Optimisation of the COMET experiment to search for charged lepton flavour violation and a new simulation to study the performance of the EMMA FFAG accelerator

Richard Thomas Patrick D'Arcy

University College London

Submitted to University College London in fulfilment of the requirements for the award of the degree of Doctor of Philosophy, 21$^{\text{th}}$ September 2012

I, Richard D'Arcy, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

The particle tracking software package, GPT, has been developed and utilised to simulate the beam optics of the EMMA injection line and ring, constructed at the Daresbury Laboratory, UK. EMMA is a proof-of-principle machine for a new type of accelerator: a non-scaling fixed-field alternating gradient (ns-FFAG) accelerator. As such the beam dynamics of the magnetic lattice require benchmarking, with GPT chosen for its space-charge self-field simulation package. Tune and time-of-flight measurements have been successfully simulated and compared to experimental data, recorded during the first few runs of the machine. Measurements confirm the successful operation of EMMA as a ns-FFAG accelerator and simulations highlight that space-charge effects are observable in the EMMA bunch-charge and energy regime.

Such accelerators have many applications within and outside particle physics, ranging from cancer therapy and accelerator driven thorium reactors to neutrino factories and muon colliders. The application of FFAGs and the design of the COMET/PRISM experiment, which is seeking to measure muon-to-electron conversion at the 10^{-18} level, is investigated. Simulations of the COMET experiment, staged in two phases, have been performed with a focus on optimisation of the stopping target design. A number of geometries have been tested, with a cone then disk structure preferred for Phase-I then Phase-II respectively. Initial data from the COMET precursor experiment, MuSIC, have also been analysed and successfully compared to simulation.

Acknowledgements

First and foremost I would like to thank my supervisor, mentor, and friend, Mark Lancaster, for a robust firsthand account of the 'Physics Skills Set 101', providing support in more ways than can be mentioned, and dropping a tonne on me in the JB. I would also like to thank Matthew Wing and Bruno Muratori for all the patience, knowledge, and supervision they have shown me over the years.

Thanks go to all the members of the UCL HEP (and CMMP) group, especially the D109 boys, for the apparently inexorable coffee breaks, Wasabi Wednesdays, and lost nights in the pub. I would also like to thank my fellow PhD students Dan, James, Tom, and Matt for all the camaraderie we've shared over the past four years. It's been a pleasure sharing this journey with you and arrant luck I've finished it with you too.

A special thank you goes to all the friends who've helped me to the end despite not having a clue what it is I've been doing all this time. There are far too many to name but special thanks must go to The Deuce, Rick Nollins, The White Rhino, SG1, The Octagon, and Noel for the indispensable opportunities of hedonism and profligacy. My bank balance may not appreciate it but I certainly do.

Finally, I would like to thank my Mum, Dad, and Smels for more than can be expressed here. Without their unwavering love and support I wouldn't be where I am today. This PhD is dedicated to them.

Contents

List of figures

List of tables

Chapter 1

The Standard Model

The Standard Model (SM) of particle physics [1, 2, 3] is a unified mathematical description of the fundamental components of matter and their interactions mediated by the strong, weak, and electromagnetic forces. It has been very successful in describing these interactions for many years, and has stood up to a number of very rigorous tests across a range of experiments.

The SM is a $SU(3) \otimes SU(2) \otimes U(1)$ gauge theory, where conventionally a Lagrangian describes the behaviour of a set of quantised fields of the twelve fermionic particles and five bosons known to exist.

This SM describes many observable properties of the Universe. All the properties of fundamental particles (with the exception of neutrino mass) and forces (excluding gravity, which is not included in the SM) are successfully described. Conservation of quantities such as energy, angular momentum, and electric charge are accounted for by the symmetries of the theory. The three forces chosen as a mathematical representation of the fields and forces of the SM account for an exceptional number of complex phenomena e.g.

- **Electromagnetism** is visible at the macroscopic level via the effects of electric and magnetic fields but also at the quantum level through the coherence of atoms and e.g. corrections to magnetic moments.

- The **weak force** dictates a number of decay processes, including the decomposition of radioactive nuclei and heavy quarks.

- **Strong** interactions are evident in the cohesion of quarks and gluons into protons and neutrons, comprising atomic nuclei.

The SM also describes the phenomenology of the extensive particle zoology currently observed. The interactions of these particles can be represented using a simple diagrammatic approach, first expounded by Feynman and Stueckelberg in 1948 [4]. In this approach complex calculations in field theory are often equivalent to a mathematical integration across a number of interaction points. An example of such a process, represented using a Feynman diagram, can be seen in Fig. 1.1. This diagram represents only the leading order term of an expansion for this process. The observable process is a sum of these higher order diagrams which generally converge (e.g. for electroweak processes) and so the calculation of the first few diagrams is sufficient. However, the SM contribution to charged lepton flavour violation (cLFV) is only non-zero at loop-level (see Fig. 1.4 for an example of a cLFV loop-level Feynman diagram). Even at loop-level the rates are extremely small, with a branching ratio (BR) of $\mathcal{O}(10^{-54})$. This topic of cLFV will be discussed further in Sec. 1.3.

The complete set of observed particles that interact in such a way, represented by the SM, can be found in Tab. 1.1. The complete phenomenology of the SM

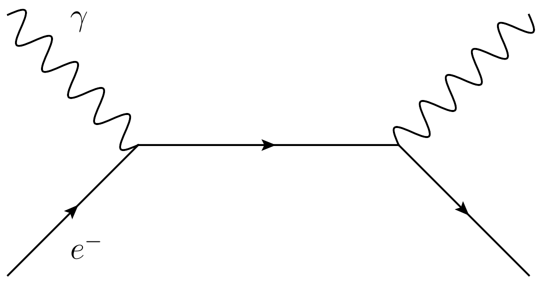

Figure 1.1: Leading order s-channel electron-photon scattering, $e + \gamma \to e + \gamma$, represented by a Feynman diagram. The positive arrow of time goes from left to right.

is far too rich to describe here in detail. However, the set of particles can be subdivided into groups: fermions (leptons and quarks) with half-integer spin, each with three generations containing two particles separated by a unit of charge, and gauge bosons which mediate the forces.

1.1 Shortcomings of the Standard Model

The SM is chiefly empirically derived and therefore rooted in experimental data. It remains a model however, and as such has its limitations. One of the most flagrant limitations is its inability to incorporate massive neutrinos. The SM defines neutrinos as having zero mass, however recent experimental data clearly contradicts this [5]. This ad-hoc extension of the SM to include neutrino masses (via Yukawa coupling to the Higgs) and neutrino oscillations is referred to as the νSM henceforth. Neutrinos are notoriously difficult to detect, as they only interact via the weak force, but recent experiments [6] have observed oscillations

Particles		**Spin**	**Charge**		
			Electric	Weak	Strong
Up-type Quarks	u, c, t	$\frac{1}{2}$	$+\frac{2}{3}$	Yes	Yes
Down-type Quarks	d, s, b	$\frac{1}{2}$	$-\frac{1}{3}$	Yes	Yes
Leptons	e, μ, τ	$\frac{1}{2}$	$+1$	Yes	No
Neutrinos	ν_e, ν_μ, ν_τ	$\frac{1}{2}$	0	Yes	No
Weak Bosons	$W^{\pm}$	1	±1	Yes	No
	Z	1	0	Yes	No
Gluon	g	1	0	No	Yes
Photon	γ	1	0	Yes	No
Higgs Boson	H	0	0	Yes	No

Table 1.1: The complete set of particles described by the Standard Model. There also exists a set of corresponding anti-fermions e.g. e^+, $\bar{\nu_e}$, etc. not denoted in this breakdown.

in the weak eigenstates that is interpreted as splittings in the mass eigenstates. This suggests the existence of at least two neutrinos with a non-zero mass. Like other fermions they could be Dirac particles, however they could be their own anti-particle and thus obey the Majorana equation [7] and not the Dirac equation.

The question of neutrino masses links to the flavour problem of the SM, of which there are many facets: undetermined fermion masses and mixing angles; CP violating phases and their origin; the observed smallness of flavour changing neutral currents. Some of these problems have been solved (e.g. the ad hoc setting to zero of CP violation in the QCD sector via the axion) but many remain outstanding.

The mechanism of electroweak symmetry breaking within the SM is an outstanding problem. Gauge-invariance dictates that the mediating gauge bosons are massless, thus operating over an infinite range. In reality this symmetry does not exist and cannot be broken explicitly without gauge-invariance being compromised. It must therefore be broken in another way and the SM accounts for this via the Higgs mechanism, which predicts the existence of a scalar boson - the Higgs particle. Recent observations of a new boson by ATLAS and cmS are consistent with the Higgs hypothesis [8, 9, 10].

The lack of a description of gravity within the SM is also a concerning limitation. At present string theory is the only known renormalisable QFT with the ability to merge gravity with general relativity, however there is no experimental evidence for it. A unification of gravity with the other three forces would require a fundamental overhaul of particle physics (e.g. an introduction of extra spatial dimensions); the three gauge interactions of the SM which define the electromagnetic, weak, and strong interactions, would likely need a new symmetry.

1.2 Beyond the Standard Model

As mentioned in Sec. 1.1, it is widely believed that new physics beyond the SM is required to accommodate non-zero neutrino masses. As neutrinos may be their own anti-particle, obeying the Majorana equation, the most elegant mechanism for generating small neutrino masses is the seesaw mechanism [11]. This requires the introduction of extra singlets to the SM Lagrangian. A consequence of this is the violation of individual lepton-flavour number, allowing processes such as $\mu \to e\gamma$ in the νSM. However, due to the small size of neutrino masses the rates for cLFV are extremely small in the νSM [12].

The gauge hierarchy problem [15] also hints at physics beyond the SM (BSM). Low-energy supersymmetry (SUSY) is one preferred candidate for a BSM solution to this problem. cLFV processes such as $\mu \to e\gamma$ are potentially amplified in such a framework, pushing the rates up to a level that may be in reach of near future experiments. In the νSM the rates for cLFV are proportional to the neutrino masses, whereas in SUSY models these processes are only suppressed by inverse powers of the SUSY breaking scale, typically at most $\mathcal{O}(1 \text{ TeV})$.

There are many different seesaw models which accommodate non-zero neutrino masses, both SUSY and non-SUSY based. The current and future experimental limits on $\mu \to e\gamma$ and $\mu N \to eN$ for a specific SUSY BSM scenario are shown in Fig. 1.2. These models demonstrate the complementarity of cLFV results to those of the LHC - in this model any future results from cLFV experiments will isolate values of the masses of the right handed boson m_{W_R} and heavy neutrino m_N, probing beyond the energy scales available to the LHC. As can be seen the current experimental limits e.g. at MEG [16] and SINDRUM [18], are beginning to probe these predictions.

An observation of just one cLFV process may indicate the mass scale of the BSM physics but it would not be enough to uncover the form of the new interaction. In fact, it would be necessary to observe cLFV in other processes to be able to

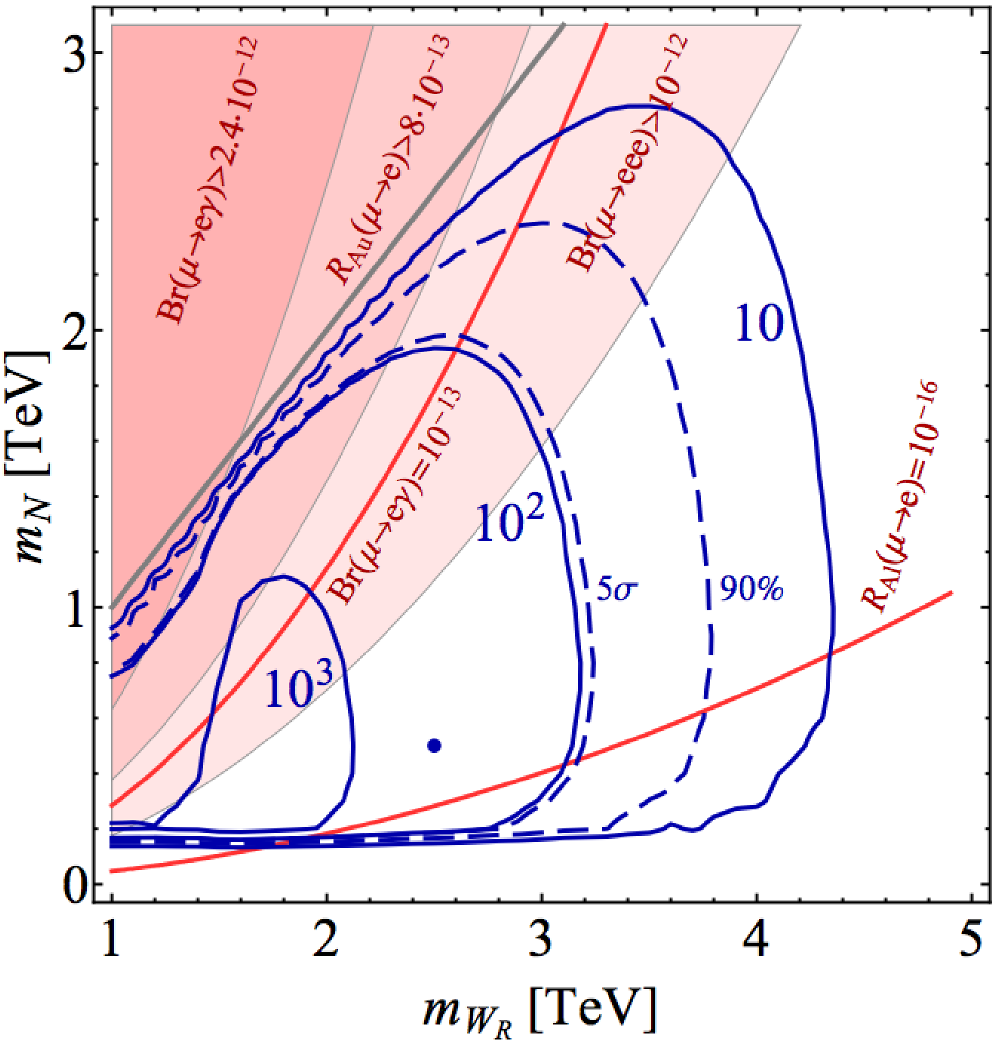

Figure 1.2: Comparison of LFV event rates at the LHC and in low energy rare decays. In order to verify the LHC sensitivity to the flavour couplings $V_{\mathrm{N_e}}$ and V_{N_μ}, the right-handed W boson mass and neutrino mass are fixed at $m_{W_R} = 2.5$ TeV and $m_N = 0.5$ respectively. This is denoted by the blue dot. The solid blue contours give the number of events for the LFV signature $e^{\pm}\mu^{\pm,\mp}$ plus jets at the LHC at $\sqrt{s} = 14$ TeV and an integrated luminosity of 30 fb^{-1}. The dashed contours define the parameter region with signals at 5σ and 90%. The shaded red areas denote the parameter regions excluded by current low energy LFV limits, whereas the red contours show the expected sensitivity of planned experiments [19].

obtain this information. Considering a generic Lagrangian for this new interaction, where Λ is the energy scale of the new interaction and κ the relative strength of two possible mechanisms (dipole and contact interactions), then Fig. 1.3 shows the experimental sensitivity in terms of these two parameters.

These parameters are related to specific BSM models e.g. for SUSY

$$\frac{1}{\Lambda^2} = \frac{g^2 e}{16\pi^2 M_{\mathrm{SUSY}}^2}\theta_{\tilde{e}\tilde{\mu}} \ , \tag{1.1}$$

where g is the SUSY interaction strength, e is the electromagnetic coupling constant, M_{SUSY} is the mass of some super-partner, and $\theta_{\tilde{e}\tilde{\mu}}$ is the mixing angle between selectrons and smuons [20]. Fig. 1.3 demonstrates that the $\mu \to e\gamma$ interaction is dominated by a dipole interaction, whereas $\mu N \to eN$ can occur via both interactions. It will therefore be necessary to observe, or place a rigid exclusion on, the $\mu \to e\gamma$ process in addition to an observation of $\mu N \to eN$ it will be necessary in order to determine the nature of the BSM model.

This demonstrates the potential reach of cLFV to energy scales far exceeding those available to the LHC, with low energy rare decays sensitive to BSM models with energy scales $\Lambda > 10,000$ TeV.

1.3 Lepton flavour violation

In the SM Lagrangian there are no interactions that change lepton flavour, contrasting the weak interactions for quarks which is based on observation. Each generation of lepton has an inherently conserved quantum number. As an example consider the electron generation, where $L_e = 1$ for e^-/ν_μ, $L_e = -1$ for $e^+/\bar{\nu}_\mu$, and $L_e = 0$ for all other particles. Processes involving this generation of lepton have $\Delta L_e = 0$ i.e. electron lepton flavour is conserved.

