are the properties of commonly produced particles in collider experi-

ments, like pions, kaons and protons. While protons are stable and have been measured to near arbitrary precision, this is not the case for the short lived pions and kaons. Their properties like mass, lifetime and decay channels, are still attached with a finite uncertainty, which should be reduced in order to improve the global fit, but also enable improved reconstruction algorithms dealing with particle decays in collider experiments. In particular the mass of the charged kaon has been a point of discussion for some decades now. In the 80s and 90s several experiments measured the mass of the charged kaon via x-ray spectra of kaonic atoms. In such an experiment, kaons from a conversion target are introduced into a material where they take the place of individual electrons in the atoms, forming kaonic atoms. Similar to electrons in regular atoms, the kaons form orbitals at discrete energy levels around the nucleus, and emit photons upon dropping to a lower level. Since these energy levels depend on the mass of the particle, different from electrons the kaons emit in x-ray, and the exact spectrum of these emissions can be used to calculate the mass of the charged kaon. According to the Particle Data Group [13], the most precise measurements were done by two groups in 1988 and 1991 [14, 15], and they disagree with each other significantly, leaving the uncertainty on the kaon mass at the level of 13 keV, compared to the uncertainty on the pion mass of 0.17 keV. Both of these values are also important inputs to lattice QCD calculations [16], including for the latest results on the determination of the hadronic contribution to the magnetic anomalous moment of the muon $g_\mu - 2$, like [17]. This result from the lattice calculation is in moderate disagreement with the leading theory calculations using the data-driven dispersion relations approach [18], but is closer to the recently updated experimental value [19] and challenges the long-standing tension between measurement and SM theory prediction on $g_\mu - 2$. The uncertainty on the lattice value is larger than the one based on dispersion, and further results from more lattice QCD groups are expected in order to resolve the disagreement of the different theory calculations. To enable more precise lattice QCD results in the future, a reduction in the kaon mass uncertainty is an important contribution.

Another test of the Standard Model is the comparison of expected and measured values of couplings. Also there, new physics can contribute to couplings between bosons and fermions at the level of higher-order loop corrections. One such collection of couplings is the CKM matrix, which summarises the quark transition factors via charged current, or in other words the couplings of the W boson to quark pairs. The unitary CKM matrix is part of the global fit and any deviations from experimental values could indicate underlying impacts from physics at higher energies. Even the unitarity of the CKM matrix itself is part of the theory and can be tested. The least well known element of this matrix is the central element V_{cs}, which is linked to the W decay into a pair of an s quark and a $\bar{\mathrm{c}}$ quark. This element is usually extracted by using the more precisely known values of the other matrix elements as well as of the decays into leptons, assuming unitarity. As an additional method, to perform a check on these assumptions, the W decays into s and $\bar{\mathrm{c}}$ can be tagged and compared to the ones into d and $\bar{\mathrm{u}}$, giving rise to an independent measurement of V_{cs}.

These questions and others can be addressed with particle colliders, providing unprecedented energies, intensities and precision. With sufficient energy and luminosity heavy new particles can be produced directly, like the Higgs boson at the LHC. But new physics also manifests itself in virtual loop corrections to effective couplings, which can be accessed at energies much below the mass of the new physics, but requires equally high luminosities and high precision in a clean environment. One accelerator which can provide these conditions is the International Linear Collider, currently under political consideration in Japan.

2.2. The International Linear Collider

The ILC [20–23], see Figure 2, is a proposed linear e^+-e^--collider, operating at center-of-mass energies of 250, 350 and 500 GeV, with a potential extension to 1 TeV. The beams have a polarisation of $\pm$ 80 % for electrons and $\pm$ 30 % for positrons, the latter being extendable to $\pm$ 60 %. The ILC is foreseen to have a peak instantaneous luminosity of $1.35 \cdot 10^{-34}\,\mathrm{cm}^{-1}\mathrm{s}^{-1}$ at a centre-of-mass energy of 250 GeV [7]. Further properties are summarised in Table 1. The full 20-year programme envisions an integrated luminosity of several ab^{-1}, with $2\,\mathrm{ab}^{-1}$ at 250 GeV and $4\,\mathrm{ab}^{-1}$ at 500 GeV being the default scenario, known as H-20 [24].

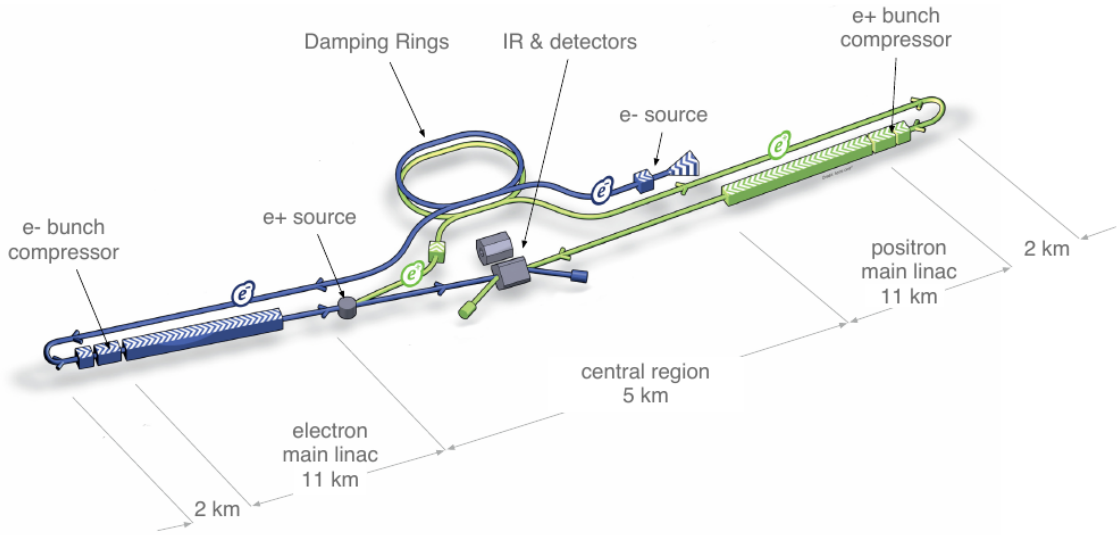

Figure 2: Scheme of the ILC, from [20].

By colliding elementary particles, the events will be rather clean compared to proton-proton collisions at the LHC, i.e. have a low fraction of low-momentum transfer collisions, only few background particles per bunch crossing, and no pile-up. This combination allows the detectors to run trigger-less, recording the detector activation of each bunch-crossing. In addition, most parts of the detector do not need to be particularly radiation-hard. ILC will operate with bunch trains: around 2000 bunches of 10^{10} electrons and positrons each collide in sequence with a spacing of about 300 ns, which in

Bunch population	N	$2 \cdot 10^{10}$
RMS bunch length	σ_z	300 µm
RMS horizontal beam size at IP	σ_x^*	474 nm
RMS vertical beam size at IP	σ_y^*	5.9 nm
Normalised horizontal emittance at IP	$\gamma\epsilon_x$	10 µm
Normalised vertical emittance at IP	$\gamma\epsilon_y$	35 nm
Electron polarisation	P_-	80%
Positron polarisation	P_+	30%
Luminosity	L	$1.8 \cdot 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$
Fraction of L in top 1% of energy	$L_{0.01}$	58%
Fractional RMS energy loss due to beamstrahlung	δ_{BS}	4.5%
Electron relative energy spread at IP	$\Delta p/p$	0.13%
Positron relative energy spread at IP	$\Delta p/p$	0.07%

Table 1: Beam parameters of the ILC 500 GeV stage, from [23].

total takes about 1 ms. Between bunch trains there is a gap of 199 ms, leading to a 5 Hz train operating mode. During most of this gap, the detector will be switched off to reduce energy consumption, and thus heat production and the necessity for cooling. With only moderate need for radiation-hardness and cooling, the detector components can be developed with more focus on precision and granularity. Colliding elementary particles also means having a fixed centre-of-mass energy, which is an invaluable constraint and allows to reconstruct missing momentum in an arbitrary direction in many physics analyses.

The main target for the 250 GeV stage of the ILC is precision Higgs physics. Through the Higgs-strahlung process a large number of Higgs particles will be produced and their properties investigated, in particular mass, coupling constants as well as visible and invisible branching fractions. It is a key feature of the ILC as a lepton collider to study the Higgs boson in a model-independent way via the Z-recoil method [25]. Running at 350 GeV will allow to scan the top-threshold and measure the top mass model-independently and with unprecedented precision. At 500 GeV W fusion will be the leading contribution to Higgs-particle production, shown in Figure 3, and the higher centre-of-mass energy will enable Higgs self-coupling processes to take place and be measured. At all energies, the large luminosity, clean conditions and unprecedented detector precision will allow for substantial improvements in electroweak precision measurements. Moreover, new physics can be detected in higher-order corrections of known physics processes, for which to be detected a high precision and background suppression is imperative. Many of these potential BSM processes again involve the Higgs boson.

There are currently two detectors proposed to operate at the ILC, one of which is the International Large Detector, ILD.

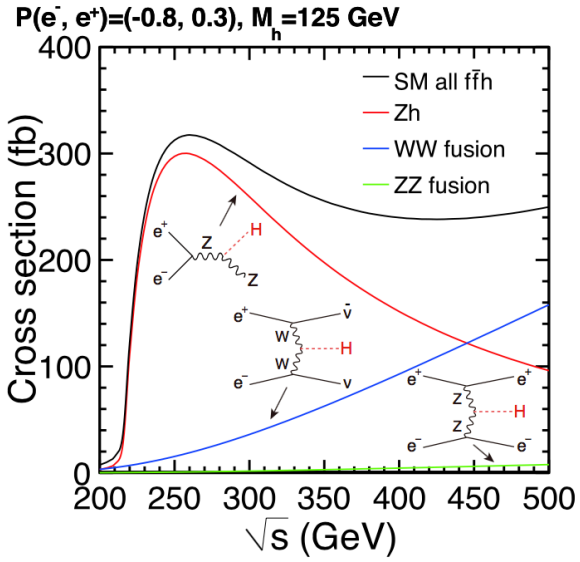

Figure 3: Higgs production cross section at the ILC for center-of-mass energies up to 500 GeV, from [21].

2.3. The International Large Detector concept

The ILD [7,26], shown in Figure 4, is a multi-purpose detector designed for collisions at the ILC. Its main targets are excellent momentum and energy resolution, effective particle flow, near full 4π-hermeticity and a very good particle identification capability. The silicon pixel vertex detector allows for vertex reconstruction up to a few µm from the IP with a track impact parameter resolution $\sigma_b < 5 \oplus 10/p \sin^{3/2}\theta$ µm as performance goal. The main tracker consists of a time projection chamber (TPC) with a silicon envelope. With a point resolution of $\sigma_r\phi < 100$ µm the TPC achieves a momentum resolution of $\sigma_{1/p_T} < 10^{-4}$ GeV^{-1}. In combination with the silicon layers which have a point resolution of $\sigma_r\phi < 5$ µm (Si pixel layer inside the TPC) and $\sigma_r\phi < 7$ µm (Si strip layer outside the TPC), respectively, the tracking system achieves $\sigma_{1/p_T} < 2 \cdot 10^{-5}$ GeV^{-1}. The tracking system is surrounded by the calorimeter system: a tungsten sandwich electro-magnetic calorimeter (ECal) with a granularity of $5\,\text{mm} \cdot 5\,\text{mm}$ and a steel sandwich hadronic calorimeter (HCal) with a granularity of $3\,\text{cm} \cdot 3\,\text{cm}$. This high granularity allows for an effective particle flow algorithm to be applied, leading to a jet energy resolution of $\sigma_{E_{jet}} < 3.5\%$ for $E_{jet} > 100$ GeV. The tracking and calorimeter system is placed fully inside of the solenoid magnet generating a magnetic field of $B = 3.5$ T. The muon system outside the solenoid serves also as tail catcher for the HCal and yoke.

ILD is optimised for the particle flow concept, which aims to reconstruct the energy of each particle and in particular its contribution to a jet in the sub-detector which is best suited due to its resolution. This means the energy of charged particles is measured via their momentum in the main tracker instead of the calorimeter, since the tracker's momentum resolution is much better than the calorimeter's energy resolution

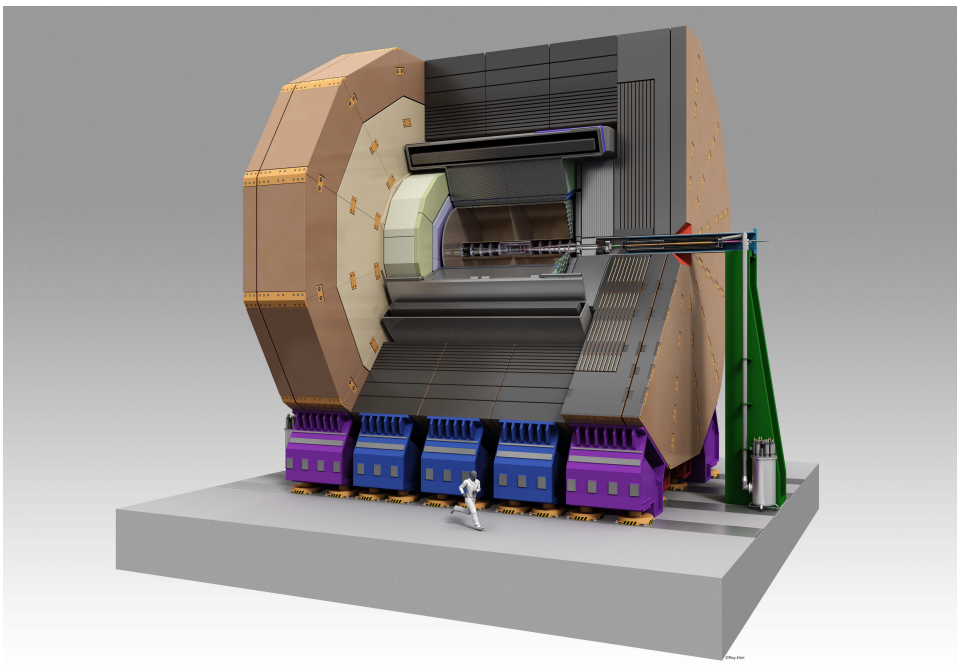

Figure 4: View of the ILD, from [7].

at typical particle energies at the ILC. Neutral particles, which do not leave a track, still need to be measured in the calorimeter, and in order to do this, their contributions need to be separated from those from the charged particles. This requires a high calorimeter granularity in order to correctly identify the respective clusters. In the optimal case, charged particles are measured in the tracker, photons in the ECal and neutrons and K_L^0 particles in the HCal, which is shown in Figure 5. This way, the uncertainty of the contribution of charged particles to the jet energy can be significantly reduced compared to classical calorimeters where the jet energy is interpreted as sum of the energy depositions in the ECal and HCal, and refined by momentum measurement, which is known as energy flow. ALEPH achieved a jet resolution of hadronic Z decays of about $\sigma(E)/E \approx 60\%/\sqrt{E/GeV}$ [27], while ATLAS [28], with an approach that is already closer to the full particle flow as described here, achieves about $\sigma(E)/E \approx 100\%/\sqrt{E/GeV}$, but with significant pile-up and at the higher average jet energies of the LHC. This results in a relative resolution of better than 10 % above a transverse jet momentum of 100 GeV for ATLAS compared to better than 3.5 % for ILD, see Figure 6.

Particle flow does not only require highly granular calorimeters, but also excellent tracking efficiency and pattern recognition, which is provided by the TPC, along with charged particle identification via dE/dx.

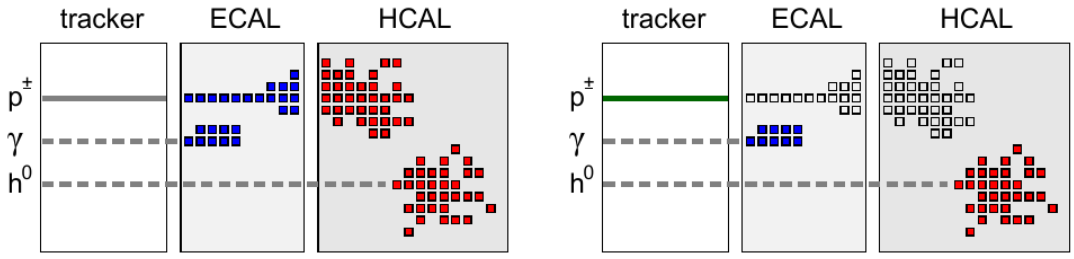

Figure 5: Illustration of the particle flow principle, from [29].

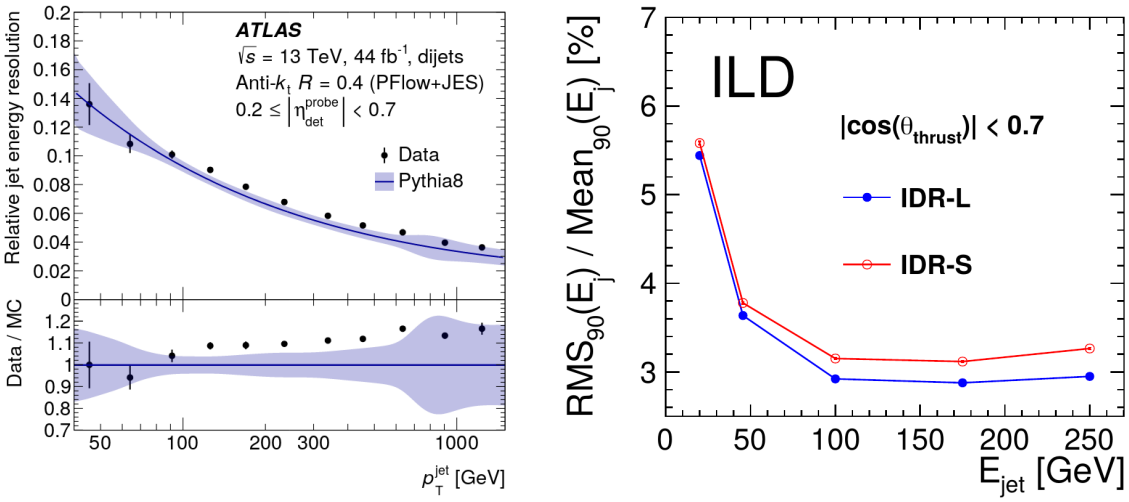

Figure 6: Jet energy resolution of ATLAS (left) from [28] and ILD (right) from [7].

3. The Time Projection Chamber

A time projection chamber is a gaseous detector for high energy particles. Particles which traverse the gas volume ionise the gas, generating ion-electron pairs. An electric field is applied to separate these pairs and to move the electrons to the readout side. Since typical ionisation densities in gas of O(10) cm^{-1} to O(100) cm^{-1} constitute a too small signal for state-of-the-art electronics, the electrons are first amplified above a segmented anode plane, nowadays mostly using micro pattern gas detectors (MPGDs). The original 3D-position of the ionisation can be deduced by combining the electron's arrival position on the anode with a time measurement, given a constant drift velocity in the gas volume. This position information is used to find and reconstruct tracks in the gas volume, and the amount of charge produced in the primary ionisation is used to measure the specific energy loss. The details of the working principle are described in this chapter.

3.1. Interactions of High Energy Particles with Matter

When a high energy charged particle traverses matter, kinetic energy of the particle is transferred to the medium, mostly via electromagnetic interaction in inelastic scattering with the electrons in the medium. Deep inelastic scattering with nuclei is rare and thus the typical energy transfer is much smaller than the kinetic energy of a particle at the GeV-scale. The interactions lead to ionisation and possible excitation of atoms and molecules in the medium, while the following de-excitation generates photons or further ionised electrons, depending on the gas constituents. Figure 7 shows the mean of the energy loss over a large range of $\beta\gamma$ of a muon. In the region of interest, between about 100 MeV and 100 GeV, the mean energy loss is described by the Bethe-Bloch formula given in Equation 1. It decreases rapidly first, reaches a minimum at about $E = 3m_0c^2$ for a particle with rest mass m_0 and then approaches a plateau, the so-called relativistic rise. A particle with an energy of $3m_0c^2$ is called a minimally ionising particle (MIP) in the narrower sense, but since the relativistic rise lies at a similar level as the minimum and significantly below the rises below $\beta\gamma = 1$ and above $\beta\gamma = 1000$, often all particles with energies in this region are called MIPs, e. g. in calorimetry. For lower energies, the shell structure of the gas atoms starts to play a role, and for higher energies, radiative effects like bremsstrahlung or Cherenkov radiation dominate the energy loss, so the energy loss is not described anymore by the Bethe-Bloch formula.

$$\left\langle -\frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left(\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right) \tag{1}$$

with:

Note that via

$$\beta^2 = \frac{1}{\frac{m^2c^2}{p^2} + 1}, \tag{2}$$

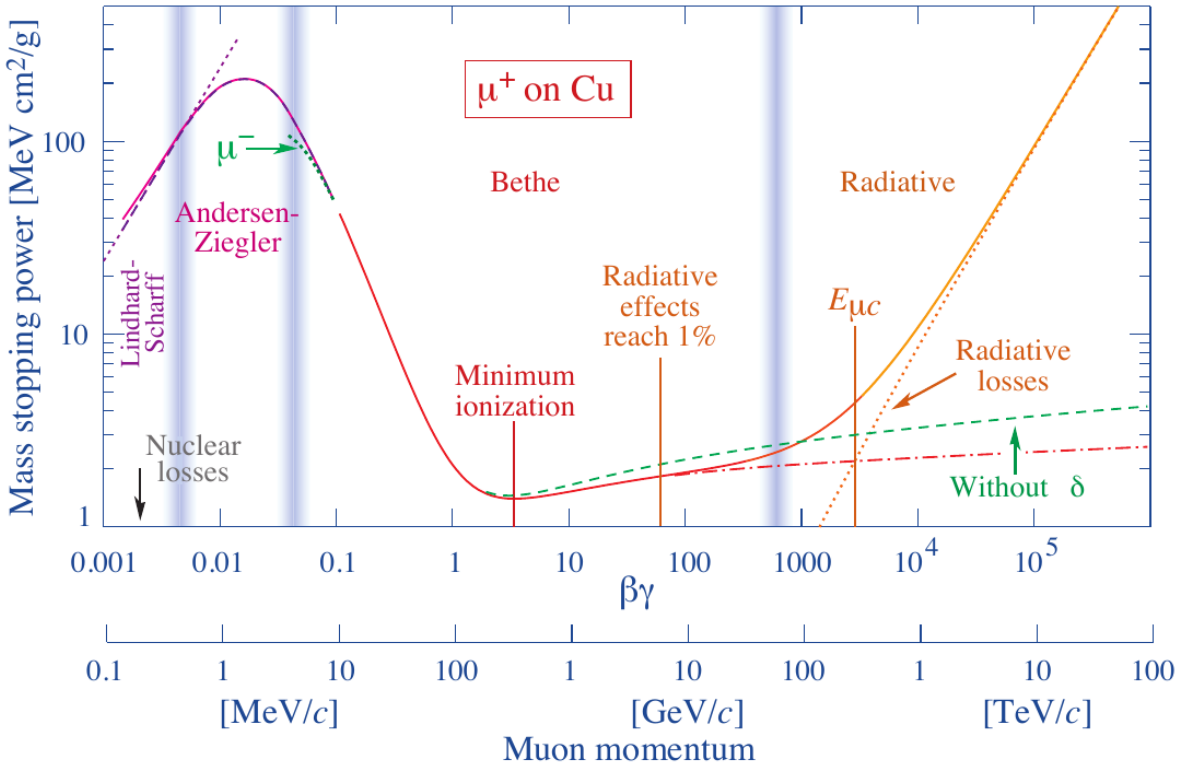

Figure 7: Mean energy loss of a charged particle from theoretical calculation, here as mass stopping power of a muon in copper, over a wide range of momenta, from [13].

z	charge of the incident particle
Z	atomic number of the absorber
A	atomic mass of the absorber
K	constant $= 4\pi N_A r_e^2 m_e c^2$
N_A	Avogadro's number $= 6.022 \cdot 10^{23}\,\mathrm{mol}^{-1}$
r_e	classical electron radius $= e^2/4\pi\epsilon_0 m_e c^2$
m_e	electron mass $= 0.511\,\mathrm{MeV}$
W_{max}	maximum possible energy transfer to an electron in a collision
I	mean excitation energy of the medium
$\delta(\beta\gamma)$	density effect correction to ionisation energy loss
β	velocity of the incident particle in units of c
γ	Lorentz factor of the incident particle $= (1-\beta^2)^{-1/2}$

Table 2: Bethe-Bloch formula with used quantities, following [13].

Equation 1 can be expressed as depending on momentum and mass of the incident particle. This allows for the principle capability to reconstruct a particle's mass and therefore species given a measurement of its momentum and energy loss in a medium.

For a given medium, the mean free path of an incident particle is a constant, and the number of interactions per depth of the medium is governed by a Poisson distribution. The energy transfer in each of these interactions, however, is characterised by a long tail to high energies, which can be approximated by a Landau distribution and is much wider than a Poisson distribution, compare Figure 8. This - theoretically infinite - tail is however cut off at the maximum energy transfer $\Delta E_{trans,max}$ in each interaction which, in the relativistic case with the energy of the incident particle $E \approx E_{kin} \approx pc$, is given by

$$\Delta E_{trans,max} = \frac{E^2}{E + \frac{m_0^2 c^2}{2m_e}}, \tag{3}$$

with the masses of the electron and the incident particle, m_e and m_0, respectively. An electron that receives enough energy in the ionisation process to cause its own short, but often visible track is called a δ-electron.

The energy transferred in the interaction does not only go into the ionisation but also into excitation, momentum transfer and radiation. Therefore, the average energy transfer necessary to generate one free electron is a medium constant, which is larger than its lowest ionisation energy.

3.2. The Time Projection Chamber

Ionising interactions as described above are utilised by a number of detector technologies, including the time projection chamber, or TPC. Its working principle is shown in Figure 9. A TPC consists of a gas-filled sensitive volume in which the ionisation happens. Electron-ion pairs are generated and separated via an applied electric field. The

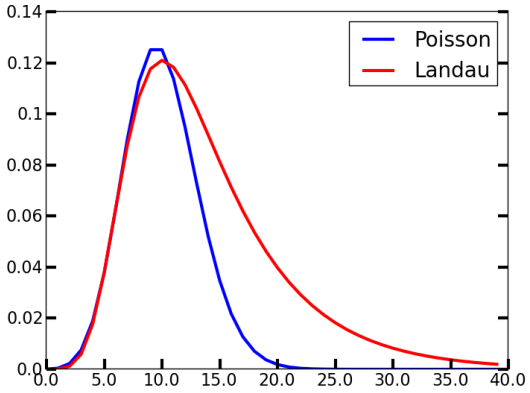

Figure 8: Examples of a Poisson and a Landau distribution, with the same most probable value and normalised to the same rise.

electrons are accelerated by the electric field towards the anode, but frequently collide with gas atoms causing them to have a constant average velocity. This electron movement is called drift. In TPCs with a drift length of more than a few cm, like the ILD TPC, a magnetic field parallel to the electric field is present. Through the Lorentz force the drifting electrons move in miniature helices around the magnetic field lines, which drastically limits the transverse diffusion. Generally, the gas of a TPC is a mixture of mostly a noble gas plus one or more organic components. The number of free electrons generated by a MIP in a gas depends on the ionisation energies of the gas molecules, dominated by the noble gas, and is a constant for the gas. This number is proportional to the gas pressure, and can be increased by operating the TPC at a high pressure. This results in typical ionisation densities of O(10) to O(100) electrons per cm track length, which is insufficient to be measured with current electronics. Therefore, the electrons need to be amplified when they arrive at the anode side.

In previous times (e.g. [31], [32]), this has been done with multiwire proportional chambers [33], with a wire-to-wire pitch of O(mm). To reduce distortions from so-called E×B-effects, when the electric field for the avalanche amplification and the magnetic field are not parallel, and to enable an increased granularity, micro pattern gaseous detectors (MPGDs) have been developed. The technology considered in this work is the gas electron multipliers (GEM) [34]. One GEM consists of a layer of non-conductive material, coated on both sides with a conductor, often copper. The foil is regularly perforated at a length scale of O(100 µm). When a high voltage is applied between the two copper-coated sides, a strong electric field is established in the holes, which is sufficient for avalanche amplification of traversing electrons. A microscopic image of a GEM foil as well as a simulation of the field and the amplification in a GEM hole is shown in Figure 10. GEMs can be stacked to increase the overall amplification, but the transport of electrons between the foils is not perfect and some losses occur [35].

Another techonology for amplification is the Micro Mesh Gaseous Detector (MicroMegas) [36]. Here, a thin metal foil is placed parallel to the anode, thus creating

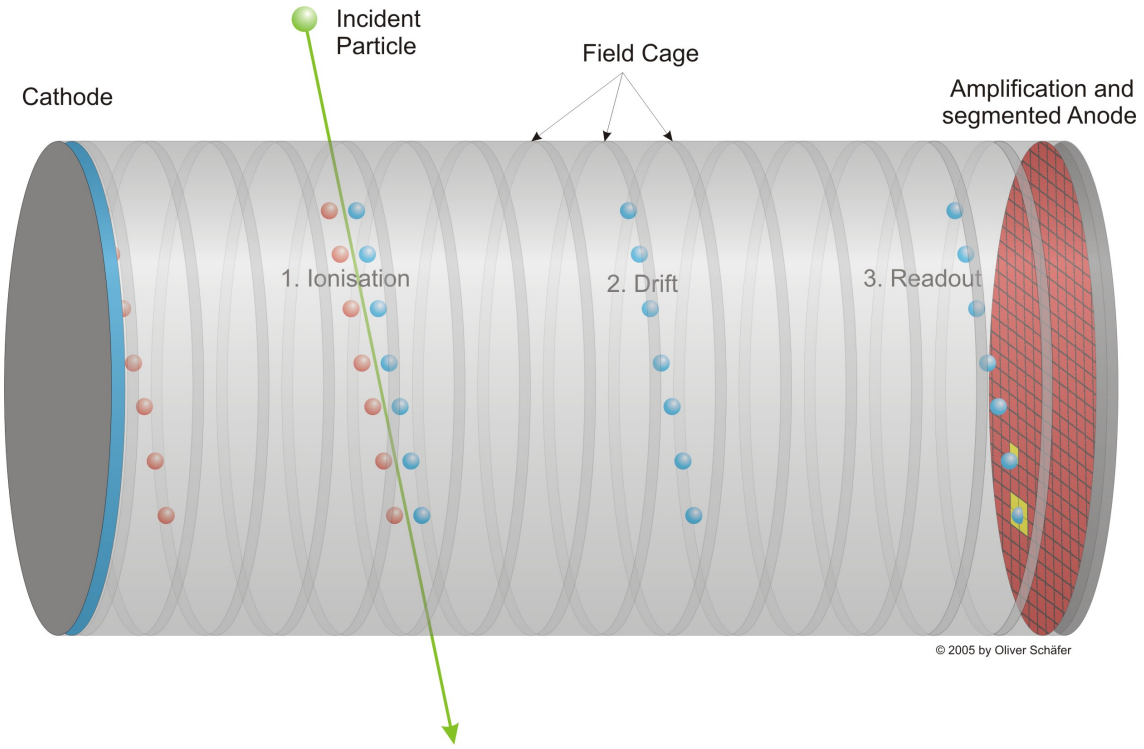

Figure 9: Schematic of the working principle of a time projection chamber, from [30].

a plate capacitor using non-conductive pillars as spacers, shown in Figure 11. Also here, the metal foil is perforated and traversing electrons from the drift region are avalanche-amplified in the region between the metal foil and the anode. This means that the necessary high voltage is applied between the metal foil and the anode, respectively the readout electronics. To avoid discharges from damaging these electronics a layer of a high-resistivy material can be applied on top of the anode. This prevents large currents, and also spreads the charge over a larger area.

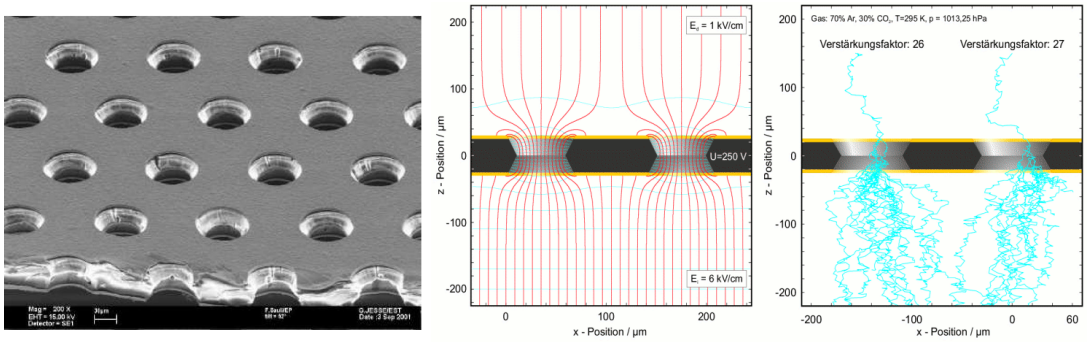

Figure 10: GEM microscopic photo from [37] and working principle.

The ILD TPC is designed as a two-sided TPC in the shape of a ring-cylinder parallel to the beam direction z. The ring structure is necessary to allow room for the beam

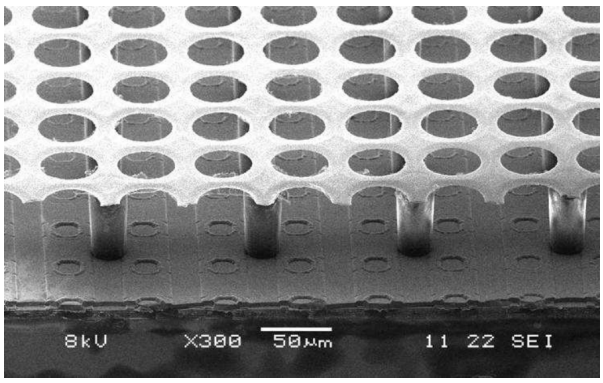

Figure 11: MicroMegas microscopic photo from [38].

pipe and the silicon vertex detector for $r < 303\,\text{mm}$. The currently proposed version of the ILD TPC [7] extends between $329\,\text{mm} < r < 1770\,\text{mm}$ and $|z| < 2350\,\text{mm}$, and is part of the ILD barrel system. There is a central cathode at $z = 0$ and each flat side of the cylinder hosts one anode at $|z| \approx 2000\,\text{mm}$. This means that both the magnetic field of the ILD solenoid and the TPCs electric field are parallel to z. The inner and outer cylinder walls at $r \approx 330\,\text{mm}$ and $r \approx 1800\,\text{mm}$ consist of a composite structure, which allows for gas tightness and mechanical stability, as well as electrical shielding. The electric field is homogenised by an integrated field cage consisting of field strips on the inside of the cylinder walls. The anode side is segmented into modules which contain the amplification and readout stage and, in the current version, have a size of about $17 \times 22\,\text{cm}^2$. The modules are placed in concentric rings around the z-axis and are positioned in an aluminium holding structure. Given the default anode readout granularity with a pad size of $1 \times 6\,\text{mm}^2$, each module contains about 6000 readout channels. The readout pads are narrow (1 mm in ϕ) and high (6 mm in r) and are aligned in 220 concentric circles (rows) around the z-axis. The default reconstruction algorithm uses a row-based pad readout. Most charged particles from the IP traverse the TPC in radial direction, generating a signal in each row. Depending on the pad response function of the amplification stage, several pads in each row are activated, and are combined into one so-called hit in that row. At least three activated pads are needed to allow for a fit of the hit center, i.e. the transverse position of the track above the anode, which determines the point resolution in r/ϕ. For the ILD TPC the target point resolution is better than 100 µm over the full drift length, while the achievable point resolution of a single active pad in a row would be $1\,\text{mm}/\sqrt{12} = 289\,\mu\text{m}$, not including effects from diffusion. GEM amplification provides a sufficiently large charge cloud to activate three pads in a row, while for MicroMegas this can be achieved with the resistive layer. The row-based hits are used for the reconstruction of tracks and the determination of their properties.

The LCTPC collaboration has developed three main amplification and readout technology options for the anodes: a GEM-based one, a MicroMegas-based one, and one with MicroMegas on a pixel chip.

The GEM-based technology has been implemented twice, once with a triple GEM-

stack [39] and once with a double GEM-stack. The former one, called 'GridGEM', will be described in the following, since it is the basis for the simulation in chapter 6. This description follows the module properties and default settings used in two successful GridGEM test beam campaigns [39, 40]. The three GEMs in the GridGEM system are based on the standard CERN GEM design, but with the same footprint as the full module. Each GEM consists of a non-conductive Kapton layer of 50 µm thickness with 5 µm thick copper layers on each side. The etched holes are double-cone shaped, have a diameter of about 70 µm and are positioned in a hexagonal pattern with a side length of 140 µm from hole to hole. The GEM foils are glued to a ceramic grid, which gives them mechanical stability and serves as a spacer between the layers, and between the third GEM and the anode. By default, the distances between two GEMs, the transfer gaps, are 2 mm, and the distance between the third GEM and the anode plane, the induction gap, is 3 mm, compare Figure 12. The module edges and the spacers cause minimal insensitive area, and the effective area coverage of the GridGEM module is about 95%. The anode plane consists of metal pads covering a printed circuit board (PCB), which are connected through the board to electronics, an adapted ALTRO system [41,42]. The pads have a pitch of $1.26 \times 5.85\,\mathrm{mm}^2$ and a metalised area of $1.06 \times 5.65\,\mathrm{mm}^2$. They are wedge-shaped to enable the positioning on concentric rows, and there are 4828 pads in 28 rows in one module. The ALTRO electronics contains a preamplifier, a shaper and a digitiser. It effectively integrates the charge of each pad including pedestal subtraction. The charge information is then used for the hit and track reconstruction.

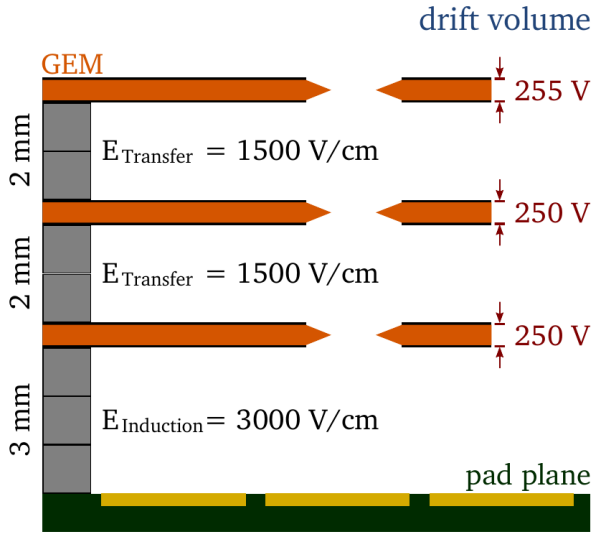

Figure 12: Scheme of a triple GEM stack as used in the GridGEM system, from [39].

The second GEM-based readout system [43] uses a double GEM stack, which requires higher amplification voltages to achieve the same gain, but is overall very similar to the GridGEM system concerning aspects relevant to this work. Another readout techology uses MicroMegas for amplification and a pad plane on a PCB [44]. Here, the MicroMegas layer is applied to the pad plane by simply laying it out over the PCB and fixing it on the

module edges. The current implementation uses a pad size of about $3 \times 7\,\mathrm{mm}^2$, which is coarser than the GridGEM, but the overall performance in terms of point resolution and specific energy loss is comparable. The double-hit resolution, to separate parallel close-by track, however, is affected by the transverse granularity, and the effect of the coarser transverse pad size is an open question.

The third technology concept, the GridPix [45, 46], differs from the previous ones mostly in anode plane: instead of metal pads on a PCB it uses the metalised openings of a pixel chip. Thus, after amplification by the MPGD the electron cloud enters the readout electronics immediately, where the signal is amplified and digitised. This allows for an anode granularity of the pixel size of the ASIC used, which is the Timepix and Timepix3 (see chapter 5.2) in recent implementations, with a granularity of $55 \times 55\,\mu\mathrm{m}$. Implementations with both GEM as well as MicroMegas amplification have been studied [47], but the GEM-based one was discontinued. To utilise the full potential of the granularity of the Timepix and to limit the necessary amplification gain, the pattern of the amplification stage was matched to the anode one, i.e. each pixel opening corresponds to one opening in the MicroMegas layer. It was found that this is best achieved by growing the amplification grid on top of the ASIC surface by microelectronic post-processing techniques, providing an exact match of the two openings. This can be seen in Figure 11, where a mesh layer has been grown on top of a Timepix chip. The avalanche process generates a narrow electron cloud that induces a signal in exactly one pixel. This signal is preamplified and digitised by applying a simple threshold counter, which allows to count the time over threshold and/or the time of arrival w.r.t. a reference time. An introduction to the Timepix ASIC and a detailed description of the digitisation process is laid out in chapter 5.2. The time-over-threshold count can be translated into an equivalent charge of the incoming signal. Thus, despite it being a digital counter, the pixel measurement provides the same qualitative data as the ADC of the pad-based readout systems, at least in their current implementation, and at a two orders of magnitude higher granularity. A qualitative difference would arise if overlapping tracks would create multiple activations of a channel in a given readout period. In that case, the pixel counter could not tell them apart and would either ignore each activation after the first one or add all up into one integrated value. An ADC, however, could in principle record all activations separately, provided a sufficient data storage is used. In the current implementations of the pad-based readout systems this is not the case and multiple activations are typically excluded in test beam analyses. Furthermore, the expected occupancy of the ILD TPC in the ILC environment is low with only few overlapping tracks in each readout window. This occupancy would be even lower with a high granularity readout like the GridPix system, and the overall loss of usable channel information due to multiple activation is expected to be negligible.

The advantage of the high granularity of a pixelised anode is a potential improvement of the track reconstruction and momentum resolution due to the higher sampling rate of the tracks. The GridPix data shows a larger effective number of reconstructed primary electrons per track length compared to the pad based systems and it was shown in simulation that the effective momentum resolution could be improved by about a factor of 1.5 [46]. The effect is even larger in the forward direction, since due to the low

occupancy of the pixel readout its momentum resolution scales with the track length, while the row-based readout one's scales only with the transverse part. With its square pads the properties of the GridPix readout do not depend on the incident ϕ-angle, other than the row-based readout. Furthermore, the measurement of the specific energy loss is expected to profit significantly from an increased granularity, which is explained in detail later in this chapter. The higher granularity can also deal with higher local hit occupancies and provide a better double-hit separation.

The disadvantage of a pixelised readout is the potential cost, the power consumption and the data rate. Both anodes of the ILD TPC have a combined area of about $19\,\mathrm{m}^2$ which needs to be instrumented. With pads of $6\,\mathrm{m}^2$ this would require about $3 \cdot 10^6$ channels. With pixels of $55 \times 55\,\mu\mathrm{m}^2 \approx 0.003\,\mathrm{mm}^2$ this would require about $6 \cdot 10^9$ channels and about 10^5 Timepix-like ASICs. This could be a significant cost aspect, but also have a larger power consumption, which in turn may require more cooling in the TPC endcap and more material in front of the endcap calorimeters. With more channels, also the data rate could be correspondingly higher in the pixel case, although zero-suppression and potential hit-combining algorithms could significantly reduce the data output. Finally, the current implementation of the GridPix system can only cover about 60 % of the anode due to the fixed square shape of the Timepix chip and the necessity for wirebond connections for power supply and data readout on the sensitive side of the ASIC. This would significantly counteract the usefulness of the aforementioned advantages of higher granularity. All of these aspects, however, are a target of ongoing development and optimisation and it remains to be seen how a full-module sized system would compare to a corresponding pad-based system.

It is the goal of this work to assess whether an intermediate system with a granularity and number of channels in between the two aforementioned options is viable and how its specific energy loss measurement and particle identification performance would turn out.

3.3. Tracking with a TPC

The arrival position at the anode together with the arrival time information allows to reconstruct individual signal positions in 3D, called hits. Charged Particles from the interaction point with a sufficiently large transverse momentum traverse the TPC and generate a hit in each row of the row-based readout, i.e. 220 in case of the ILD TPC. If a particle in ILD has a transverse momentum below about 1 GeV, its curvature in the magnetic field has a helix diameter below the radius of the outer TPC wall, causing it to curl inside the TPC. These curlers generally generate multiple half-circles of hits, each with fewer than 220 hits, but often with a much larger total number. Curlers either exit the TPC through the end plate, or they can lose all energy via ionisation or scattering in the gas volume, which causes their helix radius to shrink to a point. Similarly, 3D hits are generated in the vertex detector and the silicon external tracker, surrounding the TPC. All tracker hits are then used in a helix fit to reconstruct tracks, as laid out in [48]. This iterative fit uses a Kalman-filter and takes the energy loss in the TPC volume into account. To measure the momentum of the incident particle, the curvature of the track

needs to be extracted. Given very precise spatial endpoints of the helix fit in the silicon layers, the role of the TPC is to precisely determine the sagitta of the track, i.e. the deviation from a straight line between the endpoints. This is directly linked to the point resolution of the TPC and sets the target for the ILD TPC to be better than 100 µm over the full drift length. However, this ultimate asymptotic momentum resolution can only be achieved for very high momenta, while at momenta below about 50 GeV the impact of multiple scattering dominates the resolution, as can be seen in Figure 13a.

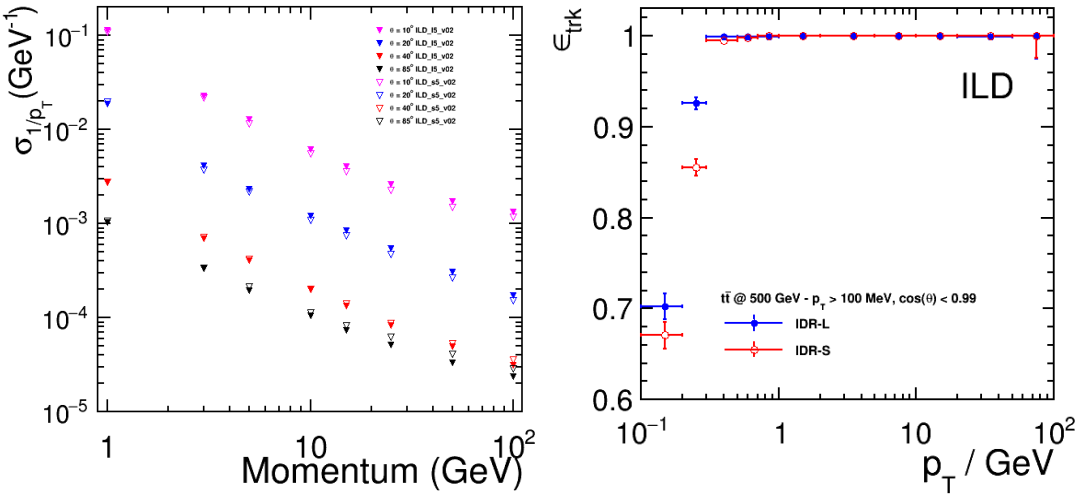

(a) (b)

Figure 13: ILD tracking performance from full-detector simulation and reconstruction. (a) momentum resolution and (b) tracking efficiency, from [7].

Compared to all-silicon trackers, the TPC registers tracks with more than an order of magnitude more individual hits and continuous tracking. At the same time a TPC adds only a comparably small amount of material to the material budget of the barrel part, and most material, including the endcaps, lies right in front of the ECal, which minimises its impact due to scattering. The pattern recognition capability not only provides an excellent track finding efficiency down to low momenta, compare Figure 13b, but also allows to detect kinks in tracks or (dis)appearing tracks, for example from in-flight decays. The latter can be used to find charged decays of neutral particles inside the TPC, so-called V0-finding, which is laid out in chapter 7.4. The central advantage of a TPC with respect to this work, however, is its capability to measure the specific energy loss dE/dx of the incident particle in the gas volume to do particle identification, which is explained in detail in the following.

3.4. Measurement of the Specific Energy Loss

As suggested by Equation 1 and Equation 2, the measurement of the specific energy loss dE/dx of an incident particle together with its momentum can be used to reconstruct its mass. The measurement resolution of the specific energy loss, the dE/dx resolution

$\sigma_{dE/dx}$, is typically only at the level of a few percent, and thus also the measurement of the mass would be of similar precision. However, the number of charged particle species which are sufficiently long-lived to leave a track in the TPC is limited and known in each experiment. In order to identify a particle, it is thus not necessary to exactly reconstruct its mass, but rather find which species it most likely belongs to. The usual way to do this is to find the closest reference line in the momentum-energy loss plane, which is displayed in Figure 14.

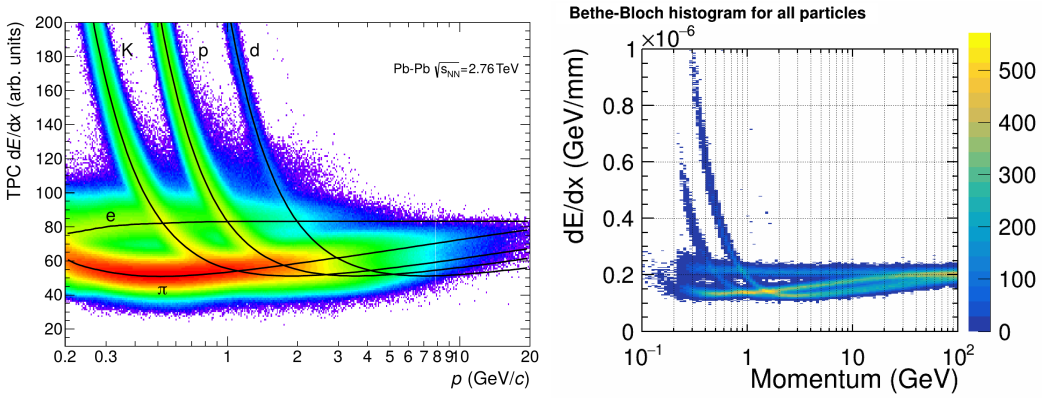

(a) ALICE data, 2014

(b) ILD simulation

Figure 14: Bethe-Bloch curves, i.e. specific energy loss vs. momentum for (a) the ALICE TPC from [49] and (b) ILD simulation with single particles. The curves for the different particle species are visible and mostly separable but also intersect in certain areas.

The specific energy loss is conventionally measured by summing up the reconstructed energy depositions dE in the tracker, which are proportional to the measured charge, and dividing them by the total flight path inside the tracker dx, which will be called the 'charge summation' method in this work. Since the energy depositions are Landau-distributed, the long tail of the distribution causes a large fluctuation of the average dE/dx when all energy depositions are simply added up. Instead, for each hit of the track a dE/dx value is determined and a fraction of the largest dE/dx contributions of these hits are rejected before summing them up. This way, the tail can be effectively, although not entirely, suppressed and the dE/dx resolution be improved, which is illustrated in Figure 15. Here, a truncation fraction of 100 % means that all hits are accepted. The more hits are rejected, the lower the impact of the Landau tail, but the fewer data is used to determine the dE/dx estimate. This leads to an optimum truncation fraction, which depends on the specific experiment and lies typically between 70 % and 80 %.

3.4.1. Charge Summation with High Granularity

How well the energy loss can be measured depends mainly on two parameters: the effective length of the sampled track, and the number of independent samples measured along track. This is expressed in the empirical formula derived by [51], p. 339, from

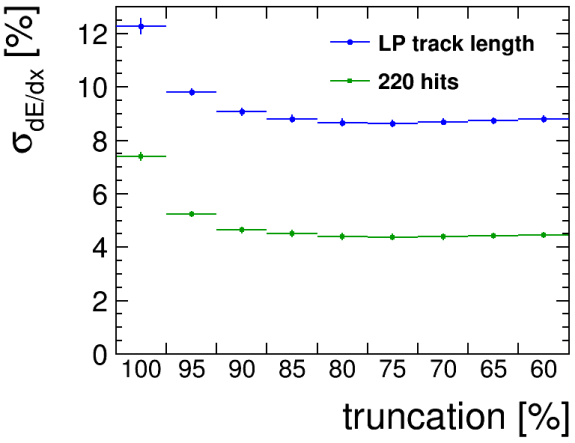

Figure 15: Measured dE/dx resolution from test beam measurements with the GridGEM system, depending on the truncation fraction. 100 % truncation means here that all hits are accepted. The plot shows the typical optimum of truncation at an acceptance fraction around 75 %. From [50].

comparing various former drift chambers and time projection chambers:

$$\sigma_{dEdx} \propto L^{-0.32} \cdot N^{-0.13}, \tag{4}$$

where σ_{dEdx} is the relative uncertainty of the energy loss measurement, L is the effective track length and N is the number of independent samples of that track. The exponent of N is quoted to be between -0.11 and -0.14, depending on the reference analysis, and in this work the average of -0.13 is used. In most use-cases when two potential systems are compared to each other, L and N are directly connected, causing Equation 4 to have limited immediate applicability. Therefore, the granularity $G = \frac{N}{L}$ is introduced and Equation 4 becomes:

$$\sigma_{dEdx} \propto L^{-0.45} \cdot G^{-0.13}. \tag{5}$$

In many cases, one compares gaseous detectors with a fixed readout structure, i.e. a constant G, but changes the effective track length, e.g. when extrapolating from a small prototype to a large scale experiment. Then Equation 5 expresses the well-known dependence of increasing precision of the energy-loss measurement with larger L, which is a little less strong than the naively expected square-root behavior. The exponent of L is estimated to be -0.37 or -0.36 in Equation 4 by [52, 53] and [54], respectively, which brings the dependence in Equation 5 closer to a square-root behavior. In this study, however, we fix the track length and vary the granularity instead. Then Equation 5 suggests, that also with increasing granularity the uncertainty on the energy-loss measurement decreases, albeit with a moderate exponent around -0.13. In [51], a corresponding behavior is shown down to a granularity of 5 mm and it is questioned whether further increase in granularity will increase the dE/dx performance. This question is addressed in chapter 6.

3.4.2. Cluster Counting

As pointed out before, the relatively large width of the Landau distribution which governs the energy deposition in each ionisation step worsens the correlation of the measured energy and the momentum of the particle. It is advantageous to count the number of ionising interactions of the incident particle instead. This is given by a Poissonian distribution with a significantly smaller width, resulting in a better correlation and particle identification power, as first mentioned in [55] and demonstrated in [56]. In Figure 16, the separation power for pion/kaon-separation depending on the cluster counting efficiency is shown compared to the conventional dE/dx determination by charge summation. In former prototype experiments, e.g. [57], a cluster counting efficiency of only 20 % was reached with Ar-based gas mixtures. Nevertheless, the resulting separation power is still better than by charge summation. Also, improved algorithms are expected to deliver a higher cluster counting efficiency. However, cluster counting can only work if a sufficient correlation between the position of the electrons of one cluster is preserved during drift. The effectiveness of one such algorithm under various conditions is investigated in a simulation study presented in chapter 6. A novel readout system capable of providing the necessary granularity is presented in chapter 5.

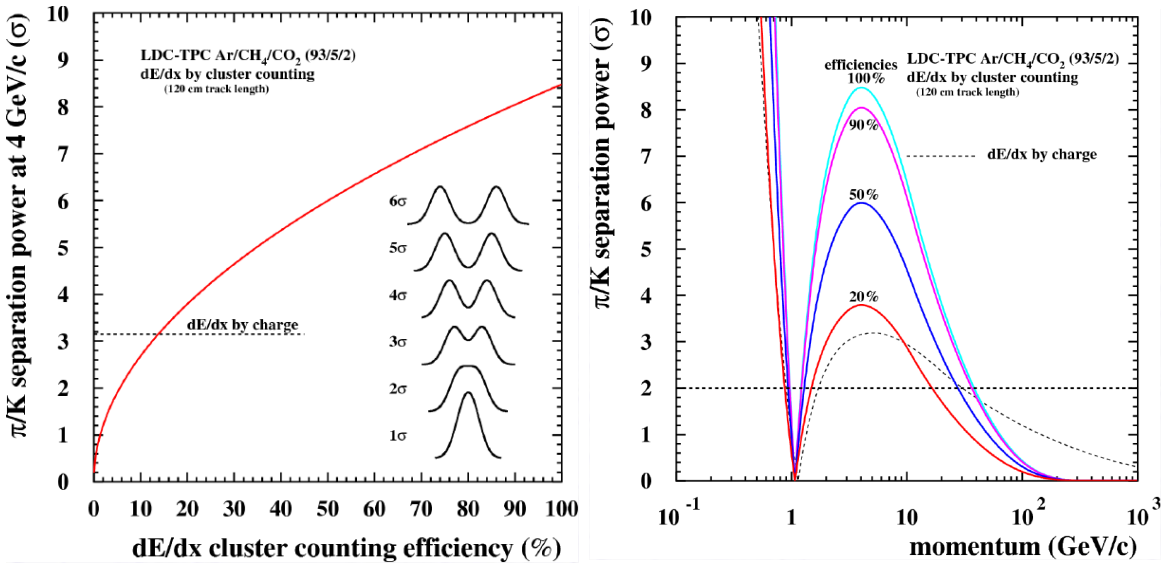

Figure 16: Pion/kaon separation power of dE/dx by cluster counting from simulation, from [56]. Left: separation power depending on the cluster counting efficiency, Right separation power vs. momentum for different cluster counting efficiencies. In both cases, cluster counting is compared to the default performance by charge summation.

4. Software Framework

Throughout this work, the software framework iLCSoft is used. The overall setup and the specific components of iLCSoft used in this work are described in this chapter. This is accompanied by the external software package Source-Extractor and its implementation into MarlinTPC as part of iLCSoft. In addition, the different Monte Carlo samples used in this work are specified.

4.1. iLCSoft

iLCSoft [58–60] is a collection of common software packages developed in the Linear Collider Community. It comprises packages to handle all tasks ranging from event simulation or conversion of measured raw data through reconstruction to analysis, mostly for e^+e^--collider experiments. This includes packages that provide data models, event handling, geometry and operating framework, but also handling of conditions data, displaying events and detectors as well as data export.
The detector geometry description is provided by DD4HEP [61] and the data format is defined by the persistency framework LCIO [62].
iLCSoft includes steering of event generation, with details given below in chapter 4.5. The Marlin [3] package provides an environment to process events through the different stages of the particle and signal propagation. Each processing step is implemented in a so-called processor. Every event is processed through a chain of these processors, which take the existing information of its predecessors, e.g. the hit information from the tracker, and add newly generated information to the event, e.g. a reconstructed track. Throughout this work, the names of Marlin processors as they appear in iLCSoft are denoted in `this font`.

4.2. MarlinTPC: Detailed TPC Simulation

The ILC software package MarlinTPC [4] provides a package for a detailed simulation, reconstruction and analysis of time projection chambers. MarlinTPC was developed to do studies on simulation and detector prototype data on a much more detailed level than is reasonably possible in a large detector study. This includes (parametrised) ionisation, drift, amplification and digitisation in the simulation part, as well as hit- and track reconstruction and analysis tools. MarlinTPC makes use of the core packages like Marlin and LCIO, but uses GEAR (GEometry API for Reconstruction) for the geometry description. It provides a very detailed simulation of a TPC, which is used for detailed studies of the performance of different TPC readout structures, but is not necessary and would not be efficient for a full ILD simulation. Instead, such a detailed simulation is run separately, analysed and then used in a parametrised way in the full detector simulation. In this work, MarlinTPC is used to simulate a highly granular readout of the TPC drift volume, which is described in 6.1.

4.2.1. Source Extractor

The simulation presented here targets a very high readout granularity with pad sizes smaller than the charge clouds generated by the GEM amplification. The resulting 2-dimensional charge patterns measured of the pad plane are similar to sky maps in astrophysics. There, often faint light sources in front of background light are investigated. This problem is comparable to finding and determining the properties of (partially overlapping) charge clouds from a highly granular GEM readout. This is depicted in Figure 17.

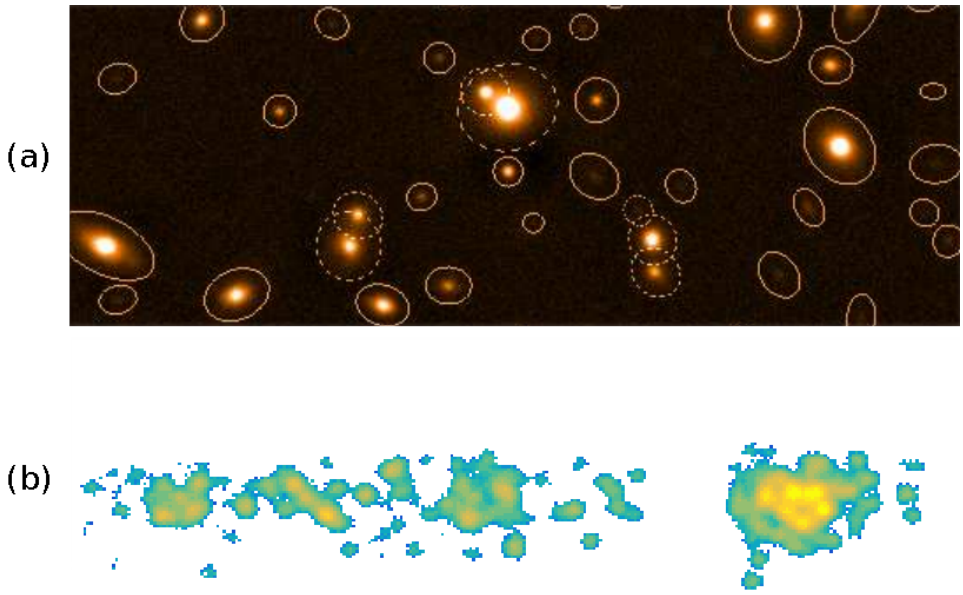

Figure 17: Illustration of (a) the data of an astrophysical sky map (from [63]) and (b) electric charge from GEM amplification picked up by a highly granular TPC readout anode. The patterns are structurally similar and a tool from astropysics, the Source-Extractor, can be used to analyse also the TPC data.

One software used in astrophysics to extract these sources from the sky maps is the program SExtractor, or Source-Extractor [63]. It was first used in the field of gas amplification studies for the reconstruction of charge depositions in a GEM-Pixel-TPC by the PixelTPC group at the University of Bonn [64]. In these studies it was implemented in MarlinTPC.
The software takes a 2-dimensional image, equivalent to a 2D-histogram, and computes the most likely position and brightness of sources, based on a provided typical shape. The software uses a multi-staged approach, iteratively separating the image in sub-sections and adding sources to the fit, which is explained in chapter 6.1.

4.3. MarlinReco: Simulation and Reconstruction of Full-ILD Events

The package MarlinReco first digitises the detector activation raw data, i.e. it creates data equal to the one later received from an actual collider, except for calibration and alignment corrections. The various detector hits created this way are then reconstructed into larger objects like tracks or calorimeter clusters, and finally combined into particle flow objects (PFOs) by the Pandora package [65], representing individual reconstructed particles as they likely were created at the IP. The energy loss measurement in the TPC and the corresponding particle identification are part of the high-level reconstruction of MarlinReco.

In this analysis, in principle the entire detector with all its components is used. However, many aspects and individual observables are not, or only indirectly, utilised in the analysis. The focus lies on the tracking system, in particular the TPC and its dE/dx resolution. V0-finding is also based on the TPC's capability to resolve in-flight decays of long-lived strange hadrons. The hit times of the ECal are used as TOF observables. Details regarding the implementation and performance of these capabilities are laid out in chapter 7. The entire tracking and calorimetric system is fundamental to the jet clustering algorithm.

4.4. Analysis software

The physics analysis parts were done on the particle flow objects, PFOs. Some scripts were implemented as additional processors and run by Marlin after the reconstruction chain, and some were implemented as ROOT scripts working on a subset of the available information to run detailed and comprehensive procedures like multivariate analyses.
The ROOT framework [66,67] provides an extensive library of tools used in high energy physics analyses, combining data structures and procedures. One such tool kit used in this work, which is worth mentioning, is the Toolkit for Multivariate Data Analysis (TMVA) [68]. In this work, the TMVA methods boosted decision tree (BDT) [69] and multi-layered perceptron (MLP) [70] were used in chapter 8.2.

4.5. MC Samples

For the analysis of full-ILD events samples from two large-scale Monte-Carlo productions, 2018 and 2020, were used. The simulation and reconstruction of these samples were done with iLCSoft v02-00-01, as laid out above in this chapter. The software Guinea-Pig [71,72] was used to simulate the beam-beam effects and the resulting centre-of-mass energy spectrum (Figure 18) based on the ILC beam parameters [23] in Table 1. To generate the hard interaction events the software Whizard [73,74], version 1.95, was used. Pythia 6.4 [75] was used to simulate the final state QCD and QED showering of generated quarks as well as muons and tauons. It was adjusted to the OPAL tune, as described in [26] sec. 2.2.1.4. The decays in the detector volume and the detector re-

sponse to the traversing particles were simulated with Geant4 [76]. The detector readout and the consecutive reconstruction of the events were done with MarlinReco, including Pandora [65] for particle flow.

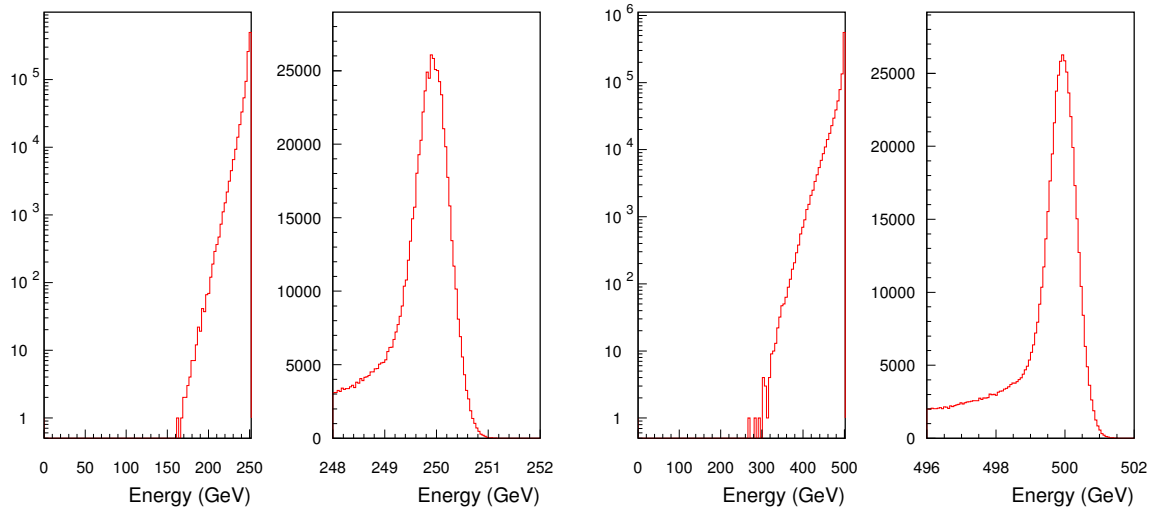

(a) 250 GeV (b) 500 GeV

Figure 18: Luminosity spectra of the ILC, calculated with Guinea-Pig, from [7]. (a) for a center-of-mass energy of 250 GeV; (b) for a center-of-mass energy of 500 GeV.

The large MC productions contain events that correspond to a large integrated luminosity of physics processes. They cover the full Standard model physics processes, as well as several dedicated BSM signatures to study which BSM cases ILC and ILD could find and differentiate. To ease usage of these samples, they are sorted into associated processes, which are characterised at the point of the final state particles "at the end of the Feynman diagram", i.e. after the matrix element calculation, but before the hadronisation, outgoing particles are picked up. They are divided by number and consistency with certain intermediate states, represented by tree-level Feynman diagrams. Two examples are "2f_Z_hadronic" and "ttbar_semileptonic". In the case of 2f_Z_hadronic the final state consists of two fermions, which are consistent to be hadronic decay products of a Z, thus a quark and its anti-quark (except for t and $\bar{t}$). In the case of ttbar_semileptonic the final state consists of six fermions, two of which must be b and $\bar{b}$, two must be quarks consistent with a W decay, thus an up-type quark and a down-type quark of the first two quark generations, and a charged and a neutral lepton with opposite lepton number (consistent with a leptonic W decay), where the charged lepton has the opposite charge as the charge-sum of the quarks from the hadronic W decay.
The cross section of the events was calculated based on a beam polarisation of $+100\,\%$ or $-100\,\%$ for each beam, and events are re-weighted in analyses to account for a realistic beam polarisation, usually $80\,\%$ for electrons and $30\,\%$ for positrons. In this work, the

re-weighting is used in chapter 8.1.
In addition, certain samples were produced for detailed study and calibration, including single-particle events. Single-particle events do not require Whizard and Pythia, instead individual generator-stable particles with defined momenta are put at the IP and their specific activation of the detector is simulated. The single-particle samples consist of electrons, muons, (charged) pions, (charged) kaons and protons (collectively addressed as "P5") with 100 k particles each, with an isotropic angular distribution and a momentum distribution that is roughly flat in $\log(p)$. In this work, these samples are used to asses the dE/dx and TOF performance of ILD.

In the 2018 MC production [77], about 150 million events with a center-of-mass energy of 500 GeV were generated and simulated with Whizard 1.95 and iLCSoft v02-00-01, including Geant4 v10.03. The number of events corresponds to roughly $500\,\mathrm{fb}^{-1}$, see [7], section 7.4.

In the 2020 MC production [78], a center-of-mass energy of 250 GeV was chosen and events were generated and simulated with Whizard 2.8.5 and iLCSoft v02-02, including Geant4 v10.04. The number of events corresponds to $1000\,\mathrm{fb}^{-1}$ for most, and $5000\,\mathrm{fb}^{-1}$ for some final states.

5. Development of a Highly Granular Readout for a TPC

In the past, several collider experiments used time projection chambers as their central tracking device. Noteworthy examples are ALEPH [79] and DELPHI [80] at LEP, STAR [81] at RHIC, and ALICE [82] at the LHC. With their relatively long readout time TPCs are well suited for collider experiments with moderate track occupancies in time, i.e. at e^+e^- colliders like LEP. At ion-ion colliders, like ALICE at the LHC, the PID capabilities of a gaseous tracker are invaluable, and the TPC is compatible with the ALICE collision rate, which is significantly reduced compared to CMS and ATLAS. For ILD at the ILC, too, a TPC is proposed as central tracker and the Linear Collider TPC (LCTPC) collaboration has been working on research and development towards its realisation. Several readout options have been developed, as laid out in chapter 3.2.

All systems have shown that they can meet not only the required point resolution, but also the collaboration's aim of a dE/dx resolution for the ILD TPC of 5 % or better. A system with an increased granularity compared to the default pad-based options, like the GridPix system, has the potential to perform significantly better than this aim due to its high granularity. However, the GridPix system comes with a number of disadvantages, which is why a novel system, called Ropperi, was devised and developed within this work.

After motivating the Ropperi system, its design and the construction of a first test system as well as an assessment of its basic properties are presented in this chapter. Based on the results, its applicability for the ILD TPC is discussed.

5.1. Motivation for a Highly Granular Hybrid Readout

To achieve the cluster counting capability introduced in chapter 3.4.2, a sufficiently high granularity of the TPC readout system is needed. Figure 19 shows a granularity comparison of the GridGEM system with the (discontinued) prototype system with GEM amplification and 8 Timepix ASICs as anode, a so-called Octoboard [83]. The electron clusters from the primary ionisation process are amplified by GEMs, creating charge clouds visible as blue blobs. Since the anode consists of 8 Timepix ASICs it has a granularity with a pitch of $55 \times 55\,\mu m^2$. The charge clouds are clearly identifiable – the granularity is even higher than needed. The green overlay of the GridGEM granularity shows that with its pads of approx. $1 \times 6\,mm^2$ a cluster identification however is not possible.

5.1.1. Concept of a Highly Granular Readout

Therefore, a novel readout structure was devised, called Ropperi (Readout Of a Pad Plane with ElectRonics designed for pIxels), that allows pad sizes of only a few hundred µm to enable cluster counting, provides a large flexibility and keeps the channel number moderate at the same time. This is achieved by reading out the pads of the anode with a highly granular digital pixel chip instead of an ADC with a much larger

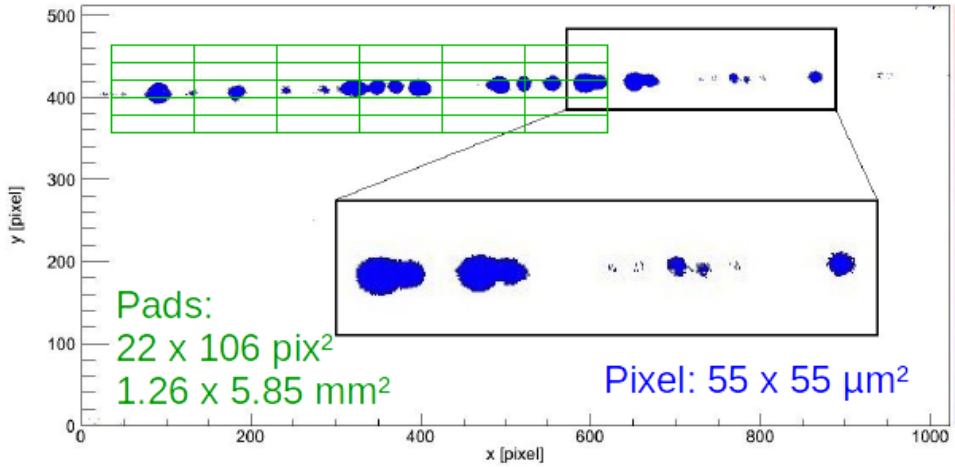

Figure 19: Event display of a track recorded with a Timepix Octoboard, from [83]. The activated pixels are shown in blue, the green overlay shows the pitch of a typical pad-based readout.

footprint. The readout pads on a printed circuit board can then be of an intermediate size between the granularity of the GridGEM and the GridPix systems, limited only by feature size limitations in PCB production. The digital pixel chip does not register the pulse shape or the integral of the charge arriving at a pad, like an ADC does, but counts the time steps the signal lies above a certain threshold. This time-over-threshold can then be transformed into an accurate measurement of the charge. The digital counting, however, does require a sufficiently low occupancy in order to avoid double hits in time, which could not be resolved, but this is easily achieved at ILD given the enhanced granularity. The pixel chip chosen for Ropperi is the Timepix ASIC, which is described in detail in the following chapter. The main reasons for this decision are twofold. On the one hand, the general properties of the ASIC in terms of granularity, expected noise, readout capabilites and readout modes were assessed to fit the Ropperi requirements. On the other hand, there are experience and ongoing work with the Timepix in the LCTPC collaboration, based on which the necessary hardware and software could be easily acquired and concrete assistance during the data taking periods was provided. The Timepix chip with a granularity of 55 µm needs to be connected to the sensitive pads of the PCB with a granularity of a few hundred µm, which requires a fan-out from the chip to the pads. The ansatz chosen for Ropperi is to use a multi-layered PCB and have this fan-out routed inside the PCB.

This leaves two major questions:
1. How can the pixel chip be connected to the PCB in a reliable and mechanically stable way? It was decided to address the connection issue with gold-stud bump bonding, which is described below. The bump-bonding was done in collaboration with the Karlsruhe Institute for Technology.
2. Can the pixel chip cope with the additional noise from capacitive coupling in the PCB's routing line length and pads connected to the pixel input?

Figure 20 shows the signal/noise ratio of a Timepix channel for a fixed input signal strength of 1000 electrons (1 ke$^-$), depending on the additional input capacitance. The additional capacitance increases the noise in the channel and reduces the S/N ratio. Routing lines in a PCB add input capacitances O(pF) per line length of a few cm, which are the typical line lengths in a board with the dimensions of the ones used in the current LCTPC readout systems (about $17 \times 23\,\mathrm{cm}^2$). While a S/N ratio of 10 is often deemed an acceptable level, at an input capacitance of 1 pF, the S/N ratio here is only about 4. However, the signal strength of one amplified electron can be much larger than 1 ke$^-$ if a stack of multiple GEMs is used. The GridGEM system operates with a total average amplification gain of about 8 k, but the gain of its GEM stack can be enhanced to about 100 k by increasing the GEM voltages. Above a gain of about 100 k, GEMs start to frequently discharge and are not considered sufficiently stable anymore. This would leave enough room to find a stable gain setting and still have a S/N ratio of 10 or more.

The system concept aims at a granularity with structures smaller than the size of GEM charge clouds in order to resolve them and reconstruct primary ionisation clusters. This means that the aforementioned signal strength of one GEM charge cloud would be divided among O(10) pads, and the S/N ratio for each pad would be reduced by that factor. At the same time, the same number of pads would be used to reconstruct the charge cloud, compensating for the charge splitting. Based on these considerations, a sufficient S/N ratio should be achievable. It is the central question of this chapter if this is possible for the developed hardware.

5.1.2. Further Advantages

Aside from the target charge cloud and cluster identification capability, the ansatz of Ropperi brings additional advantages with respect to the existing pad- and pixel based readout systems. Compared to the pad-based systems, like the GridGEM [39], Ropperi's higher granularity leads to a reduced occupancy, and potentially also to an improved double hit/track resolution, which is limited by the pad width [85]. In addition, the Timepix ASIC allows for a significantly smaller footprint of the readout electronics.

Compared to the pixel-based systems, in particular the InGrid system [86], Ropperi has a lower granularity, but still measures each GEM charge cloud with several pads, allowing for a position calculation using a fit. Instead, the InGrid system records each electron with exactly one pixel. In addition, the pixel-based anode can currently only cover approx. 50 % of the anode area in a setup with 12 octoboards [87], or up to 63 % with tightly stacked ASICs as suggested in [88], since the ASICs need some overhead area, and the module geometry of the foreseen ILD TPC is not square, as the ASICs are. The GridPix Quad has an active surface of 69 % [89] with a rectangular shape. Ropperi's separate anode PCB allows for a coverage of more than 90 %, comparable to the traditional pad-based systems. Also, it is more flexible with regard to the granularity and to the exact pad layout and PCB shape and size. These properties can be altered by a new PCB design, which is typically much simpler and faster compared to a new ASIC design.

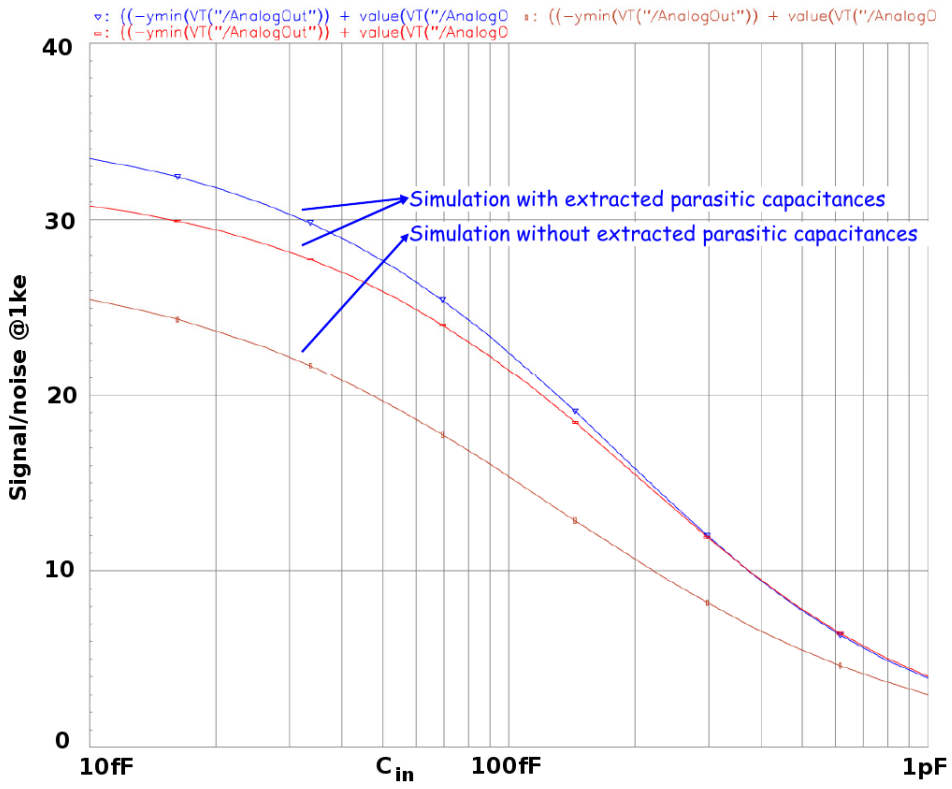

Figure 20: Simulated signal/noise ratio of Timepix channels for a fixed input signal strength of 1000 electrons, depending on the input capacitance added to the channel, from [84].

Further aspects, which would need to be taken into account in advanced TPC design considerations, are the amount of data produced and the necessary material to cover the TPC anode. Compared to a readout system with pixels and GEM amplification, Ropperi would generate less data because of its lower granularity. This is however not the case for the GridPix system, where for each electron arriving at the anode, only one channel is activated. The electronics of both pad-based and pixel-based systems are expected to undergo further development including increasing data readout efficiency. Another factor is the necessary amount of ASICs for the readout systems to cover the two ILD TPC anodes of about $19\,\mathrm{m}^2$. While one might expect a full-silicon solution to be significantly more expensive, it was estimated that the cost for pad-based and pixel-based systems would differ rather on the order of 20 %. This is because the leading cost driver in ASICs is the development, which would need to be performed for any future readout system, while the production cost would actually be moderate. A Ropperi design would require fewer ASICs than a GridPix system and be moderately cheaper in terms of silicon. While the TPC has a very low material budget in the barrel region, the readout structure has about 20 % of a radiation length in the endcap region. This material is placed directly in front of the calorimeter, so its impact is moderate, but it should still be minimised. Using fewer ASICs compared to the GridPix system, Ropperi would have a smaller electricity consumption and thus less requirement for cooling, reducing its material budget. A more detailed comparison, however, would depend on more advanced designs for all readout systems.

5.2. The Timepix ASIC

A key component in the development of a high granularity TPC readout structure is the Timepix ASIC. The chip will be described in the following, including the readout used in the scope of this study.

The Timepix [90] ASIC (*Application-Specific Integrated Circuit*) is a silicon pixel chip in a consecutive development line. It is based on the Medipix2 [91], which in turn is based on the Medipix [92]. Follow-up developments are the Timepix3 [93] based on the Medipix3 [94], and the Timepix4 based on the Medipix4, which are still in development. This development is driven by the Medipix collaborations [95] and together they form the Medipix/Timepix ASIC family [96–98].

The Timepix consists of 256×256 identical elements, called pixels, with a pitch of $55 \times 55\,\mu\mathrm{m}$, covering about $2\,\mathrm{cm}^2$ of sensitive area. In addition, on one side of this pixel matrix the periphery covers an additional area of about $0.3\,\mathrm{cm}^2$ and contains central logical units and bit storage as well as the input/output (I/O) logic and a number of wire bond pads for I/O and power supply. The schematic floorplan of the chip and the layout of an individual cell are given in Figure 21 and Figure 22. The chip as a whole and each pixel contain a number of digital registers for digital-to-analogue converter units (DACs) [99] to store operating information.

Each pixel contains an aluminium opening for the signal input, and an analogue and a digital signal processing part, the schematic of which is shown in Figure 23. The analogue

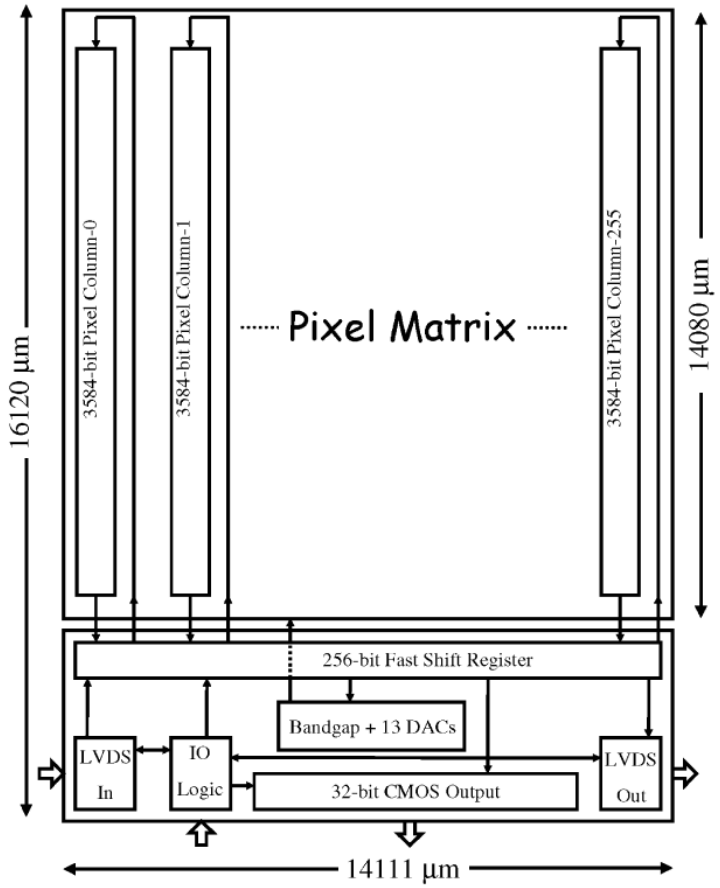

Figure 21: Schematic floorplan of the Timepix ASIC, with the pixel matrix in the upper and the periphery in the lower part, from [90].

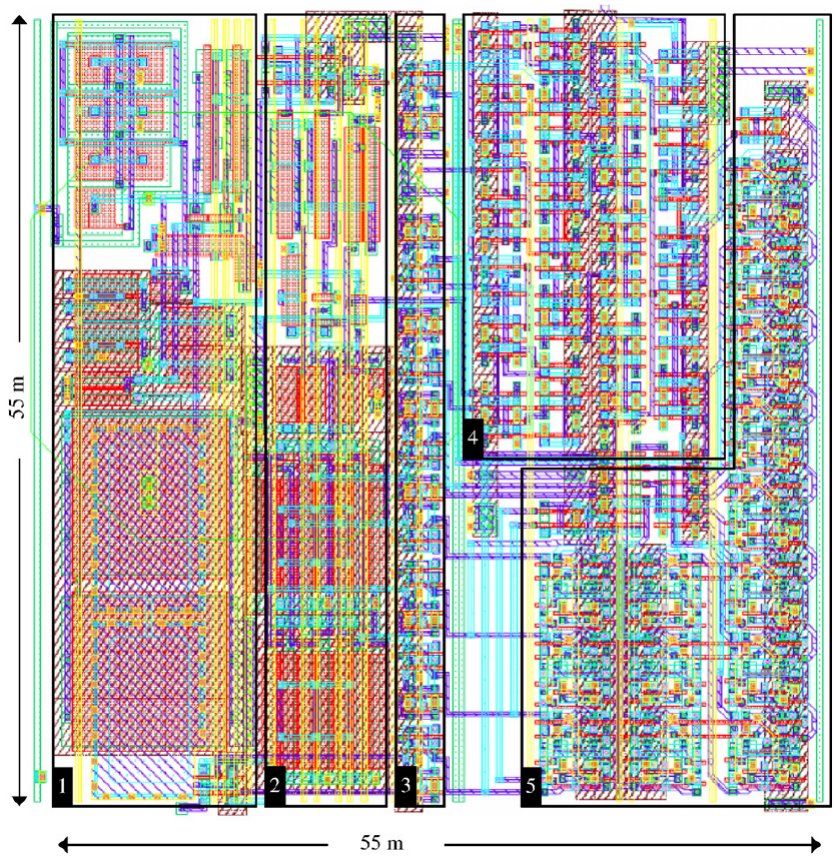

Figure 22: Timepix pixel cell layout with sections: (1) preamplifier, (2) discriminator, (3) 8-bit pixel configuration register, (4) clock buffer and synchronisation logic, (5) 14-bit shift register. From [90].

input signal goes through a preamplifier and then into a discriminator. In the digital section, the shift register counts upwards once per clock cycle, depending on the input signal strength compared to a given threshold and the selected operation mode. Three pixel operation modes are available: (i) event counting, where the counter increments each time the signal rises above the threshold. (ii) time-over-threshold (TOT), where the counter increments as long as the signal is above the threshold. (iii) time-of-arrival, where the counter increments from the time the signal (first) rises above the threshold until a global *Shutter* signal is set high; in this mode, the time information of the signal with respect to a reference time is registered.

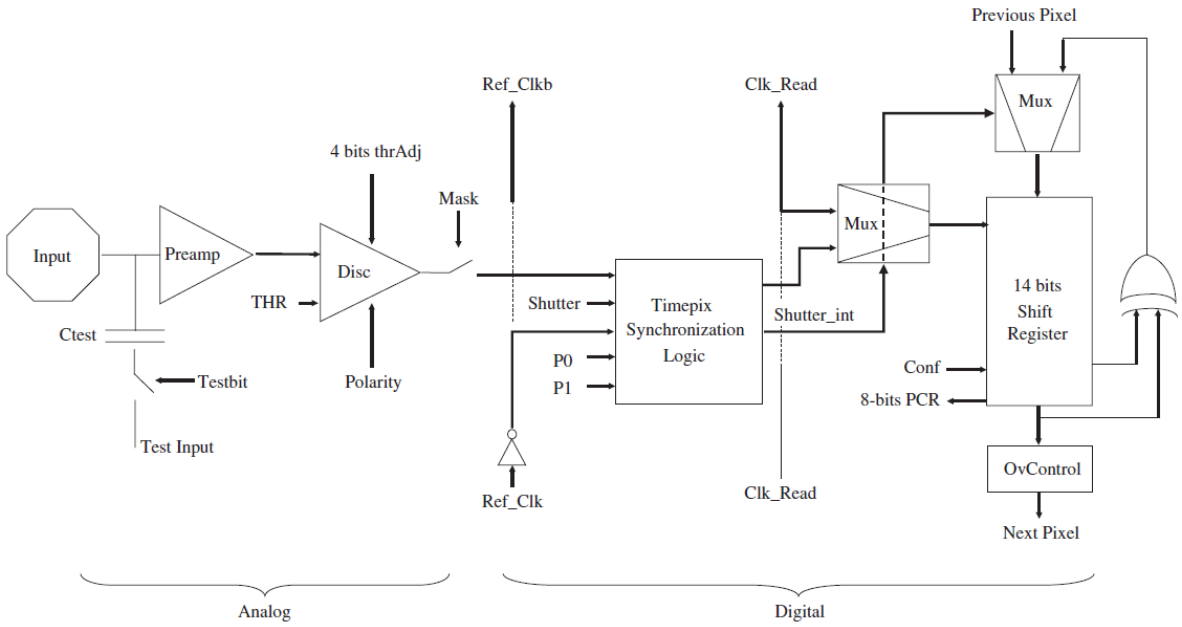

Figure 23: Timepix pixel cell schematic. From [90].

In the analogue section, charge enters the pixel through the octagonal pixel anode of about 20 µm width. This signal goes through a preamplifier that follows the Krummenacher scheme [100]. The signal is then fed into the discriminator where it is compared to a reference threshold. The 'basic' threshold is set chip-wide, but can be adjusted with a 4-bit DAC register for each individual pixel. The voltage steps of this 4-bit register can also be adjusted chip-wide in order to accommodate for individual chips with a broader or narrower distribution of necessary cell threshold deviations. A masking bit DAC allows to disable individual pixels and a polarity bit DAC to select per pixel if the discriminator should respond if the signal is above or below the threshold. The digital response of the discriminator is transferred to the Timepix Synchronisation Logic (TSL), which also receives the shutter and clock information as well as the 2 pixel operation mode selection bit DACs (P0 and P1), which can be set for each pixel individually.

Via a capacitor at the input of the preamplifier, a test pulse can be introduced into each pixel for calibration.

In data acquisition mode, the TSL interprets the signal input and generates an output

depending on the chosen operation mode. Once per clock cycle this output is transferred via a multiplexer to the 14-bit shift register, which operates as a pseudo-random counter with a maximum count value of 11810 at which it reaches the overflow state. The 14-bit shift registers of one column are connected in sequence, forming a 3584-bit column register.

The periphery contains a 256-bit Fast Shift Register (FSR), which is connected to the 3584-bit column registers of the 256 chip columns. When the chip is in readout mode, the column registers shift the data bits sequentially through the column and into the FSR, where the data is transferred either to a serial data out line, which was used in this work, or a 32-bit parallel output connector.

The FSR is also used to receive external input. This is either the values of the chip-wide DACs, in particular the chip-wide threshold value, or the DAC values of the individual pixels. The latter are used to set the matrix, i.e. to distribute the pixel bits for each pixel across the chip by shifting them through the column registers.

The clock signal used for counting as well as for shfting the register bits is an externally generated clock, which is distributed through the chip columns via minimum-sized inverter buffers. Buffering takes about 195 ps per pixel, so the distribution takes about 50 ns for the entire chip, after which a shutter signal can be sent. The maximum clock frequence is 100 MHz.

At the bottom of the periphery, the chip has 127 wire bond pads which are used to connect it to an external readout system. Only a part of these connections are necessary to run the chip, namely the supply voltage, six digital input lines, as well as the clock and the data input and output lines.

The chip operation mode is set via the six digital input lines. This is used to fully reset the chip, set the chip-wide DACs, set the matrix, start/stop the counting in the pixels, or read out the chip. The clock and the data lines are implemented as differential digital lines, i.e. they consist of two lines each, which are the exact inverse to each other with respect to ground. The desired signal is extracted by subtracting the two lines from each other. This scheme is used to improve bit quality and reliability in view of digital noise, which is of particular interest in the lines with frequently (data) or constantly (clock) changing sign.

5.2.1. Timepix Readout

In this work, the readout system connected to the Timepix and used to communicate with the chip is based on the Scalable Readout System by CERN [101–103], which was adapted for the Timepix at University Bonn [104, 105]. The setup is shown in Figure 24 including a bare Timepix ASIC glued to a carrier board. The chip is wire bonded and the lines are connected through the carrier board to an intermediate board (IMB) with initial integrated circuits to transfer the chip I/O to a VHDCI connector. This is connected via cable to an adapter card (A-Card) in the SRS front end card (FEC). The FEC has its own power supply and is connected via a network cable to a PC. The chip is powered by two DC power supply units for its digital and analogue part. The

hardware chain from the implementation by [104] used in this work contains an A-Card v4, a VHDCI cable and an IMB v10.

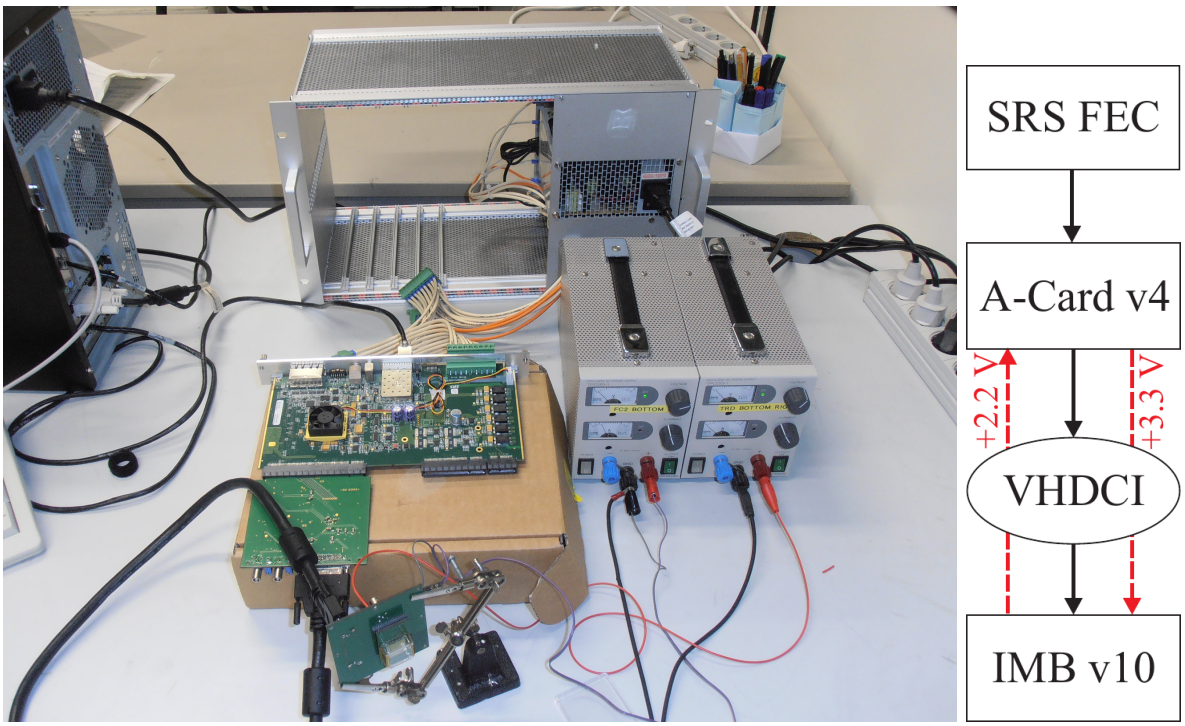

Figure 24: Setup of the Timepix readout based on Scalable Readout System [101] and the used hardware chain [104].

The data output of the Timepix is structured in frames. All pixels start and stop data taking at the same time; this fixed data taking period constitues one frame. Afterwards, the recorded data is sent out to the readout system. The length of one frame should be chosen depending on the operation mode and on the signal properties, and in view of the clock cycle $t_{clock} = 1/f_{clock}$ and the maximum register count of 11810. For example, for a measurement of the time of arrival, the data taking frame length t_{frame} should have a maximum size of $t_{frame} = 11810 \cdot t_{clock}$. If the frame length was larger, all pixels containing a hit at a time $t_{hit} \leq t_{end} - 11810 \cdot t_{clock}$ would count to the maximum number and thus be in overflow. The same is true when measuring arriving charge via the time-over-threshold in an environment where the pixels are receiving charge for a large fraction of the data taking period. If a pixel is active for more than 11810 clock cycles during the frame, its counter is in overflow, meaning the data information is pratically lost. In test beam measurements, the pixel activation fraction is low and long time frames can be chosen for the TOT mode.

The threshold level (THL) is a 10-bit DAC and can therefore have an ADC value between 0 and 1023. A step of 1 ADC value equals about 25 electrons. The pixel threshold is usually chosen in a way that it lies above the baseline including electronic noise, but below a signal peak. Generally, and in particular in the absence of a signal, this threshold can be above or below the average baseline. For a TOT measurement

without an input signal, if the threshold is above the baseline, the pixel will never register as active and its count for any amount of time will be 0, the pixel is 'silent'. If the threshold is below the baseline, the pixel registers as continuously active and it will count in each clock cycle up to the maximum value possible during the data taking frame. The electronic noise in the analogue input to a pixel, or any electrically conductive volume connected to it, will introduce a voltage fluctuation in that pixel. If the threshold is within the range of that fluctuation, the pixel will be active for a fraction of the frame time, which results in its count for that frame between 0 and the maximum. The level of these fluctuations, and thus of the electronic noise, can be assessed by applying a range of thresholds to a pixel and measuring its noise response this way. The necessary increase of the threshold to change a pixel from being permanently active to being silent is proportional to the amount of noice in that pixel's input.

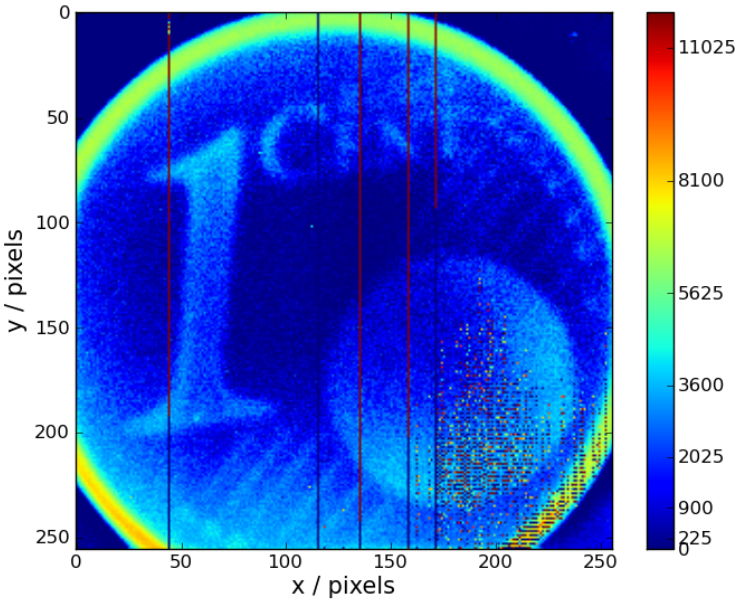

Figure 25: Example of a data frame of the Timepix. Here, a cent coin was placed on the bare chip with thin insulating foil in between. A sinusoidal AC signal supplied to the coin induced a signal in the chip's pixel via capacitive coupling.

One example of a Timepix readout result is shown in Figure 25. It contains one data frame of a TOT measurement that was taken with a cent coin placed on the bare chip and a thin insulating foil in between. The coin is supplied with a sinusoidal signal to periodically charge and uncharge it with respect to ground potential. Via capacitive coupling, this signal is transferred to the chip's pixels. The closer the surface of the coin is to the chip, the stronger the coupling and the larger the induced signal, making the surface structure of the coin visible. In addition, five vertical lines can be seen, that belong to broken pixel columns and whose pixel values only contain either the value 0 or 11810. In the lower right area of the image, a number of pixels show an incoherent noise.

These pixels are so noisy that also they are considered non-usable. A chip with these characteristics is considered low-grade and was only used for tests of the readout system, while the chips used for bonding in the following sections had at most two columns with broken pixels.

5.3. The Ropperi System

5.3.1. Design of the Prototype Board

The first board was designed with standard FR-4 as base material for the printed circuit board (PCB). On one side (side A), the sensitive pads are located, facing the drift volume of the TPC. On the other side (side B), the readout chip is bump bonded. Therefore, bonding pads are applied to side B. Lines between the sensitive and the bonding pads are routed through the PCB, requiring multiple PCB layers. Air tightness was achieved by combining separate blind and buried vias in each routing connection. Since the communication pads of the ASIC are on the same side as its pixel openings, they too are bump bonded (back) to the PCB. The electronic elements of the IMB v10 introduced in chapter 5.2.1 are placed on the PCB side B. Connectors for analogue and digital voltage supply, as well as the VHDCI connector for data transfer are applied. The principal structure of the readout system is shown in Figure 26.

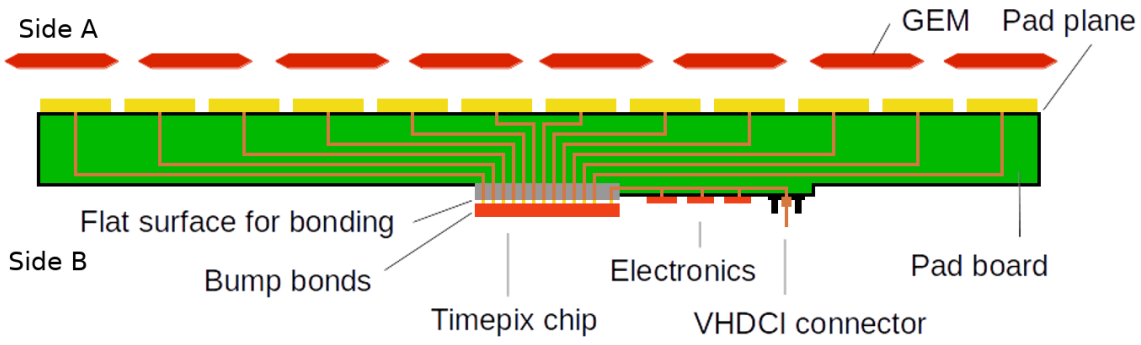

Figure 26: Principal structure of the Ropperi readout, including a potential amplification stage with a GEM.

Since the Ropperi board was intended to be used in an existing TPC prototype with $10 \times 10\,\mathrm{cm}^2$ GEMs, and since the bonding machine allowed for objects with a side length of at most 9.1 cm, the board was designed to have a size of $9 \times 9\,\mathrm{cm}^2$.

The FR-4 material of the PCB limited the minimal feature size of horizontal lines to 80 µm. The size of through-vias had to be larger than 600 µm due to restrictions on their aspect ratio. The layout of the sensitive pads in the central region of the PCB is shown in Figure 27. A large number of small pads could be connected to the bond pads with vias only (i.e. without significantly long horizontal lines in the PCB), thus sitting directly above the ASIC position. A smaller number of channels was connected to the ASIC's outer rows and columns with a significant line length. The total number of achievable connections, i.e. the number of channels, is 500.

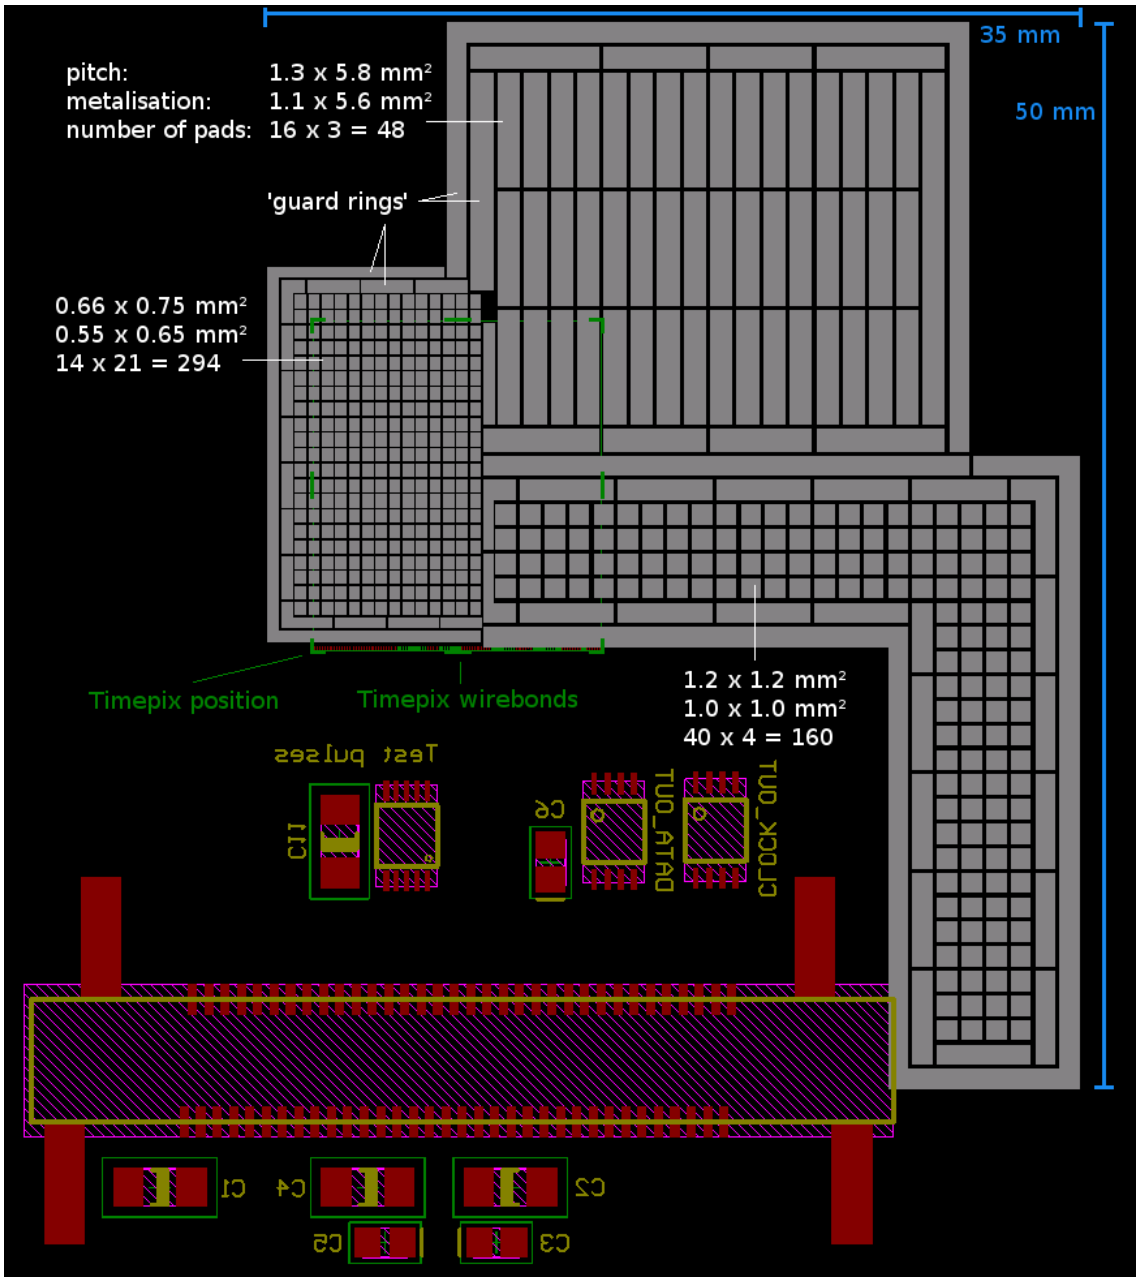

Figure 27: Layout of the active part of the ROPPERI board. The grey areas depict the metalisation with three granularities, each surrounded by two guard rings to mitigate capacitative cross-talk. The position of the Timepix ASIC, which is located on the opposite PCB side to the metalisation, is marked in green. On the side of the ASIC the initial readout electronics are depicted in red and purple.

In order to study the effect of added capacitance to the pixel readout, three different pad sizes with different line lengths are used: 294 pads, placed above the ASIC position, have the minimal achievable pitch of $660 \times 750\,\mu m^2$ (small pads) with negligible line lengths. 48 pads have a pitch of $1.3 \times 5.8\,mm^2$ (large pads), which is about the size of the GridGEM readout, as reference. To test another granularity in between these two, the remaning 160 pads have a pitch of $1.2 \times 1.2\,mm^2$ (medium pads), of which 158 are connected. These pads are placed in the shape of an L in order to test various line lengths.

The respective metalised area for each pad size is slightly smaller, about 200 µm in each direction for large and medium sized pads, and about 100 µm for the small pads. The sensitive pads are surrounded by so-called guard-ring pads, which are intended to absorb capacitative cross-talk of the sensitive pads from traversing tracks, and thus reduce noise in a test beam setup. Beyond these guard rings, the entire surface of the sensitive PCB side is also metalised. This allows to put the full PCB side A on the same anode potential in order to provide a homogeneous electric field in the detector.

A major challenge was bonding the ASIC onto the pad plane via gold stud bonding. This was done in collaboration with Michele Caselle in the bonding lab at the Karlsruhe Institute for Technology (KIT), Germany. It is difficult to bond different substrates, in particular with structure sizes of $O(\mu m)$, which - to the knowledge of the collaborators - may have been unprecedented at the level used in this work. In the gold stud bonding process, a gold wire is attached to a bonding pad with a small needle head using vibration and pressure. The wire is then cut off by a sharp metal edge, leaving a gold stud on the substrate. A schematic of the gold wire and the bonding process is depicted in Figure 28.

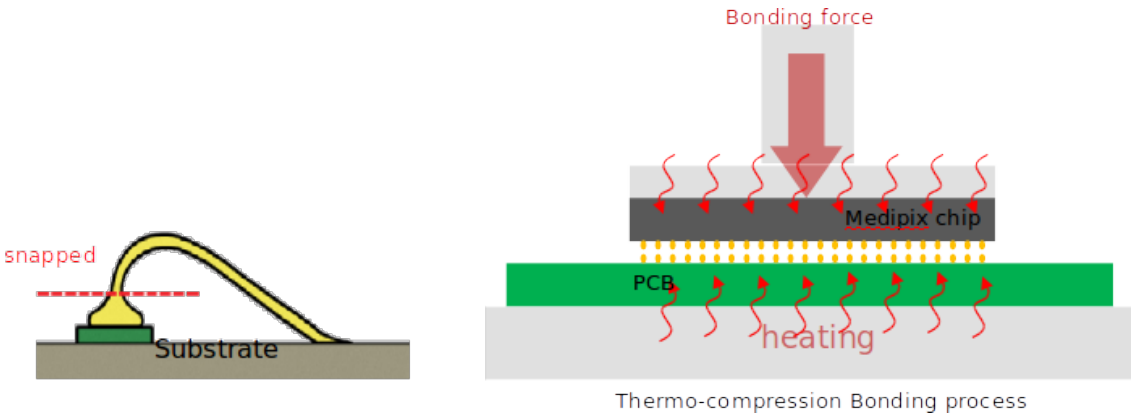

Figure 28: Principle of the bump bonding procedure. Left: bumping. A gold wire is applied with pressure to a substrate, forming a metal bump on the surface, and afterwards cut off with a sharp edge. Right: bonding. The PCB (green) and the ASIC (dark grey) with gold studs in between, applied to both substrates, are pressed together. In addition to pressure, heat is applied from both sides to achieve a stable connection between the gold studs. Images provided by M. Caselle (KIT).

Both ASIC bond pads as well as the respective metalised pads on the PCB were equipped with gold studs. The ASIC bond pads are the ocatognal aluminium openings of the chip's pixels, which are well suited for gold stud bonding. For the PCB, however, a special metalisation coating was necessary, so-called ENEPIG (Electroless Nickel Electroless Palladium Immersion Gold; see e.g. [106]). This is a surface of a few µm Ni, a few hundred nm Pd and a few 10 nm Au, which is well suited for bonding as well as soldering. For simplicity of the production process, this surface coating was applied to all areas that needed to be metalised, i.e. the bond pads, the solder pads and also the anode surface with the sensitive pads on PCB side A.

Figure 29 shows two bumped substrates: on the left, each octagonal pixel opening on the surface of the Timepix ASIC is covered by a gold stud of about 33 µm diameter. On the right, the metalised surface of the PCB is shown. The larger round areas are the blind vias, covered with metalisation, which stretches out from the via in a line to effectively reach a position which will be underneath the ASIC after bonding. At the end of this metal stretch a gold stud has been applied. The blind vias have a little visible dent from filling the via with metalisation. To avoid bonding on this potentially uneven surface, each via is connected to at least a short metal stretch to apply a gold stud, an example of which is the rightmost via in Figure 29. The pitch of the vias is about 600 µm, limiting the number of usable pixels of the Timepix. As opposed to the full pixel coverage shown in Figure 29, left, for bonding only the pixels with respective counterparts on the PCB side were bumped.

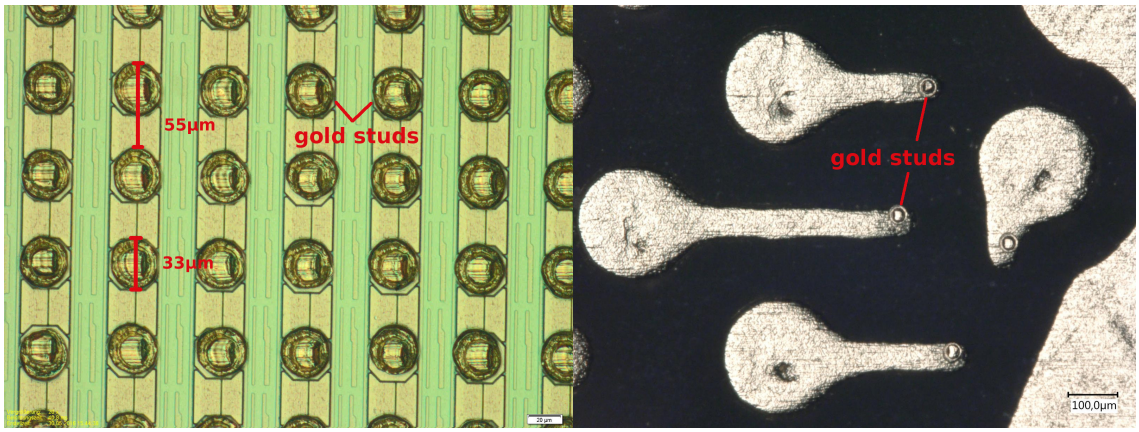

Figure 29: Microscopic images of bumped substrates. Left: Timepix surface. Right: PCB surface with metalisation. Detailed description in the text.

The fully bumped ASIC is shown in Figure 30. The upper part is the pixel matrix with bumps in about every eleventh column and fourteenth row. The outer rows and columns have a higher bump density, since additional vias in the PCB could be placed outside the area covered by the ASIC with lines reaching underneath, as shown in Figure 29 on the right. In the lower part, the wire bond pads are visible on the lower edge of the chip. Each pad has six bumps (2×3), and each counterpart on the PCB side has three bumps in order to provide a larger connection surface and more flexibility for the bonds.