However, the SM does not explain all phenomena, as described in Sec. 1.1. In particular, the observation of neutrino oscillations requires extensions to the

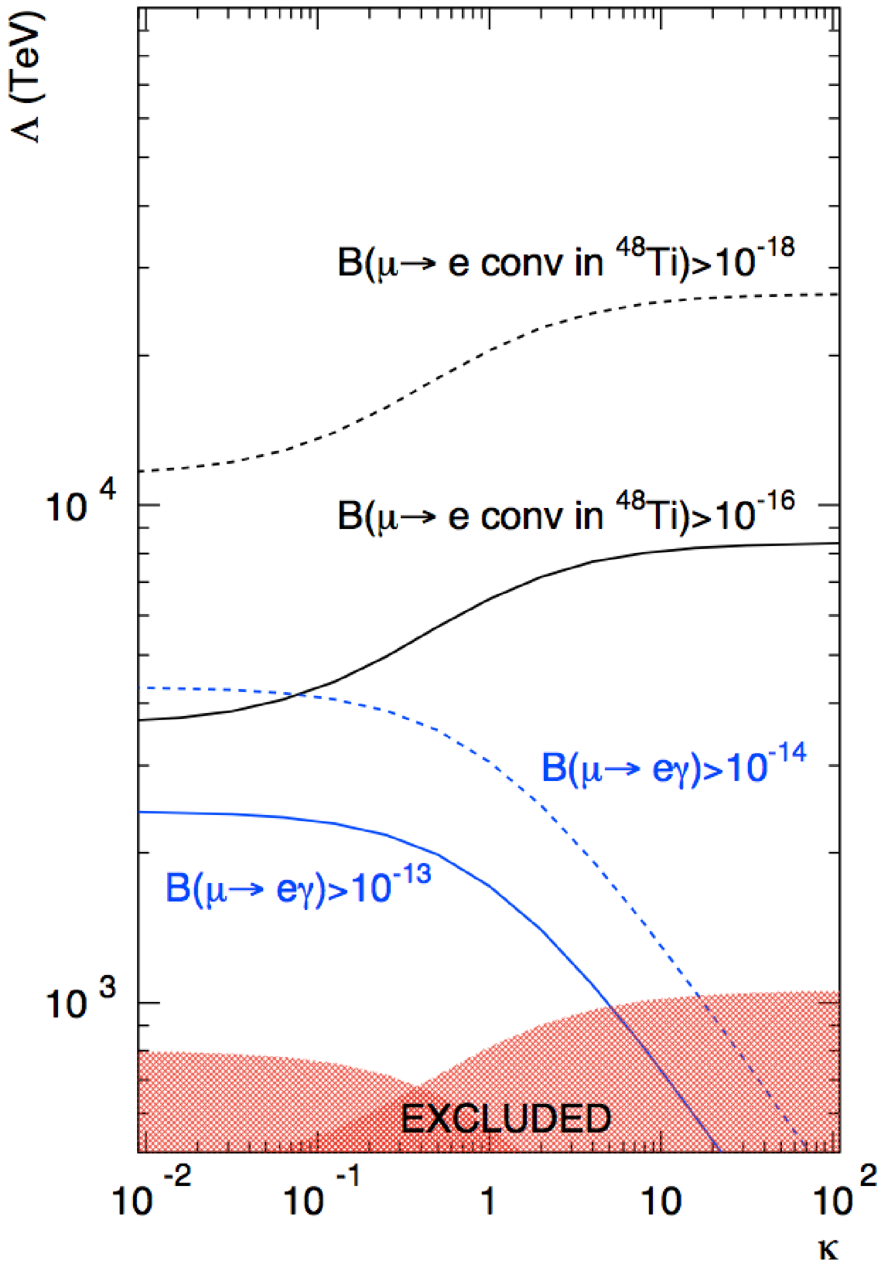

Figure 1.3: The relation of cLFV to two BSM parameters Λ and κ. In this model Λ is the energy scale of the new model and κ represents the relative strength of the process to a certain mechanism, with $\kappa \ll 1$ and $\kappa \gg 1$ indicating a fully dipole and four-fermion point interaction respectively [20].

SM Lagrangian. Such extensions result in a non-zero rate for cLFV processes, however the rate is still very small i.e. $\mathcal{O}(10^{-54})$. An example of such a lepton number violating process is

$$\mu^{\pm} \rightarrow e^{\pm}\gamma \ . \tag{1.2}$$

A Feynman diagram for this process can be seen in Fig. 1.4.

The calculation of branching ratios for these processes is determined by the neutrino masses and elements of the PMNS matrix, which are not yet all known. Using present experimental limits of the unknown parameters one can determine an upper limit of the BR in the νSM:

$$\mathrm{BR}(\mu \rightarrow e\gamma) = \frac{3\alpha}{32\pi}\left|\sum_{i=2,3} U^{*}_{\mu i} U_{ei} \frac{\Delta m^2_{1i}}{M^2_W}\right|^2 < 10^{-54} \ , \tag{1.3}$$

where $U_{\alpha i}$ are the elements of the neutrino mixing matrix and Δm^2_{1i} are the neutrino mass-squared differences. This is far too low to detect with any experiment since at least 10^{54} muons would be required to reach this branching ratio. However, when BSM theories are considered (see Sec. 1.2 for a breakdown of such models) the predictions for the rate of cLFV dramatically increases to a point where it may be possible to observe it in the near future.

These branching ratios are model dependent and the branching ratio within SUSY models can be anywhere from 10^{-14} to 10^{-25} [12]. However, any BSM results from the LHC will likely constrain such models significantly, leading to rather robust predictions of cLFV. The consistency of LHC results with cLFV will be one of the definitive and necessary checks in establishing a model of new physics beyond the SM. The LHC is a probe to high energy regimes but cannot easily measure slepton or SUSY mixing parameters [13]. Future experiments utilising cLFV and, for example, the anomalous muon magnetic moment $(g-2)$ [14] will

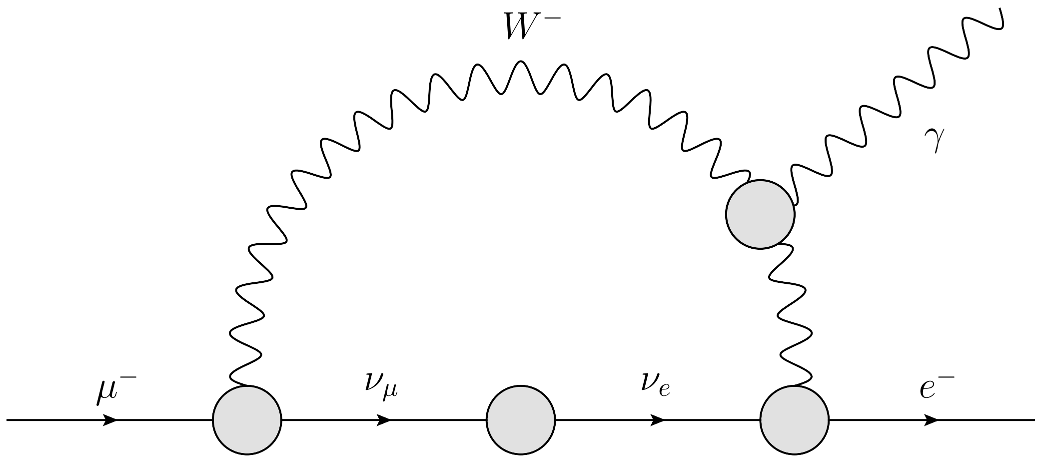

Figure 1.4: One possible charged lepton flavour violating process allowed within the SM.

therefore likely be required to isolate these parameters and resolve degeneracy between different theories of BSM physics.

1.4 Current cLFV experimental status

Since the first search for cLFV by Hincks and Pontecorvo in 1947 [17], experimental searches for cLFV have been continuously carried out with various elementary particles, such as muons and taus. The upper limits have been improved at an approximate rate of two orders of magnitude per decade, displayed in Fig. 1.5. The rate of improvement of the cLFV upper limit has slowed over the last decade but is about to take off again, as shown in Fig. 1.6 where the current, and possible near future, upper limits of various cLFV decays are shown.

The COMET experiment (described further in Ch. 3) proposes to achieve an upper limit on $\mu N \rightarrow eN$ of 10^{-16} by 2021, providing a drastic jump of four orders of magnitude on the current limit set by SINDRUM-II [18]. In order to analyse first results in the medium term, as well as gain operational experience with this type of experimental procedure, the COMET experiment will be staged with the

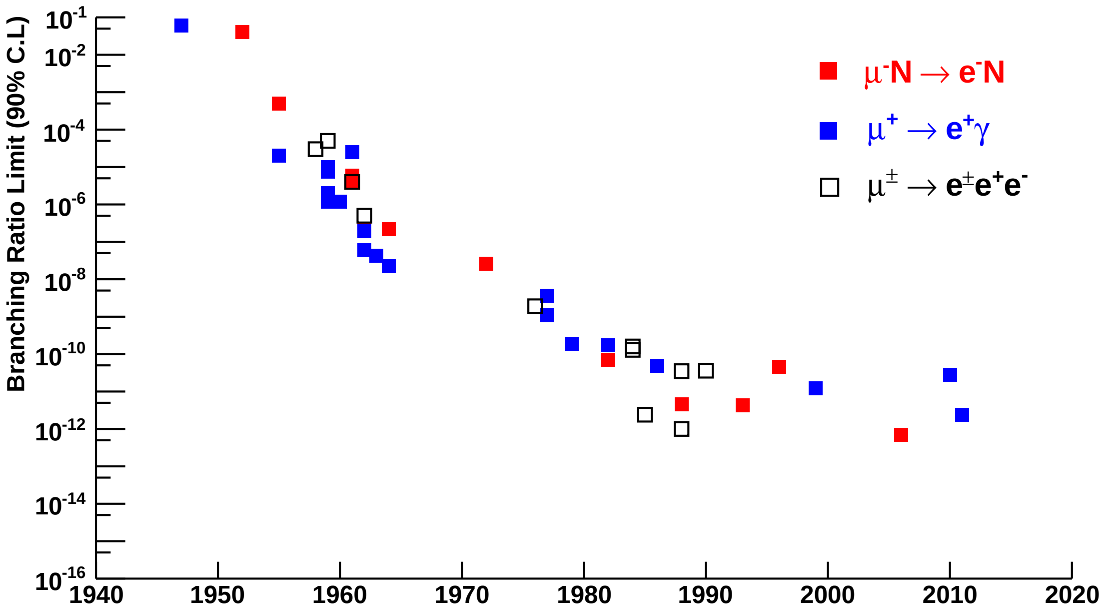

Figure 1.5: 90% confidence level upper limits on the branching ratio for three cLFV processes from 1947 to the present day.

upper limit on the BR of 10^{-14} expected to be probed by 2017. This Phase-I is described and simulated in Ch. 4.

The sensitivity of the muon system to cLFV is very high. This is owing to the large number of muons available for current experimental searches (approximately 10^{18} μ/yr proposed for COMET). The cross section of $\mu \rightarrow e\gamma$ is dependent on the lepton mass difference to the 4th power, thus changing the incident particle from a muon to tau leads to higher order rates of $\mathcal{O}(10^4)$ in the τ sector compared to the μ sector. The BR for $\mu N \rightarrow eN$ is defined as the ratio of coherent conversion to that of nuclear capture. For this reason the m_l^4 term present in both cross sections cancels, removing any intuitive suppression in the τ sector. However, the far greater yield of muons available suggests that this channel poses a deeper probe to cLFV with near-future experiments.

There are a number of channels to explore cLFV, with $\mu \rightarrow e\gamma$ and $\mu \rightarrow 3e$ the most studied [16]. The search for these processes is predicated on firing a high

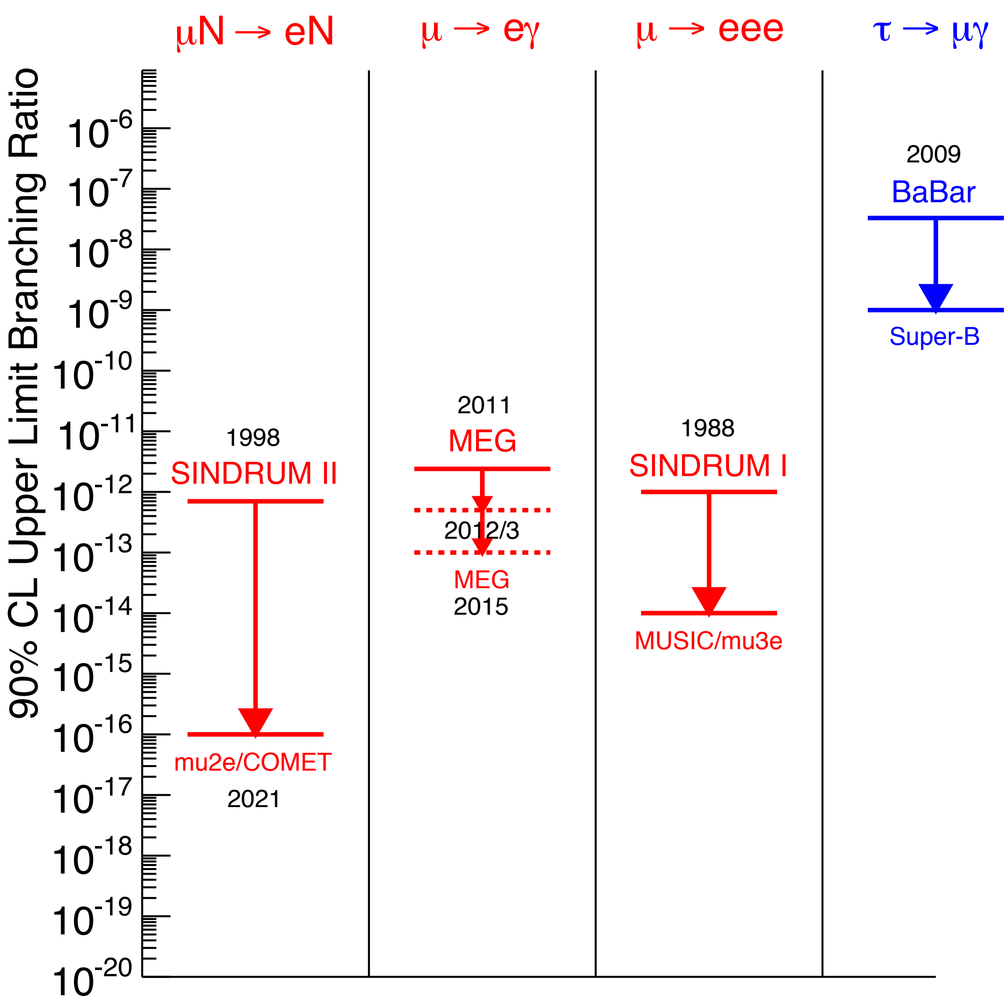

Figure 1.6: Current and future sensitivity limits for charged lepton flavour violating processes allowed within the SM.

intensity muon beam (3×10^8 μ/s at PSI [16]) onto a thin foil target, stopping as many muons as possible, then searching for the respective signal from the decay of the captured muon at rest. Both the signal and background are emitted isotropically from the target, travelling directly to the detector system. This leads to huge levels of accidental backgrounds of electrons (from non-lepton number violating processes, either in flight or at rest), therefore making it difficult to distinguish between signal and background.

The latest experimental upper limit for $\mu \to e\gamma$ at a 90% confidence level is 2.4×10^{-12}, recorded by the MEG collaboration at PSI, utilising a positive muon beam with up to 9.5×10^{15} μ/yr [16]. The experiment proposes to improve this sensitivity by two orders of magnitude through an increase in muon intensity of 100. However, the accidental background $N_{\text{acc.}}$ grows with the muon rate R_μ according to

$$N_{\text{acc.}} \propto R_\mu^2 \cdot \Delta t_{e\gamma} \cdot (\Delta\theta_{e\gamma})^2 \cdot (\Delta E_\gamma)^2 \ , \tag{1.4}$$

where $\Delta t_{e\gamma}$ is the timing resolution, $\Delta\theta_{e\gamma}$ is the anglular resolution between the electron and photon, and ΔE_γ is the photon energy [21]. For the case of $\mu \to e\gamma$ the accidental background of daughter particles therefore scales as R_μ^2. Thus at ever higher muon rates the signal becomes saturated by accidental background. For this reason other cLFV processes are considered.

The coherent neutrinoless muon-to-electron conversion process $\mu N \to eN$ represents an alternative to the aforementioned cLFV processes as it removes the necessity to detect coincidences between particles. This is achieved by searching for signal in a region boosted by nuclear recoil to a momentum range e.g. 95 MeV/c for gold, beyond the classical Michel decay limit of approximately 52.8 MeV. The SINDRUM-II experiment explored muon-to-electron conversion by firing a muon beam into a gold (and subsequently titanium) target, creating muonic atoms. No signal was detected and a 90% confidence level upper limit on the BR($\mu N \to eN$) of 3.3×10^{-13} was set [18], along with a comprehensive correlated background study. The primary source of background in the $\mu \to e$ signal momentum range is that of decay-in-orbit (DIO) electrons from $\mu^- + N(A, Z) \to e^-\nu_\mu\bar{\nu}_e + N(A, Z)$. This background, recorded by the SINDRUM-II experiment, along with simulated signal, can be seen in Fig. 1.7.

The detrimental effect of backgrounds can be mitigated by the implementation of novel beamline elements, particularly:

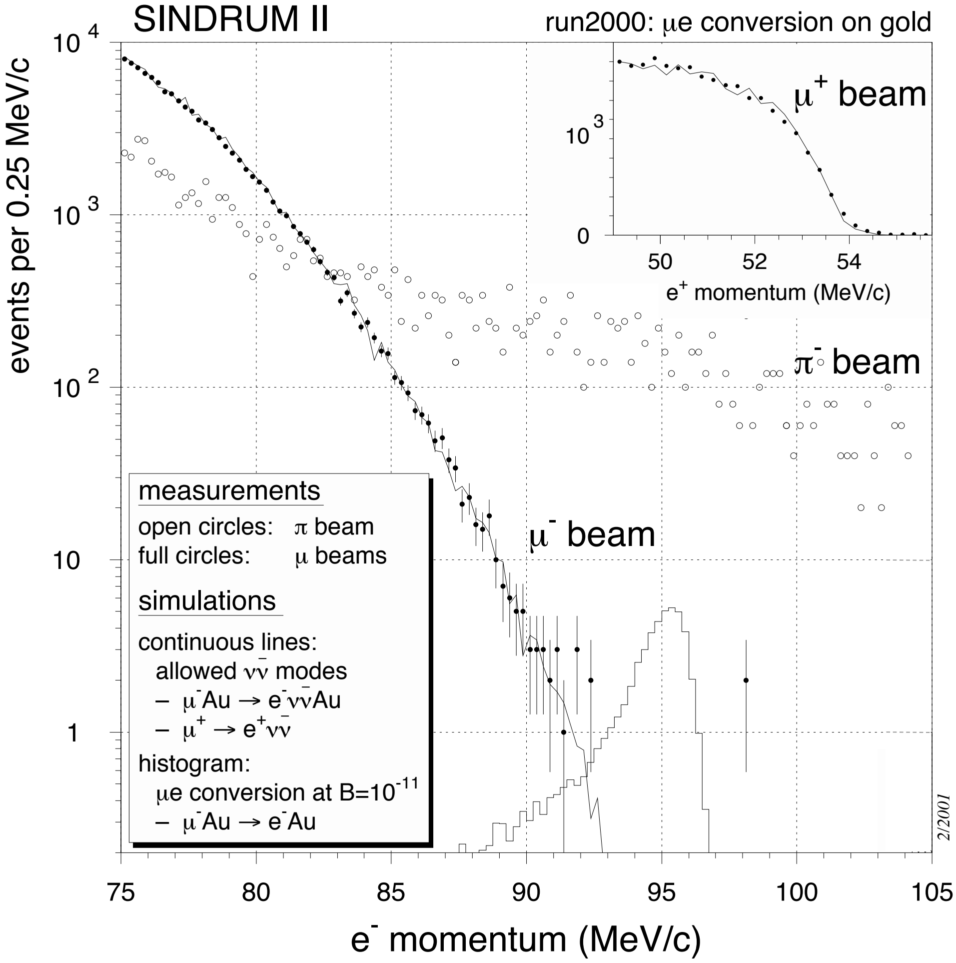

Figure 1.7: SINDRUM-II results showing the decay-in-orbit background tail at the signal region, as well as a simulated signal momentum distribution [18].