The next step is thermo-compression bonding (Figure 28, right): the PCB is placed on a heating plate in the bonding machine and the ASIC is placed precisely on it with a suction system on a lever arm, see Figure 32. High resolution cameras ensure precise placement of the ASIC down to the µm level. The result of the bonding can be seen in Figure 33, which is a side view of the PCB and ASIC at a shallow angle. On the lower part the PCB surface is visible with the metalisation structure similar to the one in Figure 29 on the right, viewed from the left. The metalisation structure is reflected in the side of the Timepix. Exactly on the boundary between PCB and ASIC, small circles can be seen which are the bonded gold studs in between the substrates. From this angle, only the gold studs on the three outermost rows are visible.

The first bonded board is shown in Figure 34. In this case, also the 'empty' surface on PCB side B is ENEPIG-coated due to production reasons.

The final production step is the equipment. A number of resistors, capacitors, integrated circuits and the VHDCI connector are soldered to the designated solder pads on PCB side B, corresponding to the functionality of IMB v10 (comp. Figure 24). The fully equipped board, in this case without the additional side B coating, is shown in Figure 35.

Production, bonding and equipping was an iterative process. Two production orders were performed, one with 3 boards and one with 20, called first and second prototype. Initially, the first board was bump bonded at KIT, then shipped to DESY, equipped and read out. This readout did not deliver usable data. Since equipping causes some thermal stress to the bonded board, it was decided to first equip the boards and then bond them, after it was clear that the equipment would fit inside the bonding machine. To avoid loss of usability during transport, it was decided to perform readouts at KIT immediately after bonding. During readout, the first prototype broke due to thermal stress and the second production was done with a different base material at a different company, which also caused the difference in ENEPIG coating to side B. This is elaborated on in the following chapter.

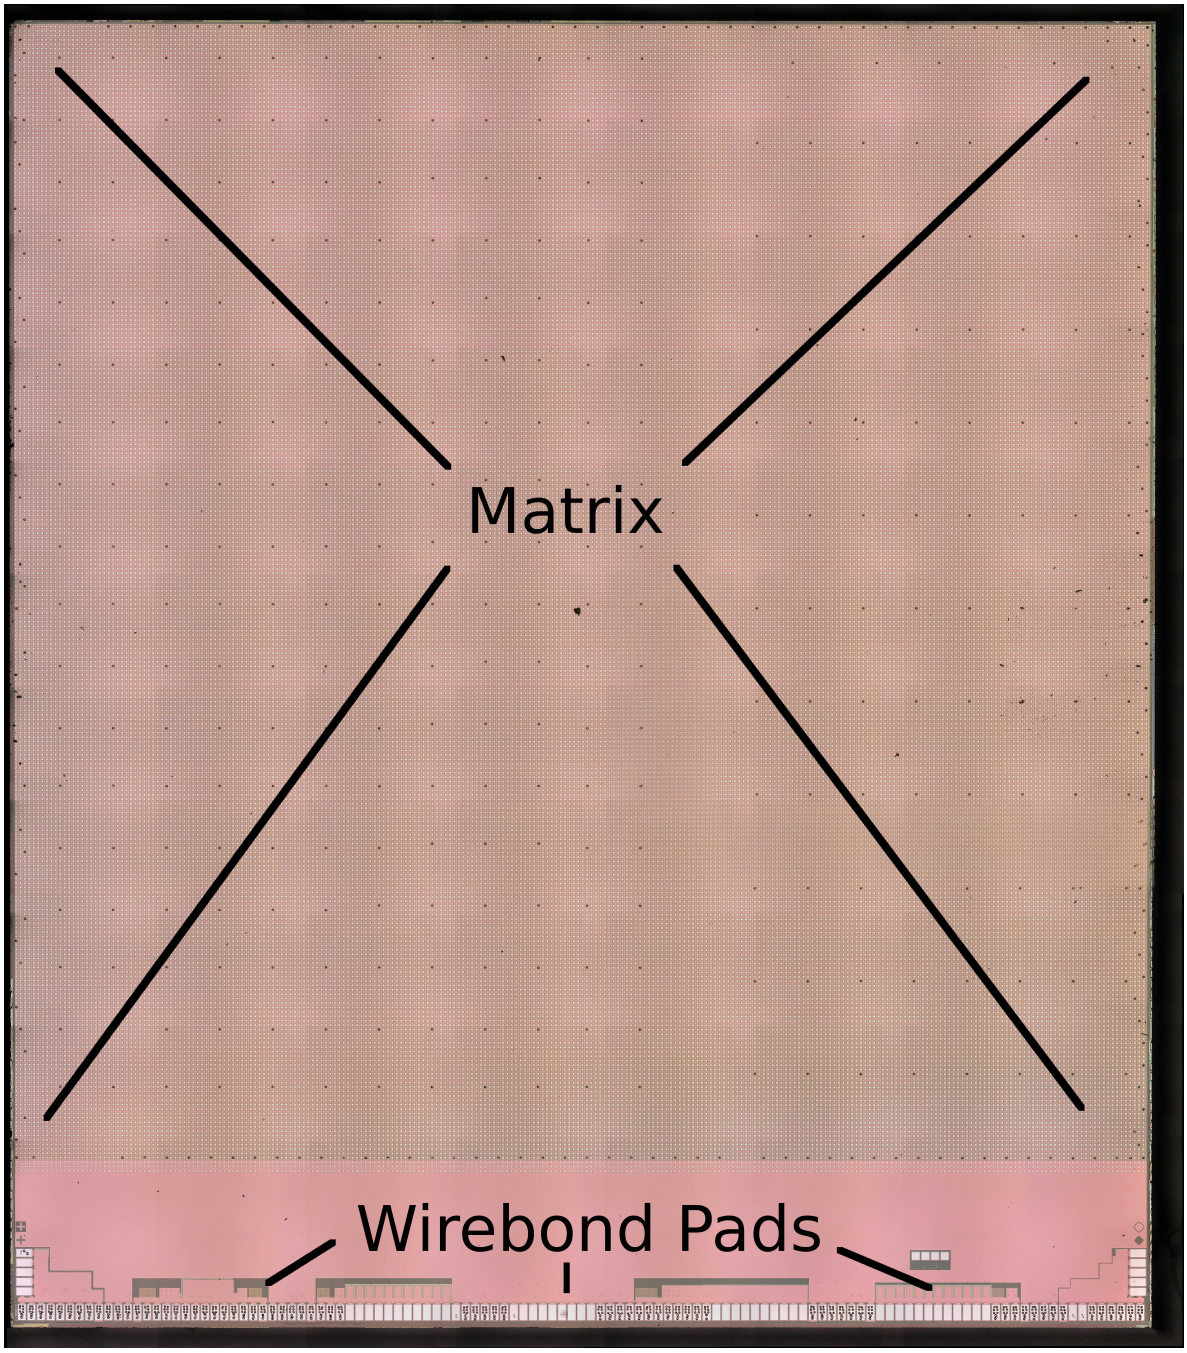

Figure 30: Microscopic image of the bumped Timepix ASIC.

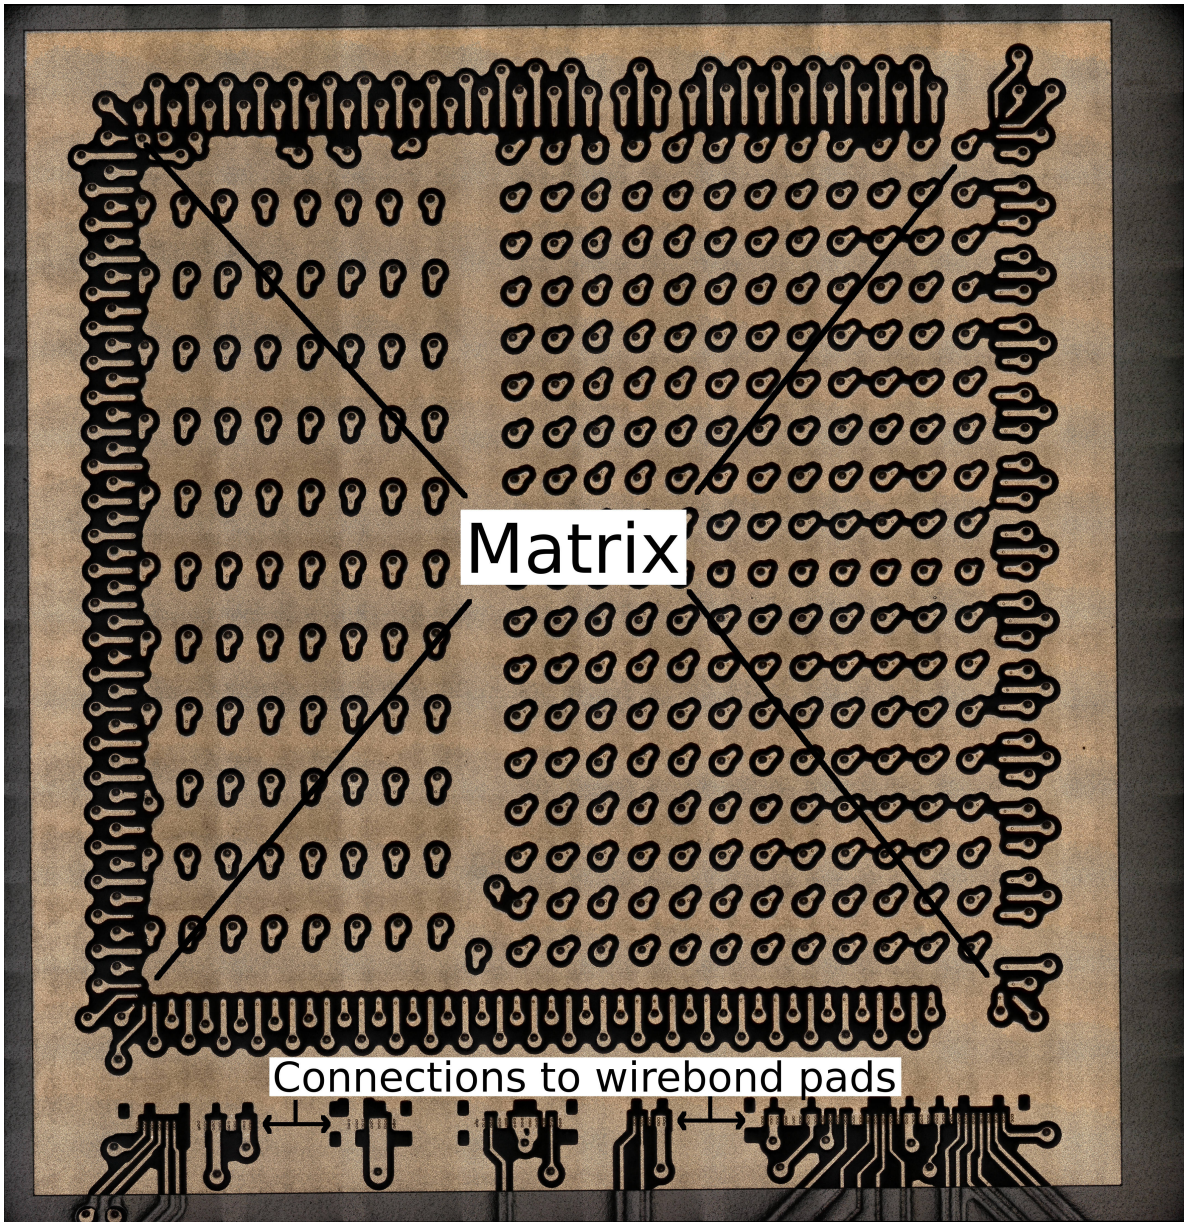

Figure 31: Microscopic image of the bumped area of the Ropperi PCB.

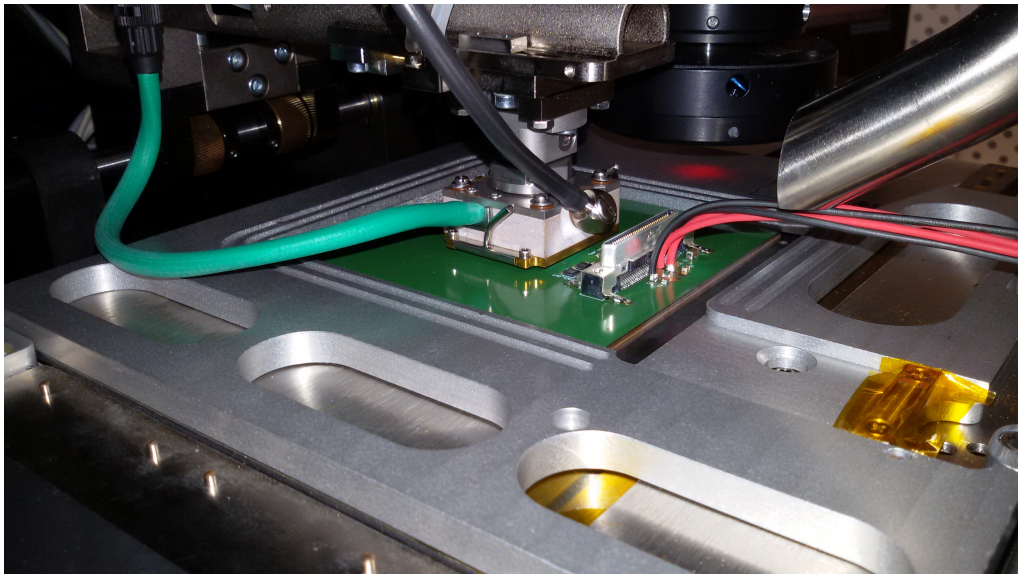

Figure 32: Ropperi board during the bonding process inside the bonding machine.

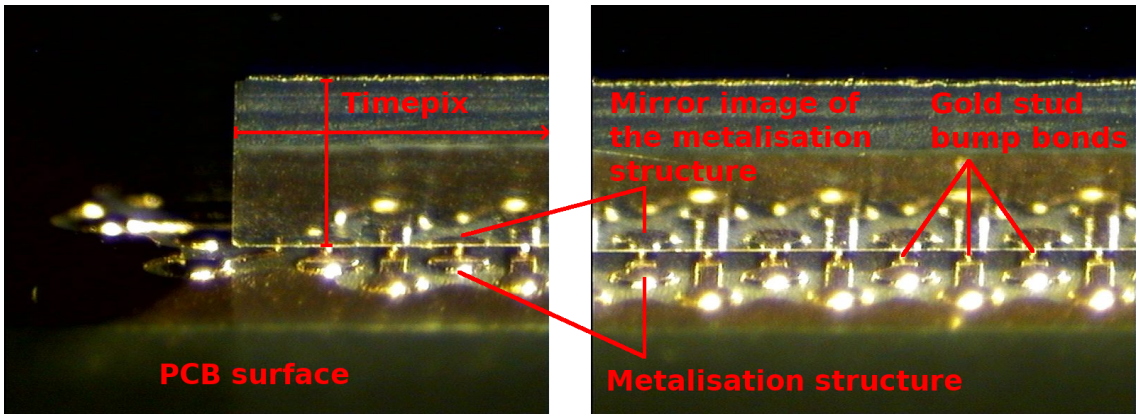

Figure 33: Microscopic side view of the bonding area of a bump bonded ROPPERI board.

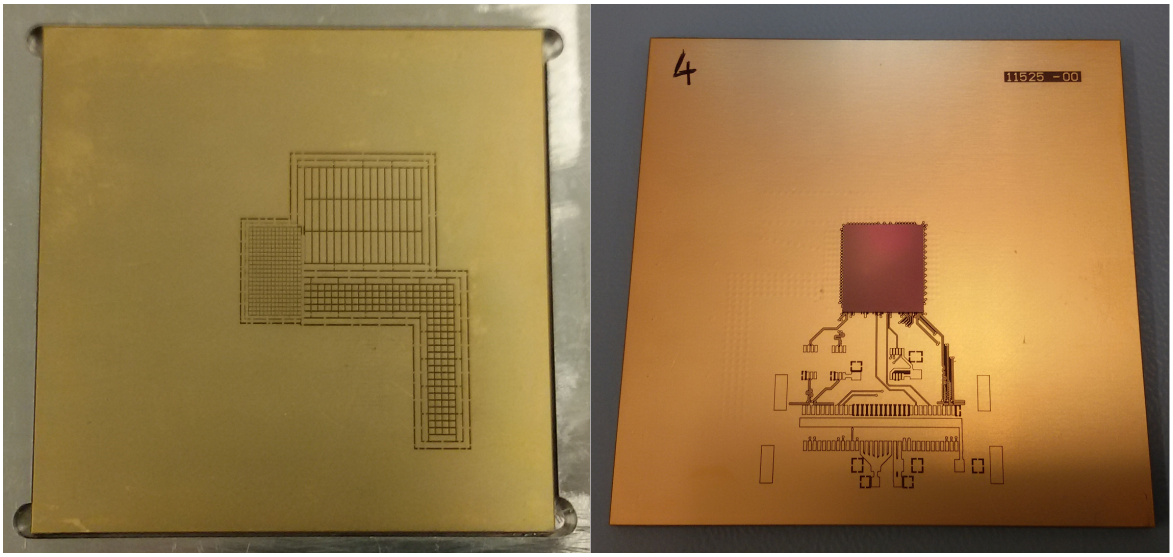

Figure 34: First bonded board. Left: side A with sensitive pads. Right: side B with Timepix (brown) bonded to the board, but without electronics elements and connector. In case, also side B was fully gold-coated.

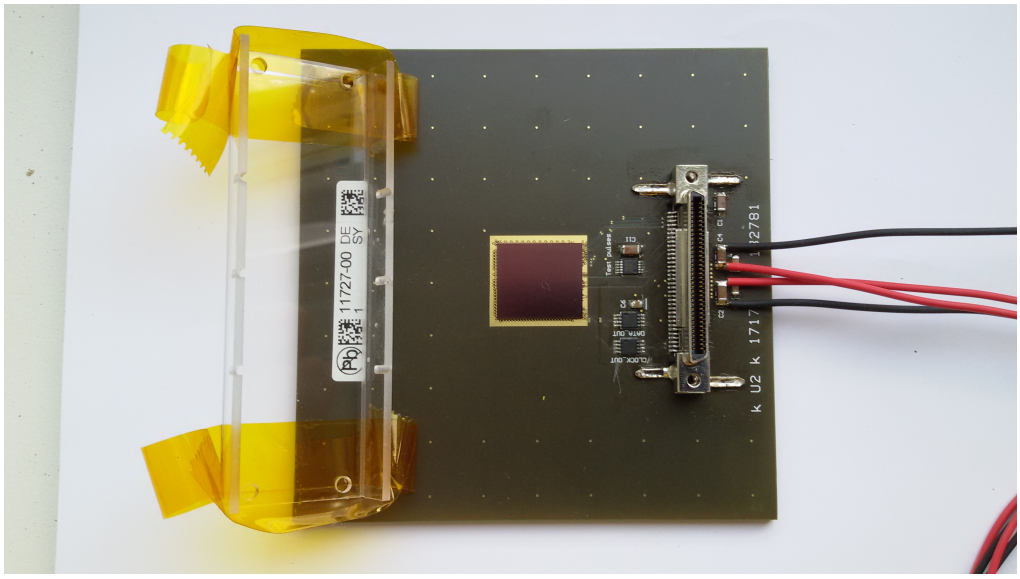

Figure 35: Fully bonded and equipped Ropperi board.

5.3.2. Basic Measurements and Conclusions from the First Prototype

Immediately after the bonding process, a first TOT readout of noise data was taken. To avoid damage by moving the bonded system, the very first measurement was done with the board lying on the bonding machine, with the sensitive pads facing its metal surface. Due to this, the readout data showed enormous noise. A very large threshold of $> 700\,\mathrm{ADC}$, compared to the usual threshold levels of unconnected Timepix chips within the given readout setup of typically $250 - 500\,\mathrm{ADC}$, was necessary to suppress the noise and have the pixels not be continuously active during the readout frames. This can be explained by the direct connection or capacitative coupling of the sensitive pads to the bonding machine, thus picking up all statistical noise stemming from the large capacitance of its metal body with respect to ground. Nevertheless, every channel of the Timepix that was supposed to be connected to a pad was noisy, all other channels were silent. This means, a connection to every target channel had been established, and cross talk between connected and unconnected pixels was not visible.

After that, the board was carefully removed from the bonding machine and a second measurement was performed on-site. The necessary threshold to suppress remaining noisy channels was around 380 ADC, i.e. much lower than before and within the usual range. It was possible to take one relevant and reasonable data frame at this threshold, then the system broke and no consistent data taking was possible anymore.

The result is shown in Figure 36. On the left, the active channels during this data frame are shown in blue, while the silent ones are white. It is noteworthy, that most channels which are connected to sensitive pads close to the Timepix chips, in particular the small ones 'on top' of the ASIC, are mostly silent, while the ones which are further away show more noise, consistent with the expected line-length dependence. The pads with larger line lengths are also the medium and large ones, which could also explain their larger noise. However, within these groups a separate line-length dependence is visible. This is quantified in Figure 37, where for all pixels the relative noise count is shown in dependence of the line length of that pixel, after histogramming and for two different bin sizes. A correlation is visible, showing very little noise at line lengths below 10 mm, and large noise up to the maximum count for larger line lengths. It is concluded that:

1. A dependence of the noise on the line length is visible.
2. The necessary threshold level to illustrate this dependence is within the usual range of $250 - 500\,\mathrm{ADC}$ for bare Timepix chips and not far above that range (compared to e.g. the necessary threshold during the readout on the bonding machine). This makes it likely that the noise level is in fact manageable and does not prohibit the intended use. Both, dependence and absolute level, need to be investigated in further detail.

After recording this single readout frame, no change in readout data was visible anymore, independent of threshold. This behaviour had been observed earlier with other boards where the bonding had failed. The board's supply voltage was switched off and back on a few minutes later. At that point, regular communication with the ASIC was possible again briefly, but it broke down before another data frame could be taken. This was repeated several times with decreasing lengths of stable communication. It was

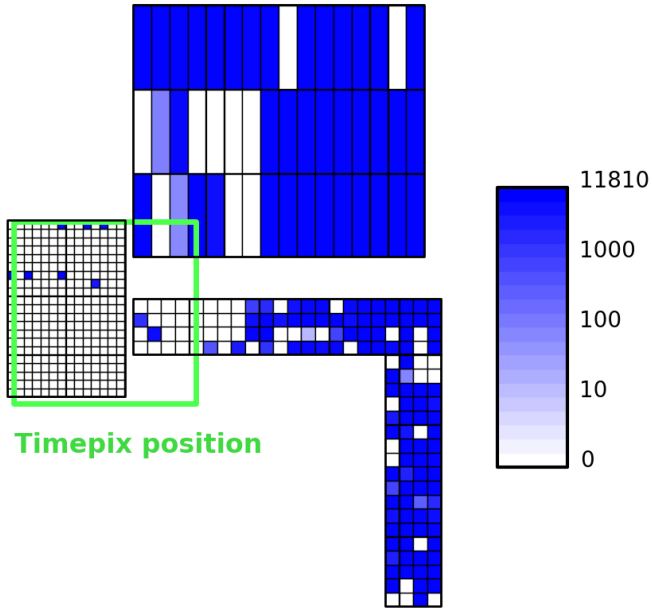

Figure 36: The single readout frame of the first successfully bonded board, with data values displayed color-coded on the corresponding readout pads.

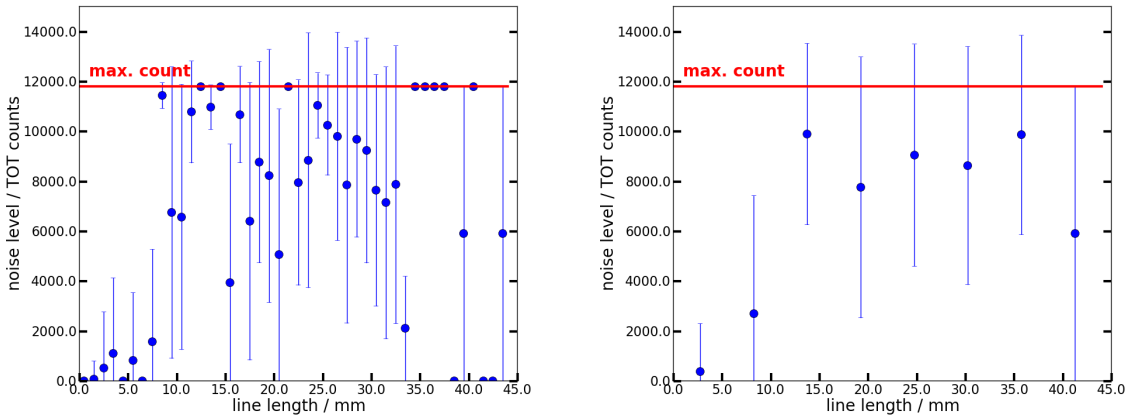

Figure 37: Dependence of the recorded noise on the line length derived the from the single readout frame of the first successfully bonded board, for two different binnings. For low line lengths the noise activity is low, for line lengths above about 10 mm the noise is close to the maximum possible count.

concluded that this time-dependent behaviour occured most likely due to temperature effects which can be (at least partially) reversed by cooling.

The PCB material FR-4 and silicon have different coefficients of thermal expansion (CTE): silicon expands approximately by 2.5 ppm/K, while FR-4 has a CTE of 16 ppm/K. At the temperature difference of about 230 °C between bonding and room temperature, the difference in CTE leads to a difference in thermal expansion of about 40 µm over the 14 mm width of the ASIC. This is large compared to the pitch of 55 µm and the size of the bonding structure of about 30 µm. Due to the resulting stress caused by the thermal expansion, enforced by the heating up during operation, it is suspected that the connections to the communication pads did break, which rendered the readout system in an uncontrollable state.

The board in question was scanned with X-rays, but the resolution of this image did not suffice to recognise if a connection was broken.

5.3.3. Basic Measurements and Conclusions from the Second Prototype

A second board generation was developed, to mitigate the issues suspected to originate from the very different thermal expansion coefficients of the chip and board material. Therefore, N7000-2 HT (see chapter A.3) was used, which has a CTE of 12 ppm/K, reducing the thermal stress. Also, the bonding temperature was reduced. With 20 boards delivered, 10 bonding attempts with different temperature profiles were done, with the maximum temperature between 120 °C and 200 °C. With several boards some level of communication could be established at least for some time, but a full data readout was again rarely possible. In the end, noise data from 3 boards could be successfully taken immediately after bonding at KIT bonding lab. In order to secure mechanical stability during shipping, underfill was applied to the bonded systems. Later in the DESY lab, a cooling element with a fan (Figure 38) was placed on the bonded ASIC with thermal paste to allow for active cooling in order to reduce the thermal stress during operation. Only one of the three boards was still operational and another noise data scan was successfully taken. After that, any further attempt to read out the bonded boards was unsuccessful. They had transitioned into a state where some communication with the ASIC was still possible, but not fully and in particular no regular data readout could be achieved. Some data was still generated, as opposed to other former attempts where there was no signal on the data line. However, the data most likely contains various bit shifts of unknown origin, compare Figure 39. No further analysis was possible.

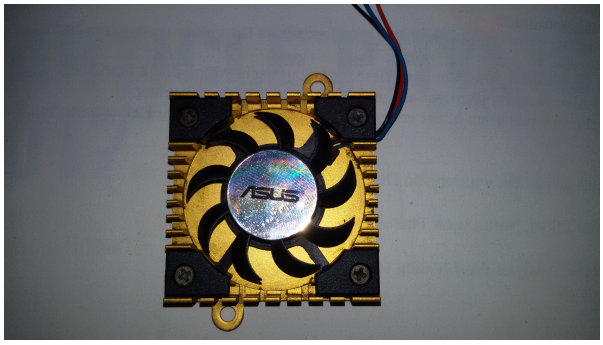

Figure 38: Active cooling element, which was placed on top of the bonded chip in order to cool it and limit the thermal stress during operation.

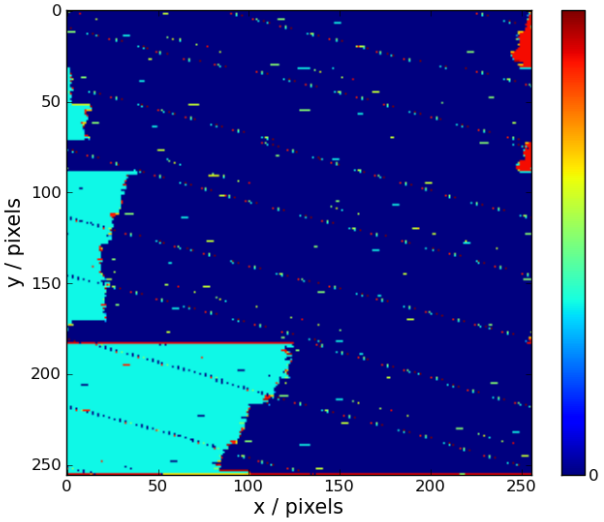

Figure 39: Timepix readout matrix after the chip connection was partially broken, likely containing various bit shifts.

5.4. Noise Analysis

Due to the limited mechanical stability of the system, is was not possible to supply the system with a defined signal to the pads and do a dedicated signal-to-noise analysis. However, right after bonding it was possible with a few of the boards to take data of random noise in the system at different threshold levels. With growing threshold, the activation level of the system declines: at very low threshold levels nearly all channels are constantly active, while at high thresholds most channels are silent. By quantifying this behaviour, it can be compared to the same data of a bare Timepix ASIC with its equivalent noise charge (ENC) of about 90 electrons. The noise drop with increasing threshold is qualitatively shown in Figure 40.

A threshold campaign was performed immediately after bonding for the three successfully bonded boards, which in the following are called boards A, B and C. For each applied threshold level (THL), 100 frames were taken with a time window of 256 times the clock frequency. The result is displayed in Figure 41 as a curve of the number of frames, in which a certain TOT count was recorded, vs. that TOT level. For very high or low threshold levels, all frames had a TOT value of 0 or 11810, respectively. In the intermediate THL region, a number of frames had a TOT value between the minimum and maximum. The values of these curves have been averaged by calculating either the arithmetic mean or the median of the noise value of the 100 frames, and are plotted against the THL in Figure 42 for an example case. Here, the noise drop is clearly visible and happens within a limited range of THL values. The width of this drop, from the largest THL with all frames at maximum TOT count (THL drop lower edge) to the smallest THL with all frames at 0 TOT count (THL drop upper edge), characterises the amount of random noise in the channel. For board C, these edges are shown in Figure 43 for one selected channel. A channel is not used further in the analysis, if the drop is not complete within the available threshold range. Unfortunately, at the time of the measurement, the scanned threshold range was chosen between THL values of 380 and 600 for board A and between 300 and 400 for board B. These ranges turned out to not cover the full drop for the majority of channels, but only 109 of 500 channels for board A and 119 for board B, see Table 3. The chosen THL range for board C was between the values of 300 and 1000; in steps of 10 between 300 and 400, in steps of 20 between 400 and 500, and in steps of 50 between 500 and 1000. This range was sufficient to cover the full drop for most channels, 422 of 500. In order to compare the scans of the Ropperi boards with a bare Timepix, a scan with the same readout settings was performed. The THL scan range and steps in this case were adjusted to be in steps of 5 from 200 to 300 and in steps of 10 from 200 to 380, covering full drops for 48754 of the 65536 Timepix channels. The chip that was used is the one displayed in Figure 25, which has some broken columns and individual pixels. These were automatically omitted in the analysis via the full-drop requirement, still leaving a sufficient number of usable channels for the analysis.

The THL drop width is different for each channel. Its abundance for all available 422 channels for the threshold campaign of board C is shown in Figure 44. This data can be subdivided depending on the expected capacitance into the 3 different pad sizes, small,

board	THL values used	number of usable channels depending on pad size (max)			
		small (294)	medium (158)	large (48)	total (500)
A	380, 400, 420, 450 500, 600	17	65	27	109
B	300, 310, 320, 330, 340, 350, 361, 370, 380, 390, 400	67	30	22	119
C	300, 310, 320, 330, 340, 350, 360, 370, 380, 390, 400, 420, 440, 460, 480, 500, 550, 600, 650, 700, 750, 800, 850, 900, 950, 1000	252	129	41	422

Table 3: Scan ranges of the threshold level ranges and numbers of usable channels for the three boards.

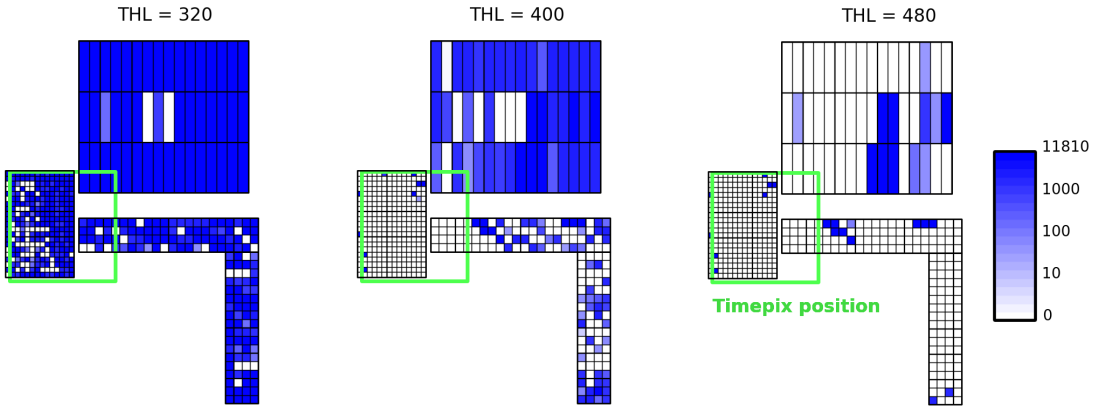

Figure 40: Example of recorded charge values dropping with increasing THL, here for board C and for 3 representative THL values.

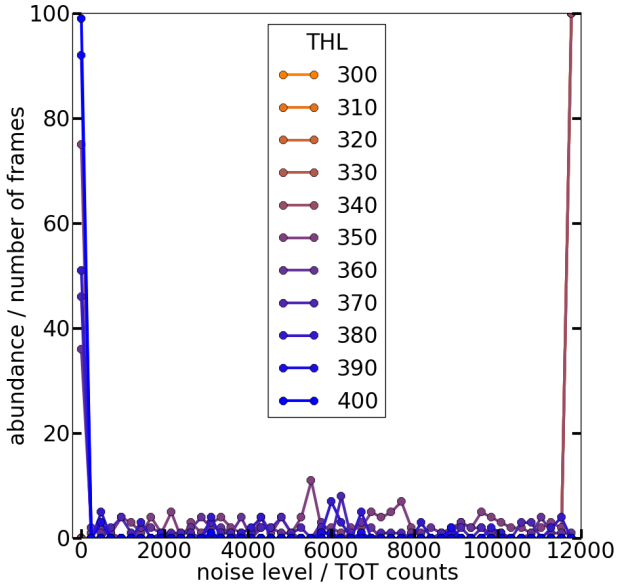

Figure 41: Abundance of recorded TOT values in 100 frames, for different THL, for one examplary channel of board B. Note that for THL values below 330 all frames are at the maximum TOT count and the curves overlap.

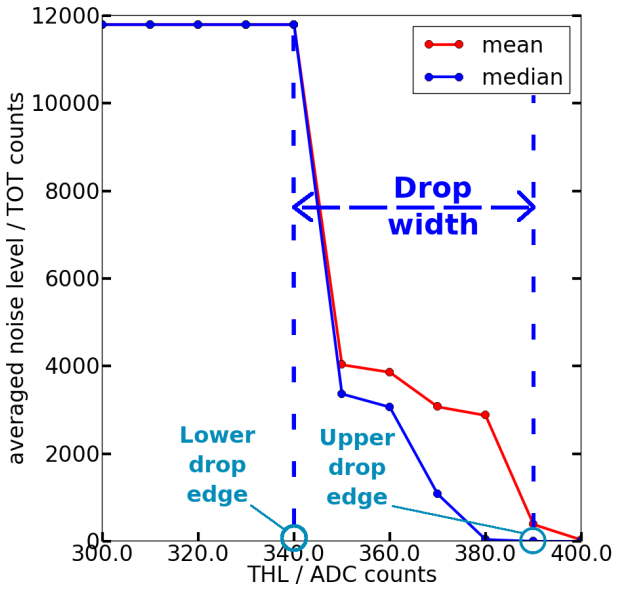

Figure 42: Averaged noise level for different THL values for one selected channel of board B. The drop with is defined as the change in THL over which the average noise level goes from the maximum count to zero, here depicted for the *median*.

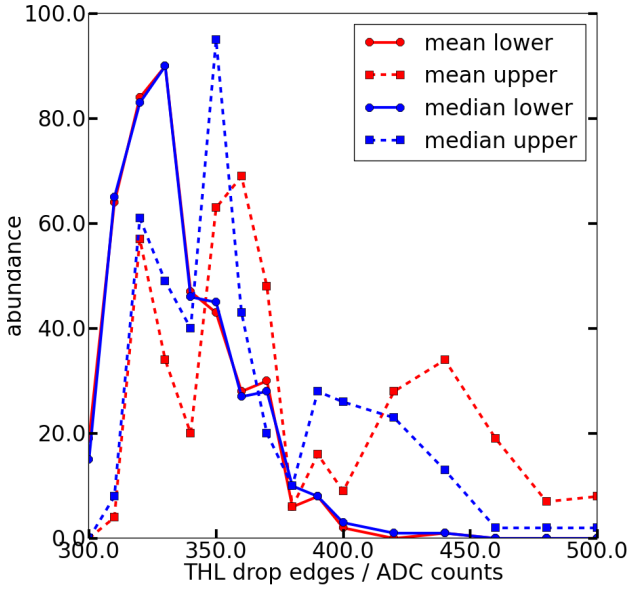

Figure 43: THL drop edges: largest THL values at which a channel was fully active (lower edge) and smallest THL value at which a channel was fully silent (upper edge), here shown for board C.

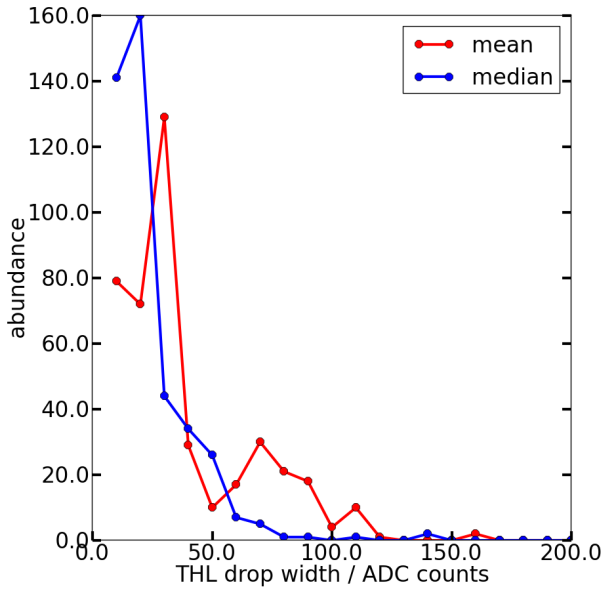

Figure 44: THL drop width: THL change required for channels to go from maximum to 0 activity, here shown for board C.

medium and large. In addition, the medium pads are split into two brackets of channels with a line length below and above 20 mm, which is half of the maximum line length. Figure 45 shows for board C the dependence of the THL drop width on the pad size and on the line length bracket. As expected, larger pads have a larger THL drop width than smaller ones, and the medium pads with a line length above 20 mm have a larger THL drop width than the ones below 20 mm. The same holds true, when the data of board C is combined per-pad-size with the limited data sets of boards A and B, shown in Figure 46. Note that the drop width can only take a value which is a possible difference between two THL values in either of the three scan campaigns. This also means that the drop cannot be smaller than the smallest THL difference, i.e. 20 ADC counts for board A and 10 ADC counts for boards B and C.

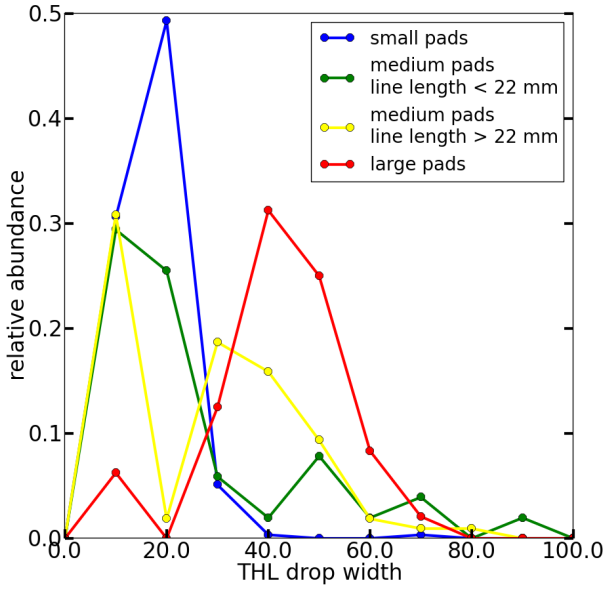

Figure 45: THL drop width of board C using median method, subdivided into pad sizes and two line length brackets for the medium sized pads.

Finally, the combined data of the three boards can be compared to data taken in the same way from a bare Timepix ASIC, which apart from its known ENC of $90-100\,e^-$ does not have any additional noise. Figure 47 shows this comparison.

The peak of the bare Timepix chip is estimated to be about 5 ADC counts, but this is limited by the scanning steps of 5 ADC counts. With an ENC of $90-100\,e^-$ and given the scaling factor of about $25\,e^-$ per 1 ADC count for the Timepix, it is likely that the peak lies rather at 4 ADC counts, which can not be resolved. The THL drop width of the small pads, which have nearly no line width, peaks at a mean value of 17.3 ADC counts with an RMS of 8.1 ADC counts. If the RMS is used as upper limit on the possible fluctuation of the mean, this is equivalent to $(432 \pm 202)\ e^-$. The drop width of the large pads peaks at about 50 ADC counts, equivalent to about $1250\,e^-$. The medium

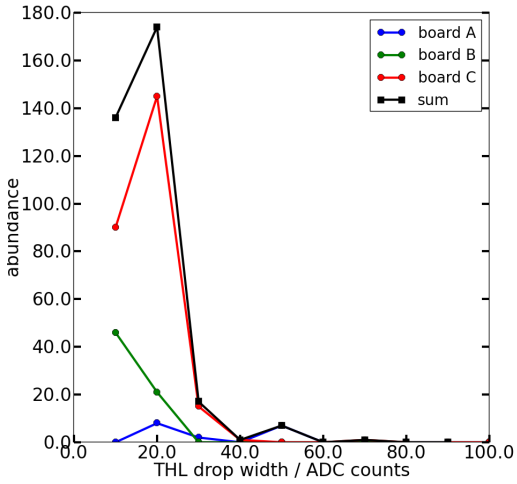

(a) small pads

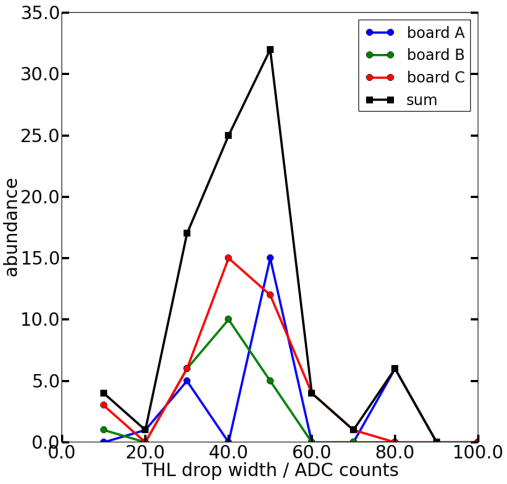

(b) large pads

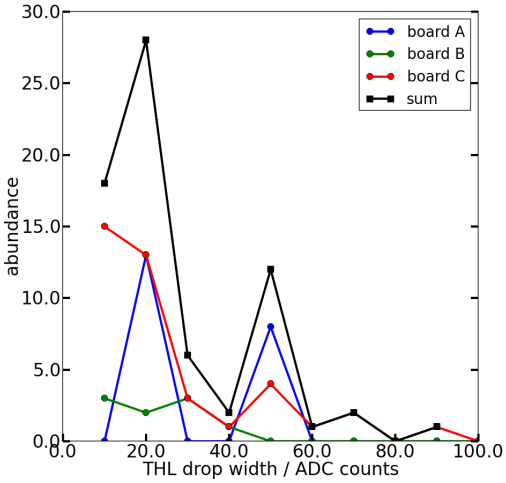

(c) medium pads, line length $< 20\,\mathrm{mm}$

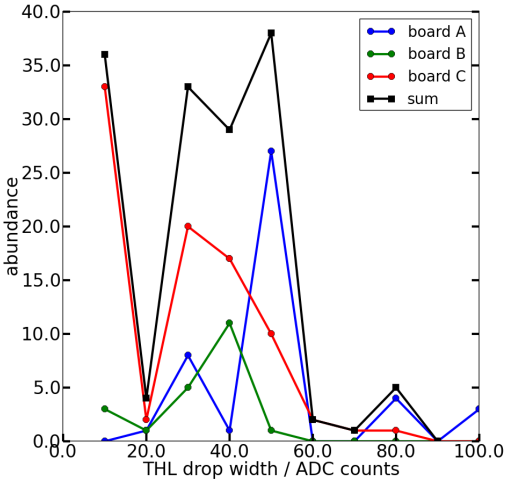

(d) medium pads, line length $> 20\,\mathrm{mm}$

Figure 46: THL drop width for all 3 boards and their combinations, for different pad sizes and line length brackets.

pads, both for short and long line length, do not have a single clear peak, but their mean drop width lies at (27 ± 18) ADC counts $= (675 \pm 444)$ e^- and (34 ± 18) ADC counts $= (850 \pm 440)$ e^-, respectively.

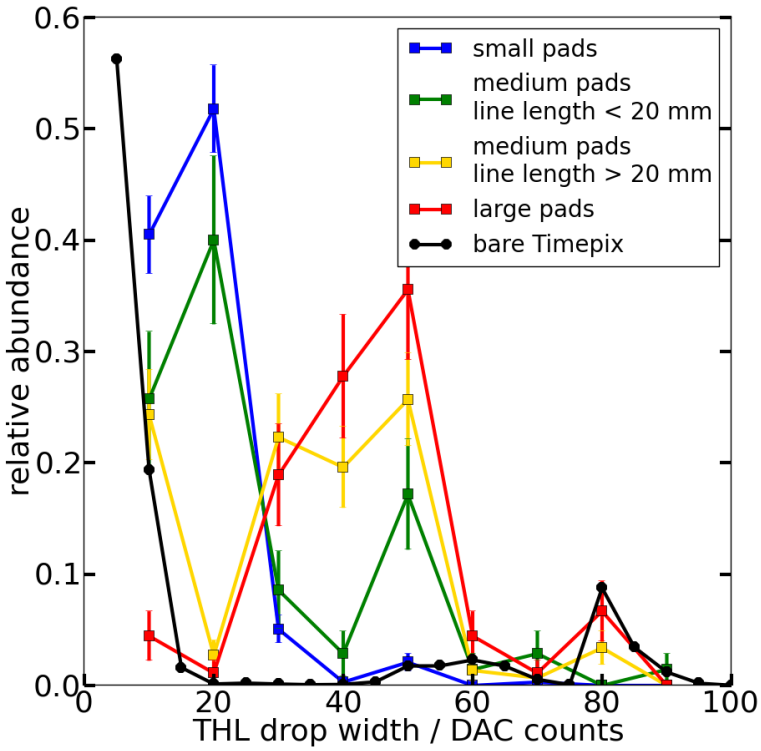

Figure 47: Comparison of the THL drop width between sums of the 3 boards divided by pad size and line lengths, and a bare Timepix ASIC.

The data points are rather widely distributed and allow only a limited amount of interpretation. This is due to the limited amount of usable data and the coarseness of the threshold campaign scans. The applied threshold steps of 10, 20 or 50 ADC counts in the three performed scans have an impact on the possible THL drop width values. This explains the peaks at 20 and 50 ADC counts, while 30 and 40 ADC counts are significantly less populated, which is particularly pronounced for the medium pads size.

Even with the somewhat coarse and limited data set, several conclusions can be drawn. As expected, the larger the pads and the longer the line lengths are, the larger is the noise and the width of the corresponding THL drop. The small pads have about 5 times the noise of a bare Timepix, while this factor is about 12 for the large pads. The statistical error on these values is in the order of 10 %, given the cumulative statistics for the different pad sizes of between 100 and 300 entries. This is smaller than any of the systematic uncertainties introduced by the scanning ranges, scanning steps and individual differences of the chips and systems. These could only be evaluated with additional data, i.e. more statistics from additional boards as well as broader scan ranges with

smaller step sizes.

To assess the possible signal-to-noise (S/N) ratio of the system, a signal from a GEM-based system is simulated, analogue to the simulation studies in chapter 6. In typical gas compositions about $10\,e^-$/mm are generated in the primary ionisation, and GEM charge clouds spread over about $1\,mm^2$ with an amplification gain of about 10 k. The pad dimensions are given in Table 4. For the small pads, on average $7\,e^-$ from pimary ionisation arrive at the pad. However, the pad only covers the charge cloud to about 70 % in either longitudinal and transversial direction of the track. Charge transfer to other pads in track direction is compensated by charge received from there, but this is not the case for the pads transverse to the track direction, and a reduction factor of 70 % has to be taken into account. This means that about 49,500 e^- arrive on the small pad with a noise of $432\,e^-$, allowing for a S/N ratio of about 115. For the medium and large pads, one pad width would fully cover a charge cloud hitting it centrally, so no transversal reduction factor has to be applied. For the large pads, it is assumed that their long side aligns with the direction of the incident particle, which is the typical case for the GridGEM system. The S/N ratio of medium sized pads is 178 for short line lengths < 20 mm and 141 for long line lengths > 20 mm. Due to their size, the large pads collect several hundred thousand electrons from an incident track, resulting in a S/N ratio of 527. The results of this assessment are summarised in Table 5. Given a signal size for each pad size, one can also calculate the allowed noise N_{crit} at which the S/N ratio is 10. This is provided in the same table and displayed in Figure 48. For each pad size, the corresponding N_{crit} is at least one order of magnitude larger than the expected signal.

	pad pitch	pad area
small pads	$0.66 \times 0.75\,mm^2$	$0.495\,mm^2$
medium pads	$1.2 \times 1.2\,mm^2$	$1.44\,mm^2$
large pads	$5.8 \times 1.3\,mm^2$	$7.54\,mm^2$

Table 4: Pad sizes of the Ropperi system.

In addition to the factors in Table 5, also the potential diffusion of primary electrons in the drift volume as well as signal induction of the GEM charge cloud into neighboring pads should be taken into account in this assessment. For the GridGEM system (chapter 3.2), the overall pad response function has a width between 0.5 mm and 1 mm, depending on the drift length [40], and the charge is typically spread over 3 to 4 pads. In the case that a track traverses a plane of pads with a width of 1.3 mm centrally between two pads, these two pads would each receive 40 % of the charge, and their immediately neighbouring pads each only 9 %. Even collecting only 9 % of the charge with medium or large pads would still be sufficient to reach a S/N ratio of more than 10. The typical row-based hit reconstruction, however, combines these 4 neighbouring pads into one hit with 90 % of the charge at only twice the noise, resulting in an adjusted S/N ratio of still more than 60 for medium and 200 for large pads.

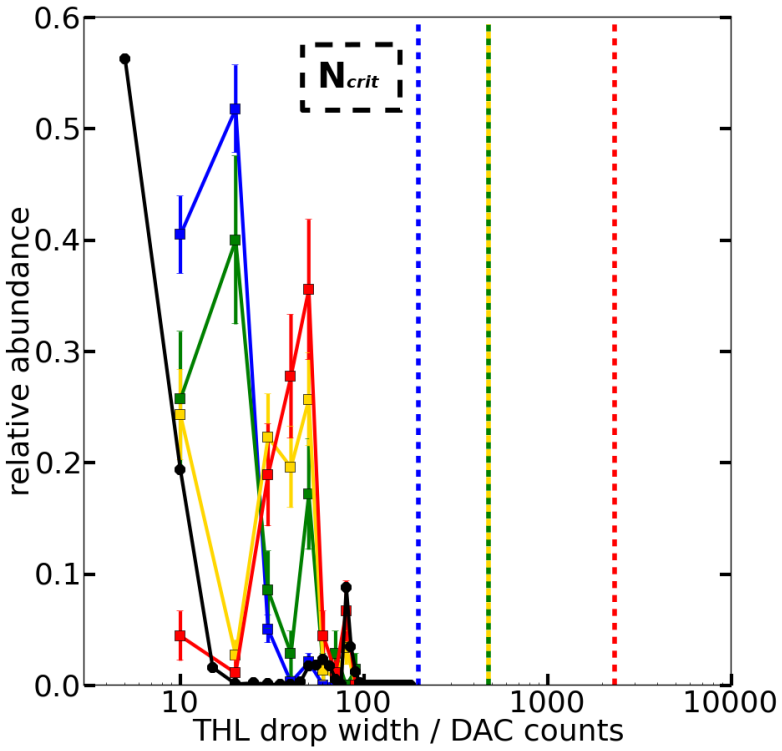

Figure 48: Comparison of the THL drop width of Figure 47 and the maximum allowed noise N_{crit} from Table 5. The colour of the vertical line corresponds to the pad size. It shows that there is at least a factor of 10 (note the log scale) between observed and allowed noise for each pad size.

	N_{el}	S / e^-	N / e^-	S/N	N_{crit} / e^-
small pads	6.6	49,500	432 ± 202	115 ± 54	4,950
medium pads I (line length < 20 mm)	12	120,000	675 ± 444	178 ± 116	12,000
medium pads II (line length > 20 mm)	12	120,000	850 ± 440	141 ± 72	12,000
large pads	58	580,000	1110 ± 368	527 ± 177	580,000

Table 5: Signal-to-noise ratio of the Ropperi system.
N_{el}: number of ionisation electrons above pad,
S: assumed signal on pad from GEM amplification,
N: estimated noise in pad,
N_{crit}: maximum allowed noise with S/N = 10.

For the small pads and pads below their size, the situation is more complicated. On the one hand, the charge of one GEM charge cloud spreads over several pads and diffusion distributes the primary electrons over more pads, which reduces the charge on any individual pad. On the other hand, in a data analysis the information of several pads can be combined into one reconstructed charge cloud, which reduces the impact of noise. As an extreme example, the signal of a single primary electron would be distributed with a PRF of 0.5 mm, and about 40 % of the charge would be distributed over an area of 0.785 mm^2, corresponding to one radial width. With a pad size of $200 \times 200\,\mu m^2$, this fraction of the charge would be distributed over 20 pads, with a combined noise of about $432\,e^- \cdot\sqrt{20} \approx 2000\,e^-$, assuming the noise level of the small pads. This indicates that a gain larger than 10 k would be required, and instead a gain of 50 k would suffice to still provide a signal of $1\,e^- \cdot 50,000 \cdot 0.4 = 20,000\,e^-$ and a S/N of 10 in this case. Such a gain value is in principle achievable with GEM amplification, but would require careful attention to avoid discharges and could generate a challenging amount of ions flowing back into the drift volume towards the cathode. The effect of varying the pad size and thus the signal distribution is presented in the detailed simulation in chapter 6.

5.5. Conclusion

The Ropperi system is an approach to combine advantages of existing TPC readout structures, in particular high granularity and a large anode flexibility and coverage, for the case of amplification with GEMs. Within this work, a collaborative effort was established to design, produce and test such a readout system. Intermediate pad sizes of about 700 µm, between conventional pad-based readouts with a size of 6 mm (in track direction) and pixel-based ones with 55 µm, were used. A PCB could be produced that contained these pads, a fan-out with vias and lines in a multi-layer structure, as well as bond pads for bump bond connections to an ASIC, which was key to the hybrid approach. The chosen ASIC was the Timepix chip, which was bonded to the PCB via gold stud bonding. The challenge to connect a silicon chip to a conventional printed circuit board may have been unprecedented, and it took several attempts until a sufficiently stable readout could be established. All channels were accessible and delivered data. Unfortunately, after all improvements which were implemented throughout the bonding campaigns, the system was still only stable for hours or, in one case, days. The most likely cause for the end of operability was breaking of the bonds due to thermo-mechanical stress which arose after bonding and during operation from the different materials and the corresponding coefficients of thermal expansion. A future development of Ropperi would need to address this issue. One possibility would be the use of an interposer, e.g. from silicon, which could buffer the different CTEs, or using a ceramic PCB with a CTE close to silicon. The first option would add significant additional material to the system, while the second was doubted to be able to provide the necessary flatness for bonding. Both options would be significantly more expensive than the PCB options used in this work. Since it was observed that during stable operation all channels were still accessible, it is concluded that - despite the three-fold redundancy - the connections to the wire bond pads broke, which are needed to communicate with the Timepix. With through-silicon vias, a technology which is available for the Timepix 4, the communication channels could be connected with wire bonds on the upper side of the ASIC, which are not sensitive to temperature fluctuations. This way, it could be tested if the bump bonds of the pixel matrix are sufficiently stable for a long-term readout. This means a time long enough to position an operational Ropperi board with GEMs in a TPC and provide a signal from incident particles, which can then be compared to the noise level in order to make a direct and full signal-to-noise measurement. In preparation for such a potential measurement, a corresponding small TPC, the Unimocs, has been set up and an interface to the Ropperi hardware has been designed and produced within this work. Another reason to target a ceramic base material for the PCB is the possible feature size. In order to utilise the full ASIC, each pixel of 55×55 µm needs to be connected and fanned out through the PCB, which comprises a challenging requirement.

Despite the limited life-time of the developed prototype, a threshold campaign was successfully conducted in order to assess the equivalent noise charge in the Ropperi system, depending on the pad size and the line length. The result shows that a typical signal which would be expected from GEM amplification would be more than a factor of 10 larger than the noise, which is deemed the required threshold. For a strong, but

achievable amplification gain of 10k the S/N ratio is typically larger than 100, which leaves room for a less beneficial situation. For example, a full ILD-module sized ($23 \times 17\,\mathrm{cm}^2$) Ropperi PCB would need to contain line lengths of 15 to 20 cm, which would be about 10 times the average length investigated with the Ropperi medium pads. Since line length increases the capacitance and noise linearly, this increase of the line length could reduce the S/N ratio from around 150 to about 15, which would still be sufficient. The target pad size of Ropperi is smaller than the tested pad sized, which would decrease the noise and thus increase the safety margin.

The overall conclusion is, that the Ropperi approach of a hybrid TPC readout with an ASIC bonded directly to a PCB and GEM amplification is viable in terms of signal-to-noise. The concrete hardware implementation, however, has proven to be a major challenge and the developed prototype did not have a sufficient stability. The issue with the current hardware approach was identified as a large difference in CTE originating from the different base matierials of PCB and ASIC and solutions were discussed. The most promising way forward is the development of a ceramic PCB with a CTE similar to silicon. Feature size limitations, flatness, mechanical machinability as well as cost are central focus points for a possible future study. Ropperi could become a 'plan B' option for a highly granular ILD TPC readout, in case the GridPix system runs into a show-stopper, e.g. in terms of data load, material budget or anode coverage.

6. Detailed Simulation of a Highly Granular Readout for a TPC

To investigate the impact of higher granularity on particle identification capabilities of a time projection chamber, a detailed simulation was carried out. Aside from granularity, a number of parameters were varied in the simulation to study their impact and optimise the PID performance.
In the following, the simulation and reconstruction chain is laid out, the role of the parameters is explained and the results on the PID performance are presented.

6.1. Simulation and Reconstruction Chain

The simulation and reconstruction chain consists of three separate major parts: simulation, reconstruction and analysis. Every step of the first two parts has a number of parameters, which affect their behaviour. In total, 12 of these parameters were used in the end to alter the setup, creating a large parameter space for scanning and optimising. The full chain and the impact of the parameters, which are denoted like this, are described in the following. An example image of several steps along the way is shown in Figure 49 and serves as reference throughout this chapter. Generally, one event of the simulation consists of one incident particle that traverses the sensitive volume and creates one image that is used for reconstruction.

6.1.1. Simulation setup

A gas volume of variable size and a readout anode on one side was simulated. It uses the following coordinate system:
The simulated particles traverse the TPC volume in x-direction. The secondary electrons drift in z-direction towards the anode, which extends in the x-y-plane. The y-direction is considered the vertical direction for purposes of event displays. The magnetic field is parallel to the z-direction. The simulated extent of the TPC volume in x is always 300 mm, which is the recorded track length. The extent in y (height) is always 256 pads of the anode, and thus depends on the simulated anode granularity. The height was chosen in a way that it contains all secondary electrons even at the smallest pad sizes and large transversal diffusion conditions simulated in this study.
As TPC gas, parameterised T2K gas was used, affecting the ionisation spectrum and the diffusion and attachment during drift and amplification. A triple GEM stack was used, as well as a simplified DAQ which transforms the charge arriving at a pad into an ADC value.

6.1.2. Detailed simulation

The simulation chain consists of several Marlin processors simulating the different steps in the detection process from the primary gas ionisation over the drift and the gas ampli-

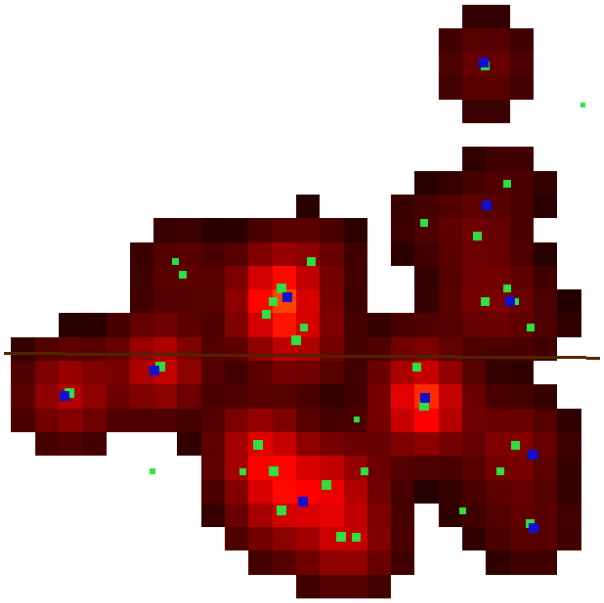

Figure 49: Example stages of the TPC-Timepix simlation and reconstruction:
The brown line shows the track of the original MC particle.
The green dots are the positions of the electrons from ionisation after drift (when they arrive at the GEMs).
The red tiles show the level of activation of the underlying pixel chip from the charge clouds created by the GEMs.
The blue dots are the reconstructed center positions of the "extracted sources" based on the pixel charge information.

fication to the signal readout and digitisation. These processors and their interplay with the simulation parameters are explained in detail in the following, and listed in Table 6. In addition, several processors are included which provide more general functionality. The `ConditionsProcessor` provides conditions data. The `TimePixMapHandlerProcessor` maps the readout channels from channel number to geometric position. The final output is provided by the `LCIOOutputProcessor`, which stores all generated collection in a .slcio-file, and the `AIDAProcessor`, which stores all root histograms generated by other processors in a .root-file. The simulation chain aims to generate the same detector readout raw data as a real detector would so that the same reconstruction and analysis chains can be used on both simulated raw data as well as (calibrated and aligned) detector raw data.

Monte Carlo particle generation:
At the start, Monte Carlo particles are generated using a particle gun contained in a separate python script. The Monte Carlo particles are stored as MCParticles and contain information about the particle, of which flavour, momentum and vertex are needed for the following steps.

Primary ionisation:
In the first step, the `PrimaryIonisationWithClustersProcessor` simulates the incident particle trajectory and the primary ionisation along its way in the sensitive volume. The processor starts from the vertex information of the Monte Carlo particle and evolves it through the TPC volume, taking into account its momentum, particle type, magnetic field and changing curvature due to energy loss. Based on parametrised ionisation tables, it generates electrons from the ionisation process along the track. This parametrisation is based on a detailed HEED simulation by [107]. For every ionisation step, the distance from the last track point and the deposited energy, given as number of secondary electrons, are chosen at random. If the number of secondary electrons exceeds a threshold, they are placed along a separate delta-electron mini-track, else they are placed at the position of the ionisation. The output of the processor are individual secondary electrons and their spatial position, which is stored in a SimTrackerHit collection. In the scope of this work, the processor was adapted to provide additional cluster information, compared to the default `PrimaryIonisationProcessor` which is commonly used in MarlinTPC simulation setups. For every ionisation step a so-called cluster is generated containing the position of the ionisation and the number of secondary electrons generated there, as a second SimTrackerHit collection. To link the individual secondary electrons to their respective cluster, each cluster receives a unique cluster ID number, which is also stored with the individual secondaries.
The processor takes the so-called cluster distance factor (CDF) as input parameter, which artificially increases or decreases the randomised distance between two consecutive clusters by the given factor. This is used to simulate different amounts of ionisation of the MCParticles in a simplified way.

Drift in the gas volume:

The `DriftProcessor` takes each of the generated secondary electrons and simulates their drift and the corresponding diffusion through the TPC drift volume. The output is a new SimTrackerHit collection with one entry per secondary electron and its position and arrival time. The z-coordinate (in drift direction) is set to the end of the drift volume, on top of the first GEM at the start of the amplification region. The x- and y-coordinates are varied randomly according to the expected diffusion, based on gas mixture, electric and magnetic field as well as the individual drift length. The arrival time contains the drift time modified by the diffusion in z-direction, but this time information is not utilised in this simulation and reconstruction chain. The DriftProcessor can also simulate attachment by randomly dropping electrons based on their drift length and the composition of the gas. This feature was not used, however.
The relevant parameters with respect to the DriftProcessor are the magnetic field (B) as well as the drift length (DL), which both affect the diffusion. The latter was however not set as a processor parameter, but is instead a consequence of the chosen geometry and the position of the initial Monte Carlo particle with respect to the sensitive volume.

Amplification:
The `GEMProcessor` simulates the amplification stage in a triple GEM stack. For each drifted electron it calculates a randomised amplification gain. This is based on the three individual GEM voltages, the electric fields between the GEMs (transfer fields) and between the third GEM and the anode (induction field), as well as the distances between the GEM foils. Again, the possibility to simulate electron attachment was available in the processor, however negligible. In addition to the amplification, the processor also changes the x- and y-coordinate of the position to the center of the nearest GEM-hole, based on a hexagonal grid with 140 µm hole-hole distance. The output is another SimTrackerHit collection, with one entry per drifted electron. In this step each entry contains the number of electrons after amplification and the new GEM-hole aligned position. Each entry thus represents the center of a charge cloud.
The relevant input parameters of the GEMProcessor are the three distances between a GEM foil and the respective next foil or the anode (Gd) and the three GEM voltages (GU). While the attachment is considered negligible during drift, the strong electric fields between the GEMs create a considerable electron attachment, which grows with increased GEM spacing. The GEM voltages determine the overall gain. In principle, both parameters can be set individually for each GEM, but were in this study generally altered as one value applied to each GEM or GEM distance.

Charge distribution:
The `ChargeDistributionProcessor` takes the aforementioned charge clouds and distributes them onto the anode. It uses a Gaussian profile for the charge cloud in x and y, and integrates over the area covered by each readout pad according to the readout geometry to sum up the charge arriving at that pad. The processor also splits the charge in time bins, which corresponds to the z-coordinate, but this is undone in the following pixel-specific digitisation step, and thus this feature is not utilised. The output is a VoxelTPC collection, specifying the charge at each pad (specified with its channel ID

according to the geometry) and each time bin.
The processor takes the granularity, respectively the pad pitch (PP), into account, though not as a processor parameter, but as geometrical information. Here, the pads were simulated to be square with the given pad pitch as width and height. For given setup, the extent of the readout plane in x-direction was always 300 mm long, corresponding to the studied track length, thus having a varying number of pads in x, depending on the chosen granularity. To limit computing memory usage at very high granularities, the extent of the readout plane in the y-direction (transverse to the track) was always 256 pads, which means a varying size with varying granularity. It was made sure that even at large diffusion values the signal was always well covered by the readout plane. The GEM distances (Gd) are also used by this processor, since they affect the diffusion in the GEM stack and thus the width of the simulated charge cloud.

Digitisation:
The `TimePixNoisyDigitisationProcessor` takes the charge in each voxel and, after summation over the time bins of each pad, converts it into a digital signal of that channel. It takes into account the partly non-linear response of the Timepix ASIC, in particular the charge threshold which must be exceeded for the channel to start counting the TOT, as well as the maximum TOT count of 11810.
Besides the minimum threshold of 500 electrons, an additional threshold (HWT) can be chosen as input parameter. As a further parameter, and novel compared to the default TimePixDigitisationProcessor, a noise value (ENC) can be chosen, which adds Gaussian noise with an adjustable width to the input charge. The output is a SimTimePixRawData collection and contains the ADC values of each active channel together with the corresponding channel ID.

Simulated process	Processor name	LCIO output collection type	Parameter(s) used
Monte Carlo particle generation	separate particle gun (python script)	MCParticle	-
Primary ionisation	PrimaryIonisation.. ..WithClustersProcessor	SimTrackerHit (2)	CDF
Drift in gas volume	DriftProcessor	SimTrackerHit	B, DL
GEM amplification	GEMProcessor	SimTrackerHit	Gd, GU
Charge distribution	ChargeDistribution.. ..Processor	VoxelTPC	PP, Gd
Digitisation	TimePixNoisy.. ..DigitisationProcessor	SimTimePixRawData	HWT, ENC

Table 6: Steps of the simulation chain, for the parameters see Table 8.

6.1.3. Reconstruction using the Source-Extractor

The reconstruction chain starts from the simulated raw data to reconstruct the tracks in the TPC. The emphasis in this work lies on the dE/dx and the cluster reconstruction to study the particle identification capabilities.
The chain consists of several processors: the `SEWriteSlcioToFitsProcessor`, the `SExtractClusterProcessor` and the `SEClusterToHitConverterProcessor` (the three SE-processors) for the actual reconstruction, together with the `AIDAProcessor` and the `LCIOOutputProcessor` to provide the data output. The key part of the reconstruction however is done outside of the ILCSoft framework by the Source-Extractor software, which is called via the `SExtractClusterProcessor`. The SE-processors, which were originally written by A. Deisting [64], were adapted for this work, to allow among other improvements for non-square images to be processed.

Data conversion:
To utilise the Source-Extractor a file format transfer is needed. This is done in the `SEWriteSlcioToFitsProcessor` which converts the SimTimePixRawData, containing the ADC values of each channel of the simulated readout anode, from the LCIO format to the FITS [108] format. This file format is widely used in astrophysics and astroparticle physics, typically to store telescope images containing light-per-channel information of recorded skymaps.

Source extraction:
The `SExtractClusterProcessor` runs the Source-Extractor, which consists of four major steps: background estimation and subtraction, application of a convolution filter, thresholding and deblending to identify sources, and measurement of source properties. The processor takes the FITS input and creates so-called catalogue files in ASCII format as output, which contain the properties of the identified sources. It performs procedural steps according to Table 7, which are explained in the following and summarised in Figure 50. Source-Extractor is steered with a file containing a large number of steering parameters. The default parameter setting can be found in appendix XY, while some were also varied as parameters of this reconstruction chain. The Source-Extractor parameters are denoted *in this way* in the following.

As first step, the Source Extractor calculates a local background noise level for finite regions of size *BACK_SIZE* of the data map. It histograms all pixel data values and iteratively rejects the extremes of this distribution, until the remaining data lies within $\pm 3\sigma$. The mean of this distribution is then the local background noise level, and is subtracted from the data.
One major difference between sky maps and the simulated TPC readout data is the level of background noise: In the TPC data there is no spatially continuous activation of a majority of the channels by random electrons, as opposed to random photons that hit a telescope. In other words, a TPC does not have an ambient electron background. Instead, in each data image, aside from the signal only few channels activate when the electronics noise exceeds the chosen charge threshold, while most channels contain no

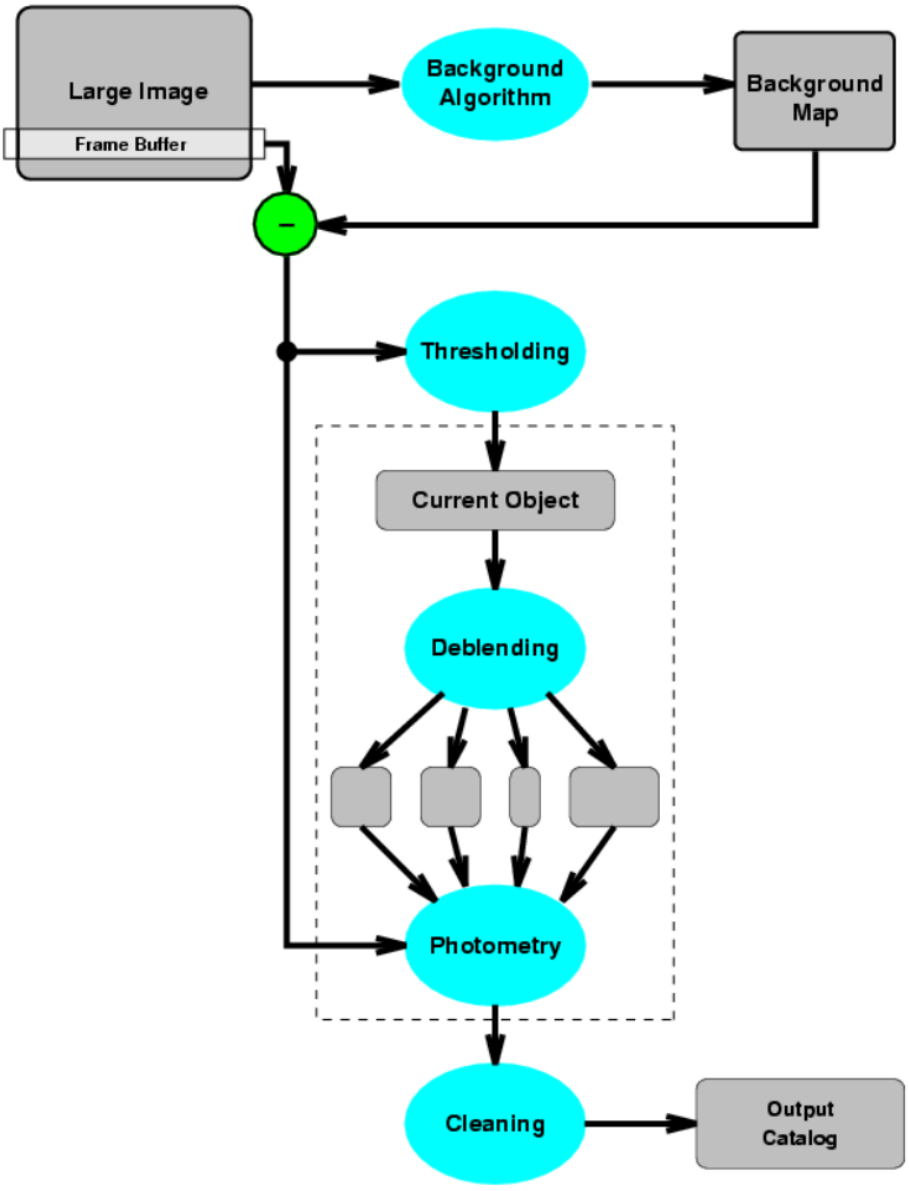

Figure 50: Principal work diagram of the SE, taken from [109], originally from a previous version of [63].

data, see also Figure 17. Thus, in this work the background subtraction was switched off by setting background mode *BACK_TYPE* to manual and selecting a subtraction level *BACK_VALUE* of 0.

Next, the algorithm applies a convolution filter to the data map. In this process, each pixel data value is replaced by a sum of contributions from itself and its neighbouring pixels within a given distance. Each of these contributions is the original value of the respective pixel multiplied by a weight specified in the convolution matrix. A simple Gaussian convolution matrix contains the values of a 2-dimensional Gaussian function, centered around the pixel whose value is to be replaced. This way, the new value contains contributions with decreasing size the larger the distance to a neighbour. This Gaussian smearing causes a blurring effect of the image, which can be useful to reduce noisy pixels, as long as the width of the Gaussian is at most as large as the characteristic sizes of the objects to be detected. In this work, in most cases a Mexican-hat filter was used for the convolution. This filter has a central positive Gaussian peak, and a surrounding broader but flatter negative part. It too is radially symmetric and follows Equation 6 with the radius r and the characteristic width of the Mexican-hat $\sigma = FWHM/2.3548$ and $FWHM$ denoting the full-width-half-maximum of the central peak.

$$M(r) = \frac{1}{\pi\sigma^2} \cdot \left(1 - \frac{r^2}{2\sigma^2}\right) \cdot e^{\frac{r^2}{2\sigma^2}} \tag{6}$$

The radial function is shown in Figure 51, and 2-dimensional examples of the pixelised filter are shown in Figure 52. The effect of the Mexican-hat filter is a smoothing by the positive part, and a contrast enhancement by the negative part on the scale of its width. It is important to choose a width of the Mexican-hat close to the size of the structures under investigating, i.e. the charge clouds arriving on the readout pads, to avoid broad blurring of the image as well as enhancement of random noise.
The size of the filter is used as variable parameter with two values in the reconstruction. The first value is the matrix size A in pixels, it must be large enough to contain the relevant structure of the Mexican-hat. The second value is the full-width-half-maximum B of the inner Gaussian, again in pixels. They are noted in the following as convolution filter (CF) "A-B".

The third major step of the source extraction is the thresholding and deblending to identify sources, or 'hits'. The algorithm goes through areas of adjacent pixels with values that exceed a certain detection threshold *DETECT_THRESH*. This value was set to 0.0001 to ensure all active pixels are used. If the number of these pixels is larger than a specified minimum number *DETECT_MINAREA* (minA) and smaller than a specified maximum number *DETECT_MAXAREA* (maxA) of active pixels for a hit, then at least one source is found, and the so-called multithresholding is applied. Between the value of the peak pixel in that area and the background value (*BACK_VALUE*), a number *DEBLEND_NTHRESH* of equidistant threshold levels is defined, here 64. The algorithm then searches for local maxima in the area, between which the pixel values are below the next lower threshold level. These local maxima are identified as separate hits (deplending), if they fulfil the minimum and maximum area condition and if their individual volume is larger than a certain fraction *DEBLEND_MINCONT* of

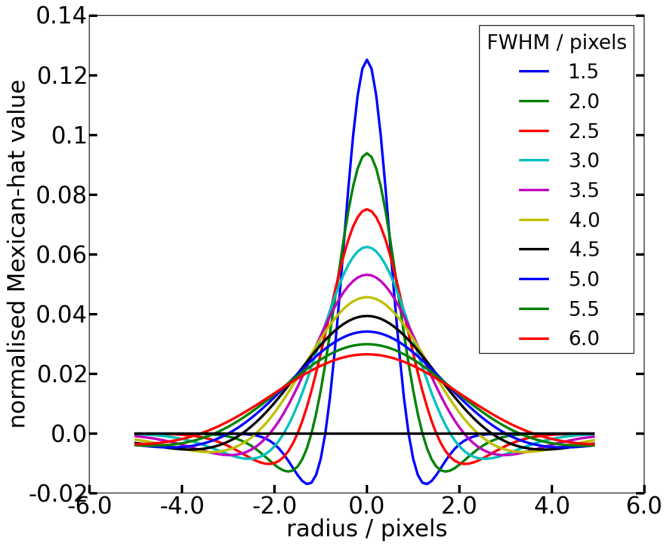

Figure 51: Radial function of the Mexican-hat potential used as convolution filter.

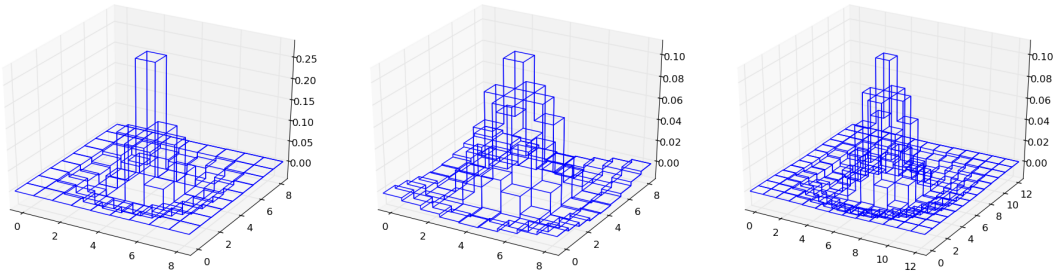

Figure 52: Examples for Mexican-hat convolution filters.
Left: 9-2.5, middle: 9-4, right: 13-4.

the combined peak volume. The deblending fraction was chosen to be 0.00000001, again to ensure all available hits are found. This is applied iteratively from the top threshold level until the background level is reached.

Finally, for each identified source several quantities are extracted and written in an output file in ASCII format, the so-called catalogue file. This file contains one line per identified source which states the calculated source parameter values. The output values chosen in this case are the position, measured as center-of-gravity, (X_IMAGE, Y_IMAGE), amplitude (FLUX_MAX) and radius (FWHM_IMAGE) of the respective source.

1. Measure the background and its RMS noise.
   This step estimates the faint background light in sky images,
   which translates to electronic noise of individual pixels.
2. Subtract background.
3. Filter, i.e. convolve with specified profile.
4. Find objects (thresholding).
5. Deblend detections by breaking detections up into different objects.
6. Measure shapes and positions.
   Here, position and radius of identified sources are measured.
7. Clean, i.e. reconsider detections, accounting for contributions from neighbours.
8. Perform photometry.
   Here, the amplitude of the identified sources are measured.
9. Classify by level of fuzziness into star- or galaxy-like.
   This step is irrelevant for this work.
10. Output catalogue and check images.

Table 7: Steps of the Source-Extractor, from [109].

Data back conversion:
The `SEClusterToHitConverterProcessor` converts the catalogue files back into the LCIO format. Each identified source is treated as one TrackerHit, which serves as basis for further reconstruction and analysis like track finding or dE/dx estimation.

6.1.4. Analysis setup

A dedicated Marlin processor, the `ClusterCountingEfficiencyProcessor`, was created to extract relevant numbers from the identified sources and the data collection at the intermediate stages of the simulation. The main observables are: the number of reconstructed sources, or hits, as listed in the catalogue file, the charge of each channel after digitisation, the number of primary ionisation clusters, and the relationship between the contribution of secondary electrons to a reconstructed hit and their primary

cluster. Further details are explained in the results chapter.

6.1.5. Parameter variation

The roles of the different simulation and reconstruction parameters are the following:
The pad size is the key parameter to vary. Throughout the simulation, square pads are used, and their size is the same as the pad pitch (PP) in x and y. It is the main task of this chapter to find the behaviour of the PID performance with changing granularity. Therefore, in the results chapter the pad size will be the abscissa in most performance plots.
The drift length (DL) is another base parameter. It determines the amount of diffusion, together with the B-field (B) in a fixed relation. For a given detector with a fixed B-field, a dependence of the PID performance on the drift length, i.e. regions of the gas volume, can correlate the applicability of PID with certain physics samples and their angular spectrum. More broadly, this can also be applied in a comparison between different detectors with different B-field and TPC sizes in drift direction. In the results, the PID performance is calculated and displayed for different drift lengths, usually in different colours.
The cluster distance factor (CDF) is used to simulate different particle species. In most cases, Monte Carlo muons were used with a CDF of 1.03 (1.19) to simulate pions (kaons), which corresponds to the working point at 3 GeV. According to Figure 53 from [107], this approach is valid, since the simulated secondary-electron spectra are the same for these species. It also protects against numerical artefacts in the simulation coming from imperfect parametrisation of the different spectra, by instead varying just the overall scale. At 3 GeV, there is a maximum in the relative ionisation between pions and kaons, compare Figure 54.

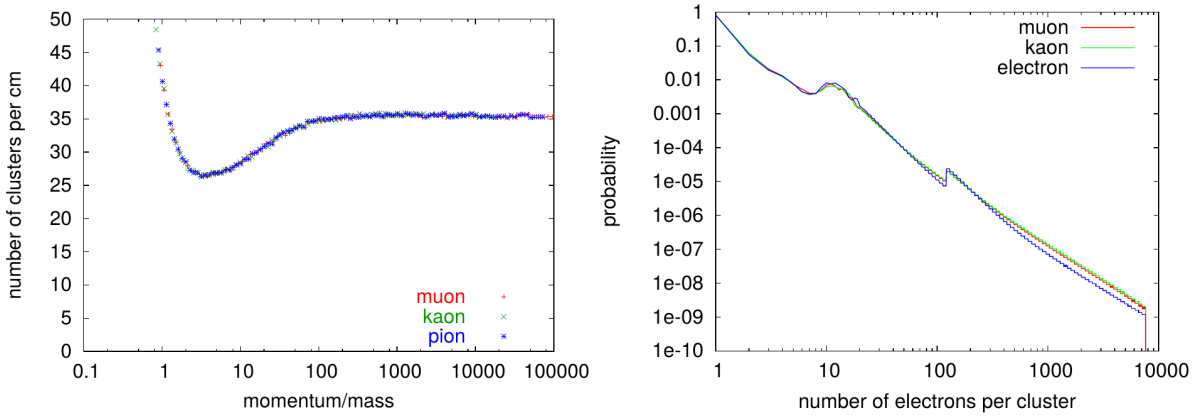

Figure 53: Plots from [107] showing the ionisation of different particle species according to HEED simulation.

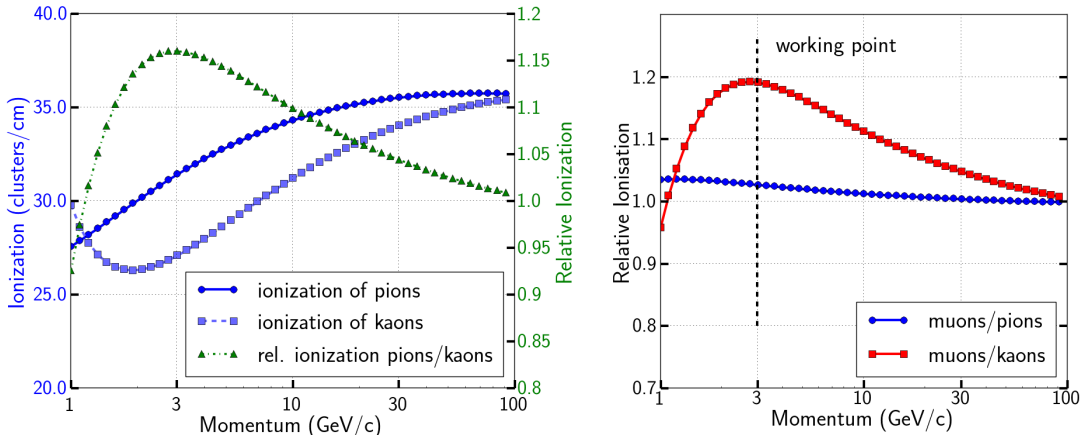

Figure 54: Ionisation densities of pions and kaons (left), and relative ionisation density w.r.t. muons (right).

The gain, determined from the GEM voltage (GU), and the convolution filter (CF) were scanning parameters used for optimisation. For each pad size, several values of these scanning parameters were simulated, and the ones resulting in the best performance were selected as result, see chapter 6.2.2. This is valid, since both GU and CF can be easily altered according to the chosen pad size in a final system. The value of GU is the voltage applied to each GEM. The CF consists of 2 parameters: the extent of the applied filter and the full-width-half-maximum of the Mexican-hat shape, both given in pixels. The minimum and maximum number of adjacent active pixels which are identified as a source in the reconstruction, MinA and MaxA, were varied to a lesser degree until a stable result was achieved, but not optimised for.
Most simulation was done without noise, but for dedicated investigations the equivalent noise charge (ENC) in electrons could be added as Gaussian width to each channel charge. The effect of noise can be combated by raising the channel threshold. This hardware threshold (HWT) was set to about 560 electrons by default in the TOT-to-ADC conversion formula, and could be raised if needed.
To reduce statistical uncertainties, a larger number of simulated events per parameter space point (NEV) could be chosen, which in turn led to a non-negligible increase in disk usage. Typically, parameters were broadly scanned with NEV= 1000 to narrow down a window, which was then studied at NEV= 10000 for more reliable results.
The default values for the altered parameters and the range in which they were varied are given in Table 8.

parameter	short form	default value	range
cluster distance factor	CDF	1	0.25-1.25
magnetic field	B	1 T	(0-4) T
drift length	DL	1000 mm	(10-2000) mm
distances between GEMs, and third GEM and anode	Gd	(1,1,1) mm	[(.1,.1,.1)-(2,2,3)] mm
GEM voltage (for each GEM)	GU	280 V	(230-300) V
pad pitch (in x and y)	PP	220 µm	(55-6000) µm
equivalent noise charge	ENC	0 e^-	(0-1000) e^-
hardware threshold	HWT	561 e^-	(561-1000) e^-
number of events (tracks)	NEV	1000	1-10000
convolution filter	CF	mexican hat with FWHM of 2.5 pads	mexican hat with FWHM of 2-7 pads
hit finding minimum number of active pads	minA	4	4-9
hit finding maximum number of active pads	maxA	1000	1000-10000

Table 8: Default values and used range of the parameters altered in the high granularity simulation and reconstruction.

6.2. Results

6.2.1. General Observables

In the first step, the cluster identification efficiency was studied. Given the high density of electrons after ionisation in the drift volume and the typical diffusion, there is a significant overlap of electrons from different clusters after drift. This effect is even increased by the defocusing in the amplification stage. About three quarters of the primary ionisation clusters contain only one electron, but the Landau tail of the ionisation process gives a long tail of larger electron numbers with decreasing probability, compare Figure 55.

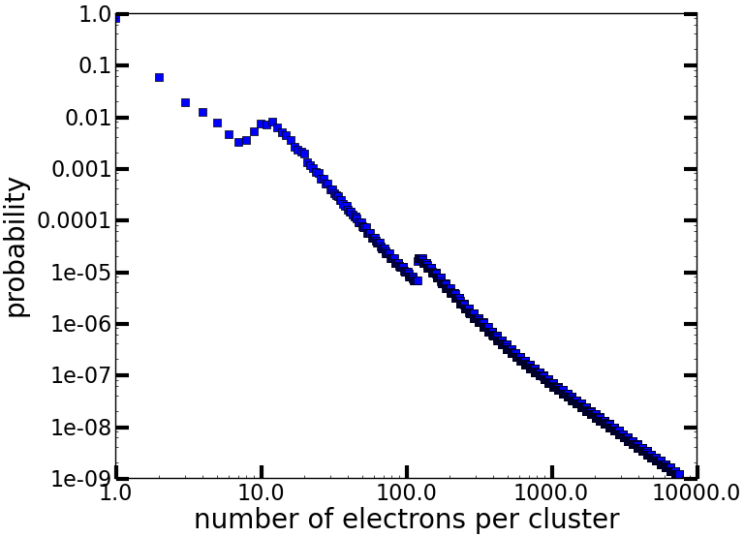

Figure 55: Cluster size probability in primary ionisation for T2K gas, analogue to [107].

On the one hand, electrons from a multi-electron cluster often drift apart and contribute to different reconstructed hits. On the other hand, reconstructed hits often contain additional electrons from these multi-electron clusters. This way, the larger clusters tend to "contaminate" clusters which otherwise contain only electrons from one cluster. Due to this, a straight-forward definition of a counting efficiency is not possible. Instead, a double-unique efficiency is used as a lower limit for the efficiency usually quoted in literature. It is defined as the number of double-unique primary ionisation clusters over the total number of clusters, where a cluster is double-unique, when its electrons contribute to only one reconstructed hit, and that hit contains contributions from only that one cluster. A specific electron from a cluster is considered to contribute to the reconstructed hit it is closest to (in the x-y-plane) when it reaches the anode. The extracted cluster counting efficiency for typical simulation conditions is shown in Figure 56. It shows that the efficiency, as defined above, reaches reasonable values only for very small pad sizes below 200 µm. It should be noted that the efficiency grows for larger drift length, since the electrons drift further apart, which reduces the aforementioned

contamination effect, especially for single-electron clusters.

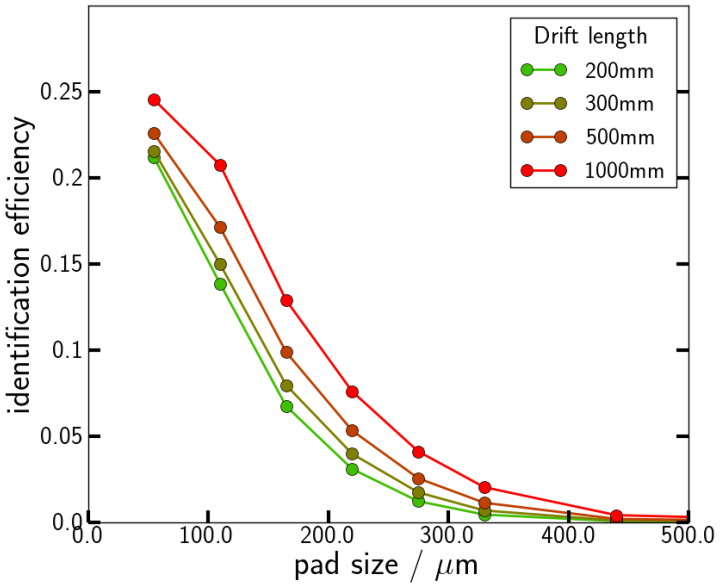

Figure 56: Cluster counting efficiency.

Secondly, the cluster counting resolution was determined, based on the number of reconstructed clusters. An example result for the default simulation setup is shown in Figure 57, with an RMS over mean of about 3%. This is rather low compared to a typical dE/dx resolution based on charge summation of about 5 % for a comparable system. The reason is a saturation effect: With the large number of clusters generated in T2K gas, there is a significant overlap between the individual GEM charge clouds. A rise in ionisation, for example by a different particle species, does not cause a proportional rise in the number of reconstructed hits. This saturation effect can be seen in Figure 58, where a low ionisation was simulated by stretching the distances between the individual ionisation interactions of the incident particle, leading to a reduced density of primary ionisation clusters along the track. While the proportionality of the number of reconstructed hits to the amount of ionisation is valid for cluster densities below about 50 % of a typical MIP, this is not the case anymore around 100 % relative cluster density. The resolution, defined as RMS over mean, is thus not a good measure for the performance of cluster counting.

This is why in the third step, the separation power S was determined. It is defined as the difference between the amount of ionisation (measured as number of hits, or as summed charge) of two different particle species 1 and 2 for a given momentum p, over the width σ of the ionisation distribution:

$$S(p) = \frac{|\mu_1(p) - \mu_2(p)|}{\langle \sigma(p) \rangle}. \tag{7}$$

Since this is a differential measure, it is unaffected by the aforementioned saturation

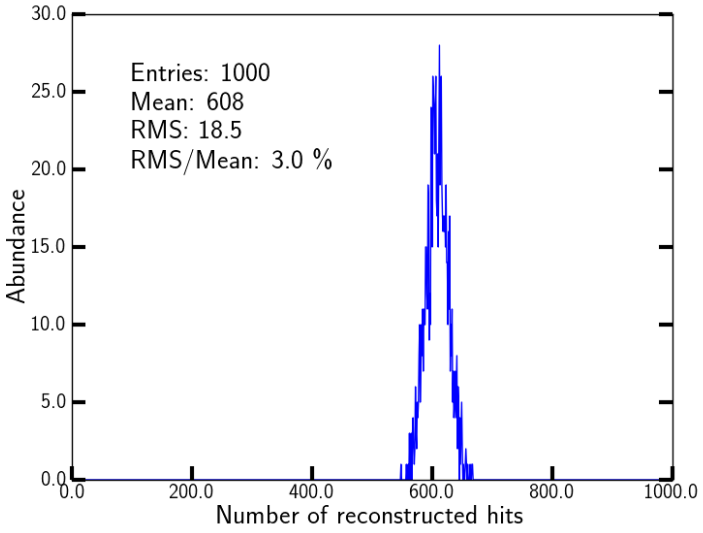

Figure 57: Resolution based on reconstructed hits from cluster counting.

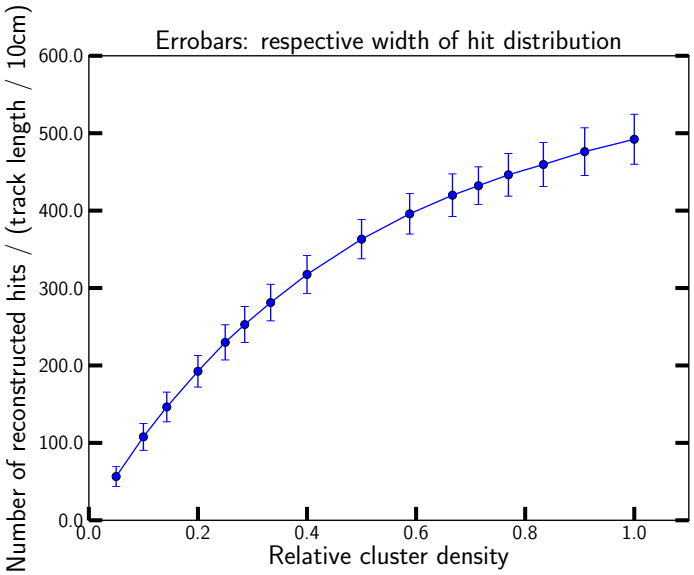

Figure 58: Saturation effect in the number of reconstructed hits from cluster counting.

effect. The width σ is usually the width of one of the ionisation distributions of species 1 or 2, depending which is considered as signal and background. In hadron identification, pions are often considered background, and thus in literature it is often chosen that $\langle\sigma\rangle = \sigma_\pi$. In the more general case here, where a separation power should be computed independently of the role of the respective species, a combined $\langle\sigma\rangle$ is used:

$$\langle\sigma(p)\rangle = \sqrt{\frac{1}{2}(\sigma_1^2(p) + \sigma_2^2(p))}. \tag{8}$$

This root-mean-square is used for this combination instead of the arithmetic mean or the squared sum to keep it comparable to the literature default. Compared to the arithmetic mean, it also takes into account the case where one of the ionisation distributions is much wider than the other, which however is generally not the case for the distributions investigated in this work.
The resulting separation power based on the number of reconstructed hits is shown in Figure 59. It drops quickly for pad sizes larger than the typical distance of electrons in a track in T2K gas. Only very small pad sizes below 300 µm provide sufficient cluster reconstruction to allow for a substantial separation power. In the range of the falling edge, the cluster counting benefits from longer drift length since the saturation effect is reduced by transversal diffusion.

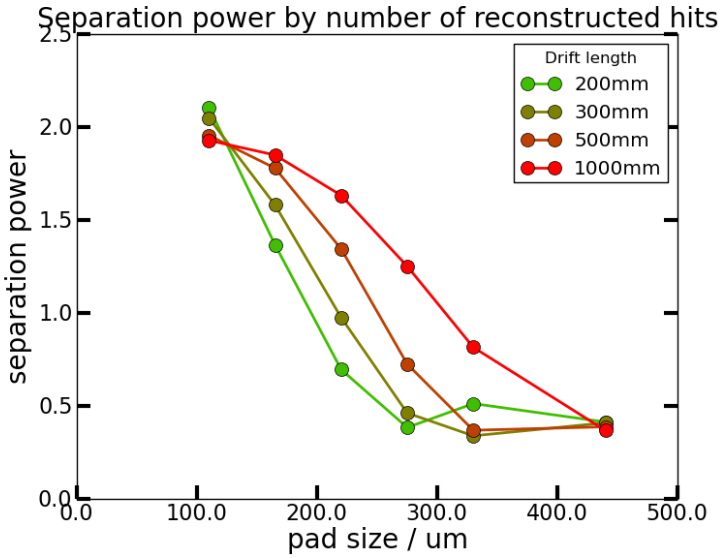

Figure 59: Cluster-counting separation power at small pad sizes. Here, the GEM voltages were, among others, a fixed parameter and not optimised for.

Figure 60 shows a comparison of this cluster-based separation power to the traditional one using charge summation. Here, no optimisation for GU and CF was applied. The resulting behaviour in the plot is largely dominated by limitations with respect to the GEM voltage: For a decreasing pad size below 440 µm, the simulated charge cloud is distributed over an increasing number of pads, reducing the charge per pad. This leads

to an increase in the fraction of pads receiveing charge below the fixed threshold level, leading to loss of information and by this to a decrease in separation power. This threshold effect can in principle be compensated by increasing the GEM voltages and by this the amplification, but the default value of 280 V is used as a maximum value, requiring reasonable stability of actual GEMs. For smaller drift distances this effect is less pronounced, since the secondary electrons experience less transversal diffusion and their charge clouds have a larger overlap with more charge per pad on average. For an increasing pad size above 440 µm, an increasing number of pixels receives more charge than its counting limit, causing an overflow and again a loss of information. This limits the achievable separation power severely, but can be compensated easily by reducing the GEM voltages.

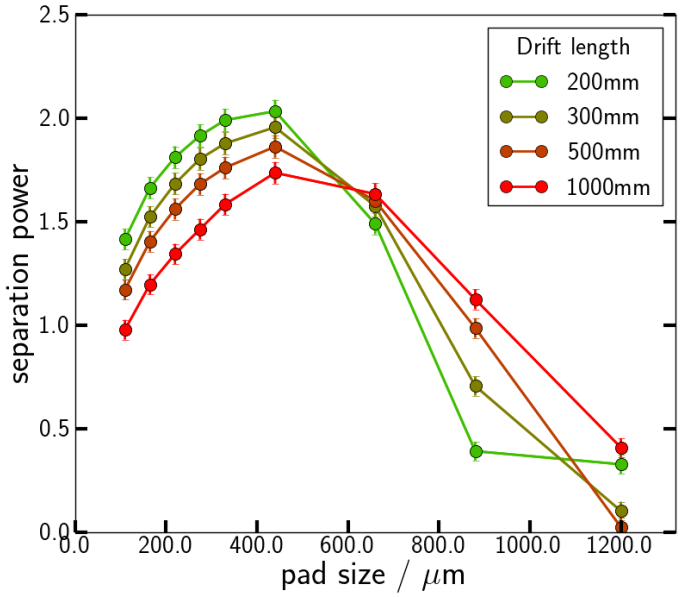

Figure 60: Charge-summation separation power over a wide range of pad sizes. Here, some parameters including the GEM voltage were fixed.

6.2.2. Paramter Optimisation

In both Figure 59 and Figure 60, GU and CF were fixed at their default values, limiting the achievable separation power. In particular a variation of the GEM voltages and thus of the gain allows for further improvement. In the simulated setup, this affects the aforementioned threshold and overflow effects of the individual data taking pixels. Figure 61 and Figure 62 show scans of the separation power based on charge summation depending on the pad size for different GEM voltages, and for a drift length of 200 mm and 1000 mm, respectively. To cover the entire range of simulated pad sizes, a log-scale was used. The threshold and overflow effects mentioned in the previous paragraph are visible for each voltage setting. Since the GEM voltage can be adjusted to the pad size,

for each pad size the maximum value among the respective curves is chosen to make a combined curve, as shown in Figure 63. The curves follow the empirical expectation well, indicated in purple. For very small pad sizes the threshold effect cannot be avoided, since a maximum GEM voltage of 280 V was assumed, which is considered a reasonable limit of GEM stability. For pad sizes between about 600 mm and 6000 mm, the optimal GEM voltage varies between 280 V and 230 V, corresponding to average gain values between 62,000 and 510, as displayed in Figure 64. Here, a reduced inter-GEM distance of 1 mm is assumed to enhance the gain, compared to 2 mm (GEM-GEM) and 3 mm (GEM-anode) gaps in the standard setup with large pads.

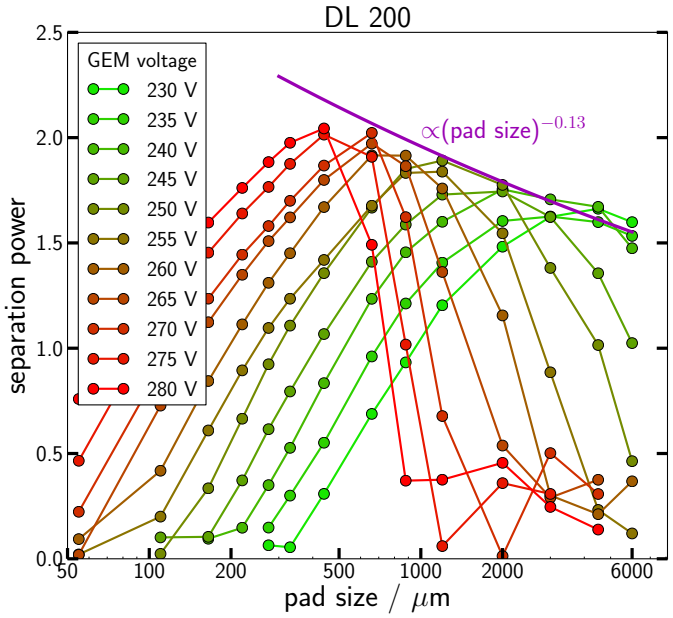

Figure 61: Charge summation separation power depending on pad size and GEM voltage for 200 mm drift length.

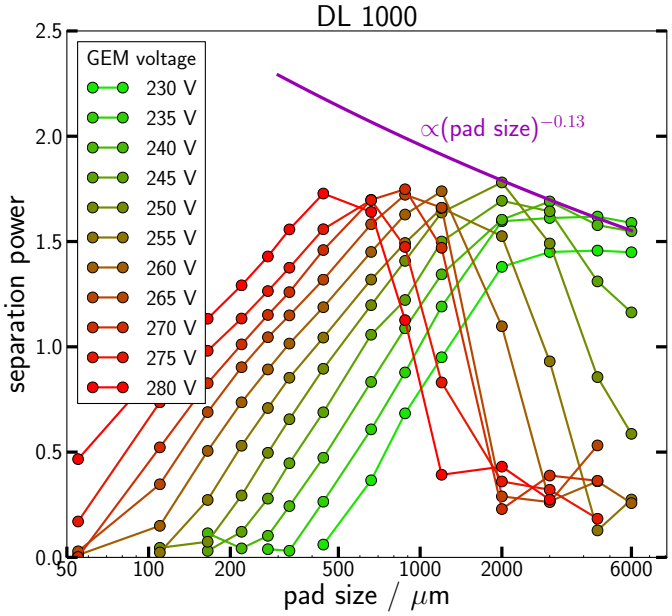

Figure 62: Charge summation separation power depending on pad size and GEM voltage for 1000 mm drift length.

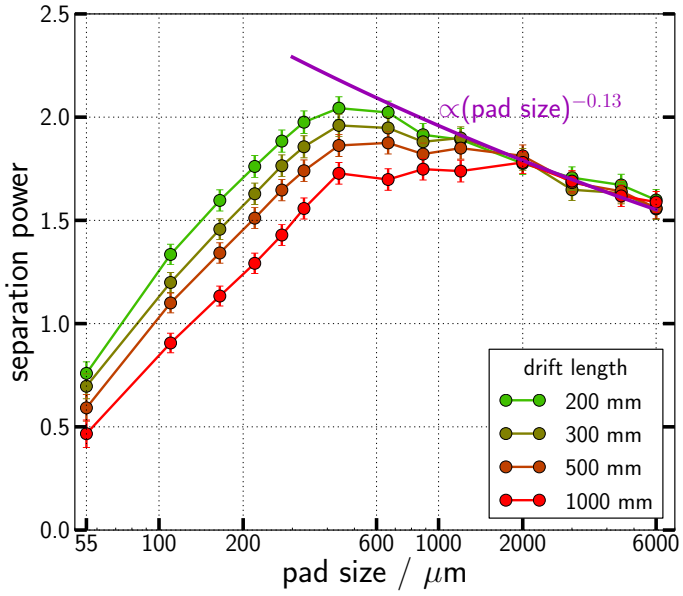

Figure 63: Charge summation separation power depending on pad size for the respective best GEM voltage setting.

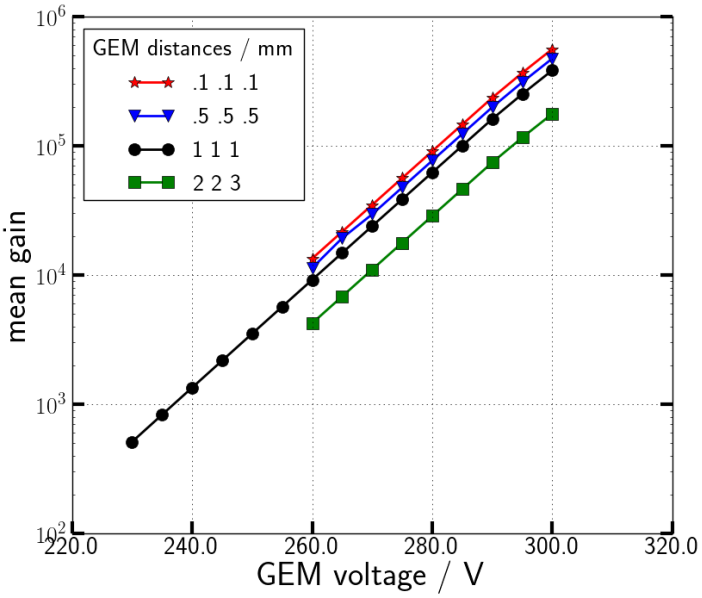

Figure 64: Average gain of the triple GEM stack, depending on the GEM voltage (same voltage for each GEM) and the distance between the GEMs and between the third GEM and the anode.

Another adjustable parameter is the convolution filter used by the Source-Extractor for finding the later hits. Since the filter size is given in matrix entries, equivalent to pads, the optimal filter should change with granularity. An example is shown in Figure 65. The filters are ordered by two numbers, indicating the matrix extent (9 or 11 pixels) and the width of the Mexican-hat shape (between 2 and 7 pixels). The same data was reconstructed with different filter sizes and the resulting separation power plotted. There is a visible maximum for a filter width of 2.5 to 3 pixels. If the characteristic width is too small, then the contrast enhancement tends to create artificial structures, which are then reconstructed as sources and counted as hits. If the width is too large, actual signal structures tend to be blurred out, reducing the number of correctly reconstructed sources. Both effects reduce the overall seapration power. Events with a short drift distance are more affected by a too large filter width than events with a larger drift distance, because they tend to keep narrow structures of clustered electrons correlated. If the filter width becomes larger than half the filter matrix extent, the negative part of the Mexican-hat is cut off and only a wide blurring remains, reducing the performance severely. To avoid that cutt-off effect, a minimum matrix extent is necessary, but aside from this, a larger matrix with the same filter width will only moderately affect the separation power. At a very small filter width, the matrix elements become increasingly discrete, or less smooth, which numerically limits the effective resemblance of the desired shape. This in turn limits the overall performance for very small filter width values even for larger pad sizes.

Figure 66 shows a similar plot, but for a pad size of 55 µm. At this granularity, the

available parameter space, in particular the optimal gain, becomes narrower and the overall separation power behaviour is more volatile. The optimal convolution filter size, however, is well visible at a filter width around 4.5 pixels. This is nearly twice the size as for the previous pad size of 110 µm, consistent with adapting the filter width to the granularity.
Examples for Mexican-hat shaped filters are shown in Figure 52.

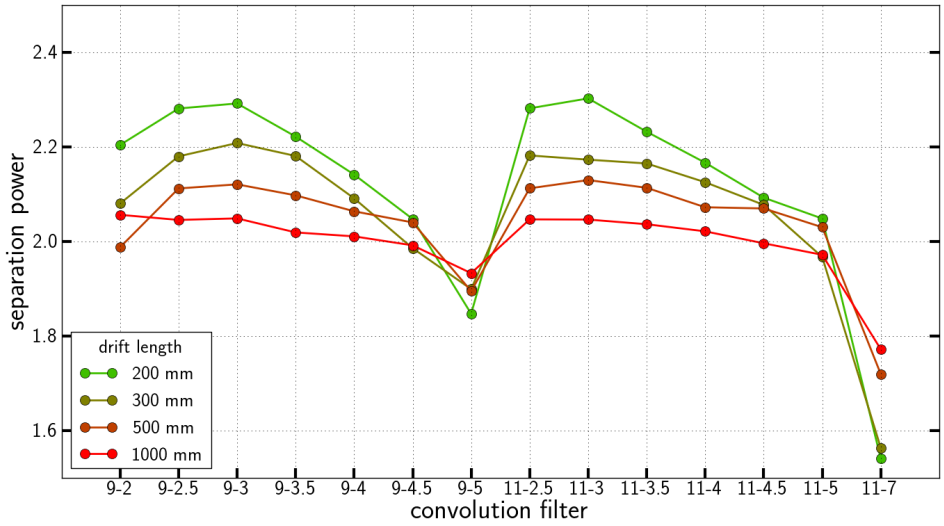

Figure 65: Scan of the convolution filter for a pad size of 110 µm.

In the following, the respective best performance with respect to GEM voltage and convolution filter is chosen for each pad size.

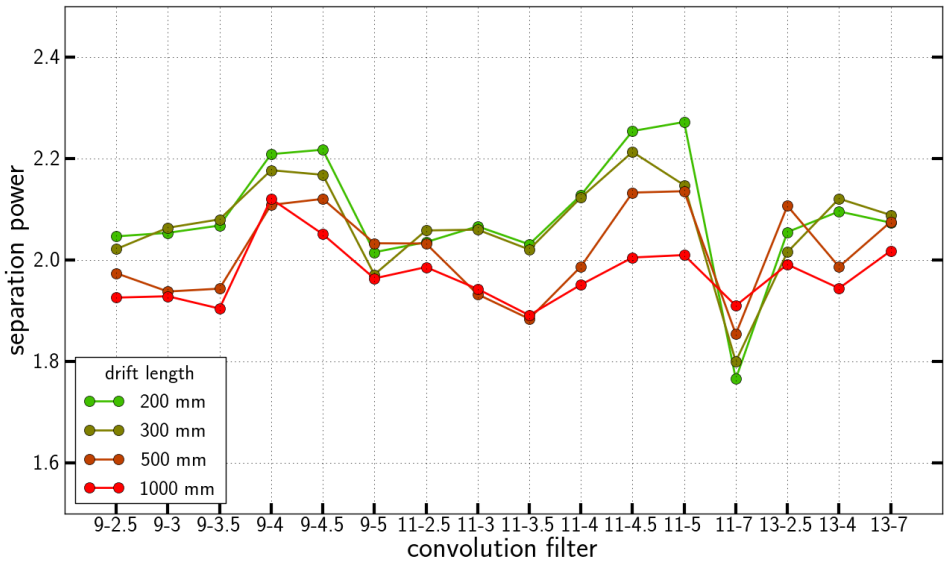

Figure 66: Scan of the convolution filter for a pad size of 55 µm.

6.2.3. Comparison and Combination

After optimisation, the separation power performance for cluster counting and charge summation is compared in Figure 67. The red and green points show the simulated data, while the corresponding lines are polynomial fits to these data to guide the eye.

In addition, test-beam based results from the existing readout systems AsianGEM, GridGEM and Micromegas as well as GridPix were extrapolated to the conditions of the simulation (in particular the track length) and added to the plot as blue squares for reference. The extrapolation of test beam results of the existing systems is done as follows.