- The prompt background arising from protons leaking between bunches may be reduced by using a pulsed proton beam, alternating particle bunches with regions of proton extinction at ($\mathcal{O}(10^{-9})$).
- Muons, typically of $p > 70$ MeV/c, responsible for decay-in-flight electrons with $p > 100$ MeV/c in the $\mu \to e$ signal region, can be eliminated by implementing curved solenoids for momentum selection.

The COMET collaboration proposes an improvement on current sensitivities of up to a factor of 10^5 through the introduction of such elements. These improvements will be outlined in Ch. 3, along with simulations of the proposed experiment.

The DIO process is an intrinsic property of the muon, with the DIO endpoint and lifetime dependent on the Z number of the stopping target material used to create the muonic atoms. As the branching ratio limit decreases the signal may move within the limits of the DIO background spectrum. It is therefore necessary to minimise energy loss on the signal by reducing background and stopping more muons. The implementation of a phase-rotating FFAG ring is proposed to decrease the mean muon momentum and RMS to a level where almost all muons are stopped, as well as providing added decay length for energetic pions. This will drastically increase the number of stopped muons per proton with a much smaller muon stopping target, thus decreasing the amount of energy loss of signal electrons en route to the detector. These plans will be included in the PRISM project, increasing the sensitivity of $\mu \to e$ conversion to $\mathcal{O}(10^{-18})$. These advances, along with their research and design, will be expounded in Ch. 3.2.

As previously mentioned, like the DIO endpoint the lifetime of the muonic atom is also an intrinsic property of the stopping material. Precise measurement of the muonic lifetime, across a range of materials, is a key aspect in setting a suitable timing window for the detector to maximise the exclusion of prompt and delayed background. The muonic lifetime is inversely proportional to Z^4, which is in turn proportional to the nuclear capture rate [22]. The nucleon charge has two effects:

- The main effect is the electric attraction: the larger Z, the closer the muon is to the nucleus, and the higher the probability for the muon to be at the position of the nucleus.

- The size of the nucleus also effectively increases with Z. The larger the nucleus, the higher the probability to capture the muon.

The MUon Science Innovative Commission (MuSIC) experiment is designed as a precursor to COMET and PRISM, used to forerun and test the new beamline elements proposed for implementation in the experimental decrease of BR($\mu N \to eN$). These elements are tested and analysed in Ch. 5.

Chapter 2

The Physics of Particle Accelerators

Since the invention of the cyclotron in the 1930s [23], particle accelerators have played a major role in the research of physics at the subatomic level. The principle objectives of accelerator physics has been to reach ever higher energies and intensities; thus advancing the depth of knowledge of elementary particle physics. There are now over 20,000 particle accelerators in the world with uses as diverse as cancer therapy and food treatment.

A novel type of acceleration proposed to perform these functions, as well as a host of experimental particle physics objectives, is that of fixed-field alternating-gradient (FFAG) acceleration. A proof-of-principle non-scaling variant of this FFAG accelerator is detailed in Ch. 6, with a demonstration of the simulation and experimental works performed. The principle of FFAGs were first demonstrated over half a century ago. However, due to the difficulties in modelling the beam dynamics their development was stunted; an issue only recently overcome with the advent of advanced simulation software and sufficient CPU power.

In order to fully understand the nature of this complex method of acceleration, a brief overview of the beam dynamics of accelerators will be presented.

2.1 Beam dynamics and beam transport

The electromagnetic lattices of all accelerators are required to achieve beam steering and confinement along a desired trajectory. A description of beam dynamics within a confined magnetic lattice is given.

2.1.1 Beam focusing

A bunch of charged particles, when left unconstrained, will ultimately diverge due to self-field effects. However, it is this intrinsic charge that allows for beam manipulation towards a desired trajectory. As such, a beam of particles may be manipulated and steered by altering the orientation and strength of an applied magnetic field. Accelerating structures are therefore formed of a lattice of interspersed magnetic elements, at regular intervals along the beamline. When a particle enters a magnetic field it experiences a force F, related to the magnetic field strength B, given a charge q and velocity v:

$$F = qvB \quad . \tag{2.1}$$

A charged particle experiencing a uniform magnetic field will therefore follow a circular trajectory. The centripetal force required to give such a circular motion is given by:

$$F = \frac{mv^2}{\rho} = \frac{pv}{\rho} \quad , \tag{2.2}$$

for a particle with momentum p travelling along a circular trajectory of radius ρ. Combining this with Eq. 2.1 gives:

$$B\rho = \frac{p}{q} \quad . \tag{2.3}$$

This quantity $B\rho$ is referred to as the magnetic rigidity and defines the magnetic layout for a charged particle with a specific momentum [23].

e.g. The LHC uses the pre-existing LEP tunnel with a circumference of 27 km. Assuming a packing fraction of 65% the bending radius is therefore 2.8 km. With an ideal beam energy of 7 TeV the required magnet strength may be calculated using Eq. 2.3, such that

$$B = \frac{p}{e\rho} = \frac{10}{3} \cdot \frac{7000}{2804} = 8.3 \text{ T} \quad . \tag{2.4}$$

The simplest type of magnetic element operating under this principle is a dipole, which serves to bend charged particles in a perpendicular direction to the motion of the beam. A uniform field is achieved by placing two coils parallel to one another, with the coils having equal coil dimensions and identical currents. If the condition of a uniform dipole field is met then the particle will travel along a circular path governed by the relation in Eq. 2.3.

2.1.2 Linear beam optics

Dipoles are the simplest magnetic element but current accelerator optics generally requires other functions such as beam focussing and steering. The field of beam optics describes the fundamental physics of maintaining a nominal trajectory (by convention along the z-axis) by utilising all these magnetic functions. By expanding the magnetic field around the nominal trajectory, the magnetic field can be regarded as a sum of multipoles, each having a different effect on the beam. The relationship between the magnetic field and the multipoles is given by

$$\frac{e}{p} B_z(x) = \frac{1}{R} + kx + \frac{1}{2!} mx^2 + \frac{1}{3!} ox^3 + ... \, , \tag{2.5}$$

where the nature of the leading order multipoles is detailed in Table 2.1. Throughout this report only the two lowest multipoles are considered, therefore the beam steering is referred to as linear beam optics, owing to the bending forces being either constant or linear throughout.

By considering the trajectory of a single particle in a co-moving Cartesian co-ordinate system $K = (x, z, s)$, illustrated in Fig. 2.1 where the particle's origin moves along the orbit of the s-axis, it is possible to derive a set of linear equations of motion as the particle travels through the magnetic lattice:

$$x''(s) + \left(\frac{1}{R^2(s)} - k(s)\right) x(s) = \frac{1}{R(s)} \frac{\Delta p}{p} , \tag{2.6}$$

$$z''(s) + k(s) z(s) = 0 , \tag{2.7}$$

where Δp describes the momentum deviation away from the nominal momentum p, $1/R$ and k are the dipole and quadrupole strengths respectively, and x'' is the second derivative of the x co-ordinate with respect to the spatial co-ordinate s.

Eqs. 2.6 and 2.7 describe the propagation of one particle, but within the field of particle acceleration it is both difficult and undesirable to create a beam with a single particle. It is therefore necessary to extend this theory to describe a composite beam of many particles and in particular to understand how the shape of the bunch evolves as it passes through the magnetic lattice. This is done by considering the limit $R \to \infty$ and $\Delta p \to 0$, with the result being Hill's differential equation of motion in a single transverse plane [24],

$$x''(s) + k(s)x(s) = 0 , \tag{2.8}$$

describing either transverse axis oscillations about the orbit; known as betatron oscillations as they were initially studied in an accelerator of that type [25] [1]. Hill's equation can then be treated as any other second order differential equation,

[1] The x-axis case is chosen for analysis from this point onwards but is applicable to both transverse axes.

Multipole	Definition	Effect
dipole	$\frac{1}{R} = \frac{e}{p} B_z$	beam steering
quadrupole	$k = \frac{e}{p} \frac{dB_z}{dx}$	beam focusing
sextupole	$m = \frac{e}{p} \frac{d^2 B_z}{dx^2}$	chromaticity/dispersion correction
octupole	$o = \frac{e}{p} \frac{d^3 B_z}{dx^3}$	field compensation

Table 2.1: The most important multipoles in beam steering/focusing, and their principle effects on the motion of the beam.

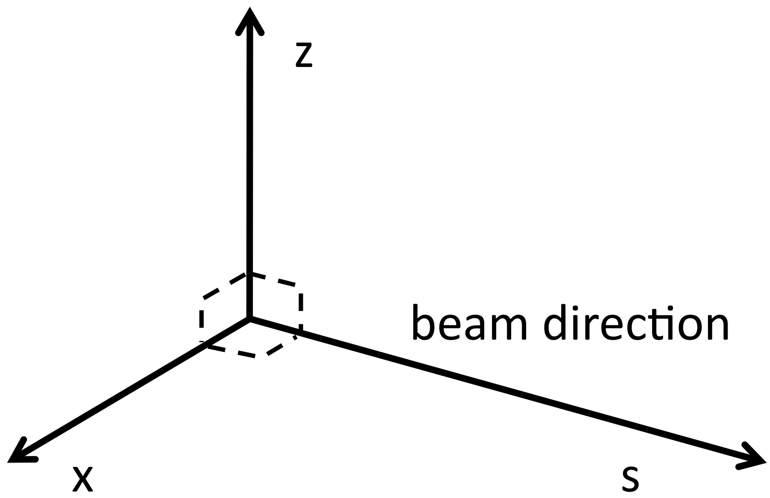

Figure 2.1: Cartesian co-ordinate system to describe the motion of particles in the vicinity of the nominal trajectory.

solved using the trial solution

$$x(s) = Au(s) \cos\left(\psi_x(s) + \phi_x\right) \ , \tag{2.9}$$

where A and ϕ_x are constants of integration. The amplitude function $u(s)$ can also be written in the form

$$\beta(s) \equiv u^2(s) \ , \tag{2.10}$$

which is the definition of the beta function $\beta(s)$. The amplitude factor A in Eq. 2.9 can also be replaced by $\sqrt{\epsilon}$, where ϵ is the emittance, the physical meaning of which will be explained later. The solution of Hill's equation can subsequently be written as

$$x(s) = \sqrt{\epsilon_x \beta_x(s)} \cos\left(\psi_x(s) + \phi_x\right) \ , \tag{2.11}$$

with the phase advancing according to

$$\frac{d\psi(s)}{ds} = \frac{1}{\beta(s)} \ . \tag{2.12}$$

The solution of the trajectory equation in Eq. 2.8 is at a maximum when $\cos\left(\psi_x(s) + \phi_x\right) = 1$. This maximum is known as the envelope function, $E(s)$, and is described by

$$E(s) = \sqrt{\epsilon \beta(s)} \ . \tag{2.13}$$

Particles in the bunch perform betatron oscillations with a position-dependent amplitude, which all lay within the bounds of the envelope function. As can be seen from the plot in Fig. 2.2, the envelope function expounds the transverse size of the beam, and if it is possible to know $\beta(s)$ at any cut along the particle's trajectory then the beam-size can be determined using Eq. 2.13. This is a fundamental quantity of any magnetic lattice as the physical beam size is limited by the magnet aperture, so if $E(s)$ exceeds this aperture the beam will be lost.

As such, the rate of change of transverse beam position, i.e. the angle, x', is given by:

$$x' = \frac{dx}{ds} = -\sqrt{\frac{\epsilon_x}{\beta_x(s)}} \left[\alpha_x(s) \cos\left(\psi_x(s) + \phi_x\right) + \sin\left(\psi_x(s) + \phi_x\right)\right] \ , \tag{2.14}$$

where $\alpha(s) = -\beta'(s)/2$.

It is not enough to simply know the size of the beam at any given point, but to know its angular divergence as well. This is essential in fully simulating the

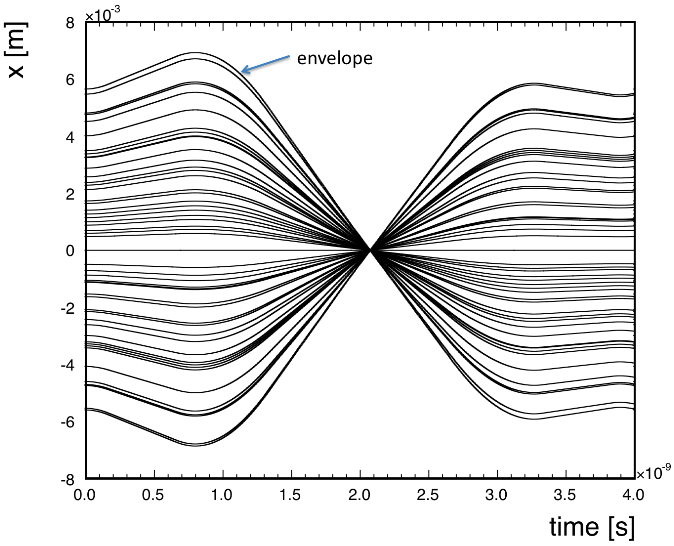

Figure 2.2: Trajectories, $x(s)$, of 50 electrons within the envelope, $E(s)$, of the beam. The beam is comprised of a combination of all the individual trajectories.

tracking of a bunch of particles through a lattice so that beam divergence, blow-up, and ultimately beam loss can be avoided. In order to describe the motion of the particle in the $x - x'$ phase space plane it is necessary to eliminate the terms which depend on the phase, $\phi(s)$. By differentiating Eq. 2.11 with respect to s, and using the trigonometric identity $\sin^2\theta + \cos^2\theta = 1$, the Courant-Snyder invariant is obtained

$$\gamma(s)x^2(s) + 2\alpha(s)x(s)x'(s) + \beta(s)x'^2(s) = \epsilon \; , \tag{2.15}$$

where

$$\begin{aligned} \alpha(s) &\equiv -\frac{\beta'(s)}{2} \; , \\ \text{and} \;\; \gamma(s) &\equiv \frac{1+\alpha^2(s)}{\beta(s)} \; . \end{aligned}$$

It is also necessary to define the relative phase advance, μ, where $\mu(\Delta s) = \Delta\psi(s)$. The functions $\alpha(s)$, $\beta(s)$, $\mu(\Delta(s))$, and $\gamma(s)$ are known as the Twiss parameters: there are two of each parameter (one for each transverse plane) and they are used to characterise the accelerator lattice. These parameters will be defined in the following section.

It is then possible to plot the phase-space of the bunch by finding the axis intersections and maxima/minima, an example of which is shown in Fig. 2.3. It is now clear that the emittance is the area of the phase-space ellipse such that

$$A = \pi\sqrt{\epsilon\beta} \cdot \sqrt{\frac{\epsilon}{\beta}} = \pi\epsilon \ , \tag{2.16}$$

where

$$\epsilon^2 = \langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2 \ . \tag{2.17}$$

Liouvilles's Phase Space Theorem states that

> "*For a nondissipative Hamiltonian system, phase space density (the area between phase space contours) is constant... for a small time increment dt.*" [26]

The condition of constant emittance is generally required by accelerators in order to maintain beam stability. This is especially pertinent to repeating lattices, such as the circular FFAG lattices of EMMA (see Ch. 6) and PRISM (see Sec. 3.2), where the beam experiences the effect of any imperfections on every revolution. The emittance of the beam is therefore considered an invariant of the particle motion throughout the entire accelerating structure. The shape of the ellipse is not constant, however, as it depends upon the amplitude function $\beta(s)$. The shape of the ellipse, therefore, represents the nature of the beam such that a beam along the x'-axis represents a completely divergent beam, whereas a beam along the x-axis represents a completely collimated beam. In modern-day circular accelerators, where the energy is not constant during acceleration, the emittance

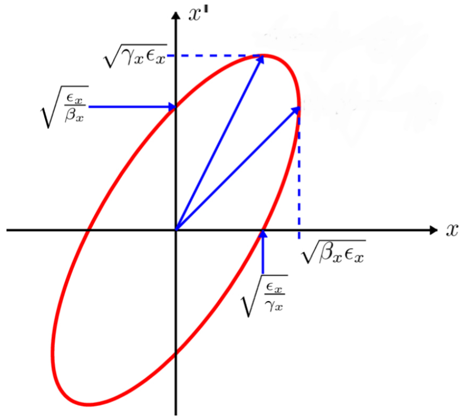

Figure 2.3: A phase-space ellipse showing the relation between the angular and spatial width of a beam of particles in the x'-x plane.

is not an adiabatic invariant and is often defined according to

$$\epsilon_N = \epsilon \times \left(\gamma \frac{v}{c}\right) \quad , \tag{2.18}$$

where here γ is the Lorentz factor (not the previously stated Twiss parameter) and ϵ_N is the invariant normalised emittance.

The phase-space ellipse can then be used to demonstrate the beam focusing abilities of magnetic elements, such as that of a quadrupole magnet. As a beam of particles passes through a magnetic lattice, a quadrupole magnet acts as a lens which focuses in one axis of the phase-space diagram but defocuses in the other. An example of this process can be seen in Fig. 2.4.

2.1.3 Transfer matrices and beam matching

In accelerator simulation and design it is useful to compare the tracking of a beam statistically (as previously described) to the intrinsic magnetic properties of the lattice. This allows direct comparison between tangible properties of the beam, such as size, to the inherent beta functions of the lattice. In order to relate the Twiss parameters at any one point to another the lattice is dissected into discrete sections with matrix transformations performed for each element e.g. drift lengths,

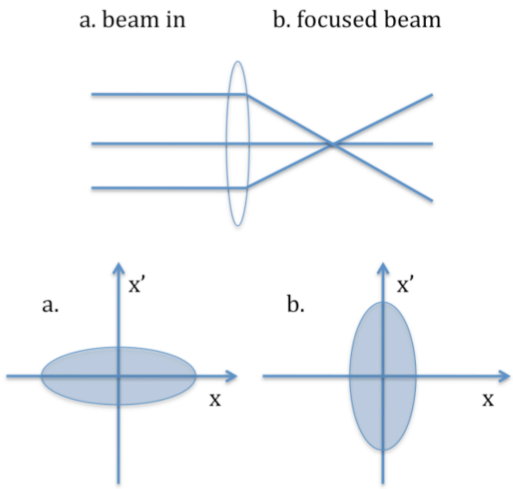

Figure 2.4: The effect of a quadrupole magnet of strength k on a beam of particles, a) the emittance before the quadrupole and b) the emittance after the quadrupole.

dipoles, quadrupoles, etc. Each of these sections can be described via a matrix transformation, translating the initial parameters at point s_1 along the accelerator to s_2. Taking the horizontal beam parameters, x and x', as an example the matrix transformation is described via

$$\begin{aligned}\begin{pmatrix} x \\ x' \end{pmatrix}_{s_2} &= \mathbf{M}(s_1 \to s_2) \begin{pmatrix} x \\ x' \end{pmatrix}_{s_1} \\ &= \begin{pmatrix} \sqrt{\frac{\beta_2}{\beta_1}}(\cos\mu + \alpha_1 \sin\mu) & \sqrt{\beta_1\beta_2}\sin\mu \\ -\frac{1+\alpha_1\alpha_2}{\sqrt{\beta_1\beta_2}}\sin\mu + \frac{\alpha_1-\alpha_2}{\sqrt{\beta_1\beta_2}}\cos\mu & \sqrt{\frac{\beta_2}{\beta_1}}(\cos\mu - \alpha_2\sin\mu) \end{pmatrix} \begin{pmatrix} x \\ x' \end{pmatrix}_{s_1} .\end{aligned} \tag{2.19}$$

The relation described in Eq. 2.19 needs to be expanded for the three spatial dimensions of the beam, spanning across a full 6D plane, to accurately describe propagation [27]. This 6×6 matrix is referred to as the R-matrix, and performs the transform

$$\begin{pmatrix} x \\ x' \\ y \\ y' \\ l \\ \delta \end{pmatrix}_{s_2} = \mathbf{R}(s_1 \to s_2) \begin{pmatrix} x \\ x' \\ y \\ y' \\ l \\ \delta \end{pmatrix}_{s_1} , \tag{2.20}$$

where l is the particle location along the longitudinal length of the bunch and δ is the relative momentum deviation from the reference particle ($\Delta p/p$).