The test beams have been performed with electron beams, resulting in a measurement of the dE/dx resolution of electrons $R_e = \sigma_e/\mu_e$, based on the mean μ_e and width σ_e of the measured distributions. This needs to be translated to the working point of this simulation study, i.e. pion-kaon separation at 3 GeV. Following the Bethe-Bloch equation and Figure 54, the mean ionisation values of pions and kaons relative to the relativistic plateau of electrons are $\mu_\pi = 0.880\mu_e$ and $\mu_K = 0.759\mu_e$. Let η be: $\mu_\pi/\mu_K = 1.16 =: \eta$. Emulating different particle species via a change in the cluster distance factor is equivalent to changing the gas pressure in the TPC's sensitive volume. Therefore, the relative dE/dx resolution R scales like $R \propto L^{-0.32}$, according to Equation 4, e.g. $R_K/R_e = 0.759^{-0.32} =: \rho_K$. Here, L is the effective track length, which is equal to the relative ionisation factors 0.880 and 0.759 above, and $\mu \propto L^1$. With $R = \sigma/\mu$ it follows that $\sigma \propto L^{0.68}$, and $\sigma_\pi/\sigma_K = \eta^{0.68}$. Then one can calculate:

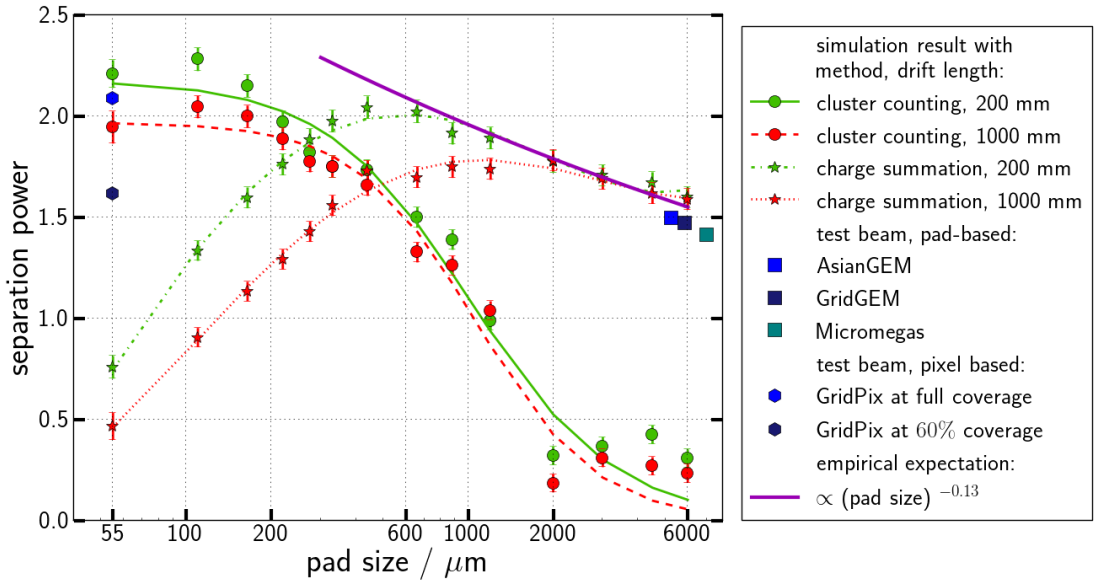

Figure 67: Combined plot for separation power by cluster counting and charge summation, depending on pad size, in comparison to test beam results of existing systems and to empirical behaviour.

$$S = \frac{|\mu_\pi - \mu_K|}{\sqrt{\frac{\sigma_\pi^2+\sigma_K^2}{2}}} \tag{9}$$

$$= \frac{\mu_K|\frac{\mu_\pi}{\mu_K} - 1|}{\sqrt{\frac{\sigma_K^2\left(\frac{\sigma_\pi}{\sigma_K}\right)^2+\sigma_K^2}{2}}} \tag{10}$$

$$= \frac{\mu_K|\eta - 1|}{\sigma_K\sqrt{\frac{(\eta^{0.68})^2+1}{2}}} \tag{11}$$

$$= \frac{|\eta - 1|}{R_K\sqrt{\frac{\eta^{1.36}+1}{2}}} \tag{12}$$

$$S = \frac{|\eta - 1|}{R_e\rho_K\sqrt{\frac{\eta^{1.36}+1}{2}}} \tag{13}$$

This can now be used to translate between the published relative dE/dx resolution values for electrons and the separation power at the working point.

The GridGEM system [40] reported a dependence of the dE/dx resolution $R_e(N)$ on the number of hits in a track N as $R_e(N) = R_e(0) \cdot N^{-k}$, shown in Figure 68, with $R_e(0) = 0.587$ and $k = 0.465$, after adjustment to the direct measurement. For the

simulated track length of 300 mm the row pitch of 5.85 mm gives $N = 51$, resulting in $R_e(51) = 9.43\%$ and via Equation 13 in a separation power of $S = 1.47$. For an extrapolation to the ILD TPC the reported resolution is $R_{e,ILD} = 4.71\%$.

The AsianGEM system [43] reported a dE/dx resolution of $R_e(26) = 13.52\%$ with a row height of 5.26 mm. Assuming a scaling with $R_e(N) \propto N^{-k}$ similar to the GridGEM with $k = 0.48$, 300 mm tracks result in a resolution of $R_e(57) = 9.28\%$ and a separation power equivalent of $S = 1.50$. The reported resolution for the ILD TPC is $R_{e,ILD} = 4.61\%$.

The Micromegas system [44] reported a resolution of $R_e(192) = 4.8\%$ with a row height of 7 mm, which is very close to the extent of the ILD TPC. An extrapolation to 300 mm tracks with $k = 0.48$ gives $R_e(43) = 9.8\%$ and $S = 1.41$.

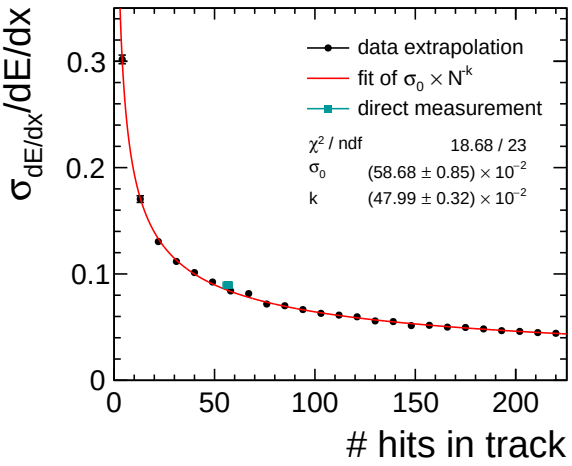

Figure 68: Extrapolation of the measured dE/dx resolution of the GridGEM system to the ILD TPC dimensions, from [40]. The fit results given in the plot change, when an exact match with the measured data point is required, resulting in a more conservative estimate as indicated in the text.

The GridPix system [46], other than the pad-based ones, reported numbers on the separation power, since also here a cluster counting algorithm was used. These numbers have been revised [110], since, and are $S_{CS} = 8.6$ for charge summation, $S_{CC} = 8.7$ for a cluster counting approach and $S_{comb} = 9.2$ for a combination of the two. This was calculated for 1 m tracks and by comparing the measured electron data to muon data emulated from the electron data. For this, 1 m of electron data was taken and the positions of the reconstructed hits on the Timepix anode were stretched by a factor of $1/0.7$ to fill 1 m, since $\mu_\mu/\mu_e = 0.7 =: \tilde{\eta}$, following the emulation conditions. Then the same reconstruction algorithms were used on these new hit positions and resulted in mean and width values for a truncated charge, in case of charge summation, and a weighted mean distance, in the case of the cluster counting approach. With these mean and width values, the separation power was calculated using the same formula as

Equation 7 and Equation 8.

In an attempt to apply the translation above to these published values in order to translate to a resolution and then back to a pion-kaon separation power it was found that the dependencies are somewhat different from the ones laid out above. It was decided within this work, to use more direct approach using the base mean and width values from the GridPix results (from [111]), in order to provide a fair comparison. The calculation, based on the measured and emulated values of the truncated charge, is laid out in the following. For simplicity, units are omitted.

$$\mu_{e,meas} = 6347, \qquad \sigma_{e,meas} = 262.9 \tag{14}$$

$$\mu_{\mu,meas} = 4284, \qquad \sigma_{\mu,meas} = 212.2 \tag{15}$$

$$\mu_{\mu,meas}/\mu_{e,meas} = 0.675 = \tilde{\eta}^{1.1}, \qquad 1.1 =: a \tag{16}$$

$$\mu_{\pi,sim}/\mu_{e,sim} = 0.759 \quad \Rightarrow \mu_{\pi,meas} = \mu_{e,meas} \cdot 0.759^{a} \quad = 4684 \tag{17}$$

$$\mu_{K,sim}/\mu_{e,sim} = 0.880 \quad \Rightarrow \mu_{K,meas} = \mu_{e,meas} \cdot 0.880^{a} \quad = 5513 \tag{18}$$

This means, if the GridPix method was applied to a target ionisation of pions or kaons at 3 GeV, then following the same dependence these are the expected mean values. Similarly for the widths:

$$\sigma_{\mu,meas}/\sigma_{e,meas} = 0.804 = \tilde{\eta}^{0.611}, \quad 0.611 =: b \tag{19}$$

$$\sigma_{\pi,meas} = \sigma_{e,meas} \cdot 0.759^{b} = 222.2 \tag{20}$$

$$\sigma_{K,meas} = \sigma_{e,meas} \cdot 0.880^{a} = 243.2 \tag{21}$$

These mean and width values can again be used with Equation 7 and Equation 8, to result in $S = 3.56$, for 1 m tracks and charge summation only. Since the same base values are not known in this work for the combined approach, a linear extrapolation is assumed, by multiplication with 9.2/8.6. In addition, the track length of 0.3 m is extrapolated with a scaling exponent of $k = -0.5$. This results in a final separation power for the GridPix system, translated to the working point in this work of $S = 2.09$. Since the current GridPix system has an anode coverage of about 60 %, in addition a reference point is included at $S = 2.09 \cdot \sqrt{0.6} = 1.62$.

The dE/dx resolution reported by GridPix is 4.1 % for charge summation and 1 m tracks. This can be extrapolated using the same scaling to the ILD TPC and a combined approach via $4.1\,\%/\sqrt{1.35}/(9.2/8.6) = 3.3\,\%$ for a full coverage, and 4.3 % for a 60 % coverage.

The references values of all systems used in Figure 67 are summarised in Table 9.

In Figure 67, at large pad sizes around 6 mm, comparable to the existing pad-based systems, the performance of the simulation with charge summation is similar to the ones observed in experiments at $S \approx 1.5$, which serves as general validation of the simulation. For decreasing pad sizes, a rise in performance is visible following a $PP^{-0.13}$-behaviour, as indicated in the plot. This agrees with the empirical observation taking into account various former experiments as summarised in [51]. The separation power via charge

Readout system	granularity	$S_{\pi/K}$	$R_{e,ILD}$
GridGEM	5.85 mm	1.47	4.71 %
AsianGEM	5.26 mm	1.50	4.61 %
MicroMegas	7 mm	1.41	4.8 %
GridPix, full	55 µm	2.09	4.3 %
GridPix, 60 %	55 µm	1.62	3.3 %

Table 9: Summary of the reference values of the different existing TPC readout system based on test beam results. The values of the separation power S at the pion-kaon working point simulated in this chapter are used in Figure 67, the relative dE/dx resolution values R for the ILD TPC are used in Figure 69.

summation has a maximum around a few hundred µm, depending on the drift length, and decreases for smaller pads due to the threshold effect explained above. For a drift distance of 1000 mm, the maximum of about 1.75 lies at a pad size of 1000 µm and 10 % above the level at 6 mm pad size. For a drift distance of 200 mm, the maximum lies between 400 and 600 mm and is about 2, exceeding the 6 mm point by about 20 %. For even smaller pads, the performance via cluster counting surpasses the one via charge summation at about 500 µm (300 µm) for a drift distance of 1000 mm (200 mm). It peaks for a pad size of 110 µm at a value of 2 for 1000 mm drift and 2.3 for 200 mm drift, exceeding the reference value at 6 mm by 20 % and 44 %, respectively. For 55 µm pads, the cluster counting performance shows a drop, and reaches a similar value as the one achieved in test beam by the GridPix system, if a full anode coverage is assumed. This drop is addressed in chapter 6.2.4.

For charge summation, the drift distance is negligible for larger pads >2 mm, but there is a significant difference for most of the simulated pad sizes, with events at short drift outperforming ones at large drift. This, however, is due to limitations in the maximum amplification and simulated electronics readout and could be at least partially compensated by using a setup optimised for the specific pad size.
In contrast, the performance for cluster counting is similar with regard to short and long drift for most simulated pad sizes, but differs for the smallest ones. Here, the necessary information of cluster correlation is actually diminished during drift, which can not be restored by choosing optimised system parameters. This is further investigated in view of the 55 µm-drop in chapter 6.2.4.

Via Equation 13 the separation power can also be converted into a dE/dx equivalent for the ILD TPC and compared to the reported or extrapolated values of the test beam systems, which is done in Figure 69. The plot does not contain additional information, but may be more convenient for the discussion of different readout options for the ILD TPC. Pad-based systems with a granularity around 6 mm report a dE/dx resolution extrapolated to the ILD TPC of about 4.7 %. With a granularity of 1 mm this could be improved to about 4.2 % via charge summation, and via cluster counting to about 3.5 % at a granularity below 200 µm. The GridPix system can achieve a dE/dx resolution

equivalent of 4.3 % with 60 % anode coverage and 3.3 % with a full coverage, at least for the so far investigated drift lengths of a few cm.

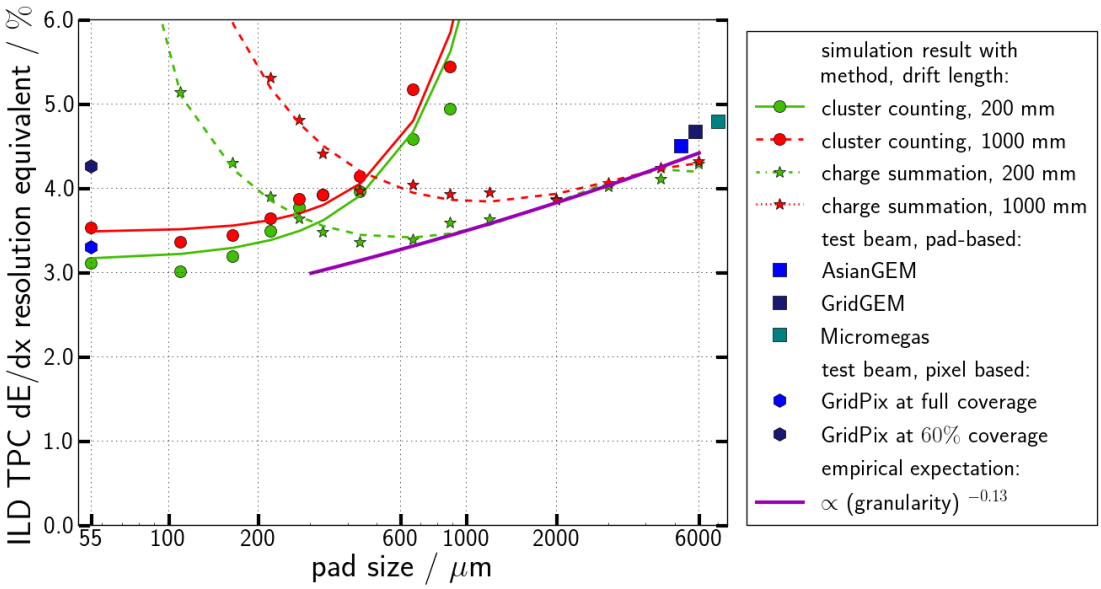

Figure 69: Combined plot of dE/dx resolution equivalent for the ILD TPC, based on separation power by cluster counting and charge summation, depending on pad size, in comparison to test beam results of existing systems and to empirical behaviour.

In a further step, one can try to combine charge summation and cluster counting for each pad size, since both algorithms can be applied simultaneously. The two estimates for the energy loss, the recorded charge and the number of hits, were linearly combined with an optimised relative factor for each granularity. The resulting combined observable was then used to measure the separation of the simulated pions and kaons. Figure 70 and Figure 71 show that for a fixed GEM voltage, there is a clear transition region visible, where both algorithms contribute to the performance. However, with a variable GEM voltage, the two algorithms have different optimal gain values at most pad sizes. Thus, the combined observable is calculated for each GEM voltage, and again the best performing voltage is chosen for that pad size. Generally, the best combination does not perform significantly better than chosing the best of either charge summation or cluster counting. This means, that the resulting combined separation power only adds a negligible improvement when compared to a system optimised for the best fitting algorithm. This is true for short and long drift distances.

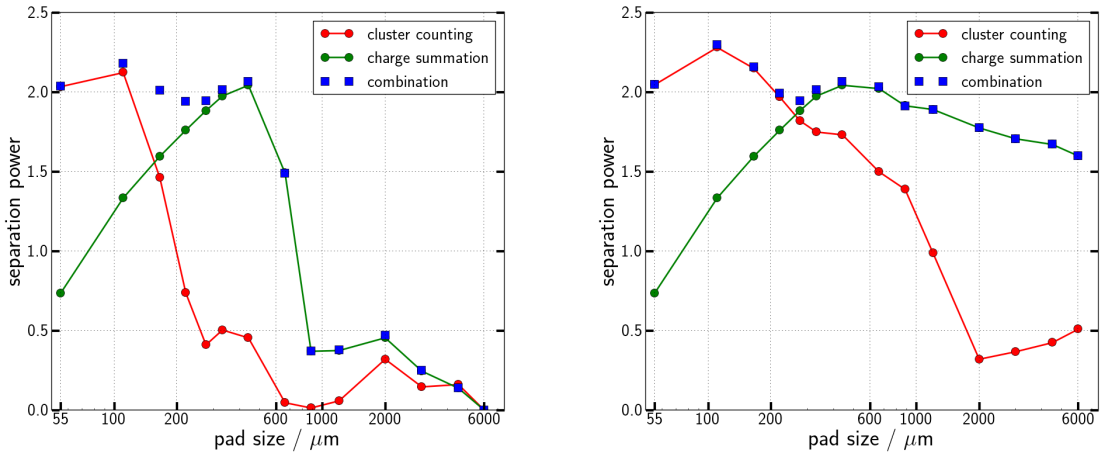

Figure 70: Combined separation power by charge summation and cluster counting, for a drift length of 200 mm. Left for a fixed GEM voltage of 280 V, right after GEM voltage optimisation, where every point reflects a different, optimal GU, but the combination was done for the same GU for cluster counting and charge summation.

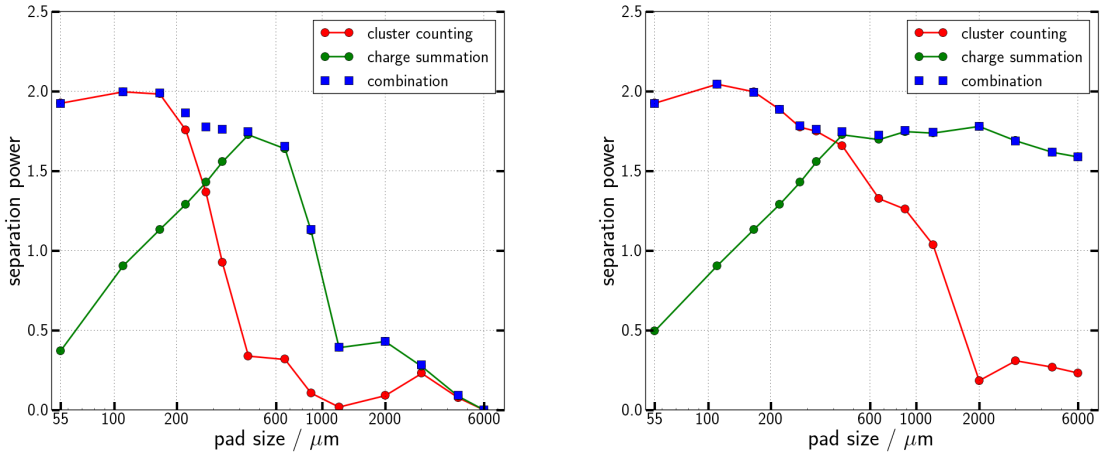

Figure 71: Combined separation power by charge summation and cluster counting, for a drift length of 1000 mm. Left for a fixed GEM voltage of 280 V, right after GEM voltage optimisation, where every point reflects a different, optimal GU, but the combination was done for the same GU for cluster counting and charge summation.

6.2.4. The 55 µm Situation

In Figure 67 the question remains why the separation power via cluster counting at 55 µm is lower than at 110 µm pad size instead of further increasing. One clear limitation is still the gain at the maximum GEM voltage: For charge summation, already at a pad size of around 1 mm the gain is insufficient to maintain the necessary charge information against the pixel threshold. This effect becomes relevant for cluster counting too, but only at a pad size below 110 µm. Allowing for a GEM voltage of 300 V, which is equal to a gain of about 384 k, raises the 55 µm point to about the level of the one at 110 µm, which also still slightly increases indicating a moderate threshold effect also for this pad size. Further studies, summarised in Figure 72, revealed another limiting aspect: the charge cloud size. If the characteristic features of the electron clusters get blurred out too much by the charge clouds in the amplification process, caused by the diffusion in the GEM stack, the increase in granularity becomes irrelevant. In simulation it is possible to reduce the charge cloud size to overcome this effect by reducing the distance between the GEMs. The simulated diffusion values are given in Table 10. By significantly reducing the transversal diffusion this way to about a third of the default value, the 55 µm point for cluster counting shifts up further and gets close to a separation power of 2.8, while the 110 µm point stays lower. This now describes one way to emulate the conditions of extreme anode and amplification granularity in the GridPix system, resulting in a performance exceeding the GridPix test beam results, instead being close to an extrapolation of the empirical expectation curve. This result indicates that the GridPix result may possibly be further improved with an optimised cluster counting algorithm, as opposed to the weighted mean distance algorithm.

It should however be noted that the simulated conditions are hardly physical, let alone achievable in a real setup with GEMs: a voltage of 300 V is not considered stable any more and would likely lead to frequent discharges in the GEMs. An inter-GEM distance of 0.1 mm would be mechanically more than challenging, in particular considering that even with the recently flatness-optimised GEM-glueing procedure, an overall flatness with an RMS of 41 µm could be achieved [40]. In addition, deformations of the GEMs can happen in the GEM stack when high voltage is applied, leading to a further change in the distances. Finally, the separation power value close to the GridPix could only be achieved with a very short drift distance. Otherwise, the diffusion during drift would again eradicate the cluster information, which the assumed extreme conditions made available. The identification of primary electron clusters can only reach the highest levels, if all three aspects that limit the spatial reconstruction are optimised: transversal diffusion in the drift volume via a short drift distance, diffusion in the amplification stage via minimised GEM distances, and reconstruction uncertainty from the anode granularity.

GEM distances Gd / mm	σ_{trans}/mm	σ_{long}/mm
2, 2, 3	0.358	0.129
1, 1, 1	0.232	0.0826
0.5, 0.5, 0.5	0.164	0.0584
0.1, 0.1, 0.1	0.0734	0.0261

Table 10: Transversial and longitudinal diffusion used in the simulated amplification stage depending on the inter-GEM and GEM-anode distances.

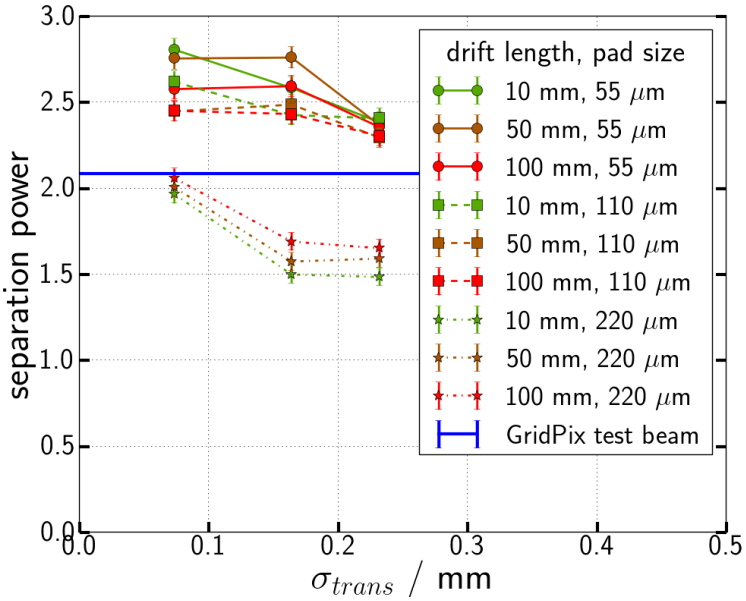

Figure 72: Separation power for extreme parameter values, depending on transverse diffusion (Table 10). Values in the legend are drift length and pad size; a GEM voltage of 300 V was used. The blue line represents the performance of the pixel-based readout.

6.3. Conclusion

6.3.1. Summary and Discussion

In order to study the dependence of the PID performance of a GEM-based TPC readout on its anode granularity, a detailed simulation chain was set up. The software package Source-Extractor was implemented in to this chain to reconstruct primary ionisation clusters. With this reconstruction tool, an effective cluster counting capability could be established. A large parameter space was scanned and the gain and reconstruction algorithm were optimised for each scanning point. The results of this simulation study span two orders of magnitude in granularity and thus manage to connect the published performances of existing pad-based and pixel-based systems. For mm-sized pads, charge summation performance increases according to empirical expectations with increasing granularity, until it becomes limited by the maximum gain and drops off. For pads below 500 µm, cluster counting becomes effective and outperforms charge summation. For the smallest simulated pads, cluster counting can perform on the same level as the pixel-based system.

The study of the situation at the smallest pad sizes does not indicate how a conventional GEM stack could be modified to enable highest performance at 55 µm granularity, but rather that the aforementioned kinds of limitations would need to be overcome to achieve this. It is also a contribution to a potential later study connecting the pad-based systems with the pixel-based ones at a more fundamental level throughout the simulation chain. Such a study would need to look further into availability and loss of spatial information and cluster correlation, and how this depends on the different technologies. If the spatial uncertainty from diffusion and granularity is kept at a minimum in each step of the signal progression, the separation power performance can even be further increased, indicating a possible path of improvement for the pixel-based system via a dedicated cluster-counting algorithm. At the same time, the simulation indicates that in an ILD-like detector the drift distance and the corresponding diffusion may fundamentally limit the potential of cluster counting to a performance which can already be achieved with a somewhat reduced granularity between 100 µm and 300 µm.

Compared to the pad-based systems, it is however clear that a significantly improved PID performance can be achieved with increased granularity, even with charge summation: from 4.7 % with 6 mm pads, to 4.2 % with 1 mm pads and charge summation, to 3.7 % with 165 µm pads, and potentially up to 3.3 % or better with GridPix, at least for short drift lengths. This motivates physics analyses which show their performance dependence on the dE/dx resolution to indicate what level of dE/dx is beneficial or even required for the ILD physics performance. To compare the default ILD case with about 4.7 % to the case of no TPC, i.e. without a dE/dx observable or a rather limited, silicon-based one, and to the case of a very optimised readout setup with a resolution in the order of 3 %, first an assessment and calibration of the currently implemented dE/dx calculation is performed in chapter 7, and one such analysis is presented in chapter 8 showing the dependence of physics observables on the dE/dx resolution.

6.3.2. Points of Improvement of the Simulation

Apart from the intrinsic limitation of a simulation regarding simplification and numerical precision, there are several points within the existing simulation and reconstruction chain that can be a target of future attempts of improvement. Generally, the simulation should in principle be compared to test beam measurements with a dedicated GEM-based high-granularity system, like a successor of the Ropperi system.

It is known that the ionisation parametrisation used by the `PrimaryIonisationWithClusterProcessor` does not entirely reflect measurements, in particular the spectrum of secondary electrons. This could affect the absolute measure of number of reconstructed clusters, as well as measured charge. This issue is compensated by using a differential measure, the separation power, as well as by stressing the relative performance behaviour with respect to the pad size and in comparison to the charge summation measure, rather than pointing to absolute measures. Nonetheless, the ionisation spectrum can impact the separation power in a non-trivial way and might benefit from a novel implementation with higher accuracy.

After an initial optimisation of many input parameters of the Source-Extractor for a default working point, only one parameter, the convolution filter, was used as a scanning tool for point-by-point optimisation. The minimum and maximum numbers of active pads to be identified as a source were also checked, but a complete scan was not possible in the scope of this thesis. Therefore, there is still a range of parameters available to be investigated in future optimisation studies.

In this study, only either the number of reconstructed sources/hits or the registered channel charge was used to calculate a dE/dx estimator. It is possible that using more information on the hit level, like width and amplitude (integrated charge), could result in a better estimator. For example, the correlation of hit size and contributing clusters could be studied to weigh large hits accordingly stronger whilst counting them.

The source-extractor convolution filter in all cases had a central Gaussian shape. Other filters, or using none at all, could be investigated. The typical shapes of anode responses to individual electrons or clusters could be quantified and used to derive an optimised convolution filter in terms of radial dependence and possibly even angular. Finally, a neural network could be applied to optimise the convolution filter, either with a dedicated tool or, as suggested in the handbook, with the EyE software [112], which performs so-called 'retina-filtering'.

Most simulation was done without channel noise, assuming it can be mitigated in the reconstruction. There are several ways conceivable to do so. Firstly, the SE has the option to scan the input image for background and then subtract it. Only the part above background is then used for source identification. Secondly, one could apply a multi-step approach: using a broad convolution filter or a large threshold value, the SE would identify only a limited number of sources associated with the track, but hardly any stemming from noise. These sources can be used for a simple track finding, e.g. by Hough transformation to identify the corridor on the image in which hits are expected. A regular source extraction can then be done on only this corridor for a dE/dx estimator.

7. Particle Identification Performance of a Full ILD Simulation

A multi-purpose detector like the ILD has more than one way to identify particles. While the previous chapters concentrated on the TPC's dE/dx measurement, it has also the capability to identify particle decays in-flight, the so-called V0-finding. With a sufficient lever arm, a measurement of the time-of-flight allows to identify low-momentum particles. This is not yet included in the ILD baseline design, but is studied here in simulation in order to assess the potential benefit of this technique. In this chapter, for each method the software implementation in the parameterised simulation of the full ILD is laid out and their respective performances are displayed, with a particular focus on the dE/dx-measurement.

7.1. Software for a Full ILD dE/dx Measurement

The main tool to calculate the specific energy loss of a track in a full-ILD simulation and reconstruction is the `Compute_dEdxProcessor`. It was created by M. Kurata, and revised and extended within this work. Its functionality will be described in the following.

The `Compute_dEdxProcessor` takes individual tracks that were found by the tracking software and calculates a dE/dx estimate for each track and attaches it to the track object. For this, the processor takes all hits associated with the track and calculates the local dE/dx of each hit by dividing the deposited energy of the hit by the cartesian distance to the previous hit. In the context of this work, alternate ways to calculate the hit dE/dx, taking curvature and missing hits into account, were implemented, but the current algorithm reflects the simulation conditions well and delivers the best performance. To make a dE/dx estimate for the entire tack, a trimmed truncated mean is calculated, with a truncation fraction of 30 % and a trimming fraction of 5 %. For tracks which revolve in the magnetic field and return into the TPC or curl entirely inside it, only the first half circle of the track is taken into account when combining hits.
The physics lists used in Geant4 work well for showers in the calorimeter, but less so for ionisation in the TPC gas. The resulting hit energy distribution is too narrow compared to beam test results, leading to a too good dE/dx resolution. Therefore, each track dE/dx estimate is smeared with a random number following Gaussian distribution, whose width is called smearing factor. This smearing factor is gauged so that the final track dE/dx distribution has a width corresponding to beam test results.

7.2. Performance of Full ILD dE/dx-Based PID

To assess the performance of the dE/dx-measurement, a new piece of software, the `dEdxAnalyser`, was implemented as a Marlin processor. It collects a large number of dE/dx-related observables, sorts and stores them, processes them to higher-level observables like separation power, and allows for a quick display of the results. Its implemen-

tation into iLCSoft should benefit and ease future related analyses. In this chapter, the single-particle samples of the 2018 production (chapter 4.5) were used, if not stated otherwise. Only tracks connected to reconstructed PFOs were used and were associated with the species of the MC particles connected to the PFO.

The first and most fundamental plot is the Bethe-Bloch curves, shown in Figure 73 for single-particle tracks. Individual particles of the five species electrons, muons, pions, kaons and protons (P5) with random momentum (approximately flat in log(p)) originating from the interaction point were shot isotropically into the detector. The plot shows the energy loss in the TPC as deposited energy per traversed distance over the particle's momentum for each track. The individual bands correspond to the five particle species, with pions and muons fully overlapping within the given resolution. Two very thin additional bands correspond to deuterium and tritium nuclei which can be produced in occasional hard scattering processes of the simulated initial particles in the the inner tracking system. The vertical width of the bands is related to the resolution, and for each bin in momentum one can define the dE/dx-resolution as $\sigma_{dE/dx} = \frac{\sigma_i}{\mu_i}$ with μ_i and σ_i being the mean and width of a (vertical) Gaussian fit applied to bin i of the individual band.

The result of this bin-wise vertical Gaussian fit is shown in Figure 74 (mean and width) and Figure 75 (constant term and χ^2/ndf) for pions. In Figure 76 the resulting dE/dx-resolution vs. momentum is plotted. Below a momentum of 1 GeV, the resolution worsens because the pion momentum is too small to reach the outer rows of the TPC. Since only the first half circle of each track is utilised in the track dE/dx calculation, fewer than the maximum number of hits are taken into account this way, which limits the resulting resolution.
To receive a single resolution number for all tracks of one species, the resolution vs. momentum plot is 'integrated' above 1 GeV by filling the energy loss of each track into a histogram after dividing it by the average energy loss in the corresponding momentum bin, as extracted from the Bethe-Bloch curve. The width of the resulting distribution, seen in Figure 77 for pions, is then a measure for the overall dE/dx-resolution of that particle species in the detector.
All curves for single particles are given in chapter A.1.

A more direct connection to the final PID is the separation power between species. It is defined analogous to Equation 7 and Equation 8 as

$$S = \frac{|\mu_1 - \mu_2|}{\langle\sigma\rangle} \tag{22}$$

with μ_i the mean energy loss for species i and $\langle\sigma\rangle$ being a mean width of the energy loss of the two species being compared, and

$$\langle\sigma\rangle = \sqrt{\frac{1}{2}\left(\sigma_1^2 + \sigma_2^2\right)} \tag{23}$$

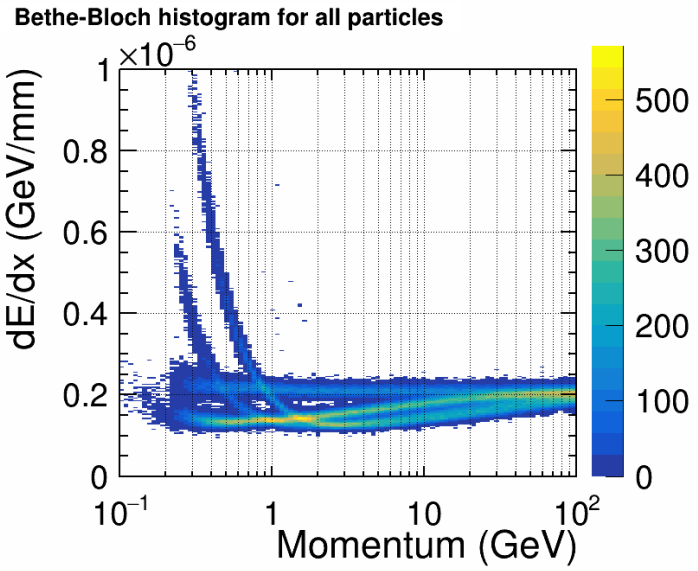

Figure 73: Bethe-Bloch curves for large ILD single particles.

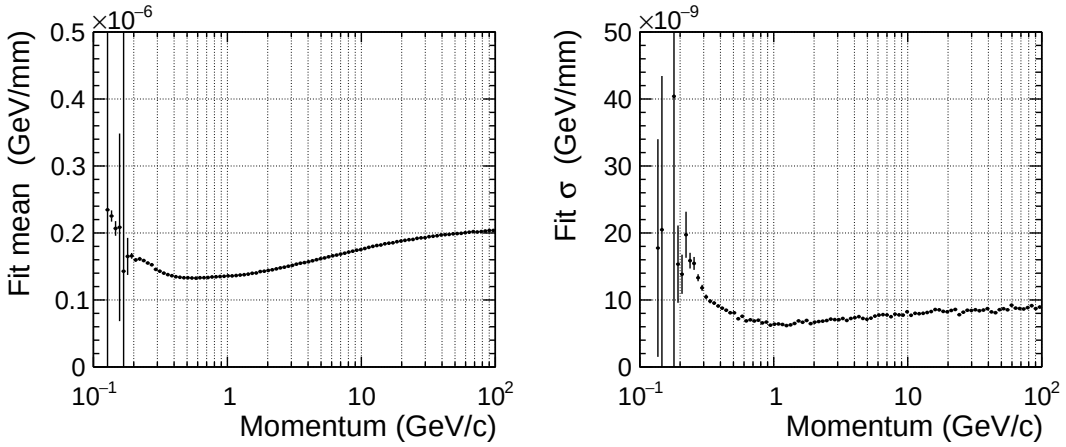

(a) (b)

Figure 74: Mean (a) and width (b) of a Gaussian fit in each momentum bin to the Bethe-Bloch curve of single pions.

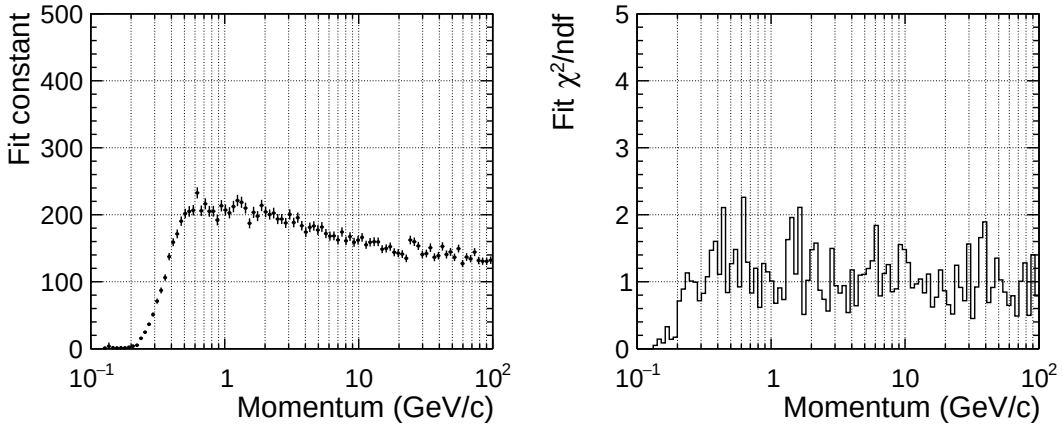

(a) (b)

Figure 75: Constant (a) and χ^2/ndf (b) of a Gaussian fit in each momentum bin to the Bethe-Bloch curve of single pions. The Constant shows that the distribution of particles used in the sample is roughly flat in $log(p)$ for momenta above 500 MeV. Above about 200 MeV the χ^2/ndf is distributed around one.

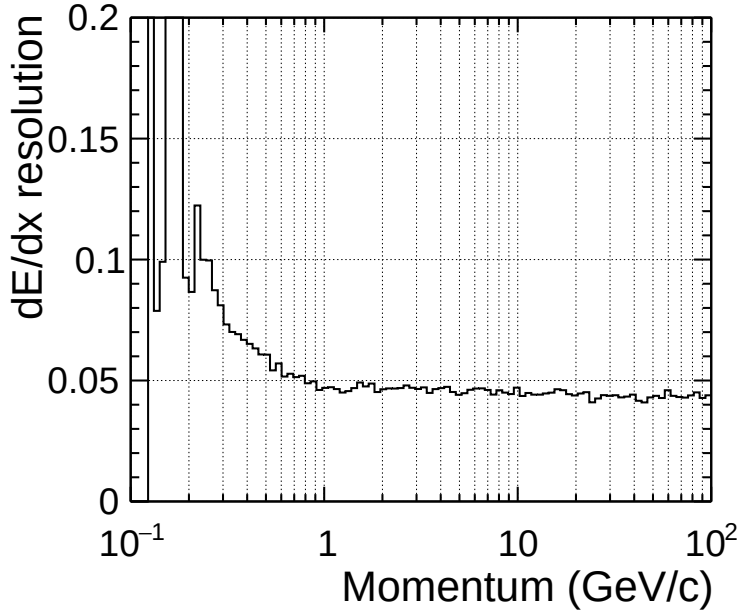

Figure 76: dE/dx-resolution of single pions over momentum. The resolution is calculated bin-wise by division of the Gaussian fit result width over mean.
For $p > 1$ GeV the dE/dx-resolution ist nearly flat and has a value close to the design value. For $p < 1$ GeV the resolution worsens because the track curls and its first half has fewer than 220 hits, which are being used to calculate the dE/dx.

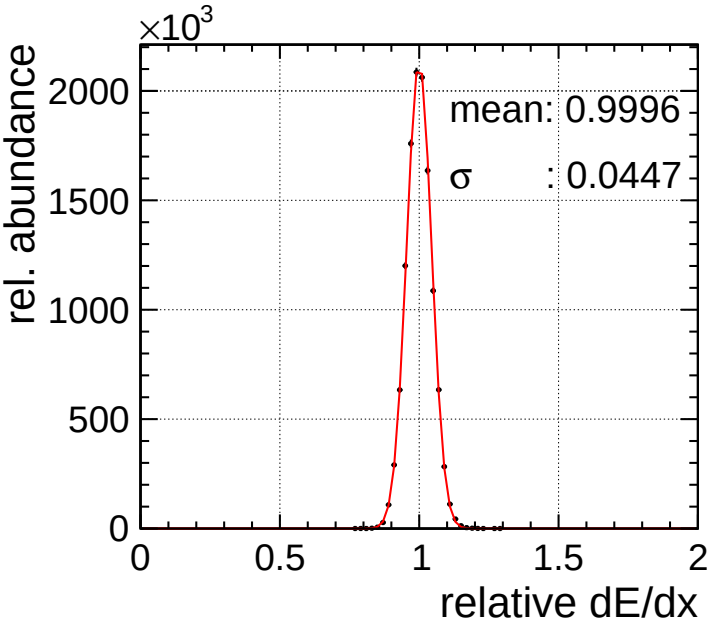

Figure 77: Combined dE/dx resolution of pions above 1 GeV momentum. A Guassian fit is depicted as a red curve, with its mean and σ highlighted.

with σ_i the width of the energy loss for species i. Since the Bethe-Bloch bands have about the same width for each species for most momentum bins, the difference between using the width of only one species and this combination of two species is mostly negligible. The quantity S is calculated for each momentum bin. As example, the separation power of electrons vs. each other species is shown in Figure 78. In the peak region the separation power reaches a value of 8 to 10, meaning in a more practical way that the two Bethe-Bloch curves are 8σ apart from each other, which leads to very small remaining cross-contamination and near perfect separation. The momentum region with $S > 3$ is often considered the area of effective separation. The separation power of all combinations of particle species can be found in chapter A.1.

Electrons and muons with a transverse momentum larger than about 1 GeV can be identified rather well based on their ECal and HCal cluster shapes. For lower momenta, electrons have a large separation power by dE/dx with respect to the other charged particles species, as displayed in Figure 78. Muons, on the other hand, have a large overlap of their Bethe-Bloch band with pions, since their masses are close to each other, and dE/dx cannot be used to effectively differentiate them. However, at low momenta their cluster signature in the forward calorimeter system can be used to identify muons (pions) with an efficiency above 75 % (50 %) for momenta as low as 300 MeV [113]. Thus, the quantities mostly used to classify a detector's PID performance and compare it to others are pion-kaon and kaon-proton separation. In the context of the IDR, the two detector options IDR-L and IDR-S (see chapter 2.3) are compared, since the different TPC sizes should have a sizable effect on the dE/dx performance. In Figure 79, prepared

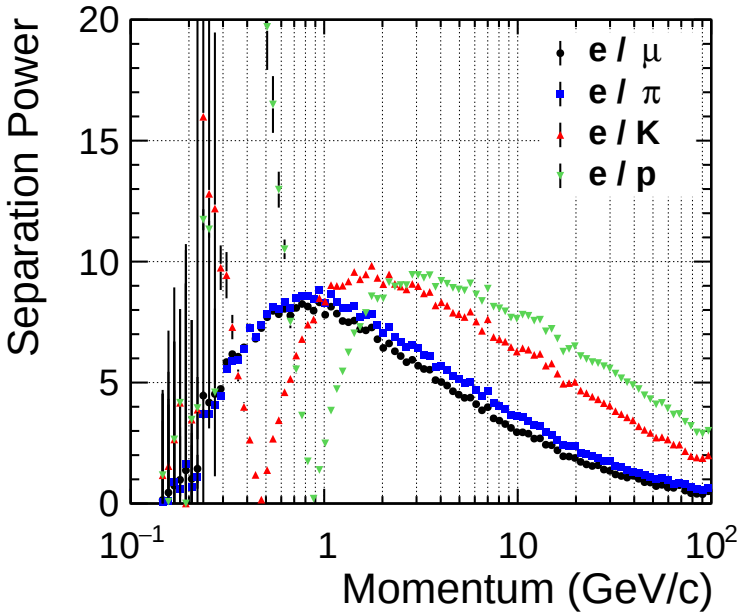

Figure 78: Separation power of electrons vs. each other species.

for the IDR, the pion-kaon separation and kaon-proton separation are compared for the two detector options. As expected, the larger TPC radius of the IDR-L allows for a better dE/dx performance. For pion-kaon separation, in the peak region at a few GeV the separation power reaches up to 3.3 for IDR-L and 2.8 for IDR-S. Similarly at higher momenta, the kaon-proton separation reaches 1.8 for IDR-L and 1.5 for IDR-S. This factor of about 1.2 between the two options stretches across nearly the entire momentum range and is non-negligible in later PID applications.

To evaluate the applicability of the single-particle studies to PFOs in full events with multiple tracks at once, the dE/dx-resolution of pions and the pion-kaon separation are compared between single particle samples and 6f-ttbar samples at 500 GeV. In $t\bar{t}$ events between 2 and 6 jets are created with a large number of hadrons originating from hadronisation. If tracks overlap, the readout channels in the overlap region are discarded for the dE/dx calculation, since the individual contributions of the tracks can not be easily interpolated or fitted. This reduces the usable number of track hits, leading to a worse dE/dx-resolution and separation power. This effect increases with larger channel foot print, since the entire area covered by one channel cannot be used in case of an overlap, even if this overlap would only actually occur in a fraction of that area. Thus, higher readout granularity helps to mitigate this performance reduction. Figure 80 shows that for ILD this performance reduction is small and both dE/dx resolution and separation power are nearly as good in $t\bar{t}$ events as in single particles. The plot also shows that the available statistics of tracks in the $t\bar{t}$ samples are much larger compared to the single particles, leading to much smaller statistical fluctuations between momentum bins.

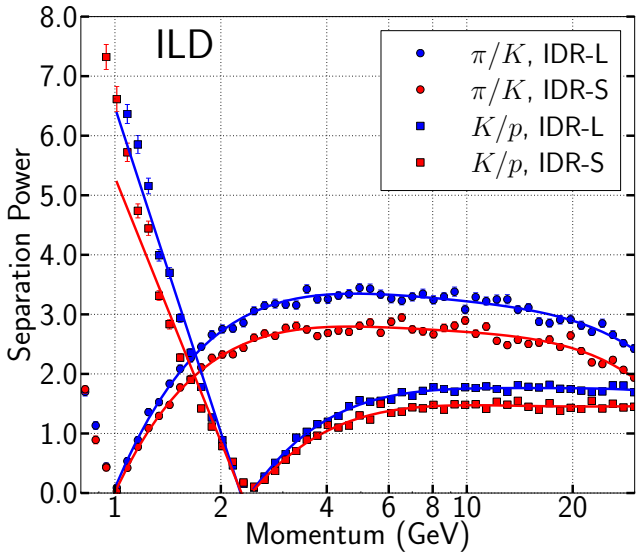

Figure 79: Pion-kaon and kaon-proton separation in large and small ILD. The continuous lines are polynomial fits over log(p) and just to guide the eye. The large ILD has an about 20 % higher separation power compared to the small ILD over the full momentum range.

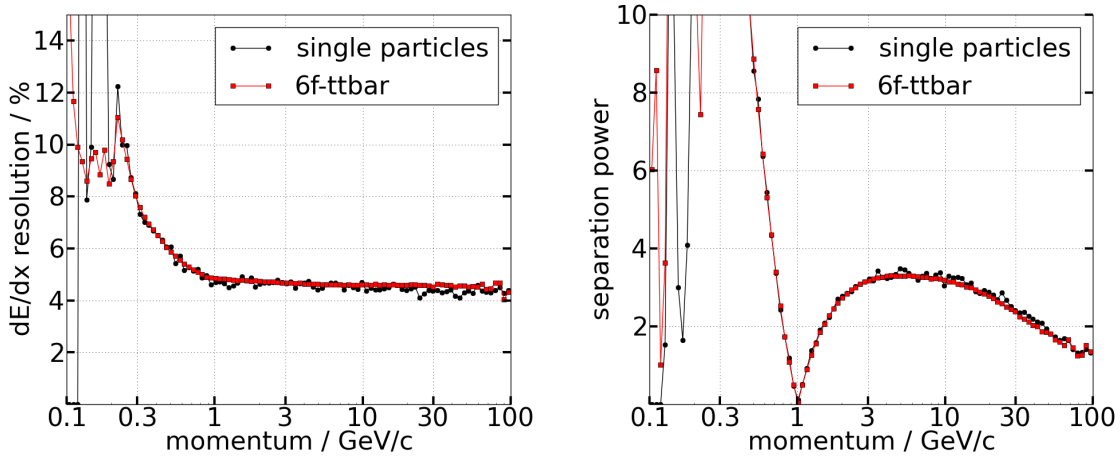

Figure 80: dE/dx resolution of pions (left) and pion-kaon separation power (right), in comparison between single particles and full $t\bar{t}$ events, for the large ILD model. The performance in complex, more realistic $t\bar{t}$ events is nearly the same as in single particles.

7.3. Calibration of the dE/dx Simulation

After a change in the Geant4 version from v10.3.2 to v10.4.3 for the 2020 MC production the hit energy distribution had changed (comp. Figure 81). The new hit energy is much closer to test beam measurements [40], but this made recalibration of the smearing factor necessary. A procedure for this was developed and is laid out in the following.

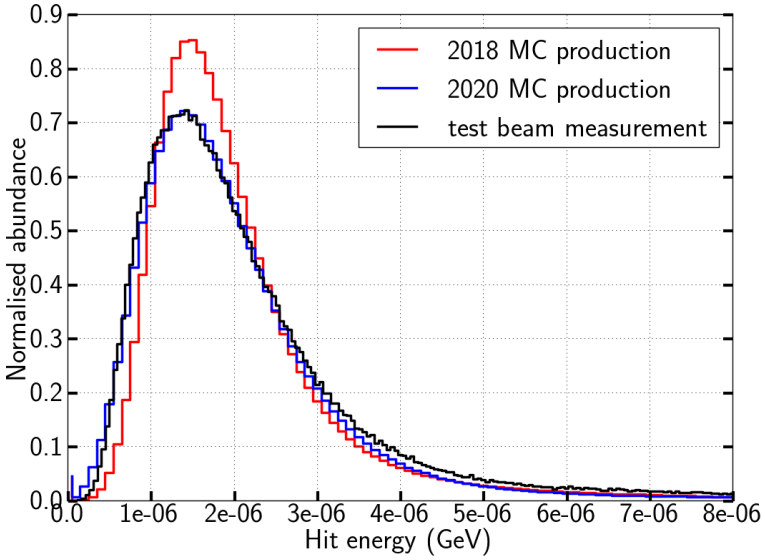

Figure 81: Hit energy distribution of TPC hits from single pions with ionisation by Geant4, comparing the 2018 and 2020 MC productions to test beam measurements [40].

By re-running the `Compute_dEdxProcessor` without a smearing factor, the intrinsic width of the track dE/dx distribution only based on Geant4 can be extracted. In the 2018 production this intrinsic width is 2.8 %, so the smearing factor was calibrated to 3.5 %, leading to a final dE/dx resolution of about 4.5 %, which is consistent with a naive combination via $\sqrt{2.8^2 + 3.5^2} = 4.48$. In the 2020 production the intrinsic width is 3.8 %, which would lead to a new smearing factor of 2.4 % with this naive approach. This adapted smearing factor led to a more accurate overall dE/dx resolution, but showed significant differences between the investigated particle species, see Figure 82. This had not been observed before, since the differences were largely smeared out by the smearing factor. However, this difference between the species is consistent with the underlying distribution when mass is considered by looking at the dE/dx resolution distribution over $\beta\gamma$ of each particle instead of its momentum, see Figure 83. All particles follow the same basic distribution which approaches a certain asymptotic value σ_0, as indicated in the figure, for large $\beta\gamma$. The rises of each curve to its low-$\beta\gamma$ end are the areas in momentum where each particle has less than 1 GeV/c momentum and curls in the TPC, leading to a drop in hit number and thus a rapid worsening of the dE/dx resolution (along with much smaller statistics). This way, within a given momentum range, each

species is sampled at a different part in the $\beta\gamma$-space, which leads to a different effective dE/dx resolution for each species.

To calibrate the simulation to beam test results the conditions of the particles investigated in simulation should be as similar to the ones in beam tests. In the beam tests of the LCTPC collaboration which are used as reference for the ILD simulation, the particles used are electrons with a momentum of 5 GeV/c, moving through the large prototype TPC parallel to the anode and cathode. These conditions were met in simulation by only using electrons with a momentum $3\,\mathrm{GeV/c} < p < 10\,\mathrm{GeV/c}$ and an angle λ relative to the cathode with $1° < |\lambda| < 10°$. Angles below 1° were excluded to avoid electrons, which all start at the IP, from interacting with the cathode. In addition, since the reference dE/dx-resolution of the beam test results is extrapolated to the maximum number of 220 hits, track with fewer than 200 hits were excluded. These cuts left about 2300 electrons of the 100 k in the single particle sample.

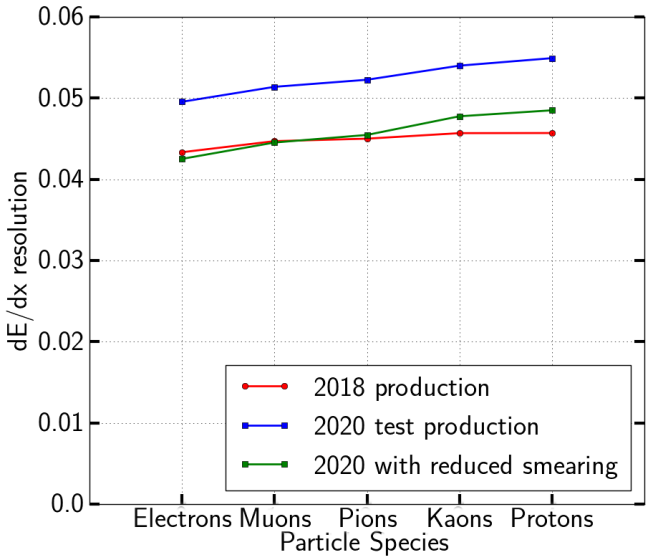

Figure 82: Effective dE/dx resolution for the P5 in single particle samples, comparing the 2018 MC production with the 2020 test production with the 2018 smearing factor and a reduced smearing factor of 2.4 %.

Given the properties of these fiducial electrons, i.e. large $\beta\gamma$ and large number of hits, they have a better dE/dx resolution than the average of all particles. This led to the observation that fiducial electrons in the 2018 production had an effective resolution of 4.3 %, and in the 2020 production with the previously proposed smearing factor of 2.4 % had a resolution of 4.2 %. This is both smaller than the values reported by the LCTPC collaboration. Instead, a larger smearing factor of 2.9 % needs to be applied in simulation in order for fiducial electrons to have an effective dE/dx resolution of 4.5 %, see Figure 84. Consequently, the final effective resolution values for the heavier particles, as displayed in Figure 85, are worse than the ones of electrons. They are also somewhat worse than the respective values from the 2018 production, but more realistic due to

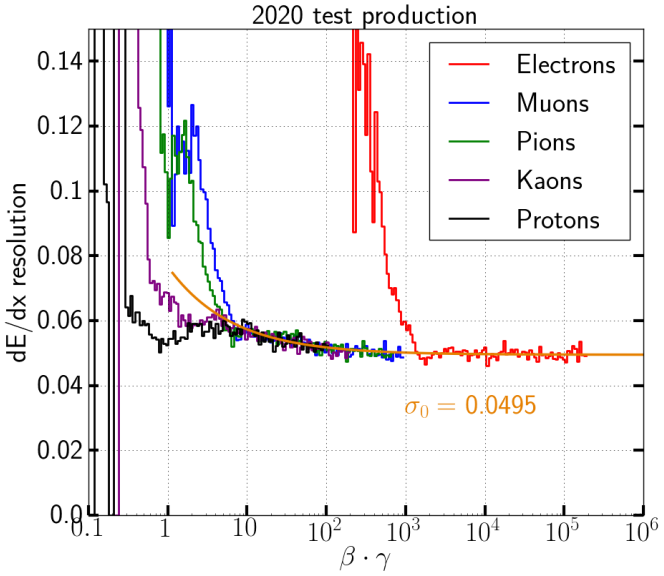

Figure 83: dE/dx resolution of the P5 vs. $\beta\gamma$ in the 2020 MC test production.

a smaller necessary smearing factor and due to the calibration with respect to fiducial electrons.

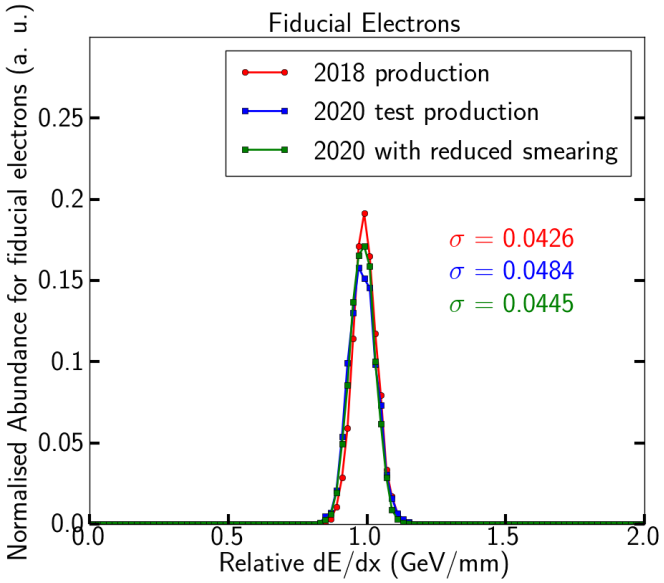

Figure 84: Track relative dE/dx distribution of fiducial electrons, comparing the 2018 MC production with the 2020 test production with the 2018 smearing factor and a reduced smearing factor of 2.9 %.