This general form of the transfer matrix can be simplified for different bunch configurations and magnetic elements. For example most bunches are Gaussian in nature, with no $x - y$ coupling, so the majority of the 36 matrix elements may be set to zero. Likewise for different magnetic elements, simplifications are possible. For example the effects of a quadrupole in the transverse $x - x'$ plane are defined by the transfer matrix:

$$R_{\mathrm{quad.}} = \begin{pmatrix} \cos(\sqrt{k}l) & \frac{\sin(\sqrt{k}l)}{\sqrt{k}} \\ -\sqrt{k}\sin(\sqrt{k}l) & \cos(\sqrt{k}l) \end{pmatrix} , \tag{2.21}$$

where l is the length of the quadrupole and k the normalised gradient, as given in Tab. 2.1. Using the thin lens approximation, where $1/kl \gg 1$, Eq. 2.21 reduces to

$$R_{\mathrm{quad.}} = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix} , \tag{2.22}$$

where the focal length of the quadrupole is defined as $f = 1/kl$. The R-matrix may then be constructed via a matrix multiplication from the definitions of the separate magnetic components, such that

$$\begin{pmatrix} x \\ x' \\ y \\ y' \\ l \\ \delta \end{pmatrix}_{s_2} = R_N R_{N-1} ... R_1 \begin{pmatrix} x \\ x' \\ y \\ y' \\ l \\ \delta \end{pmatrix}_{s_1} , \tag{2.23}$$

where N represents a discrete magnetic element of any type. It is of note that the order of multiplication is reversed, such that the first element the beam encounters is the last to undergo matrix multiplication. Using this method it is possible to track the motion of the beam through the lattice by simply knowing the initial beam parameters. The shape of the beam emittance ellipse (Fig. 2.3) may therefore also be tracked in this way.

The evolution of the optical functions through the magnet structure may be calculated using the matrix equations derived above. In addition to this, it is usually necessary to tune the optical functions to specific values at the end of the lattice, if not before. The procedure of varying magnet strengths to regulate these parameters, and so retain the beam within the magnetic aperture, is called matching of beam optics and becomes increasingly pertinent in structures with a repeating lattice, i.e. circular devices, where small fluctuations can be rapidly magnified in successive 'turns'.

Beam matching can be performed across multiple spatial and angular dimensions. Take for example the beta functions in the transverse x-plane at a point s along the accelerator. It is necessary to know the relation

$$\beta(s) = f(k_j) , \tag{2.24}$$

where k_j is the strength of the jth quadrupole. $f(k_j)$ may be any non-linear function so, therefore, it is not possible to find a general analytic solution to this relation. The function must therefore be expanded and solved iteratively. The solution will not be exact as all the higher order terms are neglected, but in general it will produce a better value than initially. This process may be repeated for all the optical parameters found in Tab. 2.2.

This type of matrix multiplication method is that utilised by the methodical accelerator design (MAD) [28] software package. It is a cornerstone of all accelerator simulation codes and one used to benchmark and characterise almost all modern accelerators, including EMMA, as detailed in Ch. 6.

2.1.4 Dispersion

The motion of particles with ideal momentum (i.e. reference momentum) has been described in previous chapters. In reality however, particles in a bunch have a momentum range relative to that of the reference particle and will therefore not behave in such a well-defined manner. Position and angle deviations are compensated for by quadrupoles but deviations in momentum are far more difficult to correct.

These dispersive deviations arise from a number of different causes e.g. energy loss from synchrotron radiation, but all result in a spreading of the momentum from that of the reference particle. This deviation is described by $\Delta p/p$.

The momentum deviation from that of the reference particle only has a significant effect on the trajectory of the particle if $1/R \neq 0$, and so it is only necessary to solve the equations of motion inside the bending magnets (i.e. when $k = 0$). The transverse displacement, Δx, due to the momentum spread is given

optical function	horizontal	vertical
beta function	$\beta_x(s)$	$\beta_y(s)$
gradient of the beta function	$\alpha_x(s)$	$\alpha_y(s)$
betatron phase	$\psi_x(s)$	$\psi_y(s)$
dispersion	$D_x(s)$	
gradient of the dispersion	$D'_x(s)$	

Table 2.2: Optical functions which can be matched by varying quadrupole strengths in a magnetic lattice.

by

$$\Delta x = x_g - x(s) = D(s)\frac{\Delta p}{p} \ , \tag{2.25}$$

where $x(s)$ is the transverse position that a particle of nominal momentum would have and x_g is the transverse position of the off-momentum particle. The dispersion function, $D(s)$, can be calculated from the lattice via a matrix-transformation method analogous to that used for the Twiss parameters [23]. It is computed independently for x and y.

The derivative of the Dispersion function, $D'(s)$, may also be calculated in a similar way. In this case the transverse position is substituted by the normalised transverse velocity, v/c.

2.2 Combined focusing and steering elements

2.2.1 Multi-quadrupole cells

An accelerator is composed of a number of discrete sections, each of which is designed to have a specific effect on the beam. Accelerators used to be built solely with dipoles, with the beam confined to a desired trajectory with a dipole's inherent weak focusing. The weak focusing of a dipole field, outlined in Sec. 2.1.1, is insufficient to provide the required focusing necessary to keep a modern, high-energy beam restricted to a desired orbit however. Quadrupoles, utilising strong focusing, are therefore required.

A quadrupole consists of four poles with hyperbolic surfaces, arranged with alternating polarity North-South-North-South, as shown in Fig. 2.5. This arrangement of poles creates a field that disappears along the beam axis but increases linearly with transverse distance, such that

$$B_x \propto y \ , \ B_y \propto x \ . \tag{2.26}$$

The magnet strength can also be expressed in terms of the normalised gradient, k, as given in Tab. 2.1. This gives rise to a dipole field perpendicular to the beam axis, the result being that particles further off-axis are steered more strongly back towards the centre of the beam-pipe. This results in an overall focusing effect on the beam in one of the transverse planes. However, this focusing in the horizontal plane, for example, leads to a defocusing in the vertical plane. It is therefore necessary to have a minimum of two quadrupoles per magnetic cell, rotated through 90° relative to each other. This basic magnetic arrangement is known as a FODO cell, where the O refers to a drift length between the two magnetic elements.

FODO cells are primarily included in a magnetic lattice as beam focusing tools. However, it is possible for the beam to see a dipole field by implementing the quadrupoles off-axis. Utilising the nature of quadrupoles to steer off-axis particles, a beam may be transversely bent by pushing the beam further off-axis. This is achieved by a physical shift in the quadrupole position. This type of off-axis FODO lattice design is a recent addition to the field of accelerator design, with a successful proof-of-principle design recently commissioned at Daresbury Laboratory, UK. The performance of this accelerator is analysed in Ch. 6.

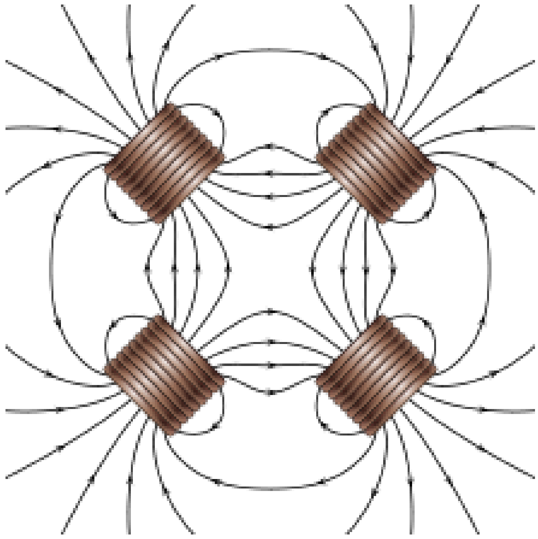

Figure 2.5: The magnet pole arrangement for a quadrupole magnet.

2.2.2 Pancake solenoids

Solenoidal fields are important in beam propagation as they create a uniform magnetic field within a closed volume. Unlike quadrupoles, the magnetic field strength is independent of transverse position within the bounds of the solenoid so a bunch will travel along the longitudinal axis without divergence. However, if a beam enters the solenoid off-axis then it experiences a kick from the B_z field. The bunch will then experience an inherent focusing effect from the solenoid, propagating along the beam axis with a transverse helical trajectory. The helical radius is dependent on the longitudinal magnetic field, such that

$$B_z \propto r \ . \tag{2.27}$$

This linear dependence of the integrated radial fields on the distance r from the axis constitutes linear focusing capabilities of solenoidal fringe fields, which is used at proton targets such as those required by the COMET experiment. An electromagnetic cascade, seemingly emerging from a point like source into a large solid angle, may be captured by placing the target in centre of a solenoid with a

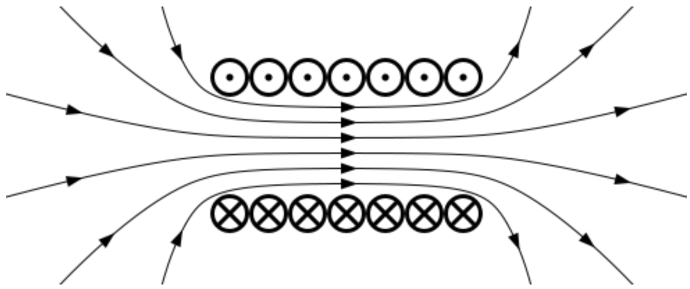

Figure 2.6: The magnetic field and pole configuration of a series of solenoid pancakes. The crosses represent the field pointing into the page and the dots the reverse.

large magnetic field. Due to the nature of the solenoidal field the radial motion of the cascade couples with the longitudinal field to transfer the radial particle momentum into azimuthal momentum. At the end of the solenoid, the azimuthal motion couples with the radial field components of the fringe field to transfer azimuthal momentum into longitudinal momentum. In this scenario a divergent beam emerging from a small source area is focused into a quasi-parallel beam of larger cross section.

As well as utilising the magnetic nature of solenoids for particle capture they are also considered for particle transport and momentum selection. One such design is based on an arrangement of thin coll pancakes, electrically and thermally connected in series along a curve. By rotating these coils in turn to form a toroid, it is possible to make the beam dispersive. This leads to a drift in the centre of the helical trajectory towards the perpendicular direction to the curved solenoid plane. The magnitude of drift, D, being given by

$$\begin{aligned} D &= \frac{1}{qB}\left(\frac{s}{R}\right)\frac{p_L^2 + \frac{1}{2}p_T^2}{p_L}\,, && (2.28) \\ &= \frac{1}{qB}\left(\frac{s}{R}\right)\frac{p}{2}\left(\cos\theta + \frac{1}{\cos\theta}\right)\,, && (2.29) \end{aligned}$$

where q is the electric charge of the particle, B is the magnetic field at the axis, and s and R are the path length and the radius of curvature of the curved solenoid respectively. p_L and p_T are the longitudinal and transverse momenta respectively, and θ is the pitch angle of the helical trajectory.

A compensatory dipole field is therefore required to correct for this drift and keep the helical trajectories on orbit. This is achieved by tilting the solenoid pancakes to create a vertical dipole field, the magnitude of which is described by

$$B_{\text{comp}} = \frac{1}{qR}\frac{p_0}{2}\left(\cos\theta_0 + \frac{1}{\cos\theta_0}\right) , \tag{2.30}$$

where the trajectories of particles with momentum p_0 and pitch angle θ_0 are corrected to be on-axis. A visualisation of this tilt in the solenoid pancakes can be seen in Fig. 2.7, thus producing a vertical dipole field.

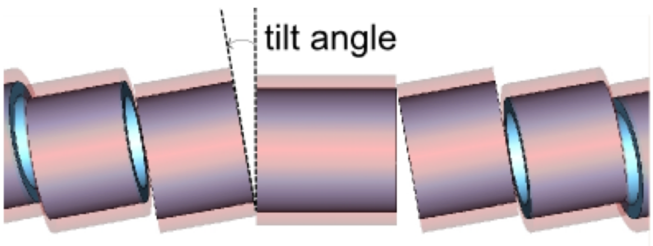

Figure 2.7: A visualisation of the curved, tilted muon transport channel of the COMET beamline demonstrating the tilt used to produce a vertical dipole field. The tilt angle shown is exaggerated for the purpose of visualisation with the angle used in simulation set to 1.43°.

This novel beamline structure is the key aspect in beam transport for high dispersive beams e.g. proton target cascades, for many cLFV experiments. The necessity, and relevance, of these pancake solenoid sections in such experiments is contextualised in Ch. 3 with the COMET experiment.

Chapter 3

COMET

The experimental search for charged lepton violation (cLFV) in the muon channel began in 1947 and has been a vibrant area of research for over 60 years. Observations of cLFV in the muon channel have primarily been sought in three processes. The 90% confidence level on the upper limit of the branching ratios for the three processes over the past 60 years is shown in Fig. 1.5. Almost 10 orders of magnitude improvement in the experimental sensitivity have been achieved.

Two major experiments, Mu2e [29] and COMET [30], are now in the final design phase and are seeking to extend the cLFV search by at least four orders of magnitude in the $\mu N \to eN$ channel.

The COherent Muon to Electron Transition (COMET) collaboration aims to observe cLFV or push the 90% confidence level limit on cLFV BR to below 10^{-16}. In subsequent sections *sensitivity* refers to the 90% confidence level (CL) on the upper limit of the BR unless explicitly stated otherwise. The collaboration is also working on a subsequent experiment (PRISM) which aims to increase the sensitivity to cLFV by a factor of 100 to 10^{-18}. This will be achieved by increasing the distance over which pions can decay and utilising the principle of phase rotation. The PRISM experiment will be discussed further in Sec. 3.2.

The aim of the COMET experiment is to search for the coherent neutrinoless conversion of muons to electrons in a muonic atom ($\mu \to e$ conversion), $\mu N \to eN$, at a sensitivity of 10^{-16}. The proposed location of the experiment is the Japanese Proton Accelerator Research Complex (J-PARC), in Tokai, Japan. As described in Sec. 1.4, a muonic atom can undergo a number of processes after cascading to its $1s$ ground state:

- Decay-in-orbit ($\mu^- \to e^- \nu_\mu \bar{\nu}_e$), where the e^- energy is similar to the Michel decay of a free muon but with a small, but significant, boost from nuclear recoil.
- Incoherent capture by a nucleus of mass number A and atomic number Z, such that $\mu^- + N(A, Z) \to \nu_\mu + N'(A, Z - 1) +$ capture products.
- The exotic coherent conversion of a muon to an electron, $\mu^- + N(A, Z) \to e^- + N(A, Z)$, under investigation by the COMET experiment.

All the nucleons participate in the latter process. As such the rate of the conversion process is enhanced by a factor of Z, and the rate of coherent conversion to incoherent capture is enhanced by a factor approximately equal to the number of nucleons.

This process of creating muonic atoms, through stopping incident muons in a solid target and allowing them to cascade to the $1s$ ground state, may be optimised by altering the design of the target i.e. the Z number or target dimensions. Along with mitigating background, optimising this stopping process is the most important experimental aspect towards providing as clean and distinct a signal as possible. These studies have been performed in simulation and are detailed in Sec. 3.1.2.

The single event signature of coherent $\mu \to e$ conversion is a mono-energetic electron emitted from the conversion. This signal electron has an energy approximately equal to the rest mass of the muon, with an additional energy shift determined by the target material used to stop and capture the muon. This energy is defined by the approximation $E_{\mu e} \approx m_\mu - B_\mu$, where m_μ is the mass of the

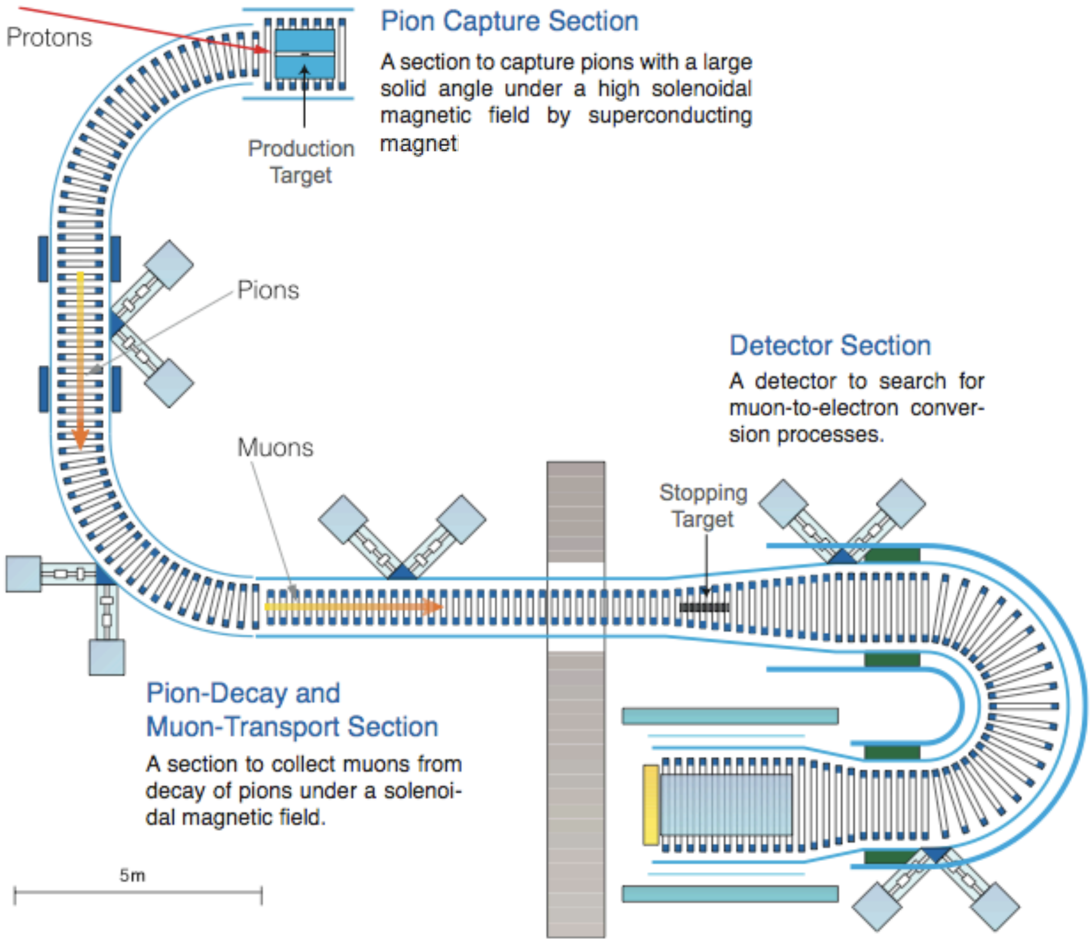

Figure 3.1: A schematic of the proposed COMET experiment, with elements of interest highlighted and detailed.

muon in its rest frame and B_μ is the binding energy of the muon in the $1s$ state. This process provides an attractive proposition for the potential improvement in sensitivity by utilising a high muon rate without suffering from accidental, coincident background, that limits the sensitivity of cLFV searches through the $\mu \to eee$ and $\mu \to e\gamma$ channels.

The main source of background in the search for $\mu \to e$ conversions arises from decay-in-orbit (DIO) electrons. The momentum spectrum of DIO extends beyond the Michel edge at $m_\mu/2$ due to a boost from the nuclear recoil, such that DIO electrons can be produced with energies very similar to those expected from the neutrinoless $\mu \to e$ conversion. In Sec. 3.1.1 the implementation of a recent calculation of the DIO spectrum into the COMET simulation is described.

There are several other potential sources of beam-related background in the signal region, such as high momentum muons decaying in flight or radiative muon/pion capture. In order to reach the proposed sensitivity of COMET, $\mathcal{O}(10,000)$ times better than that of current experimental limits, it has been necessary to introduce several design changes compared to previous experiments. These are described in the following sections. A schematic for the proposed beamline can be found in Fig. 3.1, with aspects highlighted that will be described in subsequent sections.

High intensity muon source

The total number of stopped muons needed is of the order 10^{18} to achieve an experimental sensitivity of 10^{-16}. Therefore, a high intensity muon beamline has to be constructed. Currently the most intense muon beamline is at PSI where 10^8 muons per second are produced using a number of proton targets to multiply muon yield, and using conventional magnets for capture and steering. For COMET to achieve the desired sensitivity 100 times that of PSI at least 10^{10} muons per second are required. This muon flux requires a new pion capture strategy.

Not only are higher fluxes required, but the capture of low momentum particles is essential: pions in order to decay within the decay length and muons in order to be stopped by the muon stopping target. The majority of particles created from an 8 GeV proton beam incident on a graphite target are forward moving with a mean momentum of 300 MeV/c. The backward moving pions have a mean momentum of approximately 120 MeV/c. The momentum spectra for both directions are shown in Fig. 3.2.

For muons to be stopped in the stopping target with high efficiency their momentum needs to be less than 40 MeV/c. The majority of backward moving

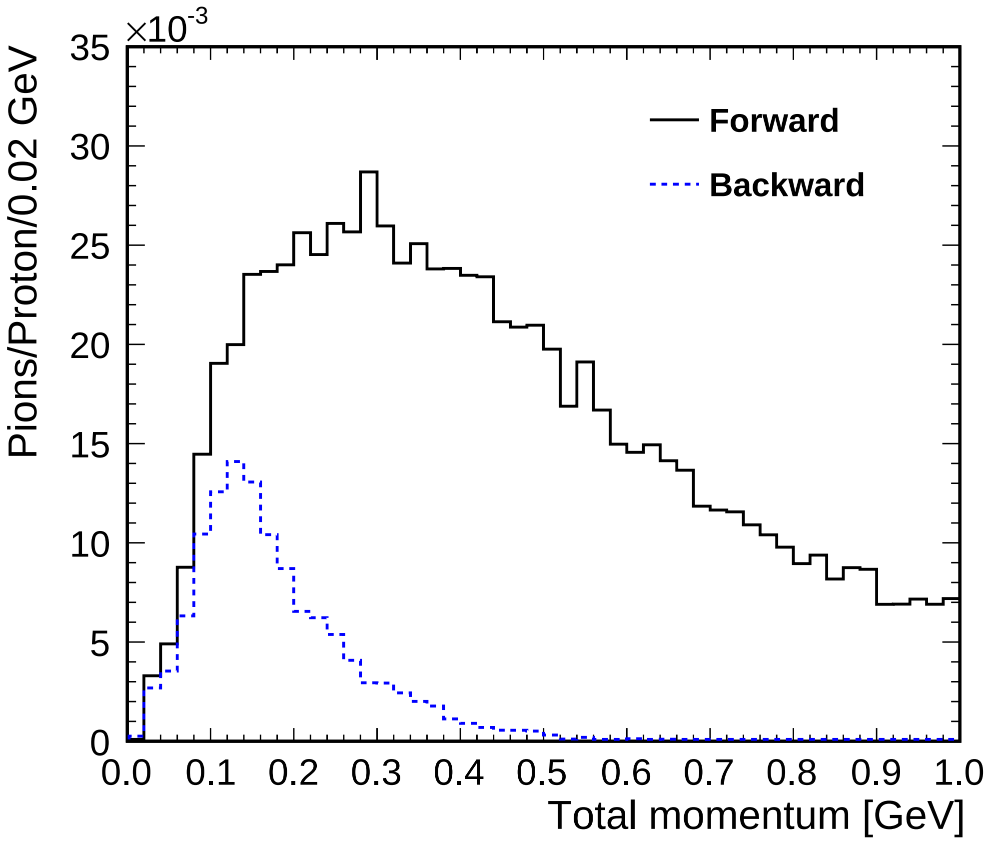

Figure 3.2: Total momentum distributions for both forward and backward moving π^- with respect to the proton stopping target. The data was produced by MARS, firing an 8 GeV proton beam on a cylindrical tungsten target.

pions satisfy this but the yield can be substantially increased by reversing the direction of the forward pions using a strong, graded magnetic field. This was proposed by Djilkibaev and Lobashev [31] in the 1980s. The idea was recently adopted for muon collider and neutrino factory R&D and studied for pion capture over a large solid angle. The superconducting pion capture solenoid system is essential for exploiting megawatt beams, with approx. 8×10^{20} protons of 8 GeV necessary to achieve the required number of muons. This type of high luminosity beam can now be produced by the intense proton sources at FNAL and J-PARC, with Monte Carlo (MC) studies predicting an $\mathcal{O}(1000)$ increase in yield on previous methods. The MuSIC experiment in Osaka, Japan represents a real study for this type of muon production, and initial results from this experiment will be described in Ch. 5.

Pulsed proton beam

There are several potential sources of electron background events in the energy region around 100 MeV, where the $\mu \to e$ conversion signal is expected. One of these is prompt electrons produced from the initial proton target interaction. They are typically produced over a time interval of approximately 250 ns. In order to suppress the occurrence of prompt beam-related background events, a pulsed proton beam is proposed. Since muons in muonic atoms have lifetimes of approx. 1 μs, a pulsed beam with short beam buckets, relative to muonic atomic lifetimes, would allow the veto of prompt beam background events by requiring the measurements to be performed in a delayed time window. An example of the time structure for the COMET proton beam can be seen in Fig. 3.3.

Pion and muon transport solenoid

The pions produced by the proton interaction with the fixed target will inevitably decay to muons. These muons are then transported with an estimated efficiency of 40% through a superconducting solenoid magnet system. Muons with high momentum ($p > 70$ MeV/c) decay to produce electron background events in the signal energy region of approximately 100 MeV and, therefore, must be eliminated with the use of curved solenoids. The centres of the helical motion of the beam particles (primarily muons and pions at this initial point in the beamline) drift perpendicular to the plane in which their paths are curved, and the magnitude of the drift is proportional to their momentum. By using this effect and by placing suitable collimators at appropriate locations, beam particles of high momentum can be eliminated. A full description of these 'pancake solenoids' can be found in Sec. 2.2.2.

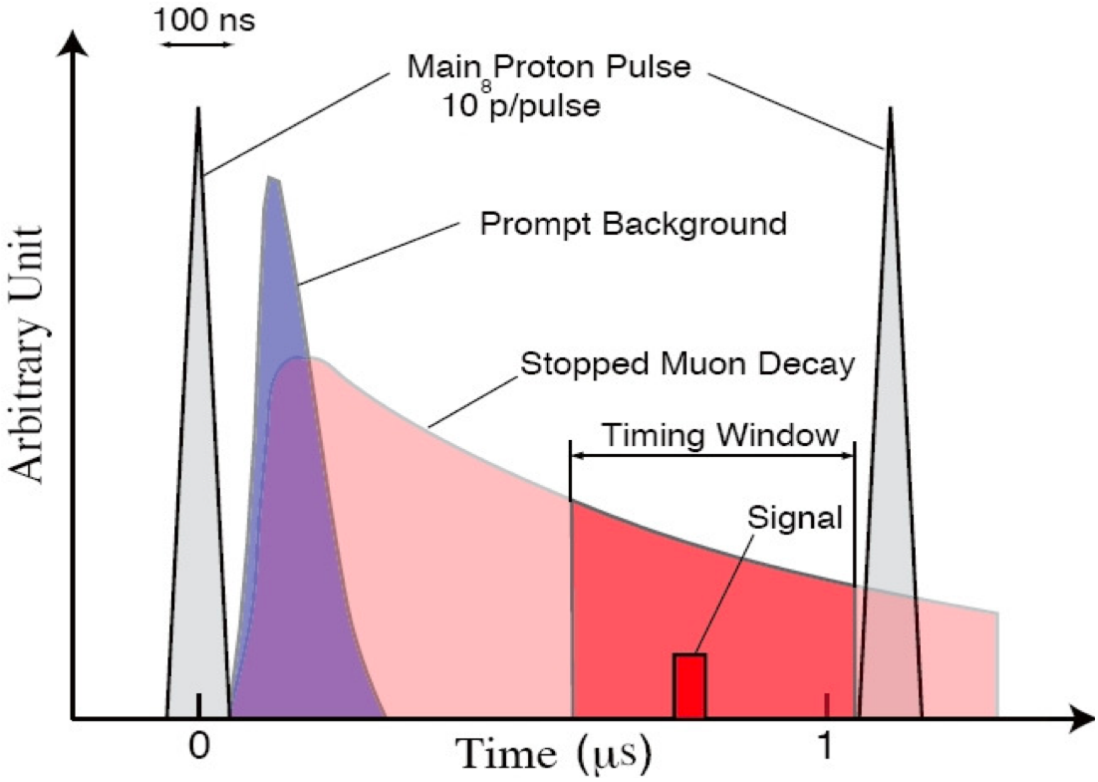

Figure 3.3: The time structure of the primary proton pulses, backgrounds, and signal for the COMET experiment.

Curved spectrometer solenoid

In order to reject electron background events and reduce the probability of false-tracking owing to high counting rates, a curved solenoid spectrometer is used to select electrons on the basis of their momenta. The principle of momentum selection is the same as that used in the transport system. However, in the spectrometer electrons of low momenta ($p < 60$ MeV/c), mostly coming from muon DIO, are removed. The tracking detector rate, composed of DIO electrons, and the secondary electrons from scattering, is expected to be less than 1 MHz. This is almost two orders of magnitude less than the expected detector rate at other proposed LFV experiments.

The curved electron spectrometer also removes high momentum protons which can be produced as a consequence of the nuclear capture process. These protons can swamp the tracker both in terms of rate and because being highly ionising they can 'deaden' the tracker response for a short period.

3.1 COMET_G4

The COMET experiment is currently in the R&D phase, with an initial Phase-I (see Ch. 4) run proposed for 2017. Comprehensive MC simulations are therefore required to model the physics effects of the magnetic lattice and particle interactions. This is necessary for the manufacturing of the experimental components, a robust estimation of background and signal rates, and optimisation of detector design.

The COMET software framework, COMET_G4, is an integrated package, comprising both GEANT4 [32] and MARS [33] components. MARS is a hadron production code that simulates the interaction of hadron beams with fixed targets, in this case a high intensity proton beam with a graphite target. A list of particle 4-vectors from the MARS simulation is stored in a file and this is read in by the GEANT4 simulation. The GEANT4 code tracks the particles and models their interactions from a point 3 cm downstream of the graphite target. In the case of COMET (and MuSIC) the pion capture system is designed to select backwards propagating particles which due to their lower momentum are more likely to stop downstream in the target region. A proof-of-principle prototype superconducting pion capture solenoid (as well as the first 36° of bend section) has already been constructed at Osaka University, Japan for the MuSIC experiment. A photograph of this, with highlighted beam components, can be seen in Fig. 3.4. This enhancement in backward going low momentum pions and muons has been demonstrated in simulation. The MuSIC experiment has been constructed in Osaka to establish the veracity of these simulations and to improve them.

The beamline is modelled in GEANT4 for its ability to simulate both magnetic elements and physical interactions. The MARS output files are automatically parsed

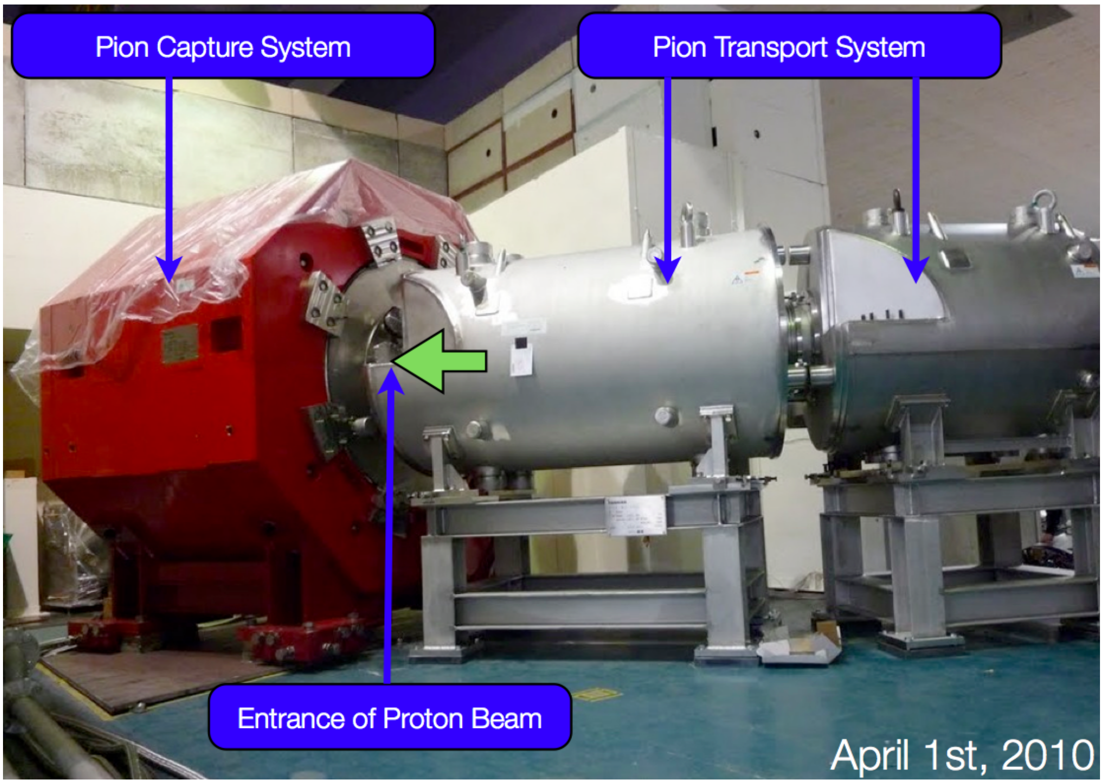

Figure 3.4: A photograph taken in April 2010 showing the MuSIC superconducting pion capture solenoid and the first 36° section of the pion transport section.

into this portion of the code, then tracked through the beamline on a particle by particle basis. The particles travel on a helical trajectory due to the magnetic field of the pancake solenoids, with momentum selection occurring in all bend sections via an induced dipole field. Due to this selection, and decay-in-flight, particles interact with beamline elements (beam collimators, the magnets themselves, etc.) where physical processes such as inelastic scattering and bremsstrahlung need to be accurately simulated; GEANT4 has an extensive set of routines, largely experimentally verified, ideally lending itself to this type of framework.

3.1.1 Software development

The COMET_G4 software has evolved substantially over the past two years and continues to be developed to provide the best possible simulation of the experiment so that robust design choices can be made. This is informed by experimental results and new theoretical developments. In the following sections the implementation of improved magnetic field maps and DIO simulation are described. Fig. 3.5 shows a computer-aided design (CAD) of the latest COMET beamline, detailing solenoid lengths and radii, as well as the direction of the proton beam and major beamline elements such as the proton target and stopping target region.

Magnetic field implementation

The magnetic field produced by the COMET magnetic lattice is a novel, and thus, crucial aspect of the experiment. Most simulation packages, `COMET_G4` included, have the ability to generate magnetic fields of some description. The original form of the package used `GEANT4`-generated fields, with basic approximations for the stopping target region (i.e. the second straight section in Fig. 3.5). These approximations were simple, linear, hard-edged models and thus not adequate to model the complex bend sections of COMET, combining both solenoid and dipole fields. For this reason a `G4beamline` [34] simulation was created to accurately navigate the particles to this second straight section. `G4beamline` is a single-particle tracking toolkit for `GEANT4`, specifically designed for the simulation of beamlines. Once this tracking was complete the distribution was parsed into the `GEANT4` model to simulate particle interactions with the stopping target.