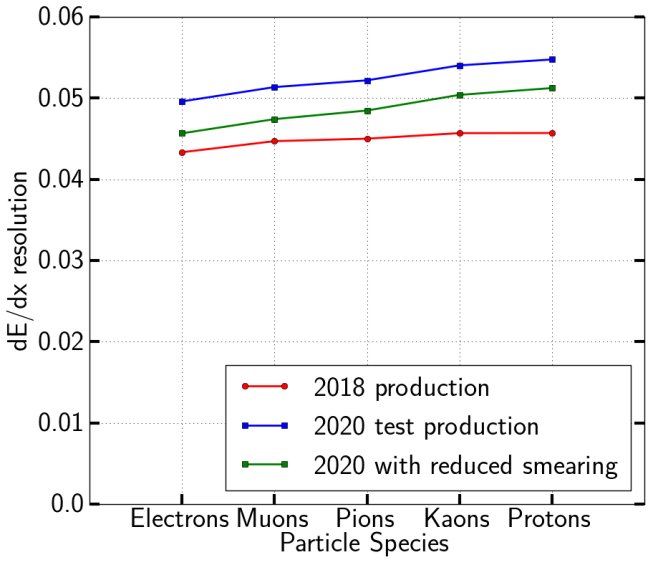

Figure 85: Effective dE/dx resolution for the P5 in single particle samples, comparing the 2018 MC production with the 2020 test production with the 2018 smearing factor and a reduced smearing factor of 2.9 %.

7.4. V0-finding PID at ILD

With its large number of track points, the TPC allows to efficiently identify decays of neutral medium-lived particles inside the tracker by their decay products. This procedure is called 'V0-finding' and applies to K^0_S and Λ, if they decay inside the tracking volume, and γ if they undergo e^+-e^--pair production. K^0_S has a typical decay length of $c\tau = 2.68\,\text{cm}$ and decays to about 70 % into charged pion pairs. Λ decays to about 64 % into a proton and a pion with $c\tau = 7.89\,\text{cm}$. Photons can interact with the material and create electron-positron pairs.

To identify these initial neutral particles, a dedicated Marlin processor, the V0Finder, creates combinations of all tracks with each other and calculates the invariant mass of the combination. If it is consistent with the masses of K^0_S or Λ within 20 MeV or a γ within 10 MeV, the two tracks are stored as the corresponding V0-particle. Since both K^0_S and Λ are strange particles, this is a way to tag strangeness.

The efficiency of tagging massive V0s depends strongly on their momentum, since a cut on the minimum distance of the V0 decay vertex from the IP is applied to suppress combinatorial background. The default value of this radial cut used in MC productions is 30 mm for K^0_S and 50 mm for Λ. This cuts away a large portion of low momentum V0 particles which decay at smaller radii, significantly reducing the overall efficiency of the processor.

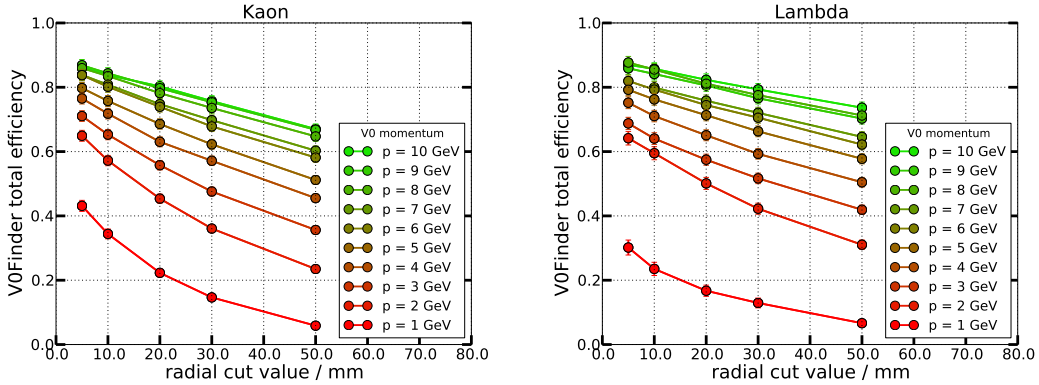

Figure 86: Reconstruction efficiency of the V0Finder for K^0_S (left) and Λ (right), for different particle momenta and depending on the chosen radial cut (minimum distance from IP).

7.5. Time-Of-Flight PID at ILD

Another way to identify particles is by reconstructing their mass from a time-of-flight measurement (TOF). For charged particles, the travel time from the interaction point to a certain detector layer which measures the arrival time, in this case the electromagnetic calorimeter (ECAL), is compared to the path length of the reconstructed track to calculate the velocity or β of the particle. The MarlinReco software contains a TOF estimator, which was implemented following [114] and [115] and will be explained along Figure 87 in the following.

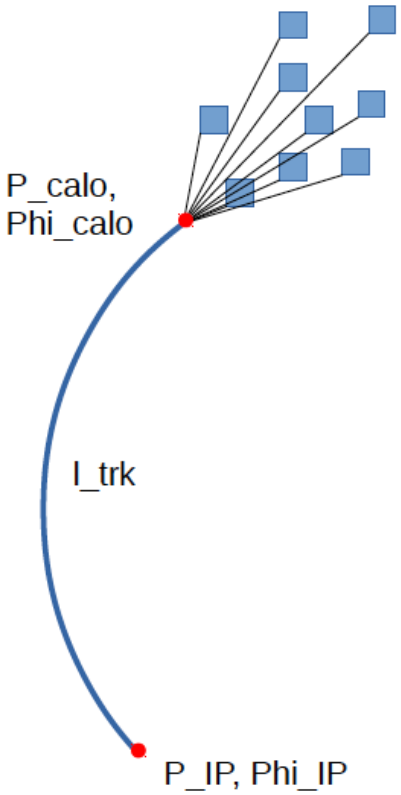

Figure 87: Scheme of the time-of-flight measurement, from [114].

For each PFO, both time difference and path length are calculated between the IP and the entry point into the ECAL P_{calo}. The path length l_{trk} is calculated based on the helix fit to the track points, which determines the point of closest approach to the IP P_{IP}, the track curvature ω, the angle at the IP and P_{calo}, Φ_{IP} and Φ_{calo}, as well as the track inclination angle λ:

$$l_{trk} = \left(\frac{\Phi_{IP} - \Phi_{calo}}{\omega} \cdot \sqrt{1 + \tan^2(\lambda)} \right). \tag{24}$$

For the time measurement, an average of up to 10 ECAL hits is used: For each of the first 10 ECAL layers, the hit closest to a linear extrapolation of the track from P_{calo}

into the ECAL is used, if one is registered in the calorimeter cluster associated with the PFO in that layer. In the simulation, the time assigned to each hit is the time since the initial hard scattering event at the IP. This hit time is smeared by adding a Gaussian with a certain width, reflecting a finite timing resolution. To estimate the time at $\mathrm{P_{calo}}$, each hit time is corrected by assuming a linear travel of a signal-creating particle with the speed of light from P_{calo} to the hit position P_{hit}. Those corrected hit times are then averaged to result in an estimate for the cluster time t_{clu} at P_{calo}.

$$t_{hit,cor} = t_{hit} - \frac{\mathrm{dist}(\mathrm{P}_{calo},\mathrm{P}_{hit})}{c} \tag{25}$$

$$t_{clu} = \sum_{i=0}^{N} t_{i,cor} \tag{26}$$

Given the path length and the time, an estimate for the speed of the PFO, respectively its β can be calculated. A PFO mass m_{PFO} can be reconstructed by taking into account the reconstructed momentum p of the PFO, which is an estimate for its momentum at the IP.

$$\beta = \frac{v}{c} = \frac{l_{trk}/t_{clu}}{c} \tag{27}$$

$$m_{PFO} = \sqrt{\frac{p^2}{\beta^2} - p^2} \tag{28}$$

Several simplifications in this procedure can introduce a systematic error to the β and mass estimates: the linear extrapolation of the track into the ECAL despite it being inside the solenoid, the correction of the hit times to the reference point using linear travel at the speed of light, and using the reconstructed momentum at the IP. Future updates to the TOF estimation software should investigate the systematic bias introduced this way and aim to remove it or correct for it. Furthermore, in the current version of the software, only the barrel ECal is used to measure the TOF of particles. This should be extended to the endcap calorimeter system.

For the IDR production of MC data, several different time resolutions were used, added as Gaussian width to each hit time. For the single particle calibration files, 10 ps, 50 ps and 100 ps estimators are available, while in the large full-SM production, 0 ps, 10 ps and 50 ps were used. Providing several estimators allows for comparison. At the time of writing of this thesis, a hit time resolution in the ECAL of 100 ps is assessed as doable, 50 ps as ambitious but achievable, and 10 ps as requiring substantial research and development and dedication. In real data, the hit time is measured with respect to the reference clock aligned to the brunch crossings. Since the foreseen bunch length at the ILC is $\sigma_z \approx 300\,\mu\mathrm{m}$, the bunch crossings take place within a ps, which is negligible compared to achievable hit time resolutions. A certain timing resolution target puts, however, contraints on the overall clock synchronisation, which also needs to be negligibly precise.

A first study of the TOF performance was done by S. Dharani in [6]. This work extended the analysis and combined the TOF and dE/dx performances.

The TOF performance of a full ILD simulation can be seen in Figure 88 and Figure 89. In Figure 88 the measured β of single particles is displayed, leading to characteristic bands for each mass, which separate well up to a few GeV. Even for single particles, spurious hits are present in the plot which mostly come from mis-measured particles that scattered in the tracker material or from back-scatter from the calorimeters. In Figure 89 the PID performance measured in π/K-separation can be seen, where it is combined with the corresponding dE/dx performance by square-root adding the separation power. For single particles, also 100 ps ECAL hit time resolution was simulated. TOF starts to work as soon as the particles have sufficient transverse momentum to reach the ECAL, which is the case at about 1 GeV for the large ILD model. At low momenta TOF is very effective and provides great separation power. This quickly reduces with growing momentum and becomes negligible at about 3.5 GeV, 4.5 GeV and 6 GeV for a time resolution of 100 ps, 50 ps and 10 ps, respectively, when requiring a separation power of at least 2. This way, TOF combines well with dE/dx-PID by covering the minimum of the dE/dx separation power where the corresponding Bethe-Bloch bands overlap. In the example of π/K, this is the case at 1 GeV, where TOF starts to be available and is most effective. TOF dominates the separation power up to a few GeV, where it becomes comparable to dE/dx and a direct combination of dE/dx and TOF information exceeds the individual contributions. TOF then quickly becomes negligible, while dE/dx continues to perform at a relevant level up to several 10 GeV. Figure 90 shows the same plot for kaon-proton separation. Here, TOF becomes more effective, since kaons and protons have larger masses and can be distinguished by TOF up to 6, 7.5 or 9 GeV with a hit timing resolution of 100, 50 or 10 ps, respectively, and again covers the blind spot of dE/dx at about 2.5 GeV. At the same time, the dE/dx separation power is effective up to several 10 GeV, but only up to $S = 1.8$.

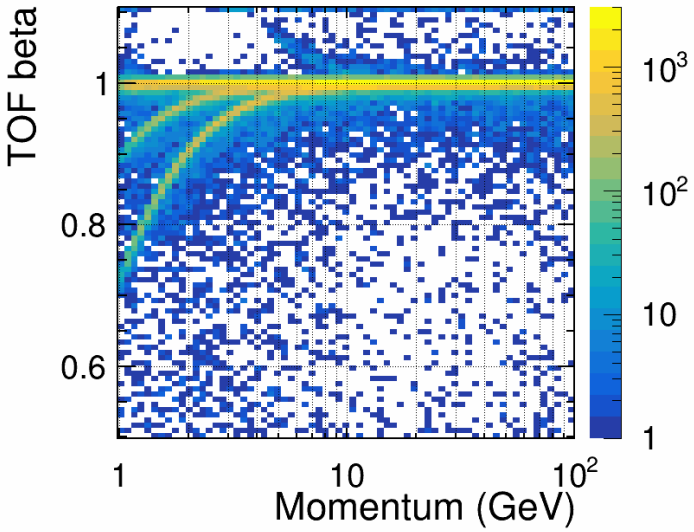

Figure 88: TOF-reconstructed β over reconstructed momentum for single particles, assuming an ECal hit time resolution of 50 ps. The three bands correspond to pions, kaons and protons, from the top, and are well separated up to a few GeV.

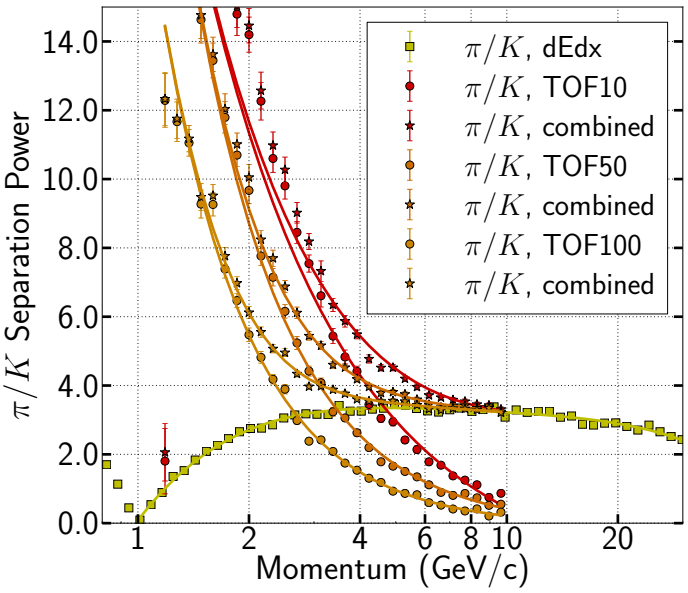

Figure 89: Pion-kaon separation in dE/dx and TOF [6] for 3 different single-hit time resolutions. The continuous lines are polynomial (dE/dx) and exponential (TOF) fits over log(p) and just to guide the eye. The plot was prepared for [7].

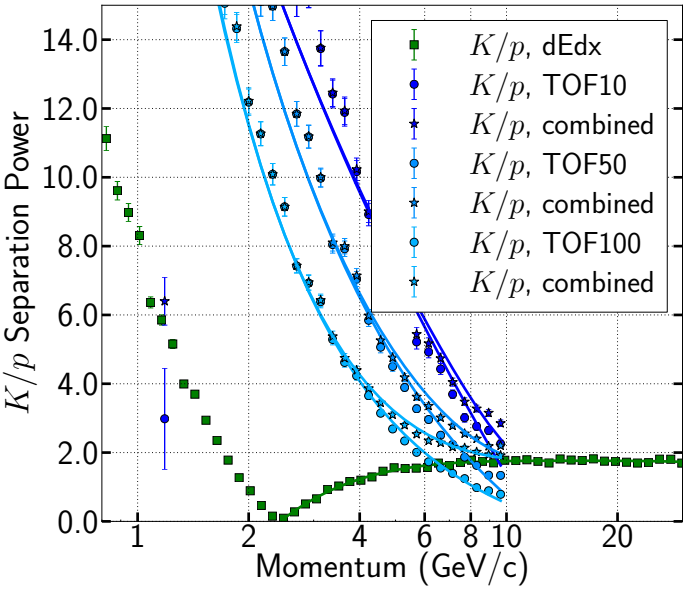

Figure 90: Kaon-proton separation in dE/dx and TOF [6] for 3 different single-hit time resolutions. The continuous lines are polynomial (dE/dx) and exponential (TOF) fits over log(p) and just to guide the eye. The plot was prepared for [7].

8. Application of Advanced Particle Identification to Physics Events in a Full ILD Simulation

After showing how an improved particle separation with increased dE/dx resolution can be achieved, now its impact on high level observables, including the flavour tag, branching fraction and reconstructed mass will be investigated.
The achievable precision of the mass of the charged kaon, whose exact value is a long-standing open question, is estimated using PID and reconstructed momentum only. This shows the prospects of a highly precise detector like ILD.
The decay of W bosons is taken as an example to show the impact of tracking-based flavour tagging on the separation of different W-decay channels compared to the conventional vertex-based tagging algorithm. The role of dE/dx in particular is tracked in a full detector simulation, leading to improved particle ID and finally to an improvement in separating different W-decay modes.

For both analyses, samples of the MC production of 2018 [77] with a centre-of-mass energy of 500 GeV were used, see chapter 4.5.

8.1. Kaon Mass Measurement Prospects

One application of the PID capabilities of the ILD is to solve a long-standing open question: the mass of the charged kaon. The current best estimate $m_K = (493.664 \pm 0.013)$ MeV published in the Particle Data Group particle physics review [13] is a combination of several measurements from the early 1990s. All the measurements were done with kaonic atoms formed by leading a slow kaon particle beam into a target material. The x-ray spectra of the kaonic atom's energy levels were observed and compared to theoretical calculations, which depend on the kaon mass. The leading two contributions to the combined estimate are two particularly sensitive energy transitions in different base elements, carbon and lead, measured by two different groups [14, 15], which disagree with each other beyond the 3σ-level, as shown in Figure 91. In the calculation of the kaon mass world average this disagreement is compensated by scaling up the error by a factor of 2.3 to the current value of 13 keV.
Using the luminosity of the ILC to produce a large number of charged kaons, and the PID capability of ILD to reconstruct them correctly and measure their mass, could help resolve the disagreement. In this chapter, the prospects of ILD to measure the charged kaon mass will be quantified.

To measure the mass of charged kaons, no specific physics process is necessary, rather an as large as possible number of kaons should be collected from any process. For this thesis, a full simulation and reconstruction including PID of physics processes at ILD was available for a center-of-mass energy of 500 GeV, see chapter 4.5. Since most kaons are produced in hadronisation, channels with hadronic final-states were selected, while many leptonic or semi-leptonic processes as well as processes with low overall cross sections were omitted in this analysis for means of practical efficiency. The samples selected are listed in Table 11. There are a number of proposed running scenarios for ILC [24] in terms of luminosity at the different center-of-mass energies. This study used the so-called H-20 scenario which at 500 GeV assumes a total integrated luminosity of $4\,\mathrm{ab}^{-1}$ over the full running time of the ILC. It is compared to a scenario with $400\,\mathrm{fb}^{-1}$, which is roughly the expected integrated luminosity before a luminosity upgrade, or in case of an energy upgrade after a luminosity upgrade the amount of data within the first one or two years of running. Each integrated luminosity is split between different beam polarisation combinations according to Table 12.

short form	final state flavours consistent with:
2f-Z-had	production of one Z boson, only hadronic decays
4f-ZZ-had	production of two Z bosons, only full-hadronic decays
4f-WW-had	production of two W bosons, only full-hadronic decays
6f-ttbar	production of $t\bar{t}$, all decay channels

Table 11: Selection from the 2018 IDR Monte Carlo production [7, 77] (see chapter 4.5) used for the kaon mass measurement.

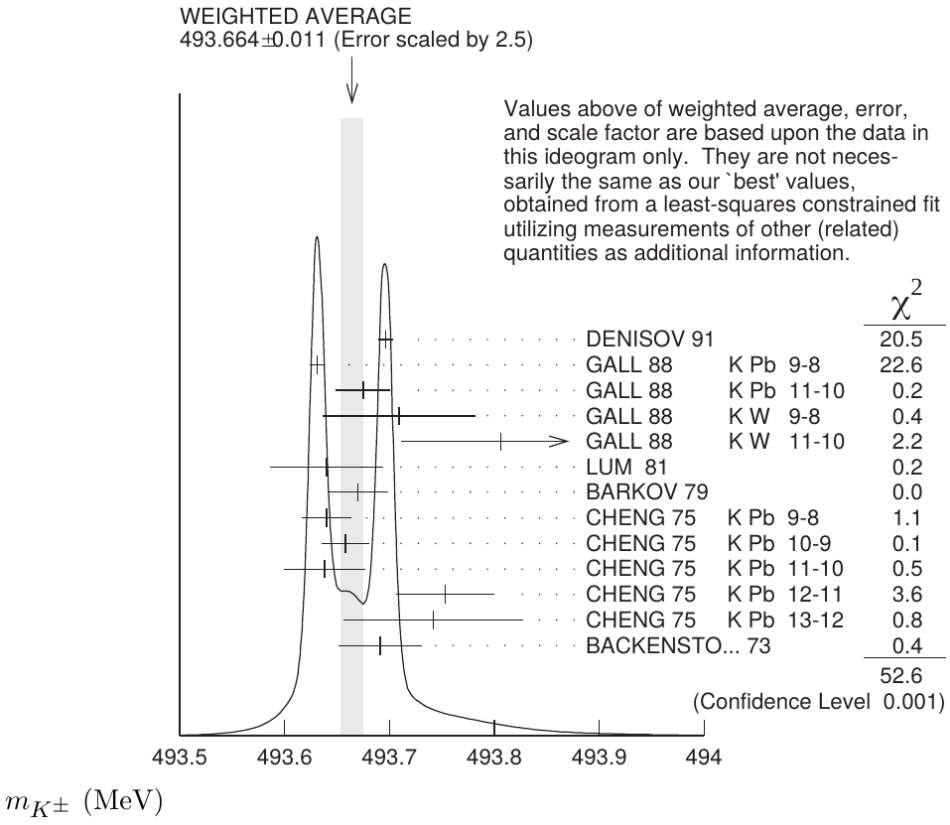

Figure 91: The contributions to the kaon mass measurement and their combination, from [13]. Note that the error on the kaon mass in this figure is taking into account only the values listed, while the world average uses more information and is slightly larger (13 keV).

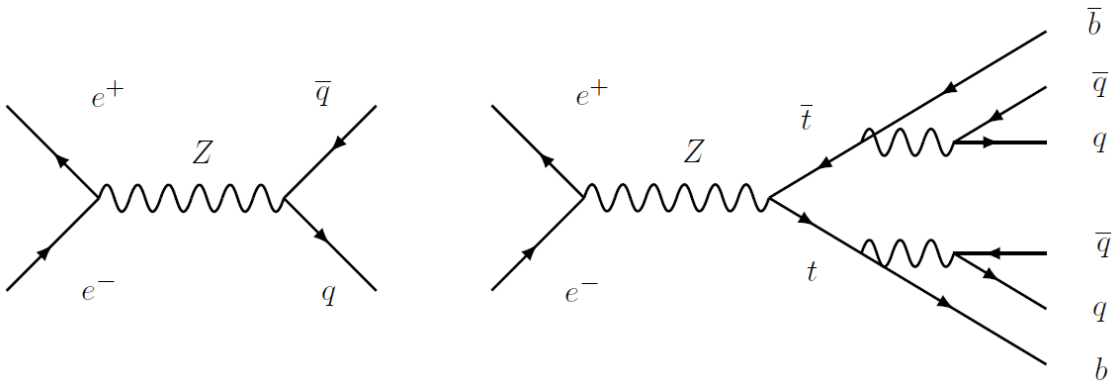

Figure 92: Tree-level Feynman diagrams of final states consistent with (left) the 2-fermion-Z-hadronic group and (right) the ttbar-production group, in this case with a fully hadronic decay.

electron beam polarisation	−80%	+80%	−80%	+80%
positron beam polarisation	+30%	−30%	−30%	+30%
$4\,\text{ab}^{-1}$ scenario	$1600\,\text{fb}^{-1}$	$1600\,\text{fb}^{-1}$	$400\,\text{fb}^{-1}$	$400\,\text{fb}^{-1}$
$400\,\text{fb}^{-1}$ scenario	$200\,\text{fb}^{-1}$	$200\,\text{fb}^{-1}$	-	-

Table 12: Split of integrated luminosities among beam polarisation combinations in the two scenarios used.

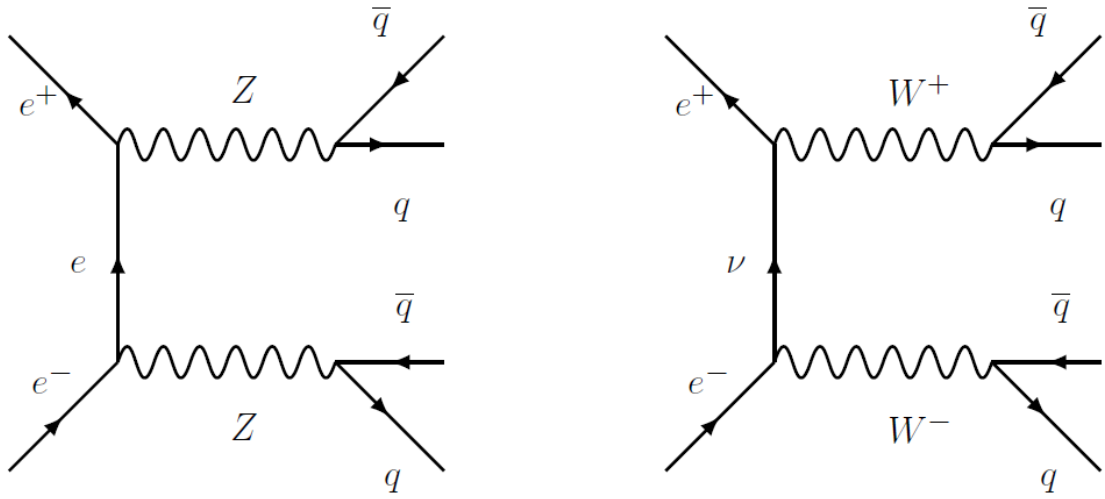

Figure 93: Tree-level Feynman diagrams of final states consistent with (left) the 4-fermion-ZZ-hadronic group and (right) the 4-fermion-WW-hadronic group.

To collect all available kaons, all PFOs of charged particles that contain a TOF measurement are considered. As layed out in chapter 7.5, only PFOs that reach the barrel ECal contain a TOF measurement in the current software version. By utilising all available PFOs, however, a large number of mis-measurements and particles not originating from the IP, like back scatter from the calorimeters, are picked up as background. Cuts on the track parameters of the reconstructed tracks can reduce this background. Figure 94 shows the TOF-bands in β-vs.-p plots, with relativistic velocity β and momentum p determined as described in chapter 7.5. The plots contain all PFOs with a TOF measurement for the $4\,\text{fb}^{-1}$ scenario and a 50 ps timing resolution and compare the situation before and after the following cuts. The bands of protons (lower), kaons (middle) and pions (upper) are clearly visible and separable, until they merge at a few GeV. The kaon band gradually merges with the pions around 3 GeV, making them less distinguishable by TOF with increasing momentum. Thus, only particles up to a certain momentum are used to reconstruct the kaon mass, with the upper momentum cut being a variable parameter in the further analysis, varying between 1.6 and 3.6 GeV. The lower momentum cut is fixed at 1 GeV, the minimum momentum necessary to reach the ECal.
Back scatter and tracks from in-flight decays often have a lot fewer than the maximum of 220 TPC hits associated with the track. The efficiency of track hit reconstruction is

large and the occupancy of the TPC is low, so most tracks that reach the barrel ECal have very close to 220 associated hits. Therefore a cut requiring the PFO to have more than 200 hits is applied. In addition, the closest approach of the reconstructed track to the IP in r/ϕ, d_0, and in z, z_0, are large, since the particles do not originate from there. The luminous region of the IP following Table 1 is smaller than 1 µm in the transverse directions and smaller than 1 mm in the longitudinal direction. Therefore, cuts of $|d_0| < 10\,\text{mm}$ and $|z_0| < 20\,\text{mm}$ are applied. Figure 94 shows that the combined cuts reduce the background by about 2 orders of magnitude, while leaving the bands of specific species intact, making them clearer identifiable.
By measuring momentum p and β of a PFO, its mass can be reconstructed via

$$m = \sqrt{p^2/\beta^2 - p^2}. \tag{29}$$

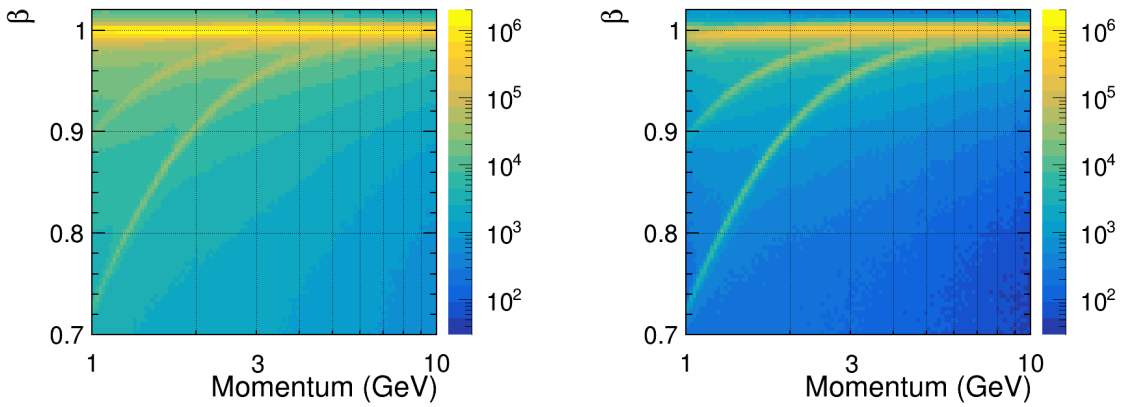

(a) (b)

Figure 94: Reconstructed velocity β vs. momentum p of charged particles calculated from TOF measurements with a time resolution of 50 ps. (a) For all charged particles, (b) after quality cuts described in the text.
All reconstructed particles from the samples listed in Table 11 that reached the barrel ECal were used here, for the $4\,\text{fb}^{-1}$ scenario. After quality cuts the bands from pions, kaons and protons can be identified more clearly.

The distribution of the reconstructed mass is displayed in Figure 95, for the $400\,\text{fb}^{-1}$ scenario and an upper momentum cut of 3 GeV. The previous pion-, kaon- and proton-bands are transformed into peaks at the respective rest masses. The three peaks reflect the masses of the pion around 140 MeV, the kaon at about 500 MeV and the proton at about 950 MeV, respectively. By adding more PID information using the track dE/dx value of the PFOs, as explained in chapter 7.2, the kaon peak can be further emphasised. A cut is added, requiring a PFO to have a track dE/dx value which is consistent with the expected mean value for single kaons within 2.5 standard deviations, see Figure 96. The mean value and width of the dE/dx-band are extracted from the single-particle samples (chapter 4.5). For particles with momenta near 1 GeV this cut does not suppress many

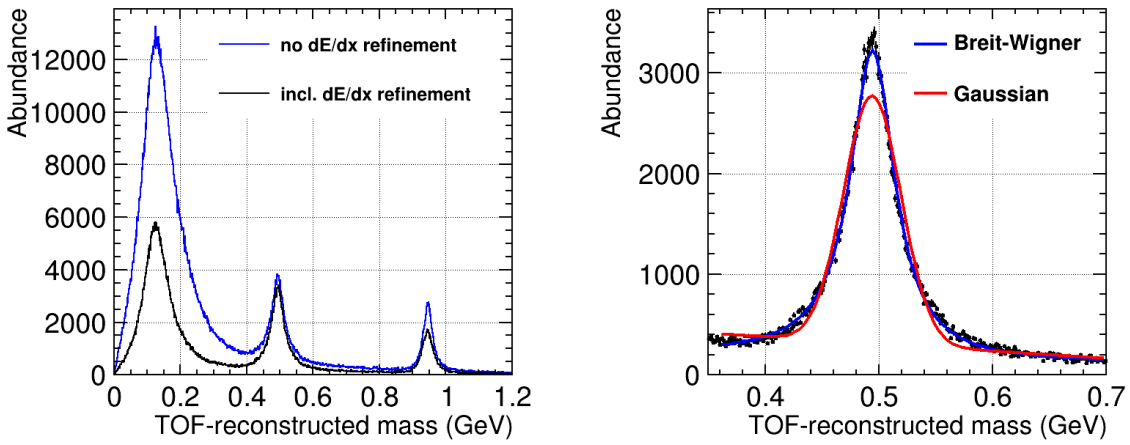

(a) (b)

Figure 95: (a) Reconstructed mass of charged particles calculated from momentum and TOF measurements. (b) Reconstructed mass of kaons with mass fit. In both plots a time resolution of 50 ps and an upper momentum cut of 3 GeV were used.

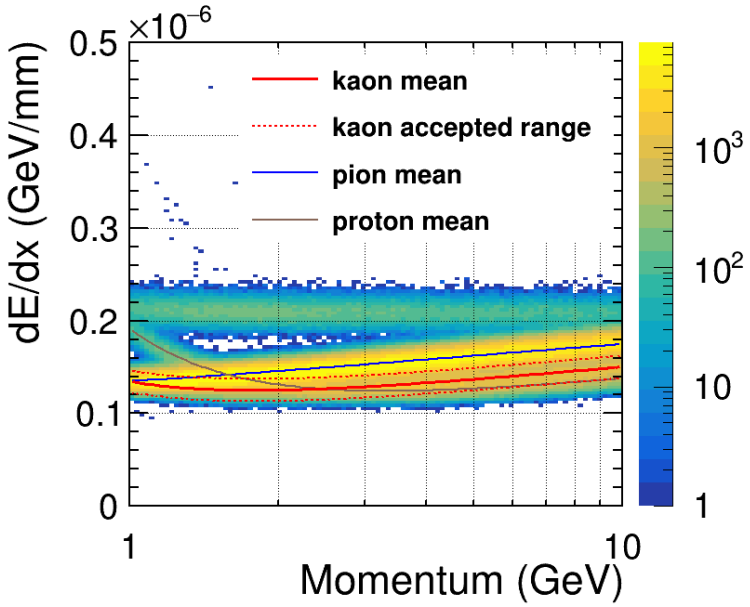

Figure 96: Bethe-Bloch bands of all PFOs in the sample used. The mean values derived from the single particle samples (see chapter 7.2) are added as lines, with the dashed lines indicating the range in which PFOs are accepted as consistent with the kaon hypothesis in this case.

pions, since pions and kaons have the same dE/dx values there, but with increasing momentum the separation between the two band grows and allows for an effective reduction of the number of pions in the sample. Figure 95, left, shows the distribution of the reconstructed mass, with and without refinement by dE/dx. The number of kaons is about the same in both cases, but the number of pions is about halved, contributing less to the continuous background. The number of protons is about halved this way as well. The crossing point between the Bethe-Bloch curves of kaons and protons is at about 2.5 GeV, however, and the reduction in protons is largest at 1 GeV. In any case, the overall number of protons is small and the separation via TOF is sufficient in the used momentum range to keep the proton background at a negligible level.
This finally leads to a fit of the kaon peak, seen in Figure 95, right. The chosen fit range is 0.36 GeV - 0.7 GeV, the limits being roughly in the middle between the respective peaks. As fit function, a Breit-Wigner with an underlying exponential was chosen, since it describes the peak shape reasonably well with a remaining χ^2/ndf of about 60. The fit function and the result are summarised in Equation 30 and Table 13. An exact model function is not known and would have to be derived from first principle, including effects of the applied cuts. Therefore a simple function with only 5 fitting parameters was selected. For comparison, a Gaussian with an exponential was used for comparison, but led to a factor 11 larger χ^2. However, both functions returned roughly the same values both for the position of the peak as well of for the error on this position. This error on the position is the achievable statistical uncertainty of the mass measurement, which is the target of this study. The position mean value itself is offset from the MC value used as input for the simulation, which points to a systematic bias.

$$BW(m) = \frac{a}{(m^2 - m_0^2)^2 + b^2 \cdot m_0^2} + c \cdot e^{d \cdot m} \tag{30}$$

m_0	=	0.494708 ± 0.000032
a	=	8.775 ± 0.031
b	=	0.04597 ± 0.00010
c	=	1456 ± 25
d	=	-0.9926 ± 0.028
χ^2/ndf	=	$2311/355$

Table 13: Fit result of the Breit-Wigner fit to the kaon mass (in units of GeV). For the 2x200 scenario, a timing resolution of 50 ps and and upper momentum cut of 3 GeV, the uncertainty on the fitted mass value is 32 keV.

The statistical error derived this way is displayed in Figure 97 and Figure 98 for various measurement conditions. They show the achieved statistical precision on the mass measurement depending on the upper momentum cut of the accepted PFO momentum window. For a larger upper momentum cut, more kaons are included and the statistical error shrinks. However, at a certain point the contamination by pions increases and with

it the statistical error. The result is plotted in different colours for the 3 different TOF resolutions under study. A clear difference is visible between 50 ps and 10 ps, but only a small one between 10 ps and the MC truth reference (0 ps). The solid line reflects the full analysis, while the dashed lines offers a comparison to a situation without dE/dx refinement. This refinement has a small but noticeable effect on the achievable resolution in this study. The effect increases with a larger upper momentum cut, consistent with the increasing pion-kaon separation power at larger momenta. Figure 97 shows the result for the 400 fb^{-1} scenario. With an ECal hit time resolution of 50 ps, a statistical uncertainty on the charged kaon mass of about 30 keV can be achieved, which would be a relevant contribution to the mass average. Figure 98 shows the result for the full ILC running time at 500 GeV and a total luminosity of 4 ab^{-1}. The achievable statistical kaon mass uncertainty is about 10 keV. Compared to the current uncertainty of 13 keV and the difference between the two leading measurements of 60 keV, this would be a decisive measurement.

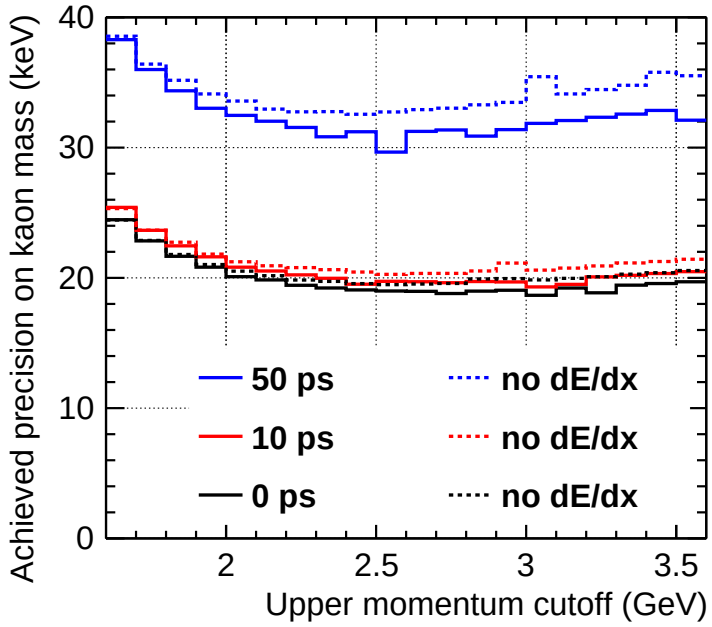

Figure 97: Statistical precision of the reconstructed kaon mass fit, for the 400 fb^{-1} scenario.

The systematic error on the kaon mass, however, which is interpreted to be the difference between the reconstructed mass and the one input into the simulation, is much larger. Figure 99 shows this comparison for 0 ps and 50 ps time resolution as well as split among the physics processes used, for an upper momentum cut of 2.5 GeV. The systematic error not only is large, O(MeV), but its sign also differs for the two time resolutions. However, the bias seems to be more consistent between the different physics

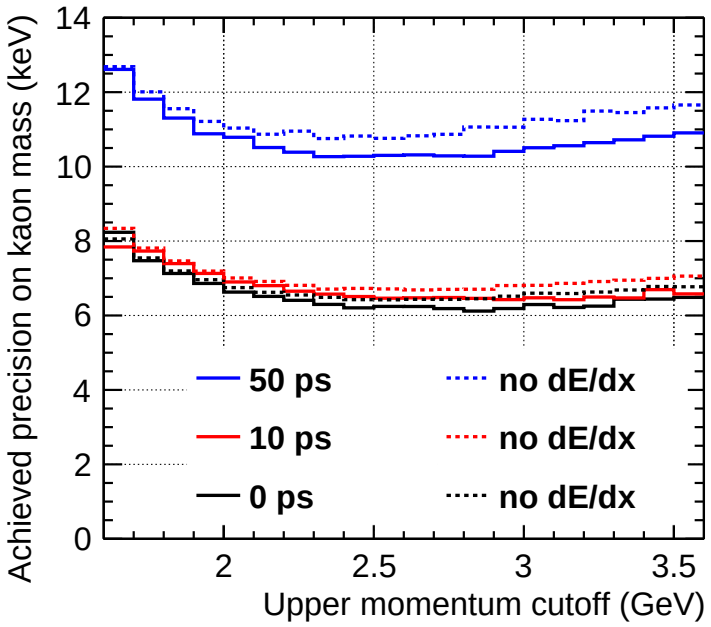

Figure 98: Statistical precision of the reconstructed kaon mass fit, for the $4\,\mathrm{ab}^{-1}$ scenario.

processes, which provide different momentum and angular spectra of kaons. This large systematic error and its dependencies are subject of ongoing and future investigation in order to reduce them and/or correct for them. To allow for a kaon mass measurement with the aforementioned precision, the systematic error needs to be reduced below the statistical one. As explained in chapter 7.5, prime targets of these studies are the usage of the reconstructed momentum at the IP despite the momentum change until impact in the ECal, and the correction of the ECal hit times to the impact point using linear propagation at the speed of light. In the end, with a data-driven approach the pion and protons peaks can be used to calibrate the mass measurement, since their masses are known to much higher precision (pions: 0.17 keV, protons: 5.8 eV, from [13]) compared to charged kaons.

With an analysis taking into account the full data set of all processes, including 250 GeV and including kaons going into endcap, the overall statistics would be even a factor 1.5 to 2 higher, which would typically lead to a further reduction in the achievable kaon mass uncertainty of about $20-40\%$. This would result in a statistical error on the kaon mass only half its current value.
In any case and even with a moderate timing resolution, ILD will be able to solve the kaon mass question with a completely independent method compared to the long-standing contradicting ones from three decades ago.

In addition to an analysis based on a TOF measurement, with dE/dx only used for a moderate refinement, the same approach could also be used based on the Bethe-Bloch

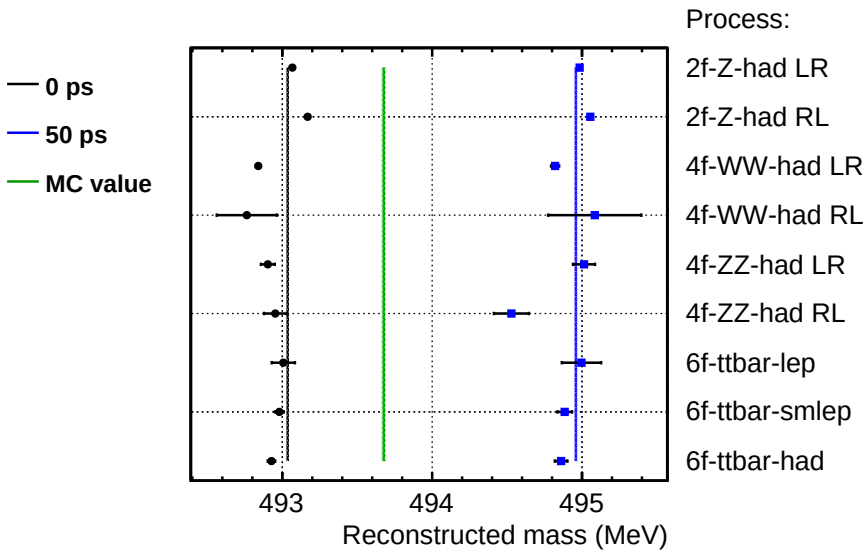

Figure 99: Simulation input and reconstructed kaon mass for 10 ps and 50 ps, divided by physics processes. The systematic bias between reconstructed and MC kaon mass is much larger than the statistical uncertainties indicated by error bars. In addition, it has a different sign for different timing resolutions.

curves directly. Figure 100 shows the corresponding curves with the kaons marked in red, where they are separable from the other species with very high purity. The position of the kaon band is linked to its mass and could therefore be used to fit the kaon mass. Compared to the method using TOF, the disadvantage would be the smaller overall number of kaons, since only a very limited momentum range would be available, all below 800 MeV, where the abundance of all PFOs is smaller than at a few GeV. However, at transverse momenta below 1 GeV the tracks curl in the TPC and could be measured much more precisely. Unfortunately, the current version of the software only takes into account the first half circle of a track, and the energy loss along a curling track is not negligible, so a dedicated treatment needs to be implemented. Conceptually, each half circle could give an independent measurement of momentum and specific energy loss and thus result in a much larger statical abundance at those low momenta. Such a dE/dx-based approach would be independent from the TOF-based one except for the aforementioned dE/dx refinement), and therefore provide additional information on the kaon mass.

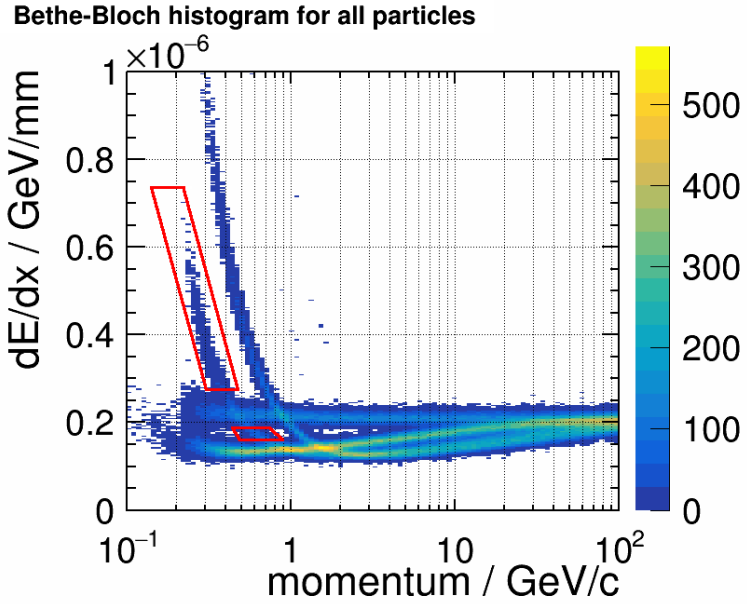

Figure 100: Bethe-Bloch curves of single particles, for ILD with a dE/dx resolution of 4.5 %. At low momenta, the kaon band is marked in red where the kaons are separable with a very high purity and the band could be used to reconstruct the kaon mass.

8.2. W-Decay Separation

One application of advanced hadron PID is the flavour tag of jets. A particular benchmark is the hadronic W decay, which takes place in about two thirds of all W decays. Following the CKM matrix, the W mostly decays via the diagonal elements of the matrix, i.e. either into a d quark and a $\bar{u}$ quark (or $\bar{d}$ and u), or into an s quark and a $\bar{c}$ quark (or $\bar{s}$ and c). The third diagonal decay, into b and $\bar{t}$ (or $\bar{b}$ and t), is impossible, since the t quark is heavier than the W boson. The off-diagonal elements of the CKM matrix are significantly smaller than one, and the corresponding decays are significantly suppressed. Decays between the first two generations, i.e. d and $\bar{c}$ (or $\bar{d}$ and c) or s and $\bar{u}$ (or $\bar{s}$ and u) are suppressed by more than one order of magnitude, and decays including a b or $\bar{b}$ by more than two orders of magnitude. In the following, explicit notation of anti-particles will be omitted in favour of briefness.

The quarks from the hadronic decay are still colour-bound and undergo hadronisation, generating on average about 20 particles, of which about three quarters are pions. During this hadronisation, the quark's initial momentum is correlated with its flavour. This way, jets from s quarks not only generate on average one more strange particle (mostly kaons) compared to d and u quarks, but this strange particle also has a large fraction of the initial quark momentum, compared to additional kaons generated from $s\bar{s}$-production in the hadronisation process. The same is true for c quarks from the W decay, which after briefly forming charmed hadrons quickly decay into s quarks, which is the dominant decay following the CKM matrix, providing more kaons. This means that W decays into an s and a c quark generate on average more strange particles and ones with higher momenta, compared to decays into a d and a u quark. Hadron PID can be used to identify these strange particles, in particular kaons, in order to separate d/u decays from s/c decays.

In fact, this method has been used in the past at LEP [116] in order to measure the hadronic W branching fractions and to determine the elements of the CKM matrix, in particular the least known element V_{cs} [117]. ILC and ILD will be able to improve on this measurement given the luminosity and the existing flavour tag capability, mostly via vertexing. Vertexing works well for tagging charm and bottom hadrons which live sufficiently long to create decay vertices well separated from the interaction point, but still indise the innermost Silicon pixel layers of the tracking system. Strange hadrons live long enough to decay either well inside the tracking system or even reach the calorimeter. Therefore, the ability to identify kaons and other strange hadrons via dE/dx, TOF or V0-finding adds valuable information to this measurement, which will be laid out in this chapter in order to show the improvement with respect to the conventional flavour tag.

The events used this analysis are taken from the most recent large-scale MC production at a center-of-mass energy of 500 GeV (see chapter 4.5). The process selected is Higgs production with the Higgs decaying into WW^*, which was selected to allow in a later advanced analysis for a clean environment, i.e. background suppression via the additional invariant mass constraint from the Higgs boson. This background can come in particular from Z decays, and the Higgs branching fraction to ZZ^* is significantly suppressed compared to WW^*. From this sample and based on MC-truth informa-

tion, only semi-leptonic W decays were selected and inter-generational decays as well as events containing a tauon were excluded. This selection was done in order to reduce possible confusion in multi-jet environments. The selection of semi-leptonic decays and the veto to tauons can be implemented in an advanced analysis, while the veto to inter-generational decays can not, and would instead need to be taken into account in a measurement of the CKM-matrix elements. This study, however, provides a differential assessment of the flavour-tag performance comparing the conventional with a new method, for both of which the selection based on true information is of similar effect.

8.2.1. Flavour Tag Performance

To investigate the improvement by tracking PID-based flavour tagging, first the performance of the conventional vertex-based method needs to be clarified. This method is integrated in the usual flavour tag by LCFIPlus [118]. It uses a number of different observables and applies a tag to each identified jet. The tag can be b, c or other (b/c/o). The used observables are derived from vertex properties as well as generic jet quantities like jet energy and number of associated tracks. They contain only basic electron and muon ID, but, crucially, no hadron ID. The derived observables are used in a BDT to acquire a good tagging performance. Events containing 6 identical quarks moving in isotropically random directions are used as training samples. Training was done centrally and the result made available together with the MC event production.

This flavour tag returns a tag value between 0 and 1 for each tagging hypothesis, τ_b, τ_c, τ_o, to each jet. To apply these hypotheses to the problem at hand, they needed to be combined to give a common tag for the 2-jet decay system. Three options to combine the tags of jet 1 and jet 2 were investigated: the sum of the tag valuess τ_{sum}, using the larger value τ_{high}, and a combination τ_{comb} according to Equation 31 from [7] (p. 114). Among these three, τ_{high} delivered the best final performance by a small margin and will be used from here on.

$$\tau_{comb} = \frac{\tau_1 \cdot \tau_2}{\tau_1 \cdot \tau_2 + (1-\tau_1)\cdot(1-\tau_2)} \tag{31}$$

For this work, the b tag is (nearly) irrelevant, since no b quarks are generated in the given statistics of the W decays, and the work aims to separate s and c quarks from d and u quarks. The c tag should identify c quarks and the o tag u, d and s quarks. A combination of the two should allow to differentiate between a d/u system correlated with two o tags and an s/c system correlated with one o and one c tag.

The most powerful feature in LCFIPlus flavour tag is the secondary vertex position and its offset from the primary vertex. Hadrons containing a b quark typically have a clearly offset secondary vertex due to their relatively large life times. On the other hand, charmed hadrons decay much quicker and it is much less likely to reconstruct a secondary vertex. The lighter quarks d, u and s form hadrons that are mostly stable in the detector, like pions and kaons, and therefore do not have a secondary vertex at

all. This leads to a non-negligible number of cases where c-quarks are tagged as a light quark with a large tag value, and vice versa.

The tagging performance of the LCFIPlus flavour tag is displayed in Figure 101 and Figure 102. The performance was evaluated using the W-decay events under investigation and comparing the resulting flavour tag to the real flavour according to the MC information of the inital quarks. The figures show the expected correlation clearly: the c tag peaks for c jets at a large tag value, while for other jets it peaks at low values and falls quickly. Applied to the combined jet systems, the c-tag curve for the d/u system follows the same shape as for the individual other jets, but the peak is shifted closer towards 0 due to the mathematical structure of the combination in Equation 31. The c-tag curve for the s/c system peaks at low values and then flattens off for higher values, crossing the curve for the d/u system at a value of around 0.25. The situation is similar for the o tag: for d, u and s jets it peaks at large likelihood values, while for c jets there is a large peak at a tag value close to 0, and a very shallow one at high values. The o-tag curve for the d/u system also peaks at high values, but again close to the maximum value of 1 as the individual contributions for the jets, while the curve for the s/c system peaks around 0 and, a bit smaller, around 1. In addition, each curve has a peak at 0, which is due to occasional mis-tagging and tagging algorithm cancellation, leading to a fixed result of exactely 0. Therefore, the 0-bin is of limited use for the further analysis.

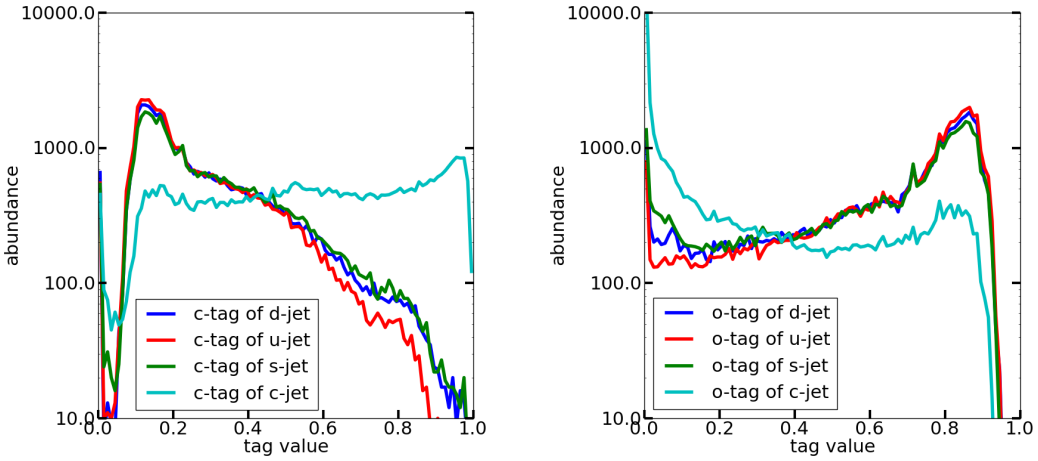

Figure 101: c and o tags for d, u, s and c jets in the investigated event sample.

The tagging performance of the LCFIPlus flavour tag can also be assessed as outcome of a BDT aiming to separate between the two decay hypotheses. Figure 103 shows the resulting ROC curves, for the different ways to combine the two jet-tag into a single system-tag value. The conventional flavour tag reaches an area-under-the-curve (AUC) value of around 0.8 as measure of its performance. This value is also used to select the most effective combination, τ_{high}, with an AUC of 0.808.

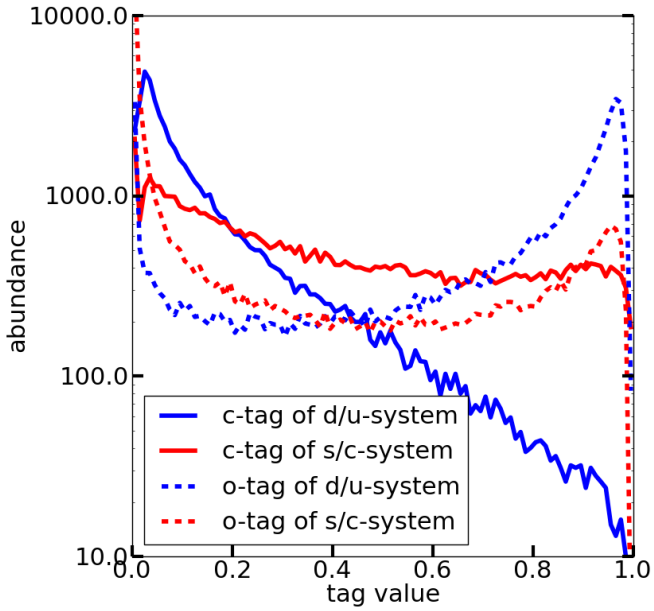

Figure 102: Combined c and o tags for d/u and s/c systems in the investigated event sample.

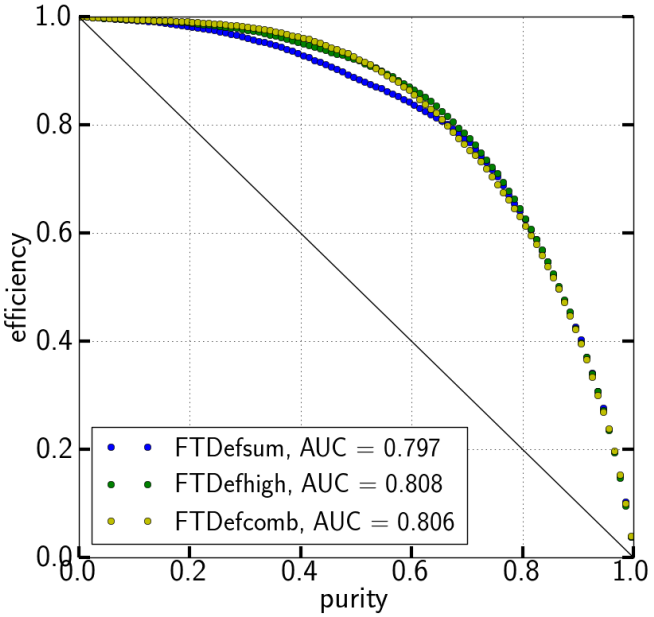

Figure 103: ROC curves of W-decay tagging by the conventional tagging algorithm, with different ways to combine the individual jet tag values.

8.2.2. Application of Tracking PID to Flavour Tag

Events from the latest MC production with a center-of-mass energy of 500 GeV were sorted for Higgs decays to WW*. The question here is: How well can we differentiate between a W decaying into a d and a u quark from a W decaying into an s and a c quark? The process is depicted in Figure 104.

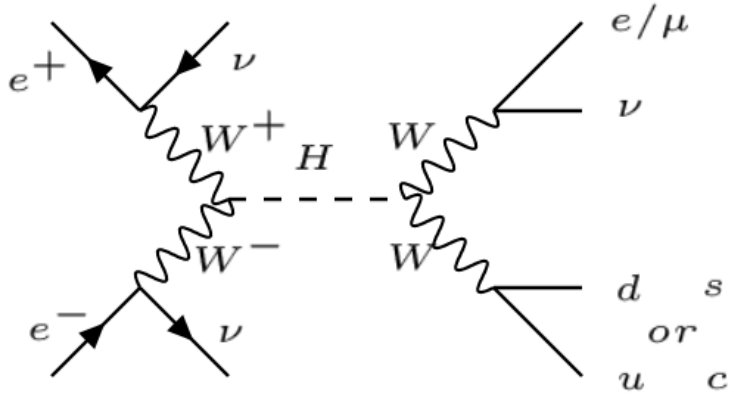

Figure 104: Leading Feynman diagram of the investigated process: At 500 GeV, W fusion dominates the Higgs production at the ILC. The chosen channel here is H to WW*, semileptonic.