As mentioned, `G4beamline` is specifically designed to model magnetic lattices and has much more sophisticated approximations of magnetic elements which complement `GEANT4`'s intrinsic physics models. Since `G4beamline` is based on `GEANT4` functionality, the code used to approximate magnetic fields was extracted from `G4beamline` and rewritten as a plugin for `COMET_G4` to define a single simulation framework.

The proposed COMET beamline is to be constructed in a stepwise manner, starting upstream with the pion capture solenoid and ending downstream with the

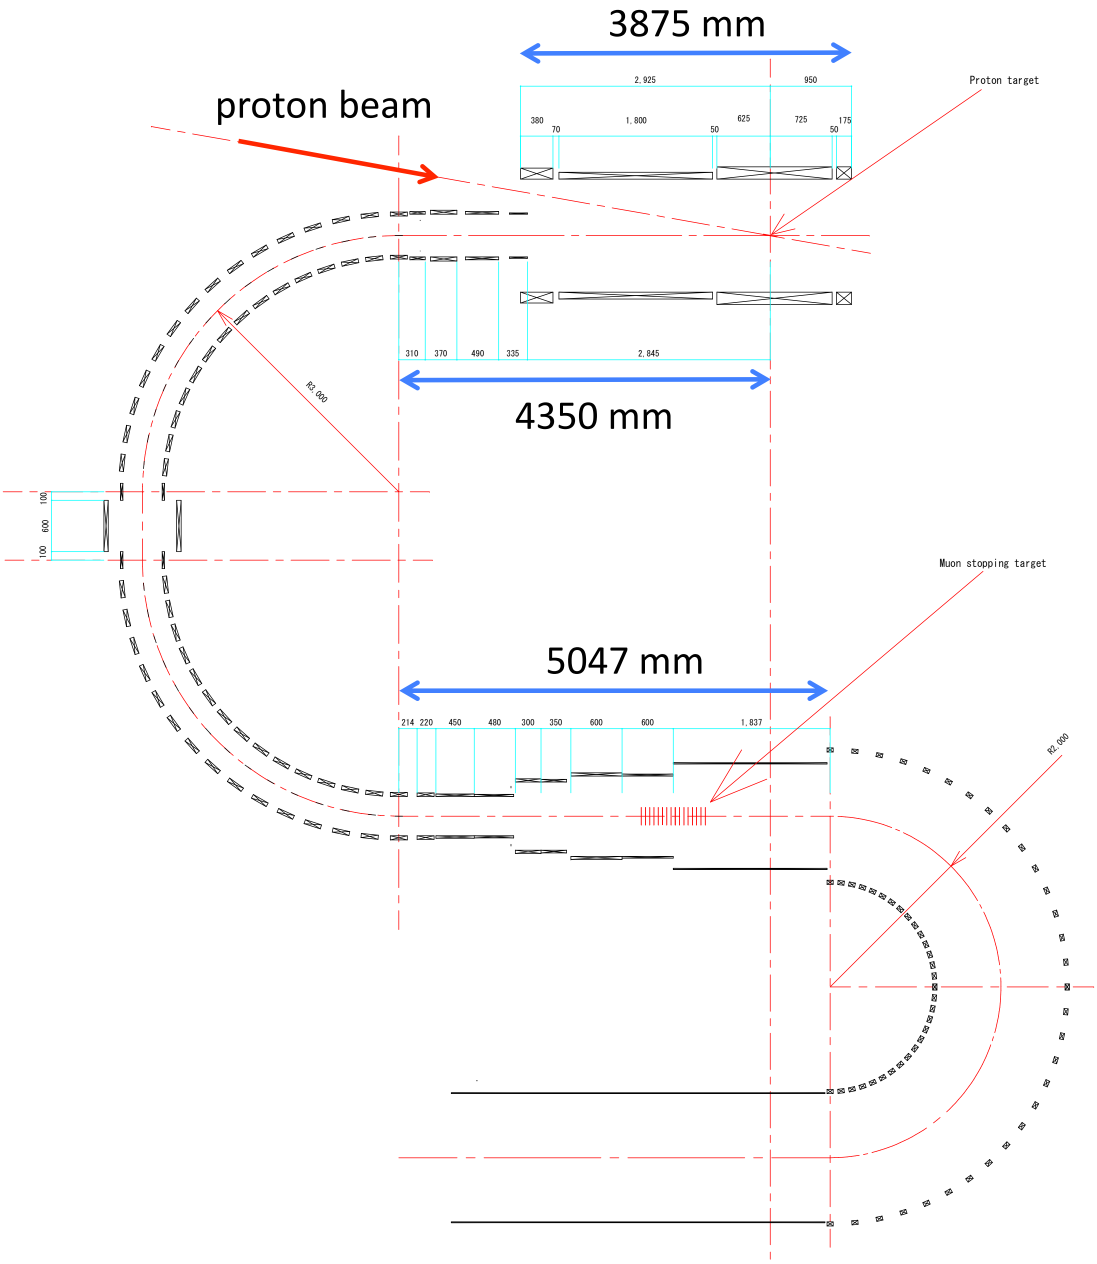

Figure 3.5: A CAD of the latest COMET beamline. Solenoid lengths, thicknesses, and radii, are shown. The direction of the proton beam as well as major beamline elements, such as the proton target and stopping target region, are shown.

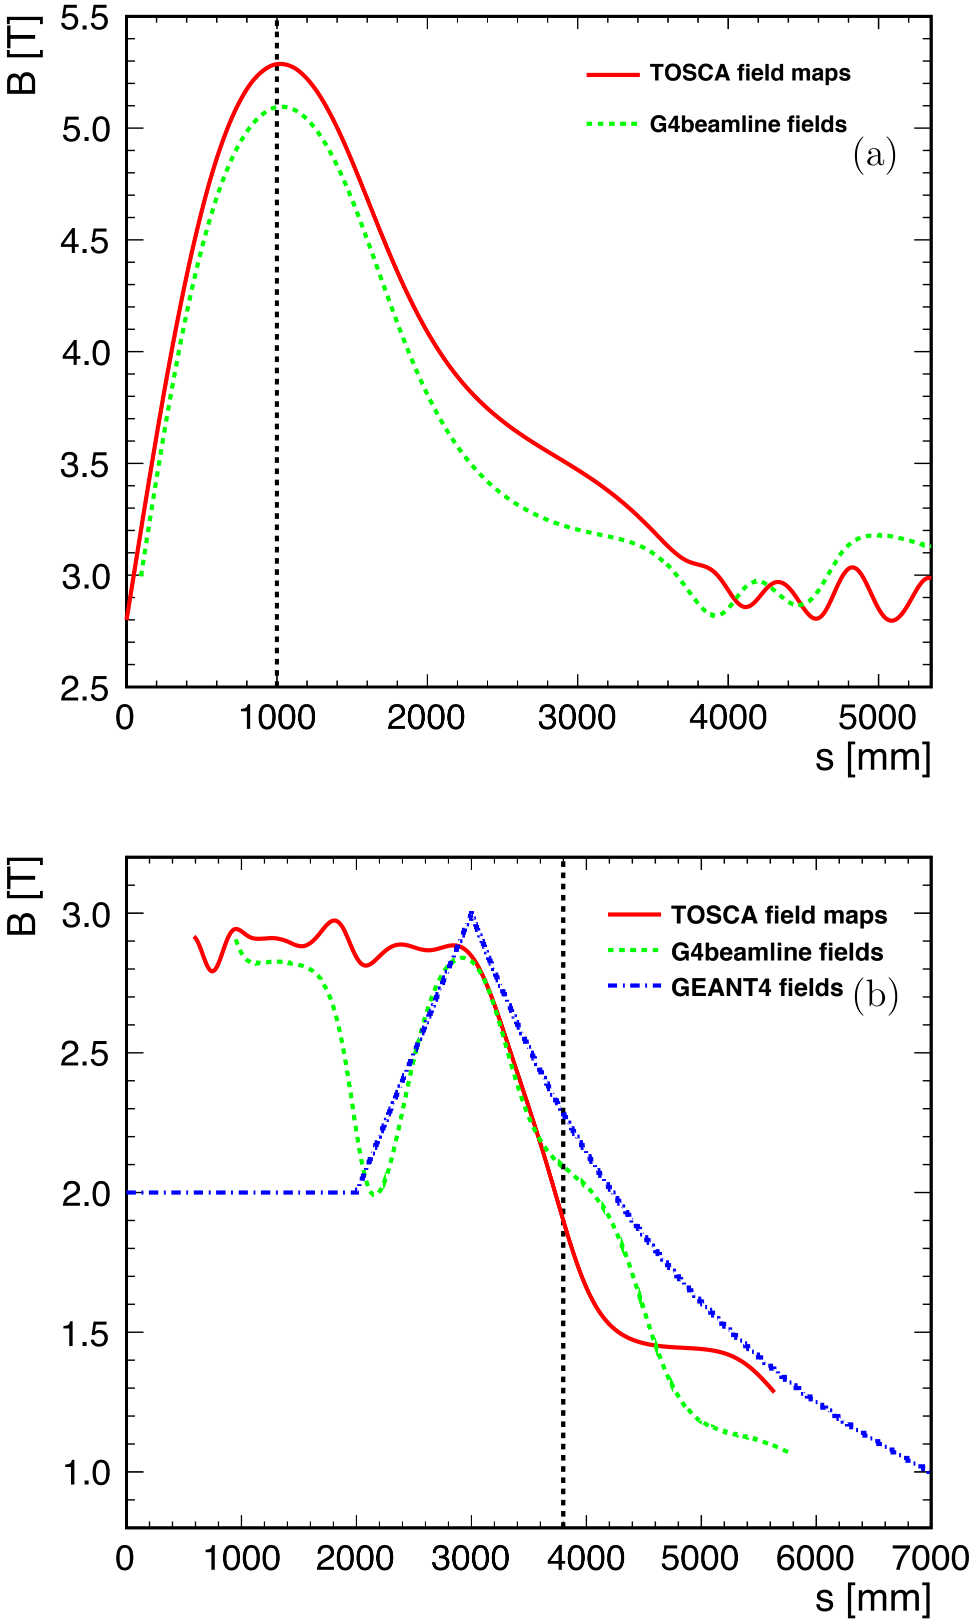

Figure 3.6: Evolution of the magnetic field for the straight sections containing a) the pion capture solenoid, and b) the stopping target region. Both plots show the total magnetic field along the beam axis, $r = 0$.

detector solenoid. As part of the procurement process for the magnetic elements of the experiment Toshiba have developed a precise simulation of the magnetic field in `TOSCA` [35] based on the exact materials and mechanical design of the magnetic beamline. They have produced a CAD file of the experimental layout (shown in Fig. 3.5) as well as a series of field maps for the separate solenoid sections. These field maps represent the most accurate field model for the intended beamline design and, as such, had to be incorporated into the `COMET_G4` framework. The field maps for the individual solenoids are concatenated, with the co-ordinate systems centred such that the origin is consistent for all the maps i.e. the proton target in the pion capture solenoid section. The `G4beamline` plugin was altered such that field maps could be parsed in a Cartesian co-ordinate format.

Fig. 3.6 shows the three magnetic field implementations: the simplified `GEANT4` model, `G4beamline` and the most sophisticated Toshiba field map modelled using `TOSCA`. The fields are shown on-axis, with $r = 0$. There is no `GEANT4` magnetic field representation for the first straight section as the pion capture solenoid was initially modelled using `MARS`. The solenoid lengths, and therefore the section lengths, have been optimised over time to provide the highest stopped muon yield. For this reason the maps are plotted around constant elements, denoted by a vertical dashed line: the centre of the proton target for the first section and the geometric centre of the stopping target for the second.

Fig. 3.6 a) shows the magnetic field strength in the first straight section, comprising the superconducting pion capture solenoid and matching solenoids. The magnetic field strength reaches an apex at the position of the proton target in order to maximise the capture of pions and muons (typically $p < 80$ MeV/c). Pions are emitted into a half hemisphere from the rear of the target and can be captured within a transverse momentum threshold,

$$p_T^{\max}(\text{ GeV}/c) = 0.3 \times B(\text{T}) \times \frac{R(\text{m})}{2}\ , \tag{3.1}$$

where B is the magnetic field strength and R is the inner radius of the solenoid bore.

Every implementation of the magnetic field in both straight sections show a graded magnetic field from high to low field strength. This gradient is implemented to mitigate the dispersion created by the isotropic particle cascade from the proton stopping target and from momentum selection in the bend sections. Liouville's phase space theorem (see Sec. 2.1.2) states that a volume in phase space (in this case the emittance) remains constant in a non-dissipative system. Knowing this, Eq. 2.18 may be manipulated to represent the proportionality

$$\epsilon \propto \frac{p_{\mathrm{T}}^2}{B} \quad , \tag{3.2}$$

where p_{T} is the transverse momentum and B the magnetic field. From this relation it is evident that if the magnetic field is reduced by one third then the transverse momentum must also decrease by a factor of $\sqrt{3}$ to maintain constant emittance. As the transverse momentum is related to the angular divergence of the beam through the relation $p_T = p \sin\phi$, a reduction in p_T will lead to a reduction in angular divergence. This is the principle of adiabatic transition and results in a 'collimation' of the beam across the two straight sections.

Thus the design of the solenoids in the stopping target region is intended to:

- maximise the geometrical acceptance for conversion electrons emitted from the stopping target disks, reflecting backwards emitted electrons using a magnetic mirroring technique, and
- maximise the transmission efficiency for conversion electrons by constraining their helical trajectories as they enter the electron transport solenoid.

The acceptance of the signal electrons is described by the formula

$$\epsilon_{\mathrm{acc}} = \frac{1 - \cos\theta^{\mathrm{crit}}}{2} \,, \tag{3.3}$$

where

$$\theta^{\text{crit}} = \pi - \sin^{-1}\left(\sqrt{\frac{B_{\text{target}}}{B_{\text{in}}}}\right) . \tag{3.4}$$

Here, B_{target} and B_{in} are the magnetic field strengths at and before the muon stopping target respectively. It is evident that the smaller the ratio between these strengths, the larger the acceptance will be, hence a negative magnetic gradient is used across this region.

The `TOSCA` field map is believed to be the most accurate but increasing the sophistication of the fields is redundant if their functionality is not also increased. Another definitive test of their effectiveness is the number of stopped muons per proton for each field map. Tab. 3.1 illustrates an increase in the yield, as desired. Both of these factors indicate the efficacy of an increase in field map sophistication.

Decay in orbit calculations

As mentioned at the beginning of Ch. 3, DIO electrons provide the main source of intrinsic physics background in the signal region. Muon DIO is a Michel decay ($\mu^- \to e^-\nu_\mu\bar{\nu}_e$) of a muon bound under the Coulomb potential in the $1s$ state of a muonic atom. The decay electron may be boosted beyond the end point of a classical Michel decay (approx. 52.8 MeV), extending to the momentum region of the $\mu \to e$ conversion signal. The end point of this spectrum depends on the

Field type	$N_{(\mu^-/p^+)}$
`GEANT4`	0.0020 ± 0.0001
`G4beamline`	0.0023 ± 0.0001
`TOSCA`	0.0024 ± 0.0001

Table 3.1: Number of stopped muons per proton for the evolving magnetic field simulations used in the `COMET_G4` software framework.

binding energy in the ground state and so depends on the element forming the muonic atom. A larger binding energy (higher Z) shifts the endpoint to a lower energy.

The DIO spectrum for gold muonic atoms, as measured by the SINDRUM experiment, was shown in Fig. 1.7. The end point of this spectrum is at a momentum of approx. 93 MeV/c, whereas the end point for aluminium (one suggested material for the COMET stopping target) is at approximately 105 MeV/c, demonstrating the far weaker binding energy than gold.

DIO occurs within the Coulomb potential of the nucleus. For this reason the nucleus recoils, sharing the final state momenta to conserve momentum. As a result, the maximum energy of DIO electrons extends to the energy of $\mu \to e$ conversion electrons. However, the intensity decreases rapidly with increasing electron energy, with the spectrum shape known to be approximated by $(E_{\mu-e} - E_e)^5$ [36]. DIO electrons may therefore be suppressed by raising the lower energy threshold of the $\mu \to e$ signal region. This comes at a cost of a reduction in signal acceptance.

The DIO spectrum has been measured and shown to agree well with theory up to energies of approx. 70 MeV [37]. Since the number of DIO events falls rapidly with energy and existing muon sources do not provide sufficient muons it has not been possible to measure the DIO spectra above 70 MeV. Therefore simulations are reliant on purely theoretical models. Watanabe *et al.* [38] derived such a model for a number of elements, valid from $1 - 100$ MeV, with the result for aluminium shown in Fig. 3.7.

The signal region for aluminium extends beyond the limit of the Watanabe calculation. For this reason another relation is required to model the DIO spectrum in the signal region. Shanker [36] derived an analytical form for $E_e > 90$ MeV,

$$N(E_e)dE_e = \left(\frac{E_e}{m_\mu}\right)^2 \left(\frac{\delta_1}{m_\mu}\right)^5 \left[D + E\left(\frac{\delta_1}{m_\mu}\right) + F\left(\frac{\delta}{m_\mu}\right)\right] dE_e \ , \tag{3.5}$$

	Hydrogen	Aluminium	Titanium	Lead
Atomic Number	1	13	22	82
Muon lifetime (ns)	2197.03 ± 0.04	880 ± 10	330 ± 7	82 ± 5
Binding energy (MeV)	0.0025 ± 0.0003	0.463 ± 0.01	1.36 ± 0.07	10.5 ± 0.2
Relative $\mu - e$ conversion BR	0.6	1	1.7	1.15

Table 3.2: Lifetimes, binding energies, and branching ratios (relative to aluminium) for a range of target material muonic atoms.

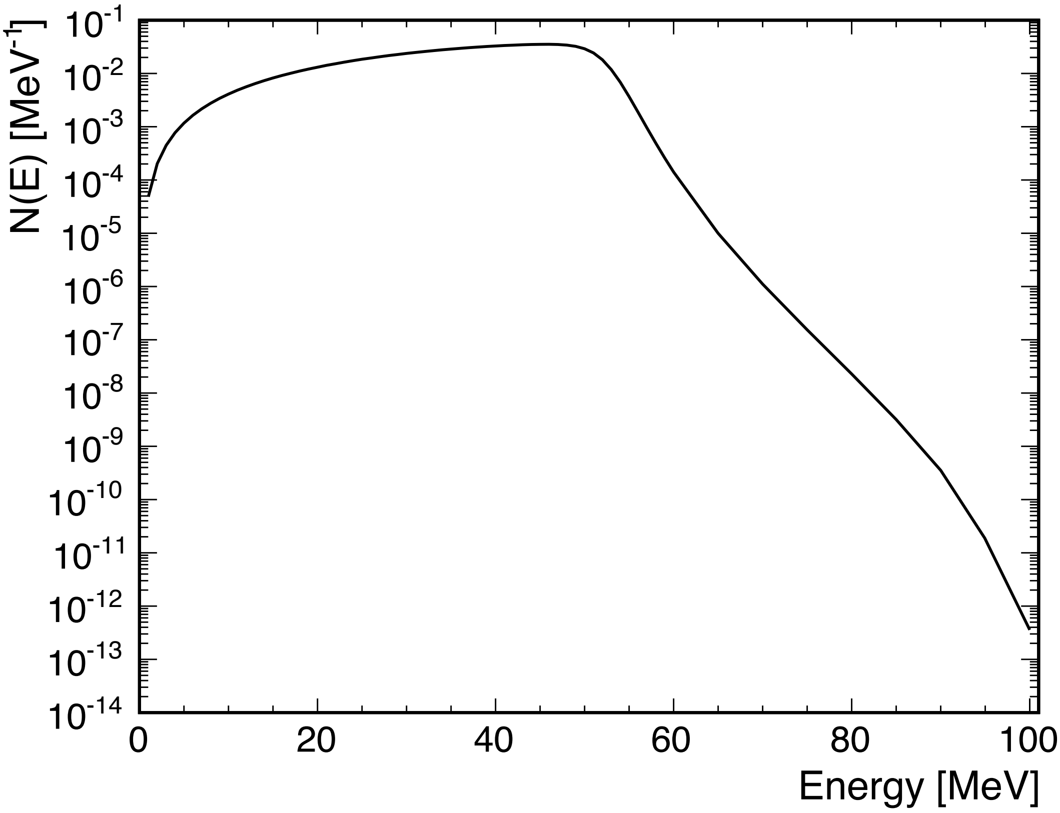

Figure 3.7: The analytical Watanabe DIO spectrum for aluminium. The calculation is only valid to 100 MeV.

where m_μ is the mass of the muon, $\delta = E_{\text{max}} - E_e$, and $\delta_1 = E_\mu - E_e - E_{\text{rec}}$. E_{max} is the end-point of the DIO spectrum, $E_\mu = E_{\text{max}} + E_{\text{max}}^2/(2M_A)$, and $E_{\text{rec}} = E_e^2/(2M_A)$ is the nuclear recoil energy. M_A is the mass of the recoiling nucleus. The coefficients D, E, and F are functions of the atomic number of the nucleus. They were calculated using a linear interpolation between the calculations for magnesium ($Z = 12$) and sulphur ($Z = 16$): values of $D = 0.388 \times 10^{-21}$, $E = 1.03 \times 10^{-21}$, and $F = 2.27 \times 10^{-21}$ are used for aluminium.

Both the Watanabe and Shanker calculations are valid between $90 < E_e < 100$ MeV. The gradients are most similar at 98.7 MeV so this is chosen as the transition point between the two calculations. The Shanker formula was then scaled at this energy in order to minimise the difference between the two calculations.

The result of this matching can be found in Fig. 3.8. The entire Shanker formula is then multiplied by this scaling factor to create a smooth transition at the crossover energy, but also to give an accurate description of the DIO spectrum in the signal region. The Watanabe-Shanker DIO spectrum is shown in Fig. 3.9.

The default stopping target material in COMET is aluminium, modelled by the Watanabe-Shanker relation up to this point. However, the target material is a property under consideration in the optimisation of the stopping target and therefore DIO spectra for different materials need to be incorporated into the COMET_G4 code. The Watanabe spectra for a range of materials [38] is shown in Fig. 3.10.

Another candidate for the stopping target material is titanium, which has a similar muonic lifetime to aluminium but also has an increased branching ratio for

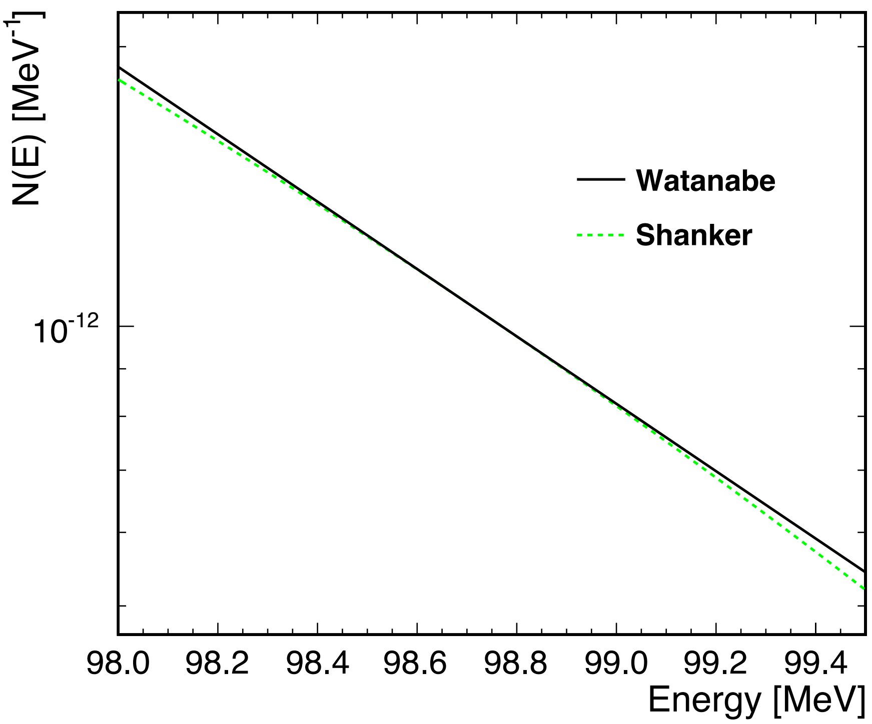

Figure 3.8: The DIO spectrum in the transition region between the Watanabe and Shanker calculations.

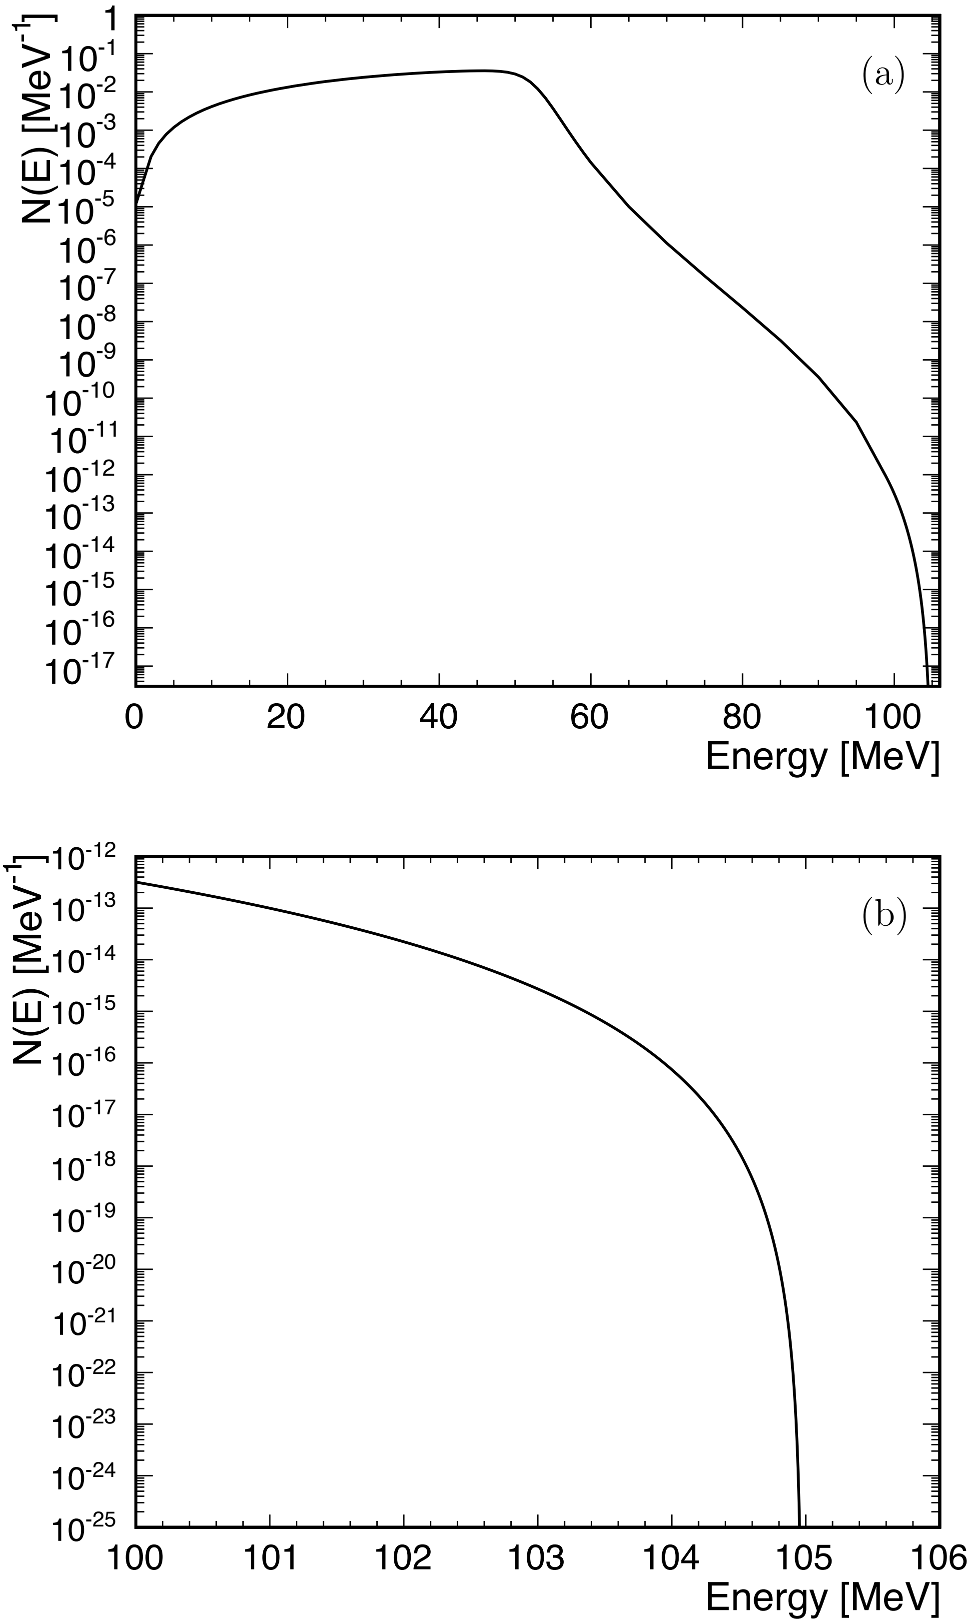

Figure 3.9: The combined Watanabe-Shanker DIO spectrum a) across the full allowed electron energy range, and b) in the $\mu - e$ signal region.

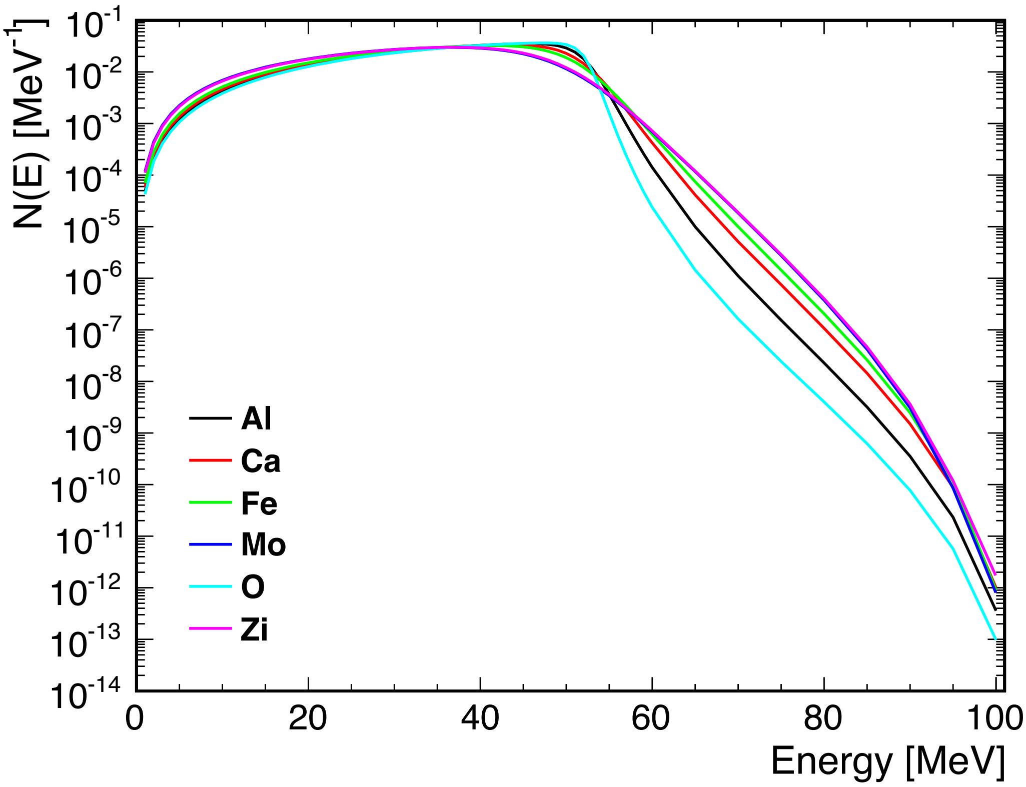

Figure 3.10: Watanabe energy spectra for a range of elements producing DIO electrons.

the conversion process (see Tab. 3.2). The Watanabe spectrum for this material is needed in order to match to the Shanker spectrum at a suitable crossover energy. However, the range of Watanabe spectra [38] do not include titanium. The spectrum for titanium ($Z = 22$) is therefore calculated by linearly interpolating between those given for calcium ($Z = 20$) and iron ($Z = 26$). Similarly, the element-dependent coefficients required by Eq. 3.6 are also linearly interpolated between the two elements expounding this region. The same method may then be applied to titanium, as for aluminium, to achieve an approximation of the DIO spectrum at its endpoint.

The equivalent spectra for titanium, compared to those of aluminium in Fig. 3.9, can be seen in Fig. 3.11. The initial differences in the Watanabe spectrum are highlighted in the first plot, up to $E_e = 98.7$ MeV where the scaled Shanker plot takes over. The second plot is the more relevant of the two, highlighting the level

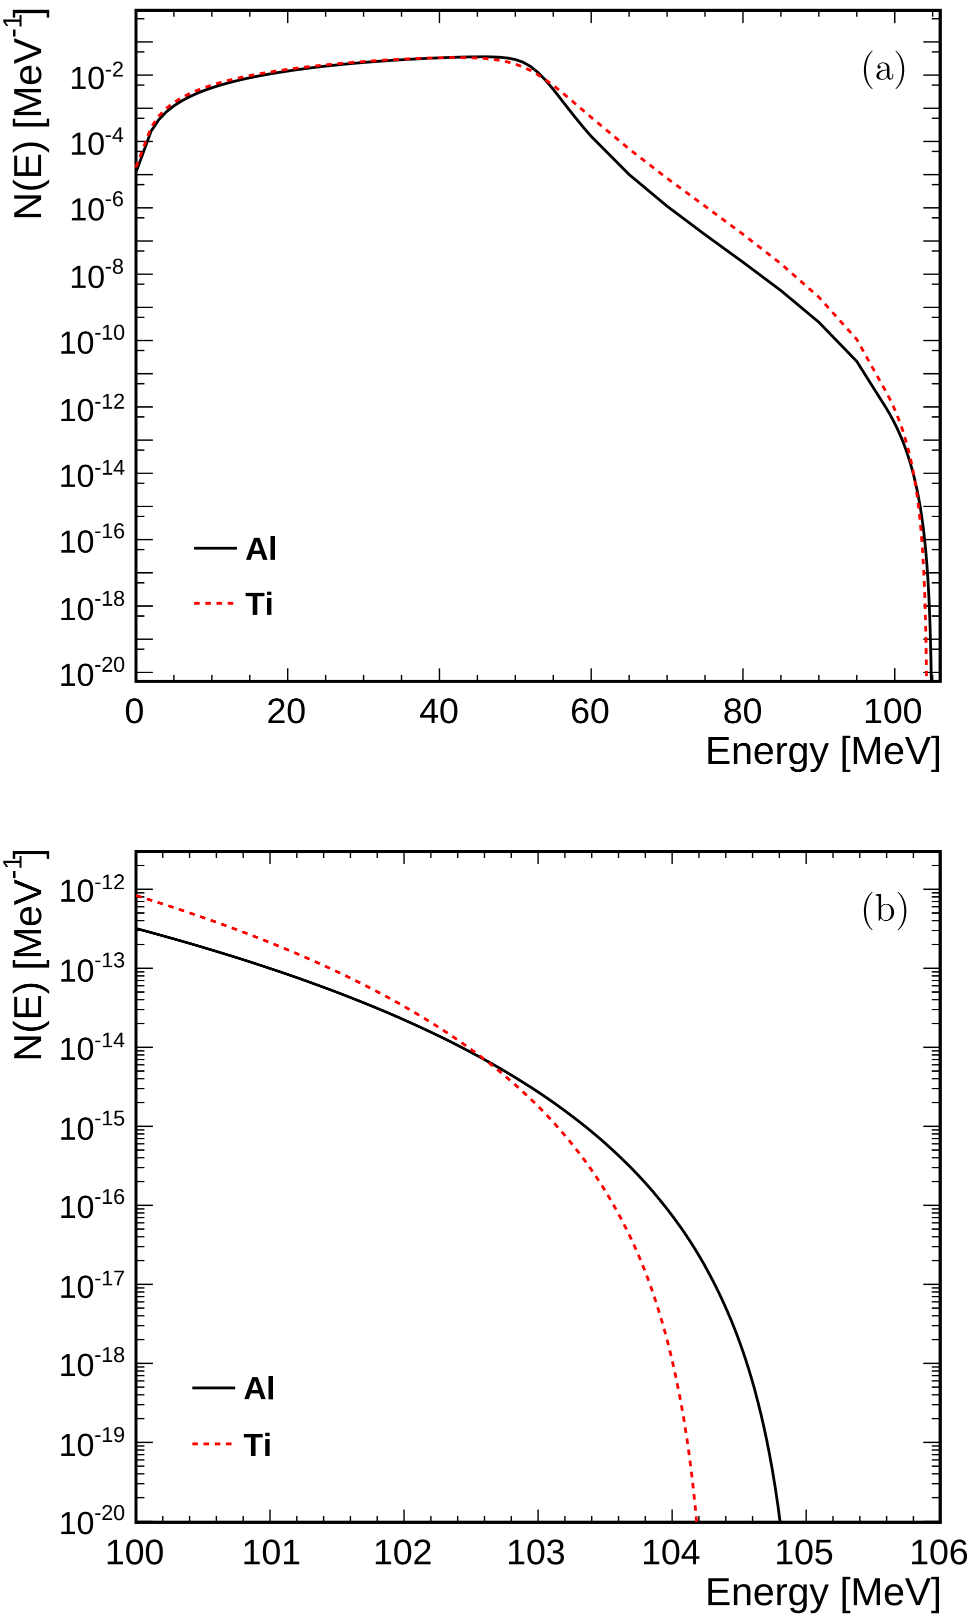

Figure 3.11: The combined Watanabe-Shanker DIO spectrum for both aluminium and titanium, a) across the entire allowed electron energy range, and b) in the $\mu - e$ signal region. The difference in end points highlights the different binding energies for the two materials.

of DIO in the signal region for the different target materials. The endpoints of the spectra differ primarily due to the higher binding energy in the $1s$ orbital of titanium. However, the shift in endpoint due to the differing material also applies to the signal conversion electrons so the more important feature is how steeply the DIO spectrum falls off in the signal region. Fig. 3.11 shows a steeper gradient for titanium, decreasing the integrated background in the signal region for this material. The sensitivity to the conversion process for aluminium versus titanium is qualitatively described in Sec. 3.1.2.