To separate the different decay channels of the W, various quantities were extracted from the MCParticles as well as from the reconstructed particles and jets. These reconstructed strange particles, $K^{+/-}$, K^0_S and Λ^0, as well as $\pi^{+/-}$ are identified with a certain efficiency and purity, and will be called PiKaLas from here. As base approach, the charged pions and kaons were identified via dE/dx-PID, fixing their reconstructed flavour. However, the relative distance to the corresponding band in the Beth-Bloch-curve can be used as a (simultaneous) measure of the pion-ness and kaons-ness of each particle. Since TOF is not (yet) part of the PID algorithm, it was used as a separate PID measure. A particle is identified with TOF by the reference curve of β vs. p it is closest to. The reconstructed particle with the largest momentum in each jet is called the leading particle, meaning each W-decay system has two leading particles, which can be identified.
Observables in a jet can then be the reconstructed number of particles of a species, the relation between two reconstructed numbers of particles of a species, and the fraction of the total jet momentum carried by particle of a species (momentum fraction). Instead of total number of pions/kaons, also a pion/kaon-ness amount can be used in each case. In addition, momentum cuts can be applied to suppress low-efficiency areas of an observable.

The 20 observables of the reconstructed jets are:

total number of each PiKaLa (4),
number fraction of charged pions / (pions + kaons) (1),
fraction of pion-ness / (pion-ness + kaon-ness) (1),
number fraction of pions / (pions + kaons) based on TOF (1),
PiKaLa momentum fraction (4),
charged pion/kaon momentum fraction for particles above 3 GeV (2),
charged pion/kaon momentum fraction weighted with pion-/kaon-ness (2),
charged pion/kaon momentum fraction based on TOF (2),
number of leading PiKaLas (3, K^0_S and Λ^0 combined).

As examples, distributions of some of these observables and the difference between d/u system (red) and from an s/c system (blue) are shown in the following figures. For reference, Figure 105 shows the true number of charged pions and kaons in the true 2-jet systems. In comparison, Figure 106 shows the abundance of the reconstructed number of pions and kaons per 2-jet system. As expected, d/u systems contain on average more pions and fewer kaons than s/c systems. While the true number of kaons in a s/c system is on average 1 larger than in a d/u system, this difference is only about half a kaon after reconstruction. In Figure 107, the fraction of jet momentum carried by kaons is seen. On the left, every reconstructed jet particle contributed to the fraction, but was weighted accoring to the distance of the measured dE/dx from the expected Bethe-Bloch curve. On the right, only particles contributed which were identified as kaons by TOF. Figure 108 shows the same momentum fraction, but for particles that were identified as K^0_S and Λ^0, respectively, by the V0Finder algorithm. All observables are shown in chapter A.2.

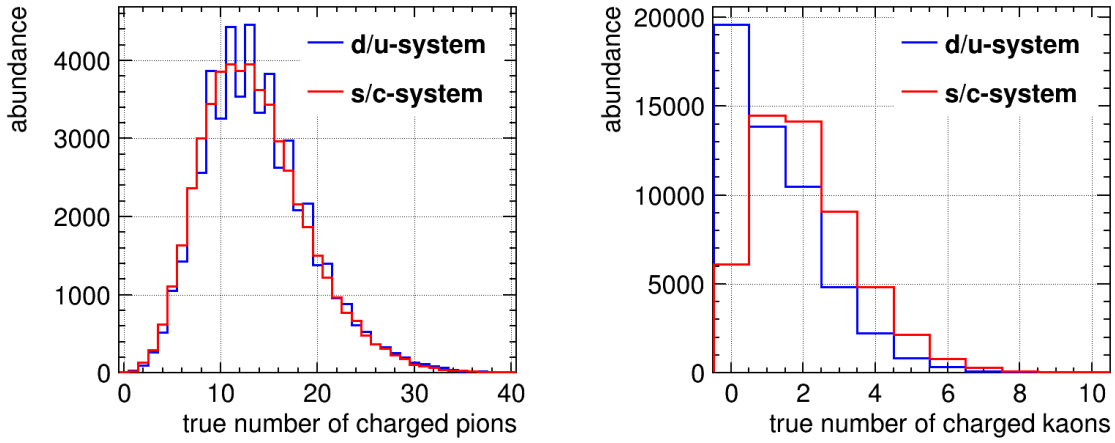

Figure 105: Abundance of MC pions and kaons, respectively, in the true jet systems.

These observables were then used as input to a boosted decision tree (BDT), which is implemented within the multivariate analysis (MVA) package in the ROOT frame-

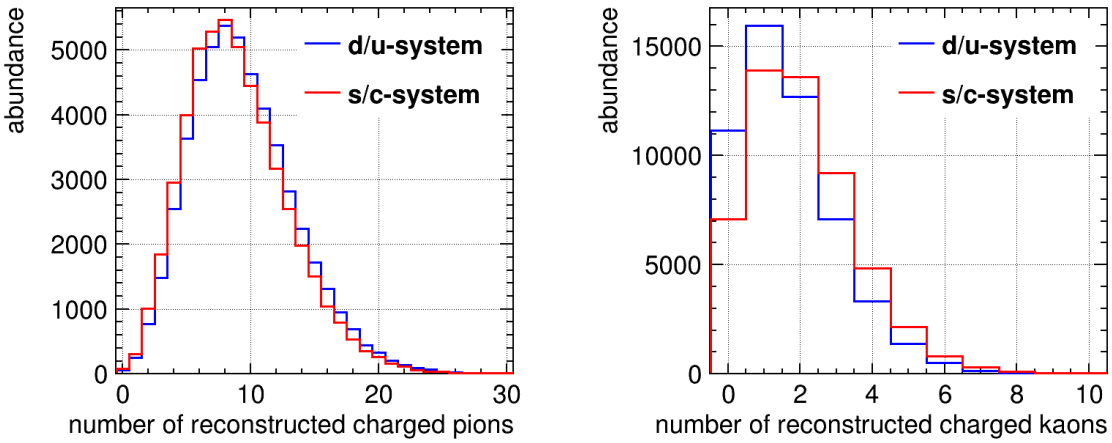

Figure 106: Abundance of pions and kaons, respectively, in the jet systems investigated, according to default dE/dx PID.

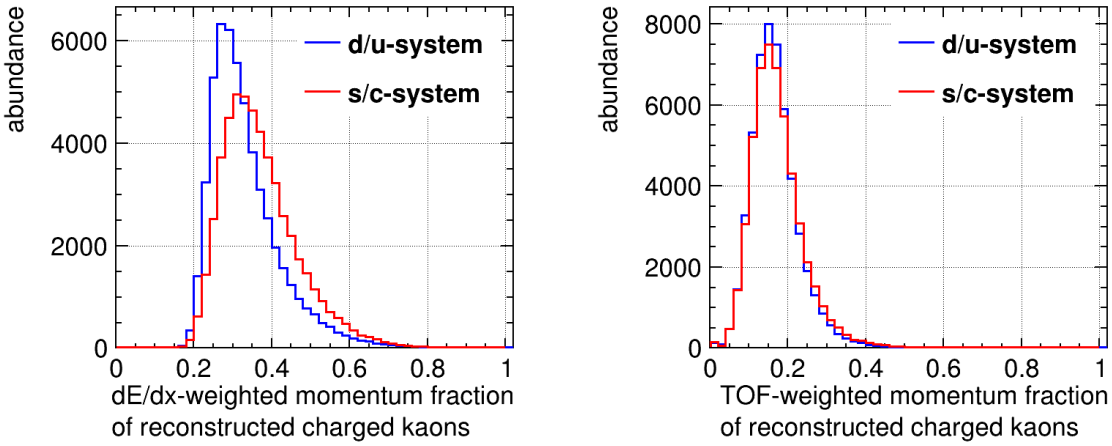

Figure 107: Abundance of jet momentum fractions of reconstructed kaons in the jet systems investigated. Left for kaon ID according to Bethe-Bloch-distance weighting, right for PID via TOF.

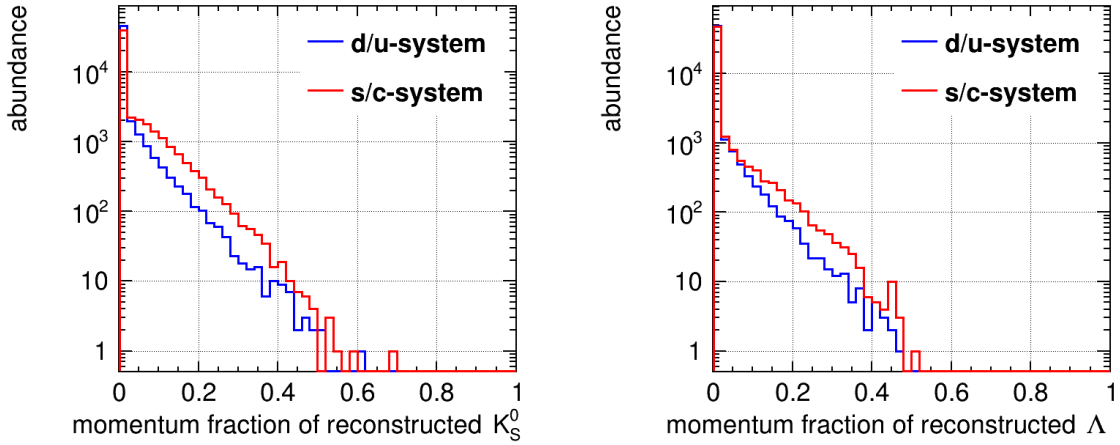

Figure 108: Abundance of jet momentum fractions of reconstructed K^0_S and Λ^0, respectively, in the jet systems investigated, with reconstruction ID according to V^0-finding.

work. In the BDT, the s/c and d/u systems were classified as signal and background, respectively. By default, half the available events were used for training 1000 trees, and half for testing the trained trees. The remaining cost function of the trained network for the training samples is lower than for the test sample, which is due to overtraining. In general, throughout this analysis the BDT parameters were chosen to keep the difference of the cost functions between training and testing sample at most at the few-percent level, which is an acceptable level, indicating very limited overtraining. The result of the BDT can be expressed as a receiver-operator-characteristic (ROC) curve. This displays resulting working points of selection purity p vs. efficiency ϵ for a given range of BDT score values. For a random selection between signal and background, i.e. no separation power, this curve is a line following $p = 1-\epsilon$ with an integral value (area under the curve, AUC) of 0.5. The larger the separation power, the further the ROC curve will stretch towards $p = \epsilon = 1$, increasing the AUC up to a value ≤ 1. Figure 109 shows an example for a ROC curve with a reduced selection of input variables. The blue points result from a BDT trained on reconstructed PFOs and only dE/dx-based PID, i.e. without V0- and TOF-based ones, and has an AUC of about 63 %, while the same BDT trained on the corresponding true particle observables from MC (green curve) covers about 75 %. In general, an AUC of around 70 % reflects only a moderate stand-alone separation power, but still reflects some valuable information which can be utilised. There is still a significant gap between reconstructed and MC-based performance which will be investigated next.

One possibility to improve the performance compared to dE/dx-only observables is to add V0 information. Since this identifies neutral particles as opposed to dE/dx, it adds nearly entirely complementary information. The result is shown in Figure 110: the AUC increases by about 3 percentage points.

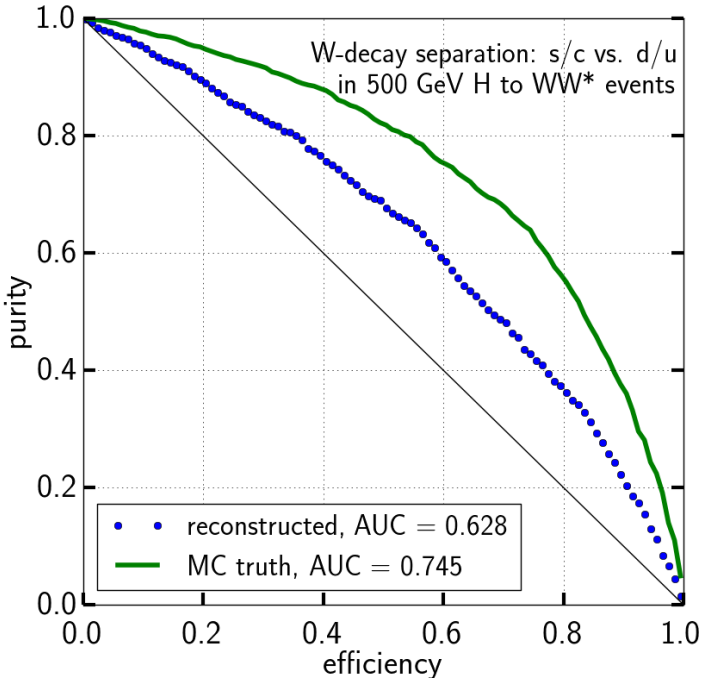

Figure 109: ROC curves for dE/dx-based PID, comparing BDTs trained on MC-based and reconstructed observables.

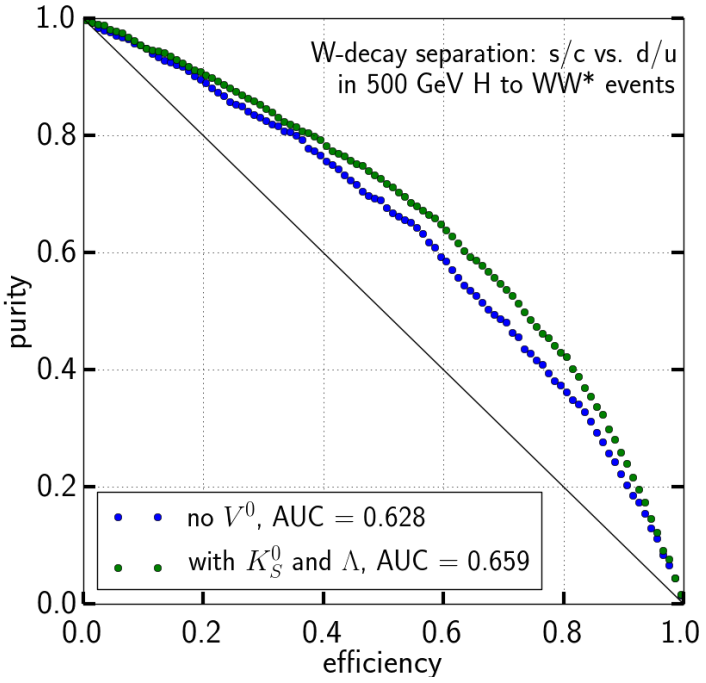

Figure 110: ROC curves for dE/dx-PID, comparing the effect of adding V0 information.

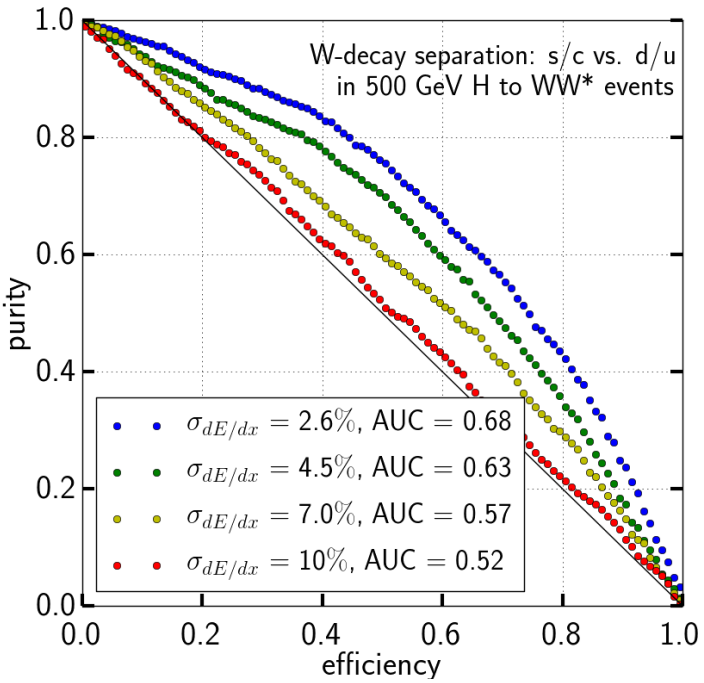

Figure 111: ROC curves for dE/dx-PID, comparing the effect of different underlying dE/dx resolutions. Here, V0 information is omitted.

To investigate the impact of the dE/dx resolution, the event files were re-reconstructed with different corresponding values. This was possible in reconstruction due to the interplay of the uncertainty of the energy loss of each TPC hit, calculated by Geant4 in the simulation, and the combination of the individual TPC hit energy values to a track dE/dx estimate by the `Compute_dEdxProcessor` in the reconstruction. The hit energy loss created by Geant4 results in an effective dE/dx resolution for tracks of about 2.6 %, which is much smaller than current extrapolations of TPC readout systems developed for ILD. These aim at a resolution of 5 % for ILD and have shown a performance of up to 4.5 % when extrapolated to ILD. This discrepancy between simulation and realistic values is compensated by the `Compute_dEdxProcessor`. After calculating a track dE/dx estimate, it applies a dedicated Gaussian smearing factor to that estimate, which is tuned to result in an overall dE/dx resolution of 4.5 %. This smearing factor can be altered to change the effective dE/dx resolution, as long as it is larger than the intrinsic 2.6 % from Geant4. In this study, the events used were re-reconstructed with no smearing and default smearing, as well as two larger smearing values, resulting in effective dE/dx resolution values of 2.6 %, 4.5 %, 7.0 % and 10.0 %. Figure 111 shows the impact of varying the underlying dE/dx resolution on the BDT AUC: at $\sigma_{dE/dx} = 10\,\%$ there is hardly any separation power left, while the AUC at $\sigma_{dE/dx} = 2.6\,\%$ is much closer the MC value.

8.2.3. Analysis Results

The performance of the flavour tag based on PID needs to be compared to the vertex-based one, in particular how much of the information gained is truely additional and can be used to improve the conventional tag.

Both LCFIPlus flavour tag and PID observables were used as inputs in a combined BDT. The corresponding correlation matrices of the 3 flavour tag observables (b-tag, c-tag and other-tag) and 6 selected PID observables, which are among the most impactful ones are shown in Figure 112. The 6 selected observables are the number and momentum fraction of both pions and kaons, and the number of K_S^0 and Λ^0. For both s/c-decay systems and d/u-decay systems there are some correlations within the selected 6 PID observables: the number and momentum fraction of kaons are correlated, as are number and momentum fraction of pions. This is not suprising, however, and it is notable that the correlation is substantial but not close to 1. The momentum fraction of kaons and the one of pions are somewhat anticorrelated, which is also to be expected. Other correlations between the 6 selected PID observables are negligible. The correlation within the vertex-based tag values differs between s/c and d/u system. For the s/c decays the tags are slightly anti-correlated, in particular o tag and b tag. This makes sense, since the b tag tends to recognise offset vertices, while the o tag avoids them. This is also true for d/u decays, but here c and b tag are strongly correlated, indicating that they react similarly to mis-identified offset vertices. Finally, the correlation between conventional and PID-based observables is negligible, indicating that PID indeed provides largely additional information for the W-decay separation.

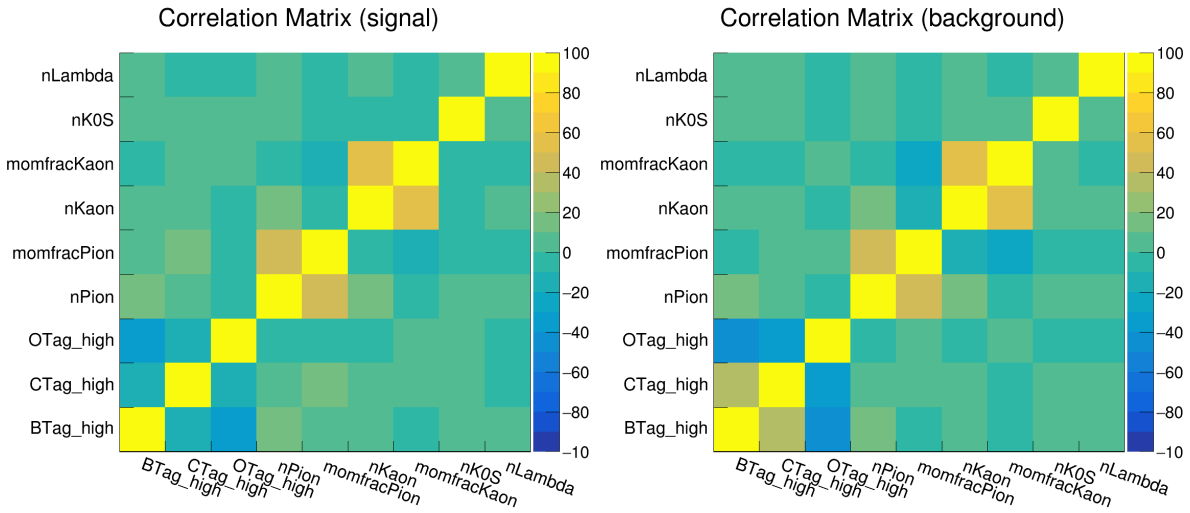

Figure 112: Correlation matrices comparing the impact of conventional tags and PID as input on W-decay separation. Left: for s/c-system; right: for d/u-system.

In Figure 113 the resulting ROC curves of the combination of conventional and PID observables are shown. The AUC of the combination is about 4 %-points larger than the conventional tag-based one alone. The difference between adding only the 6 selected impactful PID observables and all available 20 to the conventional tag information is

only about 0.7 %-points, reflecting the large degree of redundancy between the 20 PID observables. The overall increase by 4 %-points may seem small, but in particular at a large requested purity, the efficiency increases significantly. For example, for a fixed purity of 0.95, the efficiency increases by 17 % from conventional observables to adding PID-based ones, allowing more events in any physics analysis to be selected and thus reducing statistical uncertainties.

Figure 113 shows the result for a dE/dx-resolution of 4.5 %, while Figure 114 shows similar curves for the simulated 2.6 % and 7 % resolutions. Since the efficacy of the PID-based separation depends largely on the dE/dx, also the improvement of the AUC compared to the conventional method is affected. At $\sigma_{dE/dx} = 2.6\,\%$ the increase is 5.2 %-points, at $\sigma_{dE/dx} = 7\,\%$ it is 2.6 %-points. Similarly, for a fixed purity of 0.95 the efficiency increases are 20 % and 7 %, respectively.

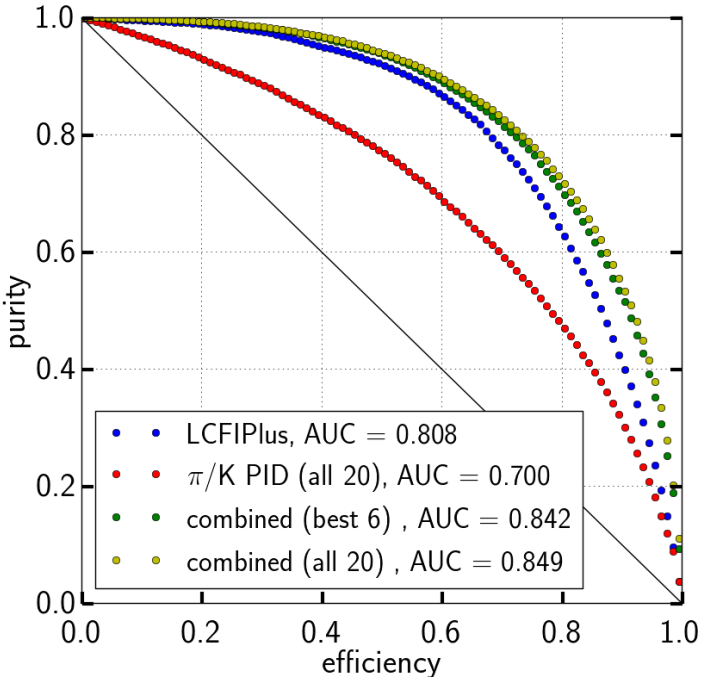

Figure 113: ROC curves for 4.5 % dE/dx resolution, showing the comparison between vertex-based flavour tag (FT) and PID-based flavour tag (pika) and their combination, once with all 20 observables and with only the 6 selected impactful ones.

These increases in efficiency can be generalised, which is displayed in Figure 115 and Figure 116, with Figure 116 showing the data in the lower right corner of Figure 115 in more detail. The higher the requested purity, the larger is the gain in efficiency through PID information. For a purity of 0.9, the efficiency increase with the default dE/dx resolution is about 10 %, for a 0.99 purity it approaches 30 %, and for even higher requested purities it quickly rises to a doubled efficiency, at which point statistical

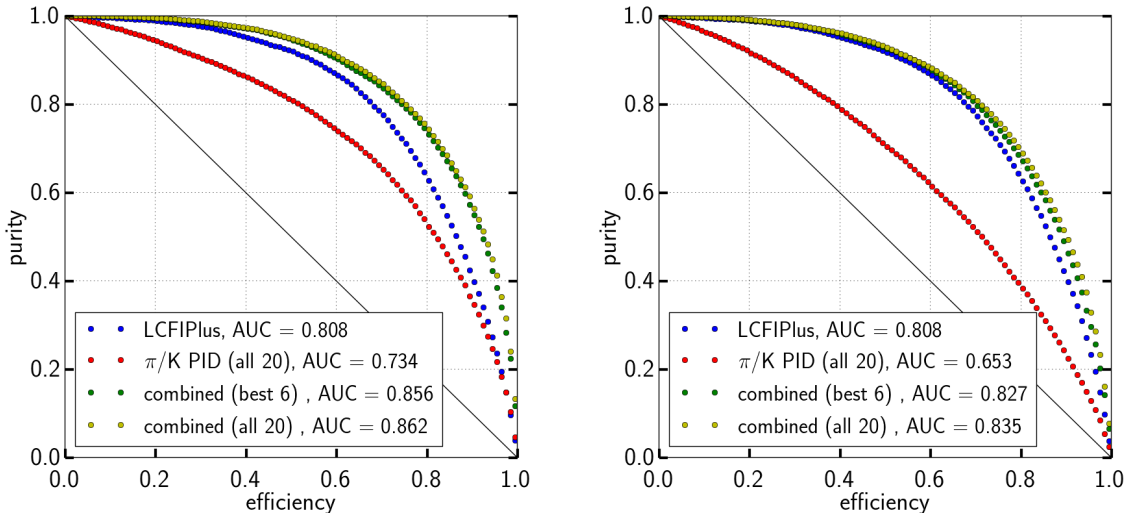

Figure 114: ROC curves for 2.6 % and 7 % dE/dx resolution. A better resolution improves the flavour tag performance correspondingly.

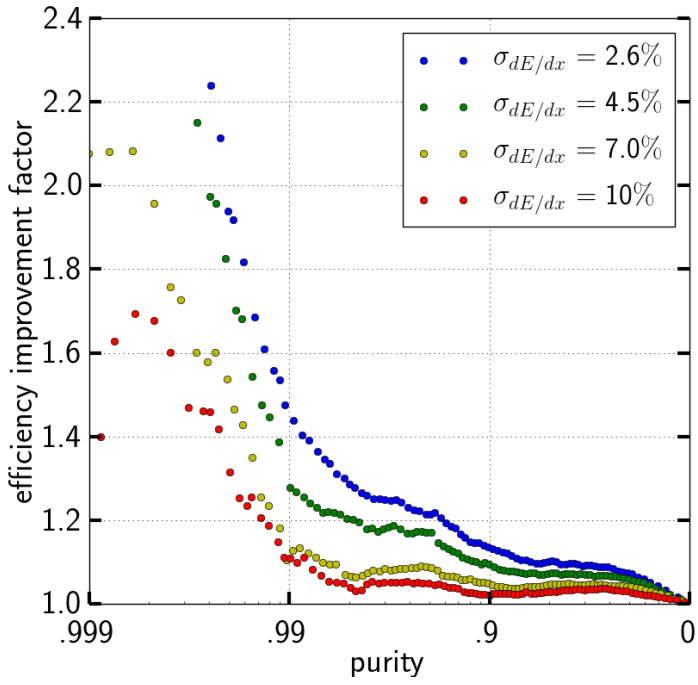

Figure 115: Efficiency improvement at high purity values for the comparison between only vertex-based flavour tag and flavour tag including PID information.

fluctuations limit the applicability of the BDT. The efficiency increase depends strongly on the dE/dx resolution: The default 4.5 % resolution outperforms the 7 % by about a factor of 2, while increasing the dE/dx resolution to 2.6 % adds another factor of 1.5.

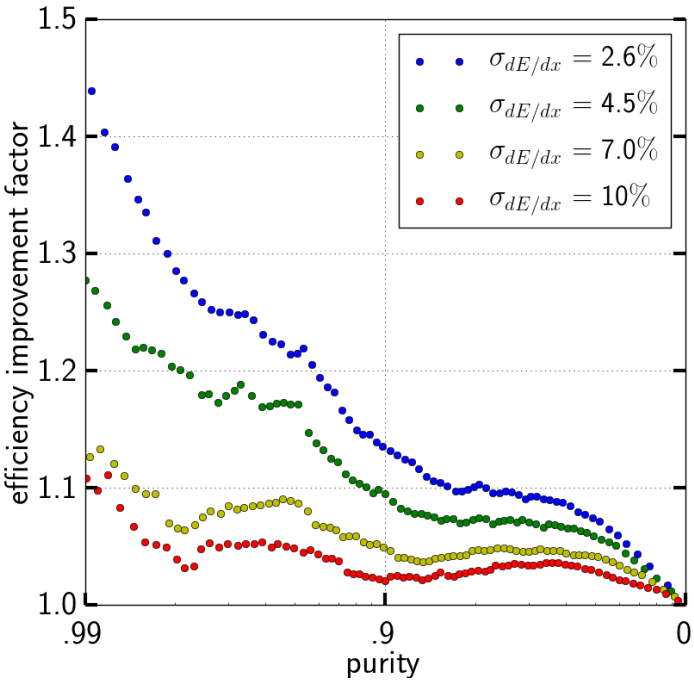

Figure 116: Efficiency improvement at high purity values, zommed into the lower right corner of Figure 115.

In order to assess the performance of the individual observables used in the BDT, a comparative analysis was executed. The calculation of ROC-curves and the corresponding AUC was done for each observable individually, i.e. only using that observable, in order to assess its individual power. In addition, the calculation was done again for each observable, but using all observables except that one, in order to assess its relevance to the final BDT performance. If observables can be identified that are entirely redundant, then they can be omitted without reducing the BDT performance. Instead, the reduction in degrees of freedom can conversely reduce the effect of over-fitting of the BDT's training sample, and even result in a somewhat better performance for the test sample. This was done by progressively eliminating the least relevant observables.

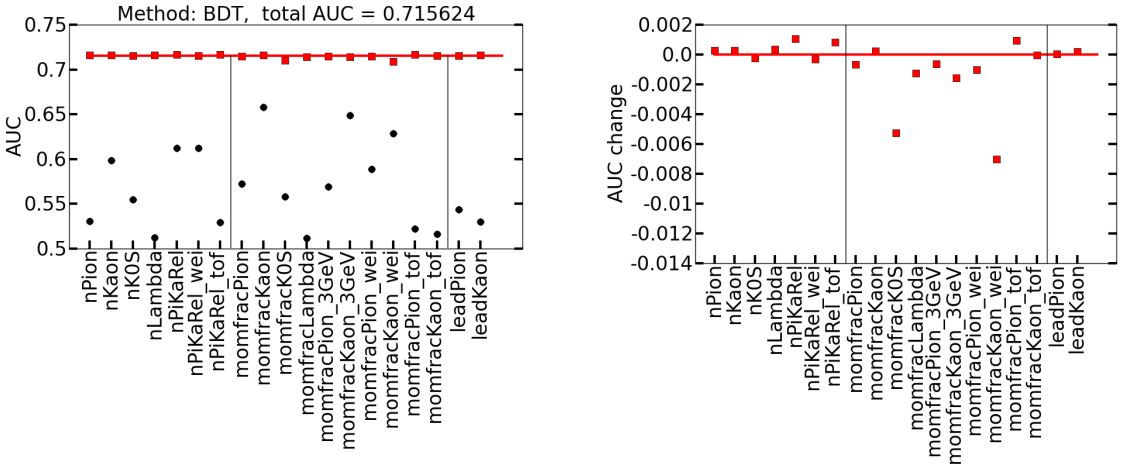

(a) (b)

Figure 117: Power and relevance of individual observables used in the BDT method. The red line indicates the BDT result as AUC including all observables. The red squares show the AUC in case this observables is omitted, indicating its relevance. The black circles show the AUC in case only this observable is used, indicating its power. (a) Absolute AUC value. (b) Difference to the reference AUC when all observables are included.

In Figure 117 the result of this assessment is shown for all observables used. The black circles indicate the power of the individual observables, and unsurprisingly, the three variations of the dE/dx-based kaon momentum fraction have the largest power. The Λ-related observables have the smallest power, due to the low number of Λ particles generated in the strange W decays. They are followed by the number relation of pions and kaons as well as their momentum fraction, when time-of-flight is used to identify them. This is due to the fact that the particles indicating the flavour of the initial quark have on average a large fraction of the overall jet momentum, usually larger than the upper limit for TOF usability at about 5 GeV. The figure also shows the relevance of the individual observables, indicated by red squares. The kaon-ness-weighted momentum fraction has the largest relevance, while the two other dE/dx-based kaon

momentum fraction observables are less relevant due to a higher degree of redundancy. The momentum fraction of K^0_S has the second largest relevance, despite it having only a moderate power. This, again unsurprisingly, shows the independence of the V0-Finding PID compared to dE/dx. Overall, the changes in AUC when removing one observable are smaller than 1 %-point, and even smaller than 2 permille for all except the two mentioned above. This is coherent with the large degree of redundancy in the set of observables.

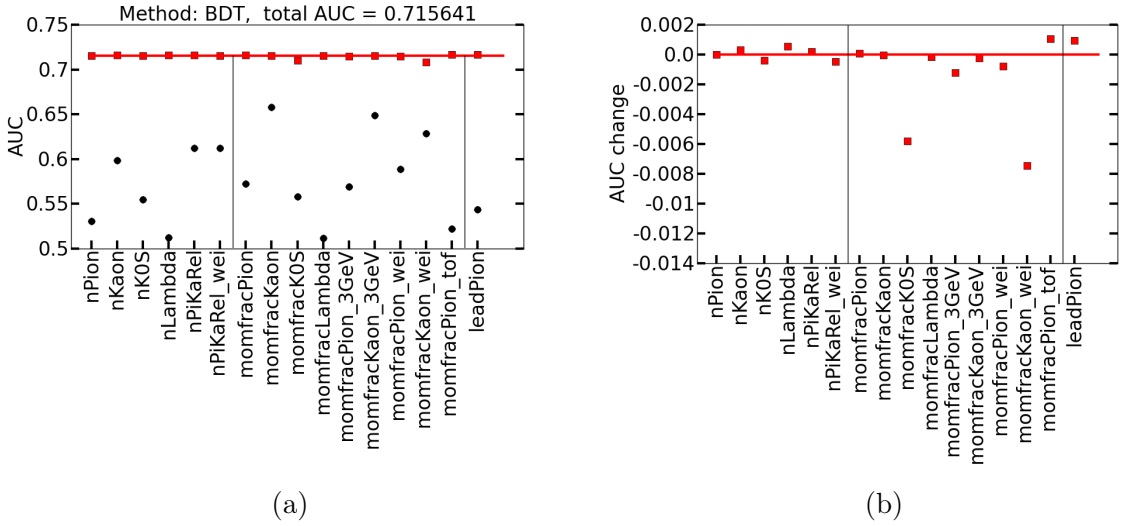

(a) (b)

Figure 118: Power and relevance of individual observables used in the BDT method, after the 3 least relevant oberservables were eliminated. The red line indicates the BDT result as AUC including all observables. The red squares show the AUC in case this observables is omitted, indicating its relevance. The black circles show the AUC in case only this observable is used, indicating its power. (a) Absolute AUC value. (b) Difference to the reference AUC when all observables are included.

The same plot is shown in Figure 118 after progressively eliminating the 3 least relevant observables, in this case the TOF-based pion/kaon number relation, the TOF-based kaon momentum fraction and the (dE/dx-based) number of leading kaons. The result in terms of power and relevance of the individual remaining observables is very similar to the one before, but despite using fewer observables, the overall AUC has not declined, but instead slightly increased by 0.2 permille. Two more variables can be excluded this way to slightly increase the AUC further, after that it decreases. Since the increase in AUC from omitting these variables is negligible, generally all variables were used in the analysis.

8.3. Conclusion

The two analyses presented in this chapter show a significant impact of PID, in particular via dE/dx, on physics observables.

With the statistics gained at the ILC and an excellent detector with timing capability, it could be possible to solve a long-standing discrepancy in the PDG book: the kaon mass. The currently best measurements are based on kaonic atoms, so the TOF-based approach serves as an entirely independent method and can narrow the statistical precision of the kaon mass to a decisive degree. In the analysis, time-of-flight is used to reconstruct the masses of any charged particle produced in the collider, and from this, kaons are both selected and their mass fitted. This way, no pre-selection is necessary for any specific physics process, and all events of the ILC can be utilised. The analysis was limited to a center-of-mass of 500 GeV, since the TOF observable was only introduced to this recent production. This makes the study likely the first one to utilise TOF at the ILC. The presented result is a statistical precision on the kaon mass of 30 keV before a luminosity upgrade and 10 keV with the full H-20 integrated luminosity, assuming a timing resolution of 50 ps for every ECal hit. With a timing resolution of 10 ps, this could increase to about 20 keV and 6 keV, respectively, clearly showing the timing dependence. With the given TOF algorithm, however, the result for 10 ps and 0 ps is nearly identical, which indicates the intrinsic limit of improvement. A better timing resolution not only narrows the mass peak, but also allows to use kaons from a larger momentum range, since they can be separated from pions up to a higher upper momentum cut, increasing the available statistics. The dE/dx value of each reconstructed charged particle was used for refinement to improve the kaon selection, which reduced contamination by pions and too allowed to use a larger momentum range. Using dE/dx refinement improved the achieved statistical error on the kaon mass by about 10 %.

Two factors limited the usable statistics in this analysis: the availability of the TOF observable in the MC productions, and the usability of TOF in the different parts of the detector. The 2020 MC production with a center-of-mass energy of 250 GeV contains the same TOF estimators as in the present analysis, and allows to add $2\,\mathrm{ab}^{-1}$ integrated luminosity to the $4\,\mathrm{ab}^{-1}$ at 500 GeV. In addition, the current TOF algorithm only assesses tracks in the barrel region. With an corresponding estimator, tracks going into the end-cap Ecal could be utilised as well, again significantly increasing the available statistics. Tracks with a momentum below about 1 GeV cannot reach the barrel ECal and instead curl in the TPC and reach the end-cap, instead. These curlers are not as abundant as particles with a few GeV, but their tracks in the TPC are much longer, which would allow for a much more precise mass reconstruction. On the other hand, the precision of their timing measurement in the end-cap ECal may be reduced due to multiple scattering in the comparably massive TPC anode, which has about 30 % of a radiation length of additional material compared to the barrel region. Overall, these two factors present a factor of 2 or more additional statistics, which would improve the statistical precision on the kaon mass by another at least 40 %.

On the other hand, the systematic error of the current algorithm is at the level of

10 MeV, three orders of magnitude larger than the statistical one. It is clear that the implementation of a TOF observable needs to be scrutinised and improved, in order to allow for a precision measurement like the kaon mass. Some work to this effect has already been done [119], reducing the systematic error to the level of about 1 MeV, equivalent to about 1 ps in relative precision, but more is clearly required. In addition, further sub-detectors should be investigated in order to assess the time-of-flight of a particle. In the barrel region, the silicon external tracker between the TPC and the ECal would be a natural candidate for a dedicated ultra-high precision timing layer. The silicon inner tracker, just inside the TPC inner field cage wall, could cover the very low momentum range below 1 GeV. In the end-cap region, first a TOF estimator based on the end-cap ECal similar to the barrel one needs to be implemented. In addition, it could be studied, if the TPC readout could reach a sufficient timing precision on the TPC hits closest to the anode to allow for a competitive TOF measurement. This would have the advantage that these tracks would not have to pass through the TPC anode material before the measurement. Reducing the absolute timing precision of the detector enough to a systematic error significantly below 1 MeV may not be feasible, but the remaining bias could be calibrated by using the corresponding peaks of pions and protons, whose masses are known to 0.24 keV and 6 eV, respectively. An actual ILD detector, however, would have more systematic uncertainties compared to the given simulation. Of particular relevance for this analysis is the homogeneity and stability of the solenoidal magnetic field. It should be noted that the NA62 experiment [120], which could do a similar kaon mass measurement as suggested in this analysis, instead inverts the process and uses kaon tracks with the current precision on their mass to calibrate their magnetic field.

It would be interesting, if other existing experiments could perform a kaon mass measurement. ATLAS and CMS at the LHC are planning to add precision timing layers with timing resolution between 30 ps and 60 ps to their detectors for the High-Lumi phase of the LHC. With the comparably large cross section of proton-proton collisions, a large number of kaons are produced, which could be used for this analysis. However, most triggers use tracks of several 10 GeV, which would be far above the upper momentum cut of TOF PID. ALICE at the LHC runs at much lower event rates, but with significantly more particles produced per event in ion-ion collisions, and it has more PID capabilities including a TPC and a dedicated TOF layer. Similarly, Belle II at the SuperKEKB accelerator has a central drift chamber, a TOF layer and and time-of-propagation counter, adding to the pion-kaon separation. At the same time, as a b-factory Belle II is operating with e^+e^- collisions at a center-of-mass of only 10.6 GeV, the $\Upsilon(4S)$ resonance. Compared to the ILC, the higher cross section at this energy allows Belle II to collect 50 ab^{-1} over its running time and the generated hadrons are at sufficiently low momenta to allow TOF PID. This raises the question if a similar kaon mass analysis could be performed with already existing or soon-to-be-taken data, and if not, what improvements to the detector would be needed to succeed.

The second analysis focused on the effect of dE/dx on the flavour tag. In a differential analysis, the added benefit of tracking-based PID information on the flavour tag

in hadronic W decays compared to the default mostly vertex-based flavour tag with LCFIPlus was studied. The capability to separate 2-jet systems where the W decayed into a d- and a u-quark from those where it decayed to an s- and a c-quark was assessed via a boosted decision tree and measured as area under the curve of a ROC-curve. Since s/c-systems have on average more and higher momentum kaons than d/u-systems, PID information weighted with measured momentum in 18 different observables could be used to differentiate the two cases. The additional information came mostly from dE/dx-related quantities, with V0-finding delivering a smaller but independent contribution from neutral strange particles. TOF, however, was of negligible effect in this case, since the correlation with the jet flavour lays mostly in high-momentum pions and kaons, which can not be separated from each other with TOF. Usage of the PID information showed a moderate increase in the BDT AUC compared to the well-established LCFIPlus performance. This increase was, however, significantly larger when high purities are needed in an analysis that uses the flavour tag, allowing to use several 10 % more data, depending on the purity level. This performance improvement also shows a strong dependence on the applied dE/dx resolution in the TPC. For a requested purity of 99 % for example, the default 4.5 % resolution increases the efficiency by about 27 %, while a 2.6 % resolution would achieve about 44 % and a significantly reduced dE/dx resolution of 7 or 10 % would only add 10 % efficiency, mostly coming from the V0-Finding which is unaffected by the dE/dx resolution.

An analysis that uses flavour tag to specifically separate hadronic W decays is the measurement of V_{cs}, the central element of the CKM quark-mixing matrix. The values of the matrix have been determined to the level of about 10^{-4} - in a global SM fit and with the contraint of three-generation unitarity. If this contraint is fulfilled is not evident a priori, and a violation could indeed indicate new physics. Therefore, an independent measurement of the individual elements of the CKM matrix is a necessary cross-check. While most elements have been measured to the sub-% level, V_{cs} and V_{tb} are only known to an accuracy of about 2 %. One example of a direct measurement of V_{cs} was performed by the DELPHI experiment at LEP [117], which in 1996 used about 100 W pairs to determine $|V_{cs}| = 0.94^{+0.32}_{-0.26}(\text{stat.})\pm 0.13(\text{syst.})$ via only the flavour tag method, i.e. even independent of the W branching fraction into leptons. The ILC is expected to produce about 60 million W pairs in the H-20 scenario, with a production cross section of about 15 pb and an integrated luminosity of 2 ab^{-1} at 250 GeV, and 7 pb and 4 ab^{-1} at 500 GeV, and an additional about 20 million single W at 500 GeV with a production cross section of 5 pb. This gain in statistics of nearly 6 orders of magnitude will allow for an unprecendented precision on the statistical uncertainty of the V_{cs} measurement below the per-mille level. With these large statistics a reduced selection efficiency can be tolerated and it would be beneficial to chose a flavour tag working point at a high purity level to more effectively suppress background from alternative W decay modes, which is where the dE/dx PID is most impactful.

9. Conclusion

In this work, a bridge was built from hardware development to detailed and more general simulation and finally to physics observables, demonstrating why advanced PID is desirable in a Higgs factory, how this can be implemented in software in particular via dE/dx, how dE/dx PID depends on the readout in a TPC, and how an improved TPC readout system can be realised.

The Ropperi system shows an ambitious path forwards to a novel detector concept that combines flexibility and granularity in a highly integrated way. The bonding between ASIC and 'standard' PCB is a novel process at the attempted bond sizes and still provided a temporary readout allowing to assess the signal-to-noise level in the system, which is in the order of 100 for typical signal sizes. In terms of TPC readout, the detailed simulation makes clear: granularity is king. Decreasing the mm-sized pads of pad-based readouts already improves the dE/dx PID, while application of cluster counting requires a granularity below 300 µm, but enables further improvement. One such cluster counting algorithm, utilising astrophysics software, was adapted in this work and applied for the first time to a broad range of granularities spanning two orders of magnitude from pads to pixels. Should the GridPix system run into an unexpected show-stopper, then a next generation of Ropperi could fill the gap and provide this level of granularity.

For ILD, this allows to formulate a more ambitious aim for the dE/dx resolution: 3 % instead of 5 %! This would also be a significant step forwards compared to DELPHI's 5.7 %, ALEPH's 4.4 % [121] and currently ALICE's 5.2 % in pp collisions and 6.5 % in Pb-Pb collisions [49]. The software tools to calibrate and monitor this were developed and provided to the collaboration. Beyond dE/dx, there are more tools for PID, like V0-finding in the TPC and TOF outside of it. The physics benefits of these PID capabilities are laid out in two example studies. With an excellent timing layer, the kaon mass can be measured and a decades old discrepancy in the PDG book resolved - provided the systematics can be kept in check. Improved measurements on 'basic' quantities are increasingly needed as input to simulations - the kaon mass is a crucial input to lattice QCD calculations, which are now tied to more conventional methods of precision QCD and will proceed to improve with growing CPU power. As particle physics community we need to measure the necessary input quantities as boundary conditions for theory calculations in time, to enable theory to indicate the path to new physics via high precision. The W-decay analysis shows the benefit of advanced PID for the flavour tag. Not only enhances the inclusion of tracking-based PID quantities into the existing flavour tag its performance, but with a highly granular TPC readout and its optimised dE/dx resolution this can be improved even further. For an analysis using the flavour tag at a high-purity working point, this dE/dx improvement alone could provide as much as 13 % more usable data, the equivalent of 1.5 years additional running time of the ILC after the luminosity upgrade in accordance with the H-20 scenario.

This leaves one central conclusion: use the detector as well as possible! Collisions are expensive and should be analysed to the best of our abilities. Optimal usage of collision information includes usage of a TPC in ILD in the first place, and in particular striving for an optimised high-granularity TPC readout. V0-Finding, as a feature of the

excellent pattern recognition of the TPC, should be revised and optimised with respect to its current implementation in iLCSoft. It extends to further PID layers, namely TOF via excellent timing in an existing sub-detector or a dedicated timing layer. Also entirely new ideas, like a time-of-propagation layer in the barrel region between TPC and ECal, could be implemented in simulation to study their potential additional benefits. The benefits of optimised PID were shown in this work and already several others for the ILD. Similar to the W-decay analysis presented here, the quark couplings to the Z can be measured via hadronic Z decays [122]. Kaon identification is used in measuring the forward-backward asymmetry of hadronic Z decays [123, 124]. PID of any charged particle should be applied to its momentum fit by using its actual rest mass instead of the 'pion hypothesis' to improve momentum resolution and consequently jet energy resolution [125]. Even in Supersymmetry analyses PID can be crucial, for example to identify low-momentum decay products in low-Δm SUSY [126]. Countless analyses have used PID in the past, in particular at the LEP experiments, and will provide inspiration to ILD analyses, while young physicists will think of new ways to utilise all information extracted from the collisions - if we provide it.

A. Appendix

A.1. dE/dx PID Plots

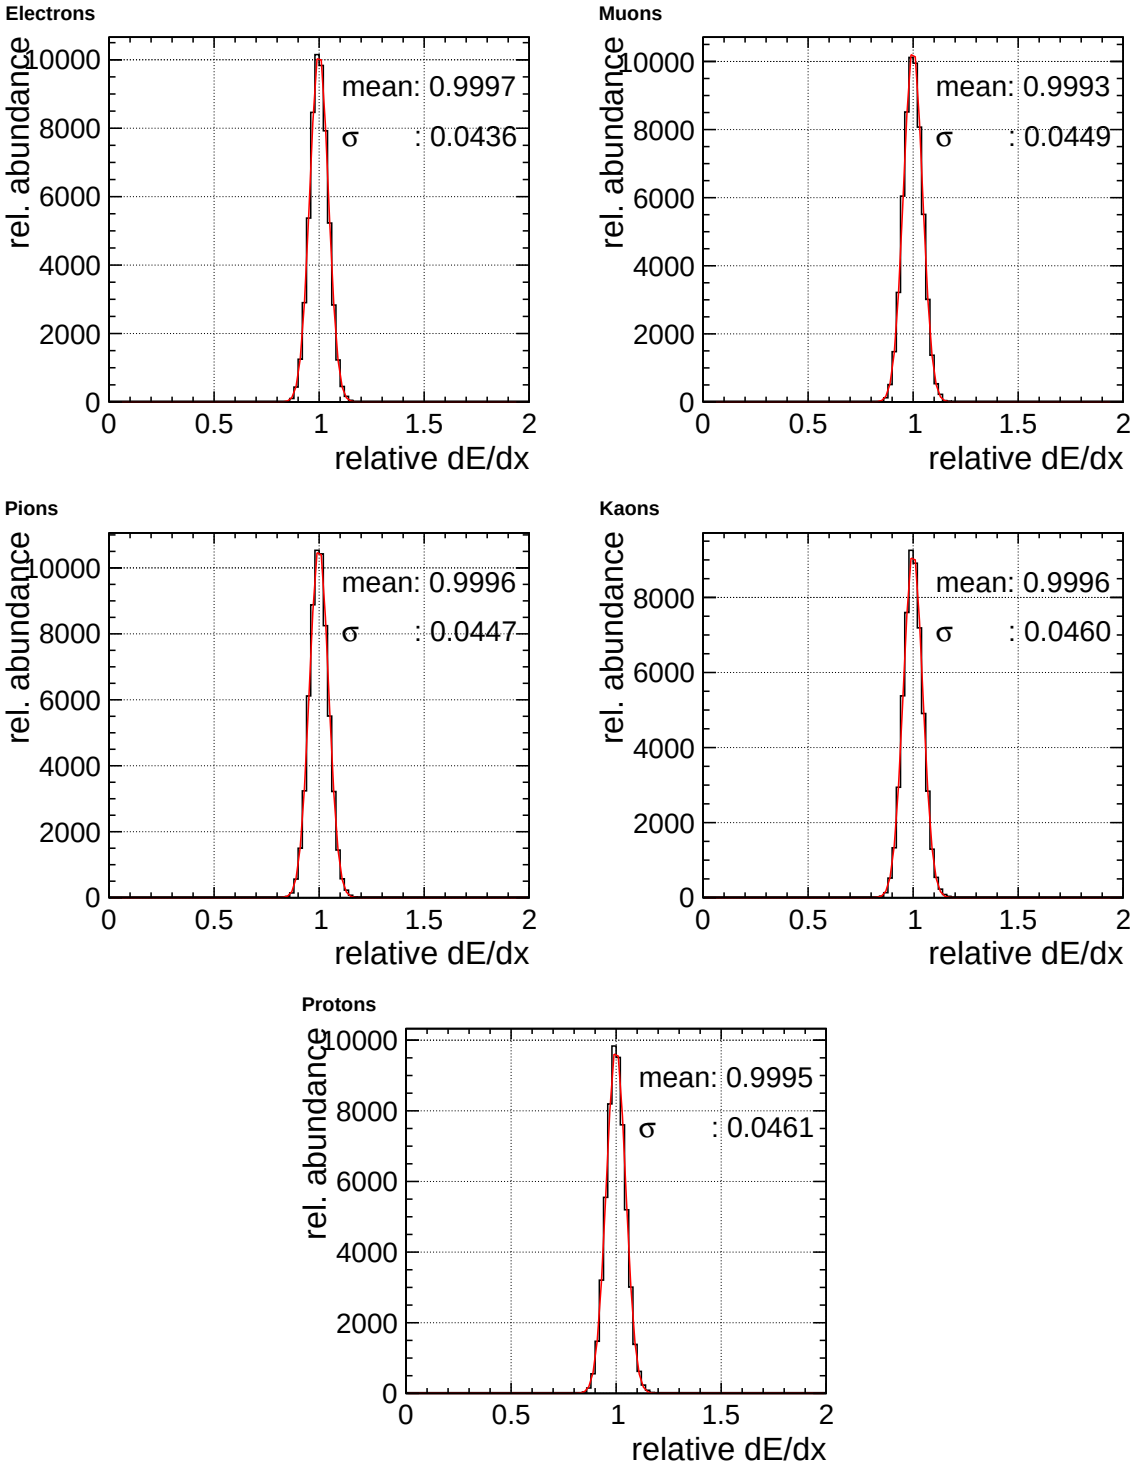

Figure 119: Combined dE/dx resolution of single particles above 1 GeV momentum. A Guassian fit is depicted as a red curve, with its mean and σ highlighted.

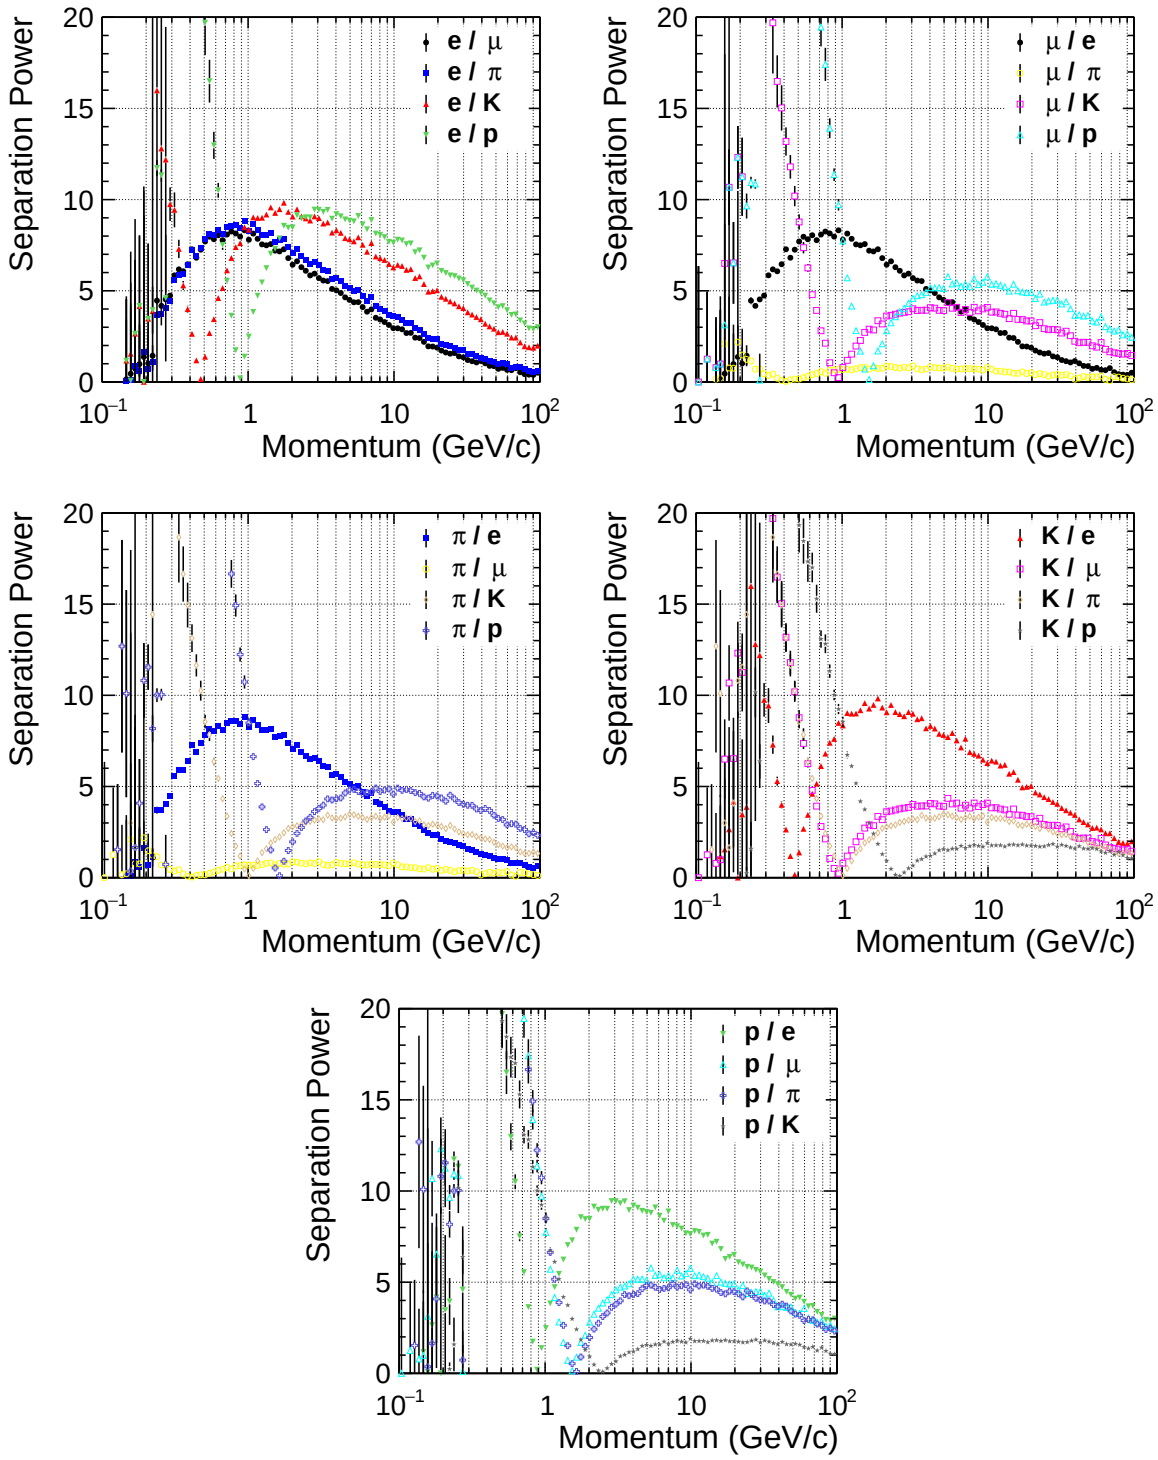

Figure 120: Separation power of each species vs. each other species, derived from single particles.

A.2. W decay Observables Plots

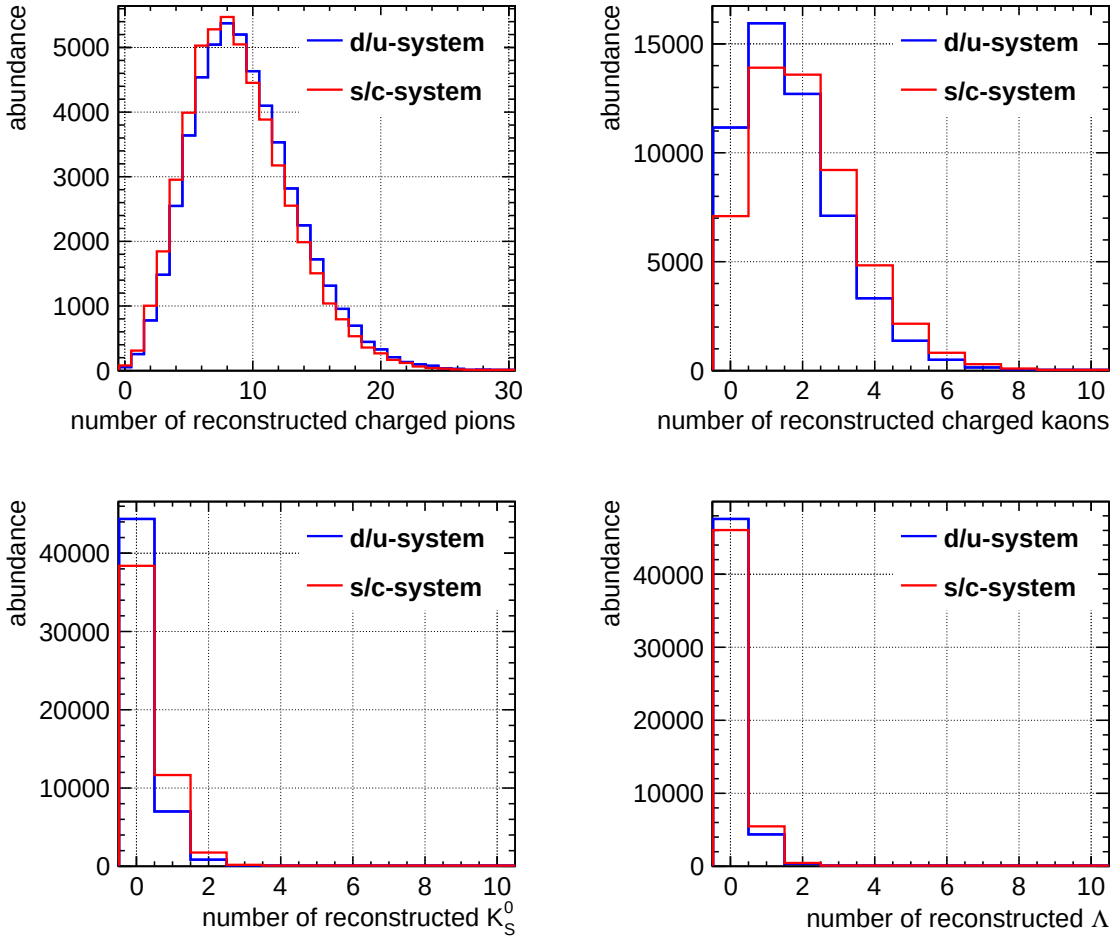

Figure 121: Reconstructed jet observables: numbers of PiKaLas.

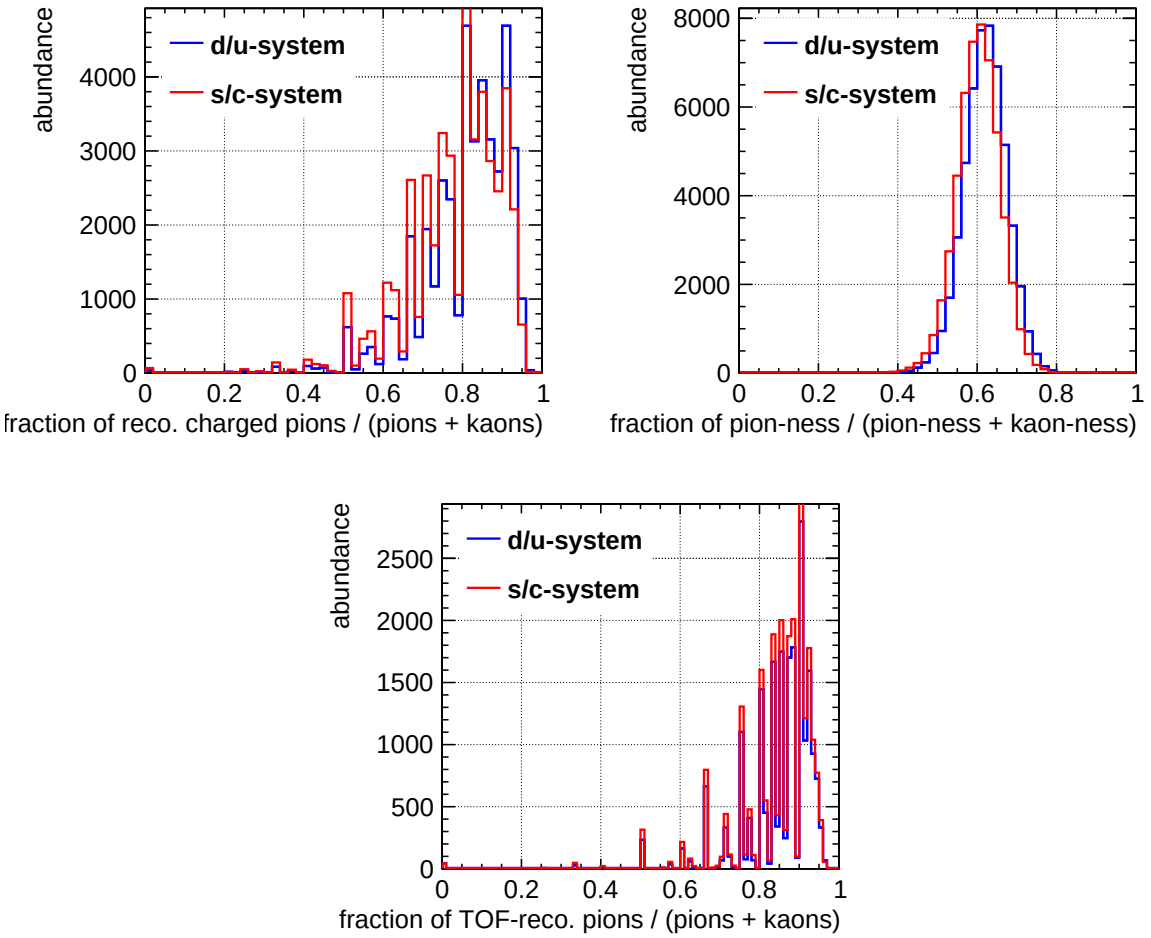

Figure 122: Reconstructed jet observables: (weighted) fraction of pions and kaons.

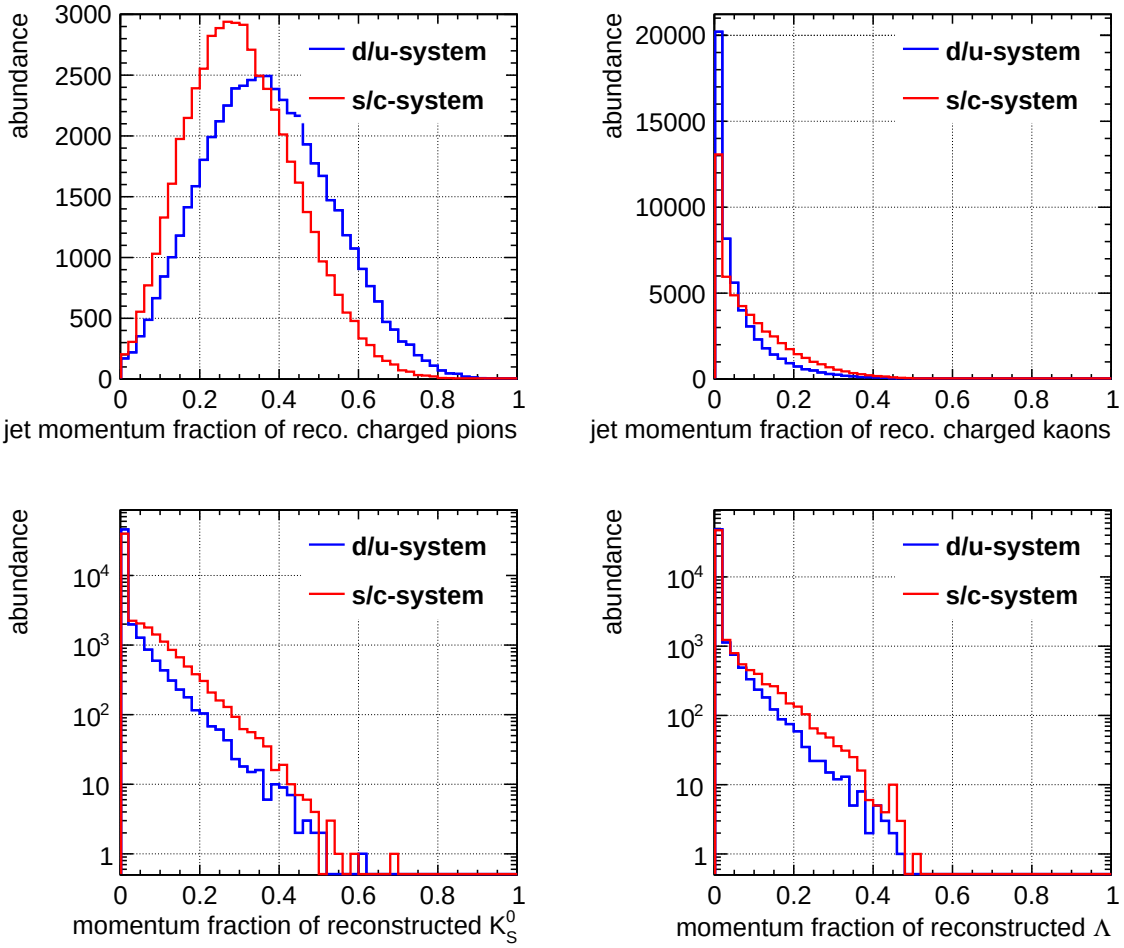

Figure 123: Reconstructed jet observables: jet momentum fraction of PiKaLas.

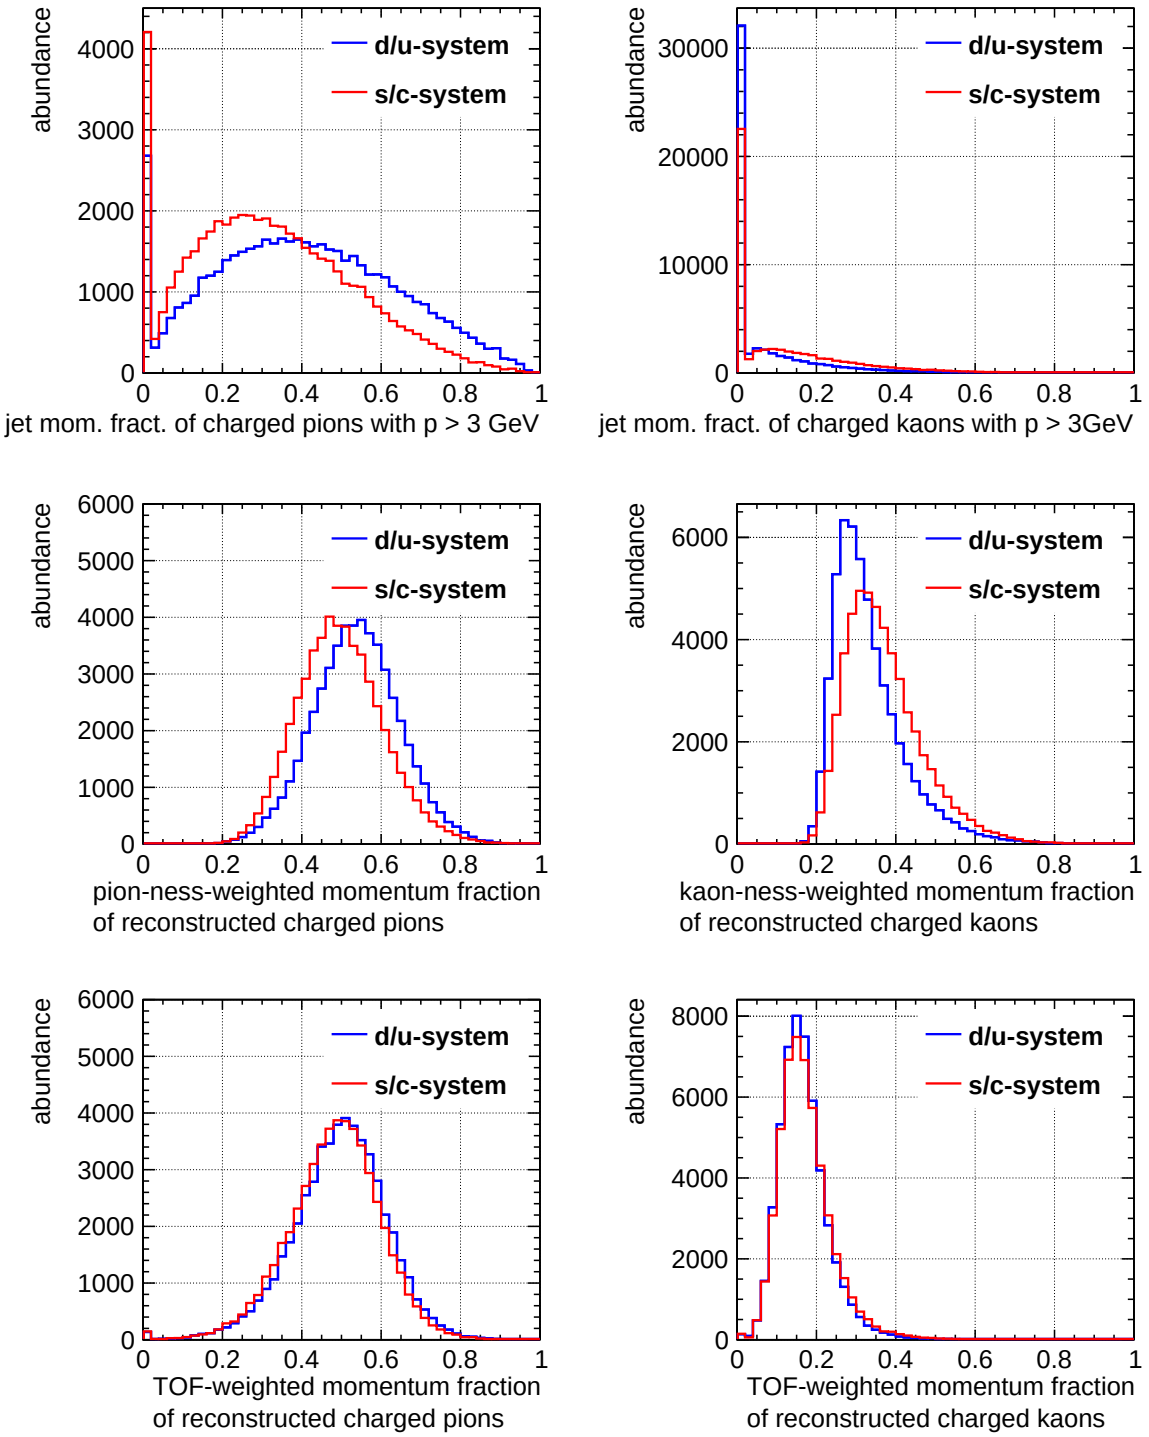

Figure 124: Reconstructed jet observables: jet momentum fractions above 3 GeV, weighted by pion-/kaon-ness, and based on TOF.

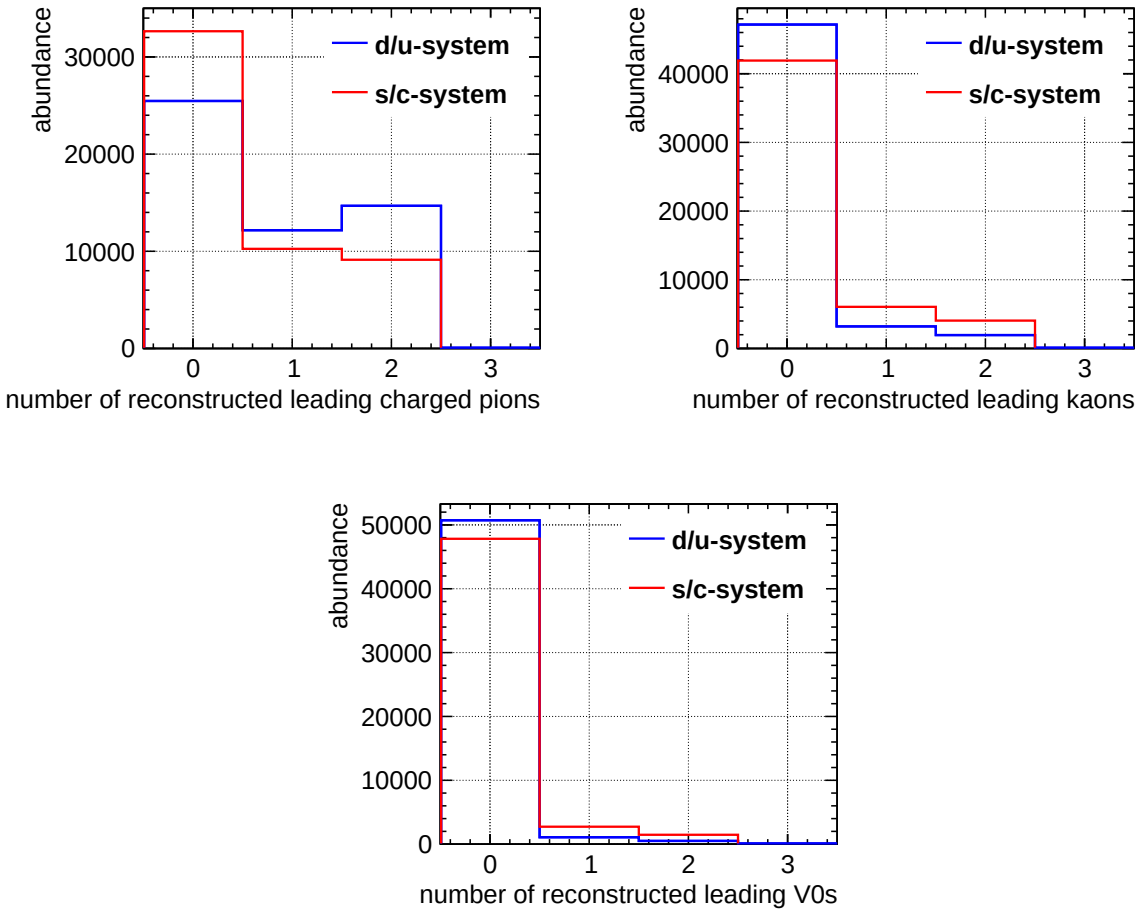

Figure 125: Reconstructed jet observables: number of leading pions, kaons and V0s.

A.3. Datasheet: N7000

Nelco® N7000-2 HT Laminate
Nelco® N7000-3 Prepreg and Laminate

Toughened Polyimide Prepreg & Laminate

The Nelco N7000-2 HT laminate and N7000-3 prepreg are a series of toughened polyimide material for use in high-reliability multilayers. This combined resin system provides excellent thermal performance, improved processing characteristics and is exceptional for use in a wide variety of applications that include fine geometry multilayer constructions and extreme reliability requirements.

Key Features

Polyimide Resin Chemistry
- Robust thermal stability and reliability
- Toughened resin system
- High temperature tolerance and chemical resistance

Lead-free Assembly Compatibility
- Withstands multiple thermal excursions
- Tg 260°C by DSC
- T_{260} >12 minutes
- Low Z-Axis CTE
- Designed for use in severe conditions

Supports Current and Previous Military and Industrial Standards
- N7000-2 HT and N7000-3 meet IPC-4101/40 and /41
- Complies with the old GIJ and GIL military specifications

Reliable Plated-through Holes
- Low Z-Axis CTE and toughened polyimide chemistry providing good dimensional stability

Reliable Processing
- Improved fracture resistance compared with traditional polyimide systems
- Reduced cure time compared to other traditional polyimide systems

And Much More
- Vacuum laminated
- Available in a wide variety of constructions, copper weights and glass styles
- All Nelco materials are RoHS compliant

Applications

- Fine-Line Multilayers
- Backplanes
- Surface-Mount Multilayers
- BGA Multilayers
- Direct Chip Attach
- Underhood Automotive
- Burn-in Boards

Global Availability

Nelco Products, Inc. (California) - Americas
+1.714.879.4293
Neltec, Inc. (Arizona) - Americas
+1.480.967.5600
Nelco Products Pte. Ltd. - Asia Pacific
+65.6861.7117
Neltec, S.A. - Europe
+33.562.98.52.90
www.parkelectro.com
info@parkelectro.com

Park's UL file number: E36295

PARK
ELECTROCHEMICAL
CORP.

Nelco® N7000-2 HT / N7000-3

Toughened Polyimide Laminate & Prepreg

Mechanical Properties	U.S. Units		Metric Units		Test Method
Peel Strength - 1 oz. (35 micron) Cu					
After Solder Float	7.5	lb / inch	1.31	N / mm	IPC-TM-650.2.4.8
At Elevated Temperature	7.0	lb / inch	1.22	N / mm	IPC-TM-650.2.4.8.2a
After Exposure to Process Solutions	7.0	lb / inch	1.22	N / mm	IPC-TM-650.2.4.8
X / Y CTE [-40°C to +125°C]	9 - 12	ppm / °C	9 - 12	ppm / °C	IPC-TM-650.2.4.41
Z Axis Expansion [50°C to 260°C]	<2.5	%	<2.5	%	IPC-TM-650.2.4.24
Young's Modulus (X / Y)	3.1 / 3.3	psi x 106	21.1 / 22.2	GN / m2	ASTM D3039
Poisson's Ratios (X / Y)	0.146 / 0.153		0.146 / 0.153		ASTM D3039
Thermal Conductivity	0.45	W / mK	0.45	W / mK	ASTM E1461
Specific Heat	1.06	J / gK	1.06	J / gK	ASTM E1461
Electrical Properties					
Dielectric Constant (50% resin content)					
@ 1 GHz (RF Impedance)	3.8		3.8		IPC-TM-650.2.5.5.9
@ 2.5 GHz (Stripline)	3.5		3.5		IPC-TM-650.2.5.5.5
@ 10 GHz (Stripline)	3.5		3.5		IPC-TM-650.2.5.5.5
@ 10 GHz (Split Post Cavity)	3.5		3.5		
Dissipation Factor (50% resin content)					
@ 2.5 GHz (Stripline)	0.015		0.015		IPC-TM-650.2.5.5.5
@ 10 GHz (Stripline)	0.015		0.015		IPC-TM-650.2.5.5.5
@ 10 GHz (Split Post Cavity)	0.009		0.009		
Volume Resistivity					
C - 96 / 35 / 90	10^7	MΩ - cm	10^7	MΩ - cm	IPC-TM-650.2.5.17.1
E - 24 / 125	10^7	MΩ - cm	10^7	MΩ - cm	IPC-TM-650.2.5.17.1
Surface Resistivity					
C - 96 / 35 / 90	10^7	MΩ	10^7	MΩ	IPC-TM-650.2.5.17.1
E - 24 / 125	10^7	MΩ	10^7	MΩ	IPC-TM-650.2.5.17.1
Electric Strength	1200	V / mil	$4.7x10^4$	V / mm	IPC-TM-650.2.5.6.2
Dielectric Breakdown	>50	kV	>50	kV	IPC-TM-650.2.5.6
Arc Resistance	100	seconds	100	seconds	IPC-TM-650.2.5.1
Thermal Properties					
Glass Transition Temperature (Tg)					
DSC (°C)	260	°C	260	°C	IPC-TM-650.2.4.25c
TMA (°C)	250	°C	250	°C	IPC-TM-650.2.4.24c
Degradation Temp (TGA) (5% wt. loss)	376	°C	376	°C	IPC-TM-650.2.4.24.6
Pressure Cooker-60 min then solder dip					IPC-TM-650.2.6.16
@288°C until failure (max 10 min.)	Pass		Pass		(modified)
T260	12+	minutes	12+	minutes	IPC-TM-650.2.4.24.1
Chemical / Physical Properties					
Moisture Absorption	0.35	wt. %	0.35	wt. %	IPC-TM-650.2.6.2.1
Methylene Chloride Resistance	<0.50	% wt. chg.	<0.50	% wt. chg.	IPC-TM-650.2.3.4.3
Density [50% resin content]	1.70	g / cm^3	1.70	g / cm^3	Internal Method

Park Electrochemical Corp. is a global advanced materials company which develops and manufactures high-technology digital and RF/microwave printed circuit materials and advanced composite materials, parts and assemblies. The company operates under the Nelco®, Nelcote® and Nova™ names.

All test data provided are typical values and not intended to be specification values. For review of critical specification tolerances, please contact a Nelco representative directly. Nelco reserves the right to change these typical values as a natural process of refining our testing equipment and techniques.

PARK
ELECTROCHEMICAL
CORP.

References

[1] U. Einhaus and J. Kaminski. *ROPPERI - A TPC readout with GEMs, pads and Timepix*, 2017. http://arxiv.org/abs/1703.08529.

[2] U. Einhaus, J. Kaminksi, and M. Caselle. *ROPPERI - A TPC readout with GEMs, pads and Timepix*, 2018. http://arxiv.org/abs/1801.07178.

[3] F. Gaede. *Marlin and LCCD: Software tools for the ILC.* Nucl. Instrum. Meth., A559:177–180, 2006. http://dx.doi.org/10.1016/j.nima.2005.11.138.

[4] J. Abernathy et al. *MarlinTPC: A Marlin based common TPC software framework for the LC-TPC-TPC Collaboration.* LC Notes, 2007, http://bib-pubdb1.desy.de/record/83142.

[5] *Code base of MarlinTPC.* https://svnsrv.desy.de/desy/marlintpc. [online; accessed 12-September-2021; access authorisation required].

[6] S. Dharani. *Particle Identification using Time of Flight in the International Large Detector ILD.* Bachelor thesis, Universität Leipzig, 2018. https://web.physik.uni-bonn.de/group/view.php?&group=1&lang=de&c=t&id=28.

[7] The ILD Collaboration. *International Large Detector: Interim Design Report*, 2020. http://arxiv.org/abs/2003.01116.

[8] I. Flegel, M. Fleischer, and T. Konstandin, editors. *Particle physics 2020: highlights and annual report; Editorial deadline: 15 February 2021.* Deutsches Elektronen-Synchrotron, Hamburg, 2021.

[9] U. Einhaus. Charged Hadron Identification with dE/dx and Time-of-Flight at Future Higgs Factories. In *Proceedings of The European Physical Society Conference on High Energy Physics — PoS(EPS-HEP2021)*, volume 398, page 760, 2022.

[10] G. Aad et al. *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC.* Physics Letters B, 716(1):1–29, 2012. http://dx.doi.org/https://doi.org/10.1016/j.physletb.2012.08.020.

[11] S. Chatrchyan et al. *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC.* Physics Letters B, 716(1):30–61, 2012. http://dx.doi.org/https://doi.org/10.1016/j.physletb.2012.08.021.

[12] *Standard Model of particle physics.* https://en.wikipedia.org/wiki/Standard_Model. [online; accessed 27-August-2021].

[13] P. D. Group, P. A. Zyla, et al. *Review of Particle Physics.* Progress of Theoretical and Experimental Physics, 2020(8), 08 2020. http://dx.doi.org/10.1093/ptep/ptaa104. 083C01.

[14] K. P. Gall, J. P. Miller, F. O'Brien, B. L. Roberts, D. R. Tieger, G. W. Dodson, M. Eckhause, J. Ginkel, P. P. Guss, D. W. Hertzog, D. Joyce, K. J. R., C. Kenney, J. Kraiman, W. C. Phillips, W. F. Vulcan, R. E. Welsh, R. J. Whyley, R. G. Winter, R. J. Powers, R. B. Sutton, and A. R. Kunselman. *Precision Measurements of the K^- and Σ^- Masses.* Phys. Rev. Lett., 60:186–189, 1988.

[15] A. S. Denisov, A. V. Zhelamkov, Y. M. Ivanov, L. P. Lapina, P. M. Levchenko, V. D. Malakhov, A. A. Petrunin, A. G. Sergeev, A. I. Smirnov, V. M. Suvorov, and O. L. Fedin. *New measurements of the mass of the K^- meson.* JETP Lett., 54:558–562, 1991.

[16] S. Aoki et al. *Review of lattice results concerning low-energy particle physics: Flavour Lattice Averaging Group (FLAG).* The European physical journal. C, Particles and fields, 77(2)(112), 2017. http://dx.doi.org/10.1140/epjc/s10052-016-4509-7.

[17] S. Borsanyi, Z. Fodor, J. N. Guenther, C. Hoelbling, S. D. Katz, L. Lellouch, T. Lippert, K. Miura, L. Parato, K. K. Szabo, F. Stokes, B. C. Toth, C. Torok, and L. Varnhorst. *Leading hadronic contribution to the muon magnetic moment from lattice QCD*, 2020. http://arxiv.org/abs/2002.12347.

[18] T. Aoyama, N. Asmussen, M. Benayoun, J. Bijnens, T. Blum, M. Bruno, I. Caprini, C. Carloni Calame, M. Cè, G. Colangelo, and et al. *The anomalous magnetic moment of the muon in the Standard Model.* Physics Reports, 887:1–166, Dec 2020. http://dx.doi.org/10.1016/j.physrep.2020.07.006.

[19] B. Abi et al. *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm.* Phys. Rev. Lett., 126:141801, Apr 2021. http://dx.doi.org/10.1103/PhysRevLett.126.141801.

[20] T. Behnke, J. E. Brau, B. Foster, J. Fuster, M. Harrison, J. McEwan Paterson, M. Peskin, M. Stanitzki, N. Walker, and H. Yamamoto, editors. *The International Linear Collider Technical Design Report - Volume 1: Executive Summary.* 2013. ILC-REPORT-2013-040.

[21] H. Baer, T. Barklow, K. Fujii, Y. Gao, A. Hoang, S. Kanemura, J. List, H. E. Logan, A. Nomerotski, M. Perelstein, M. E. Peskin, R. Poeschl, J. Reuter, S. Riemann, A. Savoy-Navarro, G. Servant, T. M. P. Tait, and J. Yu, editors. *The International Linear Collider Technical Design Report - Volume 2: Physics.* 2013.

[22] C. Adolphsen et al. *"The International Linear Collider Technical Design Report - Volume 3.I: Accelerator R&D in the Technical Design Phase".* 2013.

[23] C. Adolphsen et al. *"The International Linear Collider Technical Design Report - Volume 3.II: Accelerator Baseline Design".* 2013.

[24] T. Barklow, J. Brau, K. Fujii, J. Gao, J. List, N. Walker, and K. Yokoya. *ILC Operating Scenarios*, 2015. http://arxiv.org/abs/1506.07830.

[25] H. Li, R. Pöschl, and F. Richard. *HZ Recoil Mass and Cross Section Analysis in ILD*, 2012. http://arxiv.org/abs/1202.1439.

[26] T. Behnke, J. E. Brau, P. N. Burrows, J. Fuster, M. Peskin, M. Stanitzki, Y. Sugimoto, S. Yamada, and H. Yamamoto, editors. *The International Linear Collider Technical Design Report - Volume 4: Detectors*. 2013. ILC-REPORT-2013-040.

[27] D. Buskulic et al. *Performance of the ALEPH detector at LEP*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 360(3):481–506, 1995. http://dx.doi.org/https://doi.org/10.1016/0168-9002(95)00138-7.

[28] ATLAS Collaboration. *Jet energy scale and resolution measured in proton–proton collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector*. The European Physical Journal C, 81:689, 2021. http://dx.doi.org/10.1140/epjc/s10052-021-09402-3.

[29] N. Feege. *Low-energetic Hadron Interactions in a Highly Granular Calorimeter*. Dissertation, Universität Hamburg, Hamburg, 2011. http://wwwalt.physnet.uni-hamburg.de/services/biblio/dissertation/dissfbPhysik/___Volltexte/Nils___Feege/, Dissertation, Universität Hamburg, 2011.

[30] O. Schäfer. *Ein Monitorsystem für Gasbasierte Detektoren am International Linear Collider (ILC)*. Diploma thesis, Universität Rostock, 2006.

[31] G. Dellacasa et al. *ALICE time projection chamber: Technical Design Report*. Technical design report. ALICE. CERN, Geneva, 2000.

[32] H. Wieman et al. *STAR TPC at RHIC*. IEEE Trans. Nucl. Sci., 44:671–678, 1997. http://dx.doi.org/10.1109/23.603731.

[33] G. Charpak and B. Southworth. *New particle detectors—The multiwire proportional chamber and the drift chamber*. The Physics Teacher, 15(1):26–31, 1977. http://dx.doi.org/10.1119/1.2339524.

[34] F. Sauli. *GEM: A new concept for electron amplification in gas detectors*. Nuclear Instruments and Methods, 386(2):531 – 534, 1997. http://dx.doi.org/10.1016/S0168-9002(96)01172-2.

[35] K. Zenker. *Studies of Field Distortions in a Time Projection Chamber for the International Linear Collider*. PhD thesis, Universität Hamburg, 2014. https://bib-pubdb1.desy.de/record/207651, DESY-THESIS-2014-044.

[36] Y. Giomataris, P. Rebourgeard, J. Robert, and G. Charpak. *MICROMEGAS: a high-granularity position-sensitive gaseous detector for high particle-flux environments.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 376(1):29–35, 1996. http://dx.doi.org/https://doi.org/10.1016/0168-9002(96)00175-1.

[37] Gas Detector Development Group at CERN. 2021. https://gdd.web.cern.ch/gem.

[38] M. Chefdeville. *Development of Micromegas-like gaseous detectors using a pixel readout chip as collecting anode.* PhD thesis, Universiteit von Amsterdam, 2009. https://cds.cern.ch/record/1169417.

[39] D. Attié et al. *A Time Projection Chamber with GEM-based Readout.* 2016. http://dx.doi.org/10.1016/j.nima.2016.11.002.

[40] D. Attié et al. *Double/hit separation capability and dE/dx resolution of a Time Projection Chamber with GEM redout.* in prep., 2021.

[41] R. Esteve Bosch, A. Jimenez de Parga, B. Mota, and L. Musa. *The ALTRO chip: A 16-channel A/D converter and digital processor for gas detectors.* IEEE Trans. Nucl. Sci., 50:2460–2469, 2003. http://dx.doi.org/10.1109/TNS.2003.820629.

[42] L. Jönsson. *Front-end electronics for the TPC in ILD; a status report April 2014.* LC Notes, LC-DET-2014-007, 2014, http://flc.desy.de/lcnotes.

[43] A. Shoji. *Measurement of dE/dx resolution of TPC prototype with gating GEM exposed to an electron beam.* 2018. http://arxiv.org/abs/1801.04499.

[44] P. Colas. *A new scheme for Micromegas TPC readout: the encapsulated resistive anode with reverse grounding.* 2019, https://indico.cern.ch/event/757322/contributions/3387077/.

[45] J. Kaminski, Y. Bilevych, K. Desch, C. Krieger, and M. Lupberger. *GridPix detectors – introduction and applications.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 845:233–235, 2017. http://dx.doi.org/https://doi.org/10.1016/j.nima.2016.05.134. Proceedings of the Vienna Conference on Instrumentation 2016.

[46] C. Ligtenberg, K. Heijhoff, Y. Bilevych, K. Desch, H. van der Graaf, M. Gruber, F. Hartjes, J. Kaminski, N. van der Kolk, P. Kluit, G. Raven, L. Scharenberg, T. Schiffer, S. Schmidt, and J. Timmermans. *Performance of a GridPix TPC readout based on the Timepix3 chip.* 2019. http://arxiv.org/abs/1902.01987.

[47] P. Colas, A. Colijn, A. Fornaini, Y. Giomataris, H. van der Graaf, E. Heijne, X. Llopart, J. Schmitz, J. Timmermans, and J. Visschers. *The readout of a GEM or*

Micromegas-equipped TPC by means of the Medipix2 CMOS sensor as direct anode. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 535(1):506–510, 2004. http://dx.doi.org/https://doi.org/10.1016/j.nima.2004.07.180. Proceedings of the 10th International Vienna Conference on Instrumentation.

[48] F. Gaede, S. Aplin, R. Glattauer, C. Rosemann, and G. Voutsinas. *Track reconstruction at the ILC: the ILD tracking software*. Journal of Physics: Conference Series, 513(2):022011, jun 2014. http://dx.doi.org/10.1088/1742-6596/513/2/022011.

[49] ALICE Collaboration. *Performance of the ALICE experiment at the CERN LHC*. International Journal of Modern Physics A, 29(24):1430044, Sep 2014. http://dx.doi.org/10.1142/s0217751x14300440.

[50] Ties Behnke and Ralf Diener and Ulrich Einhaus and Uwe Krämer and Paul Malek and Oliver Schäfer and Mengqing Wu. *Recent Performance Studies of the GEM-based TPC Readout (DESY Module)*, 2020. http://arxiv.org/abs/2006.08562.

[51] W. Blum, W. Riegler, and L. Rolandi. *Particle Detection with Drift Chambers*. Springer-Verlag Berlin Heidelberg, 2008.

[52] I. Lehraus. *Progress in particle identification by ionization sampling*. Nuclear Instruments and Methods in Physics Research, 217(1):43–55, 1983. http://dx.doi.org/https://doi.org/10.1016/0167-5087(83)90108-4.

[53] M. Hauschild. *Gaseous Tracking and dE/dx at Future Colliders*, talk presented at the Universität Bonn Seminar über Teilchenphysik, 2007, https://web.physik.uni-bonn.de/events/events.php?id=92.

[54] M. Hauschild. *Progress in dEdx techniques used for particle identification*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 379(3):436–441, 1996. http://dx.doi.org/https://doi.org/10.1016/0168-9002(96)00607-9. Proceedings of the Sixth International Conference on Instrumentation for Experiments at e+ e- Colliders.

[55] V. Davidenko, B. Dolgoshein, V. Semenov, and S. Somov. *Measurements of the relativistic increase of the specific primary ionization in a streamer chamber*. Nuclear Instruments and Methods, 67(2):325–330, 1969. http://dx.doi.org/https://doi.org/10.1016/0029-554X(69)90467-4.

[56] M. Hauschild. *dE/dx and Particle ID Performance with Cluster Counting*, talk presented at the ILC Ws. Valencia, 2006, http://ific.uv.es/~ilc/ECFA-GDE2006/.

[57] A. Bamberger, K. Desch, U. Renz, M. Titov, N. Vlasov, P. Wienemann, and A. Zwerger. *Readout of GEM detectors using the Medipix2 CMOS pixel chip*.

Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 573(3):361–370, Apr 2007. http://dx.doi.org/10.1016/j.nima.2006.12.049.

[58] T. Behnke and F. Gaede. *Software for the International Linear Collider: Simulation and reconstruction frameworks.* Pramana, 69:1089–1092, 2007. http://dx.doi.org/10.1007/s12043-007-0233-z.

[59] *Homepage of iLCSoft.* https://ilcsoft.desy.de/portal/. [online; accessed 06-April-2021].

[60] *Code base of iLCSoft.* https://github.com/iLCSoft. [online; accessed 06-April-2021].

[61] M. Frank, F. Gaede, M. Petric, and A. Sailer. *AIDASoft/DD4hep*, Oct. 2018. http://dx.doi.org/10.5281/zenodo.592244. webpage: http://dd4hep.cern.ch/.

[62] F. Gaede, T. Behnke, N. Graf, and T. Johnson. *LCIO: A Persistency framework for linear collider simulation studies.* eConf, C0303241:TUKT001, 2003. http://arxiv.org/abs/physics/0306114. SLAC-PUB-9992, CHEP-2003-TUKT001.

[63] E. Bertin and S. Arnouts. *SExtractor: Software for source extraction.* A&AS, 117:393–404, June 1996. http://dx.doi.org/10.1051/aas:1996164.

[64] A. Deisting. *Rekonstruktion von Ladungsdepositionen in einer GEM-Pixel-TPC mit einem Algorithmus aus der Astronomie.* Bachelor thesis, Universität Bonn, 2012. https://web.physik.uni-bonn.de/group/view.php?&group=1&lang=de&c=t&id=28.

[65] M. Thomson. *Particle flow calorimetry and the PandoraPFA algorithm.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 611(1):25–40, Nov 2009. http://dx.doi.org/10.1016/j.nima.2009.09.009.

[66] R. Brun and F. Rademakers. *ROOT — An object oriented data analysis framework.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 389(1):81–86, 1997. http://dx.doi.org/https://doi.org/10.1016/S0168-9002(97)00048-X. New Computing Techniques in Physics Research V.

[67] *Code base of ROOT v6.08/06.* https://github.com/root-project/root/releases/tag/v6-08-06. [online; accessed 22-April-2021].

[68] A. Hocker et al. *TMVA - Toolkit for Multivariate Data Analysis.* 3 2007. http://arxiv.org/abs/physics/0703039.

[69] B. P. Roe, H.-J. Yang, J. Zhu, Y. Liu, I. Stancu, and G. McGregor. *Boosted decision trees, an alternative to artificial neural networks*. Nucl. Instrum. Meth. A, 543(2-3):577–584, 2005. http://dx.doi.org/10.1016/j.nima.2004.12.018.

[70] C. Peterson, T. Rognvaldsson, and L. Lonnblad. *JETNET 3.0: A Versatile artificial neural network package*. Comput. Phys. Commun., 81:185–220, 1994. http://dx.doi.org/10.1016/0010-4655(94)90120-1.

[71] D. Schulte. *Study of Electromagnetic and Hadronic Background in the Interaction Region of the TESLA Collider*. Dissertation, Universität Hamburg, Hamburg, 1997. https://inspirehep.net/literature/888433.

[72] D. Schulte. *Beam-beam simulations with Guinea-Pig*. eConf, C980914:127–131, 1998, https://inspirehep.net/literature/470649.

[73] W. Kilian, T. Ohl, and J. Reuter. *WHIZARD—simulating multi-particle processes at LHC and ILC*. The European Physical Journal C, 71(9), Sep 2011. http://dx.doi.org/10.1140/epjc/s10052-011-1742-y.

[74] M. Moretti, T. Ohl, and J. Reuter. *O'Mega: An Optimizing Matrix Element Generator*, 2001. http://arxiv.org/abs/hep-ph/0102195.

[75] T. Sjöstrand, S. Mrenna, and P. Skands. *PYTHIA 6.4 physics and manual*. Journal of High Energy Physics, 2006(05):026–026, May 2006. http://dx.doi.org/10.1088/1126-6708/2006/05/026.

[76] S. Agostinelli et al. *Geant4—a simulation toolkit*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 506(3):250 – 303, 2003. http://dx.doi.org/https://doi.org/10.1016/S0168-9002(03)01368-8.

[77] A. Miyamoto and H. Ono. *ILD MC production for detector optimization*, 2019. http://arxiv.org/abs/1902.02516.

[78] H. Ono and A. Miyamoto. *Status of ILD new 250 GeV common MC sample production*. https://indico.cern.ch/event/995633/contributions/4259664/. [contribution to LCWS2021; online; accessed 15-April-2021].

[79] M. De Palma et al. *ALEPH: Technical Report 1983*. 4 1983.

[80] W. Bartl et al. *DELPHI: Technical Proposal*. 5 1983.

[81] J. W. Harris. *The STAR experiment at the relativistic heavy ion collider*. Nucl. Phys. A, 566:277C–285C, 1994. http://dx.doi.org/10.1016/0375-9474(94)90633-5.

[82] S. Beole et al. *ALICE technical design report: Detector for high momentum PID*. 8 1998.

[83] M. Lupberger. *The Pixel-TPC: first results from an 8-InGrid module.* Journal of Instrumentation, 9(01):C01033, 2014, http://stacks.iop.org/1748-0221/9/i=01/a=C01033.

[84] X. Llopart. *TIMEPIX Manual v1.0*, 2006.

[85] O. Fedorchuk. *Investigations of the long-term stability of a Gas Electron Multipliers and double hit resolution for the highly granular Time-Projection Chamber.* Dissertation, Universität Hamburg, Hamburg, 2019. https://ediss.sub.uni-hamburg.de/handle/ediss/8760, Dissertation, Universität Hamburg, 2019.

[86] M. Lupberger, J. Bilevych, K. Desch, T. Fischer, T. Fritzsch, J. Kaminski, K. Kohl, M. Rogowski, J. Tomtschak, and H. van der Graaf. *InGrid: Pixelated Micromegas detectors for a pixel-TPC.* PoS, TIPP2014:225, 2014.

[87] M. Lupberger. *Preliminary results from the 160 InGrid test beam*, talk presented at LCTPC-WG Meeting, 2015, http://agenda.linearcollider.org/event/6873/.

[88] J. Timmermans. *Progress with pixelised readout of gaseous detectors*, talk presented at ALCPG Eugene, 2011, http://pages.uoregon.edu/jimbrau/LC/alcpg11/.

[89] C. Ligtenberg, Y. Bilevych, K. Desch, H. van der Graaf, M. Gruber, F. Hartjes, K. Heijhoff, J. Kaminski, P. Kluit, N. van der Kolk, and et al. *Performance of the GridPix detector quad.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 956:163331, Mar 2020. http://dx.doi.org/10.1016/j.nima.2019.163331.

[90] X. Llopart, R. Ballabriga, M. Campbell, L. Tlustos, and W. Wong. *Timepix, a 65k programmable pixel readout chip for arrival time, energy and/or photon counting measurements.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 581(1):485–494, 2007. http://dx.doi.org/https://doi.org/10.1016/j.nima.2007.08.079. VCI 2007.

[91] X. Llopart, M. Campbell, R. Dinapoli, D. San Segundo, and E. Pernigotti. *Medipix2: A 64-k pixel readout chip with 55-/spl mu/m square elements working in single photon counting mode.* IEEE Transactions on Nuclear Science, 49(5):2279–2283, 2002. http://dx.doi.org/10.1109/TNS.2002.803788.

[92] M. Campbell, E. H. M. Heijne, G. J. Meddeler, E. Pernigotti, and W. Snoeys. *A Readout chip for a 64 x 64 pixel matrix with 15 bit single photon counting.* IEEE Trans. Nucl. Sci., 45:751–753, 1998. http://dx.doi.org/10.1109/23.682629.

[93] T. Poikela, J. Plosila, T. Westerlund, M. Campbell, M. D. Gaspari, X. Llopart, V. Gromov, R. Kluit, M. van Beuzekom, F. Zappon, V. Zivkovic, C. Brezina, K. Desch, Y. Fu, and A. Kruth. *Timepix3: a 65K channel hybrid pixel readout*

chip with simultaneous ToA/ToT and sparse readout. Journal of Instrumentation, 9(05):C05013–C05013, may 2014. http://dx.doi.org/10.1088/1748-0221/9/05/c05013.

[94] R. Ballabriga, J. Alozy, G. Blaj, M. Campbell, M. Fiederle, E. Frojdh, E. H. M. Heijne, X. Llopart, M. Pichotka, S. Procz, L. Tlustos, and W. Wong. *The Medipix3RX: a high resolution, zero dead-time pixel detector readout chip allowing spectroscopic imaging.* Journal of Instrumentation, 8(02):C02016–C02016, feb 2013. http://dx.doi.org/10.1088/1748-0221/8/02/c02016.

[95] *Homepage of the Medipix Collaborations.* https://medipix.web.cern.ch/. [online; accessed 14-June-2021].

[96] R. Ballabriga, M. Campbell, and X. Llopart. *Asic developments for radiation imaging applications: The medipix and timepix family.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 878:10–23, 2018. http://dx.doi.org/https://doi.org/10.1016/j.nima.2017.07.029. Radiation Imaging Techniques and Applications.

[97] R. Ballabriga, M. Campbell, and X. Llopart. *An introduction to the Medipix family ASICs.* Radiation Measurements, 136:106271, 2020. http://dx.doi.org/https://doi.org/10.1016/j.radmeas.2020.106271.

[98] X. Llopart. *The Timepix family ASICs*, 2017. https://indico.cern.ch/event/627245/contributions/2676748/attachments/1522549/2379247/Timepix_Vertex2017_XL.pdf. Talk at: The 26th International Workshop on Vertex Detectors [online; accessed 14-June-2021].

[99] F. Corsi, R. Dinapoli, P. Lamanna, and C. Marzocca. Design and characterization of a DAC for the slow control of the pixel chip. In *6th Workshop on Electronics for LHC Experiments*, pages 560–564, 9 2000.

[100] F. Krummenacher. *Pixel detectors with local intelligence: an IC designer point of view.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 305(3):527–532, 1991. http://dx.doi.org/https://doi.org/10.1016/0168-9002(91)90152-G.

[101] S. Martoiu, H. Muller, and J. Toledo. Front-end electronics for the Scalable Readout System of RD51. In *2011 IEEE Nuclear Science Symposium Conference Record*, pages 2036–2038, Oct 2011.

[102] S. Martoiu, H. Muller, A. Tarazona, and J. Toledo. *Development of the scalable readout system for micro-pattern gas detectors and other applications.* Journal of Instrumentation, 8(03):C03015–C03015, Mar 2013. http://dx.doi.org/10.1088/1748-0221/8/03/c03015.

[103] M. Lupberger. *The Scalable Readout System*, 2018. https://indico.cern.ch/event/722363/contributions/3031273/attachments/1692586/2723568/EIC_Tracker_Wokrshop_MLupberger_SRS_3.pdf. Talk at: EIC Tracker Workshop [online; accessed 30-June-2021].

[104] M. Lupberger, K. Desch, and J. Kaminski. *Implementation of the Timepix ASIC in the Scalable Readout System.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 830:75–81, 2016. http://dx.doi.org/https://doi.org/10.1016/j.nima.2016.05.043.

[105] M. Lupberger. *The Pixel-TPC: A feasibility study.* PhD thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, 2016. https://hdl.handle.net/20.500.11811/6760.

[106] K. Johal, S. Lamprecht, and H. Roberts. *Electroless nickel/electroless palladium/immersion gold plating process for gold-and aluminum-wire bonding designed for high-temperature applications.* 06 2018, https://www.researchgate.net/publication/228881159_Electroless_nickelelectroless_palladiumimmersion_gold_plating_process_for_gold-and_aluminum-wire_bonding_designed_for_high-temperature_applications.

[107] A. Münnich. *Simulation Studies for a High Resolution Time Projection Chamber at the International Linear Collider.* PhD thesis, RWTH Aachen, 2007. http://publications.rwth-aachen.de/record/52595, RWTH-CONV-114808.

[108] D. C. Wells, E. W. Greisen, and R. H. Harten. *FITS - a Flexible Image Transport System.* A&AS, 44:363–370, 1981, https://ui.adsabs.harvard.edu/abs/1981A%26AS...44..363W/abstract.

[109] B. W. Holwerda. *Source Extractor for Dummies*, Manual. http://mensa.ast.uct.ac.za/~holwerda/SE/Manual.html.

[110] C. Ligtenberg. *A GridPix TPC readout for the ILD experiment at the future International Linear Collider.* PhD thesis, Vrije Universiteit Amsterdam, Amsterdam, 2021. in prep.

[111] C. Ligtenberg. *Private conversation, email.*

[112] E. Bertin. Mining pixels: The extraction and classification of astronomical sources. In A. J. Banday, S. Zaroubi, and M. Bartelmann, editors, *Mining the Sky*, pages 353–371, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.

[113] H. Sert. *Light Higgsinos at the ILC: Precision Measurements and Detector Requirements.* PhD thesis, Universität Hamburg, Hamburg, 2015. http://dx.doi.org/10.3204/DESY-THESIS-2016-001.

[114] F. Gaede. *TOF for PID*, Presentation, 2018. https://agenda.linearcollider.org/event/7760/contributions/40943/attachments/32768/49846/gaede_ild_tof.pdf.

[115] N. Weinhold. *Private conversation.*

[116] E. Fernandez. *Collider Physics.* PoS, silafae-III:008, 2000. http://dx.doi.org/10.22323/1.005.0008.

[117] P. Abreu et al. *Measurement of |V(cs)| using W decays at LEP-2.* Phys. Lett. B, 439:209–224, 1998. http://dx.doi.org/10.1016/S0370-2693(98)01061-2.

[118] T. Suehara and T. Tanabe. *LCFIPlus: A Framework for Jet Analysis in Linear Colider Studies.* Nucl. Instrum. Meth., A808:109–116, 2016. http://dx.doi.org/10.1016/j.nima.2015.11.054.

[119] B. Dudar, J. List, U. Einhaus, and R. Ete. *Prospects of fast timing detectors for particle identification at future Higgs factories*, 2021. http://arxiv.org/abs/2105.12495.

[120] E. C. Gil, E. M. Albarrán, E. Minucci, G. Nüssle, S. Padolski, P. Petrov, N. Szilasi, B. Velghe, G. Georgiev, V. Kozhuharov, and et al. *The beam and detector of the NA62 experiment at CERN.* Journal of Instrumentation, 12(05):P05025–P05025, May 2017. http://dx.doi.org/10.1088/1748-0221/12/05/p05025.

[121] H. J. Hilke. *Time Projection Chambers.* Rep. Prog. Phys., 73:116201. 62 p, Oct 2010. http://dx.doi.org/10.1088/0034-4885/73/11/116201.

[122] P. Malek. *PhD Thesis in preparation.* Dissertation, Universität Hamburg, Hamburg, 2021. Dissertation, Universität Hamburg, 2021.

[123] S. Bilokin, R. Pöschl, and F. Richard. *Measurement of b quark EW couplings at ILC*, 2017. http://arxiv.org/abs/1709.04289.

[124] A. Irles-Quiles, R. Poeschl, and F. Richard. *Determination of the electroweak couplings of the 3rd generation of quarks at the ILC.* PoS, EPS-HEP2019:624, 2020. http://dx.doi.org/10.22323/1.364.0624.

[125] Y. Radkhorrami. *Heavy Flavour Jet Reconstruction with the ILD Detector Concept*, 2019. https://indico.desy.de/event/24272/contributions/53215/attachments/34350/42953/Yasser_Radkhorrami_Reconstruction_of_Heavy_Flavours_at_ILC.pdf. Talk at: 13th Annual Meeting of the Helmholtz Alliance 'Physics at the Terascale' [online; accessed 12-August-2021].

[126] K. S. Sasikumar. *Photon-Photon processes at the International Linear Collider and BSM signatures with small mass differences.* Dissertation, Universität Hamburg, Hamburg, 2020. https://ediss.sub.uni-hamburg.de/handle/ediss/8876, Dissertation, Universität Hamburg, 2020.

Acknowledgements

I would like say thanks to a number of people who have all contributed to this work in one way or another.

My primary gratitude goes to my supervisors, my professor Johannes Haller and my day-to-day supervisor Ties Behnke, who always had an open ear and supported me over these last years, providing direction for the work and space for me to decide how to execute it.

Furthermore, my thanks goes to Ralf Diener and Jenny List for their supervision and lots of feedback on the TPC and the analysis part, repectively, as well as the occasional more personal exchange.

I am also grateful to my collaborators in the Ropperi project, Jochen Kaminski, Michele Caselle and Markus Gruber.

I would like to thank all of the FLC group, now FTX, the group leaders, the technicians, the postdocs and my fellow PhD students. The atmosphere always was and is one of mutual support and companionship, and I could not have done it without you guys! I can only fail in an attempt to name everyone, so my thanks goes out to all of you for making DESY a home.

Nevertheless, I would like to thank Oliver and Remi for their software support and their endurance in keeping up with my occasional IT frustration.

A special thanks goes to my proof-readers, Ralf and Jenny, was well as Elke, Markus, Möller and Hannes.

I'd like to thank coffee for keeping my body up and music for keeping my spirit up.

My thanks goes out to all my friends outside of physics, again too many to mention and do justice. We spent many a day and even more nights with games and laughter, with glasses and fun. We spent weeks of lockdown online together in solitude, keeping each other sane. Thanks for all of you!

Schlussendlich geht mein Dank an meine Familie, für meinen Ruhepunkt zu hause und die stete, unbedingte Unterstützung. An meinen Bruder, meine Mutter und meine Oma.

Eidesstattliche Versicherung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Hamburg, September 2021 Ulrich Einhaus