Since the completion of this work a single calculation valid across the whole momentum range, and incorporating nuclear recoil, was published [39]. A comparison between this relation and the Watanabe-Shanker spectrum in the signal region for $\mu \to e$ conversion is shown in Fig. 3.12. As can be seen the relations are similar, converging towards their endpoints. The differences between these spectra are not expected to alter the qualitative conclusions in the following section as any minor discrepancies will be lost due to smearing from detector resolution, etc. In Sec. 3.1.2 a measurement window of $103.5 - 105.2$ MeV is defined. In this region the two DIO definitions agree to within 2%.

3.1.2 Stopping target optimisation

The `MARS` code is used to simulate the cascade of particles produced by 100×10^6 protons incident on the proton stopping target. A solid angle of 180^o behind the proton target is used to select the backwards moving particles. Of these selected particles there are approximately 3.1×10^6 μ^- and π^-. This distribution is then parsed into `COMET_G4` where it is tracked around the first 180^o bent solenoid and through the straight section until the stopping target section. The number of muons at this point is approx. 240×10^3, giving a transmission efficiency of 0.0024 μ^- per proton.

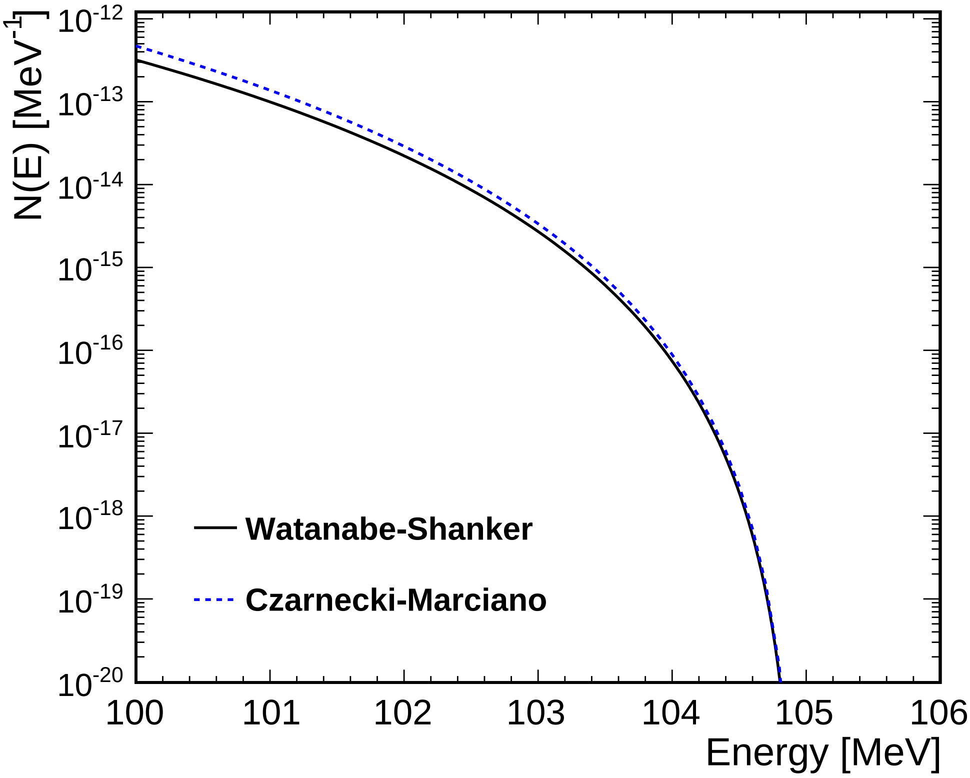

Figure 3.12: Comparison between the Watanabe-Shanker energy spectrum for DIO electrons and the new Czarnecki-Marciano relation.

The COMET stopping target is designed to stop as many of these μ^- and provide an ideal medium for those muons to convert to signal electrons. For that reason the stopping target design must be optimised to:

- maximise stopped muons,
- minimise energy loss of conversion electrons,
- minimise the number of electrons produced by muon decay-in-orbit (DIO).

These constraints are at some level mutually exclusive as using more material stops more muons but also increases the energy loss for conversion electrons and the smearing. The default stopping target design is taken from the MECO experiment [40], composed of 17 aluminium disks of 200 μm thickness, separated by 50 mm. The target parameters considered for optimisation are summarised in Tab. 3.3.

Fig. 3.13 [41] demonstrates the stopping power of copper as a function of the momentum of incident muons. The majority of μ^- and π^- produced via the interaction of the 8 GeV proton beam and tungsten target are high momentum ($p > 300$ MeV/c), as demonstrated by the results of MARS simulations in Fig. 3.2. Muons are close to minimum ionising at this momenta, and so would only be stopped by a large amount of material. Muons with $p < 45$ MeV/c lose energy far more rapidly and can be stopped by a thin material. It is thus essential to ensure that the yield of $p < 45$ MeV/c is maximised.

The effect of changing the stopping target parameters in Tab. 3.3 can be significant. For example, Fig. 3.14 shows the effect on the electron momentum by reducing the disk spacing to zero. The stopping target has the same amount of material but, as the signal electrons are emitted isotropically from the muonic atoms, the lack of spacing between the disks prevents the signal from escaping the stopping target volume without interacting with further material. Despite the number of electrons in the signal region being approximately equal, the peak height is drastically reduced and shifted down in momentum.

The metric against which the stopping target will be optimised is the fraction of muons stopped by the target compared to the total incident on the target, integrated over the momentum range $0 < p < 90$ MeV/c. An example of this can

parameter	value
disk radius	100 mm
disk thickness	200 μm
number of disks	17
spacing	50 mm
material	aluminium

Table 3.3: Default parameters of the muon stopping target.

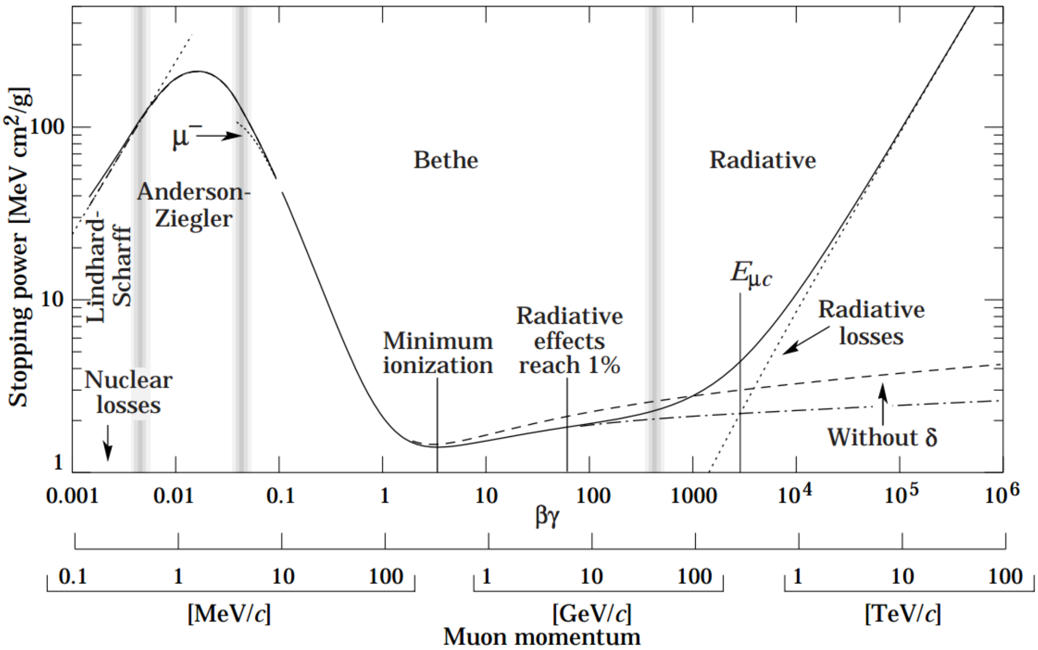

Figure 3.13: Stopping power ($-dE/dx$) for muons as a function of $\beta\gamma = p/Mc$ in copper over nine orders of magnitude in momentum. The muon momentum range expected at the COMET stopping target is $0 - 90$ MeV/c [41].

be seen in Fig. 3.15, with the green histogram displaying the momentum of the muons stopped by the target. This case gives a stopping efficiency of 0.46 for the default stopping target configuration. Error bars are included on this plot but are too small to observe due to the high statistics involved. This is also the case for Figs. 3.17 and 3.20 as the same initial data set was used for all simulations.

As mentioned, it is not enough to simply stop the muons but also ensure that the resulting signal electrons do not encounter enough further material to significantly reduce their momenta below the DIO upper edge i.e. approximately 104.0 MeV/c for aluminium at a branching ratio of 10^{-16}. Energy loss smears both the signal and background events and ultimately the metric of interest in optimising the stopping target is the highest possible sensitivity to the conversion process i.e. the enhancement of signal over background.

Several different ways have been used to quantify the sensitivity of a search, which makes it sometimes difficult to compare them. In particular, two different

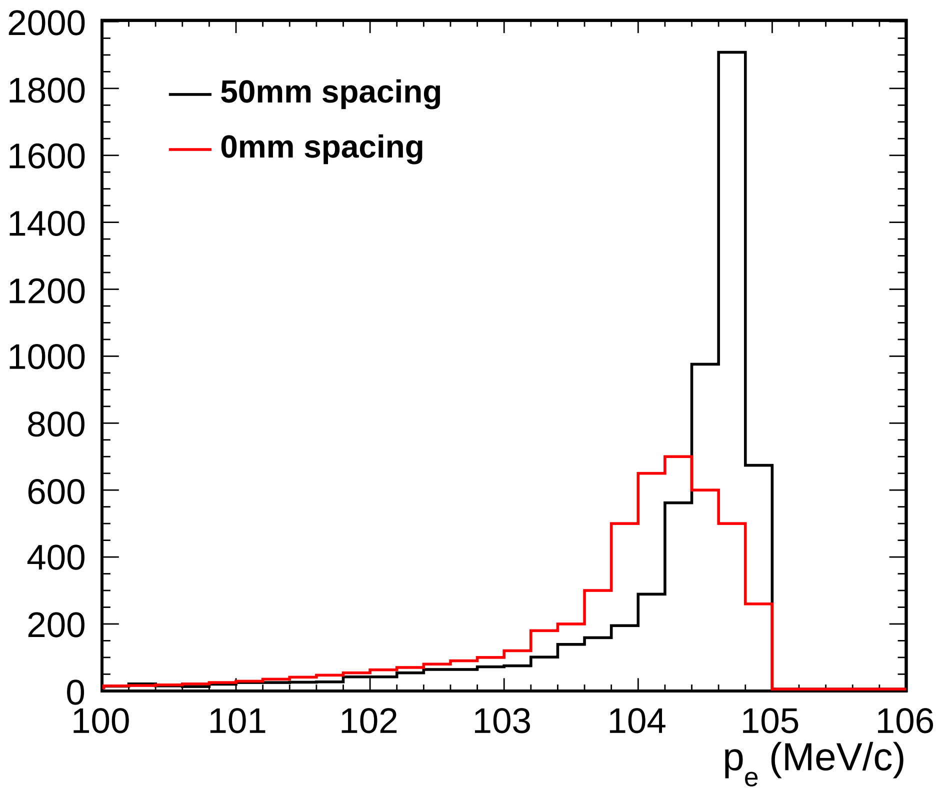

Figure 3.14: μ-e signal for the default COMET stopping target configuration and for the same configuration but with zero spacing between the disks. A beam of 10000 μ^- (of $p = 40$ MeV/c) are simulated in each case.

sensitivity figures are often quoted, one that is relative to the potential for actually making a discovery, and another to characterise how strong a constraint is imposed on the unknown phenomena if no evidence is found for a deviation from the standard theory. Correct statistical practice requires a decision (before the experiment occurs) on appropriate values for the significance level α and confidence level CL. Assuming their values are given the region of the parameters m for which the power of the chosen test is greater or equal to the confidence level chosen for the limits in the case of no discovery is,

$$1 - \beta_\alpha(m) > CL \; , \tag{3.6}$$

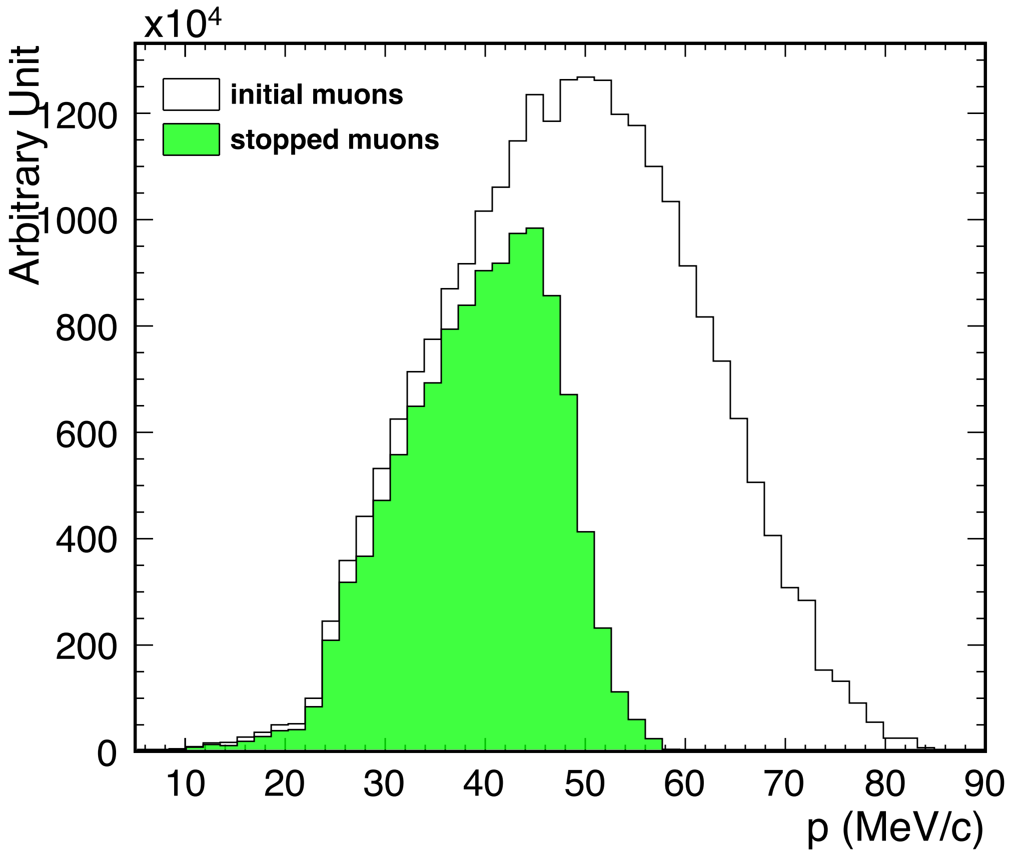

Figure 3.15: Momentum distribution of the muons incident on the stopping target. The momentum spectrum of those stopped is also shown, with a stopping efficiency of 0.46. Error bars are included on this plot but are too small to observe due to the high statistics involved.

where $1 - \beta$ is the power function describing the probability that a discovery will be claimed.

In particle physics a single-event sensitivity is typically quoted to give an estimated sensitivity of the threshold at which a signal may become apparent. This is for the case where background is zero so may be used for a discovery but not as an exclusion limit. Therefore the 90% confidence level upper limit on the branching ratio is typically quoted as a statement of exclusion when no significant signal is seen. A single metric is required combining both information, however. Such a metric has been proposed by Punzi [42] whereby a single BR is quoted at which a signal is established with 90% confidence with 3σ significance or is excluded at 90% confidence (BR[Punzi]). This naturally takes into account signal and background so is therefore more appropriate than a single-event sensitivity with large backgrounds.

The BR[Punzi] is evaluated using 3×10^{18} muons incident on the stopping target which is the value expected for two years COMET running [43]. Backgrounds from radiative muon and pion capture, neutrons, and cosmic rays are neglected in this analysis since DIO is the dominant background. An event selection efficiency of 0.096 [43] is assumed which includes the effects of timing cuts, trigger and reconstruction efficiency and the acceptance of the electron spectrometer and detectors. The conversion and DIO electrons are smeared with a series of Gaussians and exponential decays to simulate the effects of energy loss from interactions with the stopping target and detectors. Timing cuts are also implemented to remove prompt background: transit times to the muon stopping target, time position in the pulse, muonic atom lifetime, etc. have all been simulated. The minimum time at which signals will be recorded is particularly sensitive to a change in target material, with the relative muonic lifetimes for a range of materials detailed in Tab. 3.2.

Fig. 3.16 is a reconstruction of the e^- momentum entering the tracker. Depending on their momentum, the signal and background may interact with further material in the stopping target region, leading to energy loss. The electrons then traverse the curved electron spectrometer solenoid and detector solenoid. The detector contains an electron calorimeter and tracker, both providing media of interaction. The resolution effects from the tracker will smear the electron energy, either up or down. A smearing with an RMS of 150 keV is added to each DIO and signal event to account for the tracker resolution [43]. This information, and the other values necessary to calculate the BR[Punzi], can be found in Tab. 3.4. This yields a BR[Punzi] of 5.9×10^{-16}, 1000 times better than achieved by SINDRUM-II.

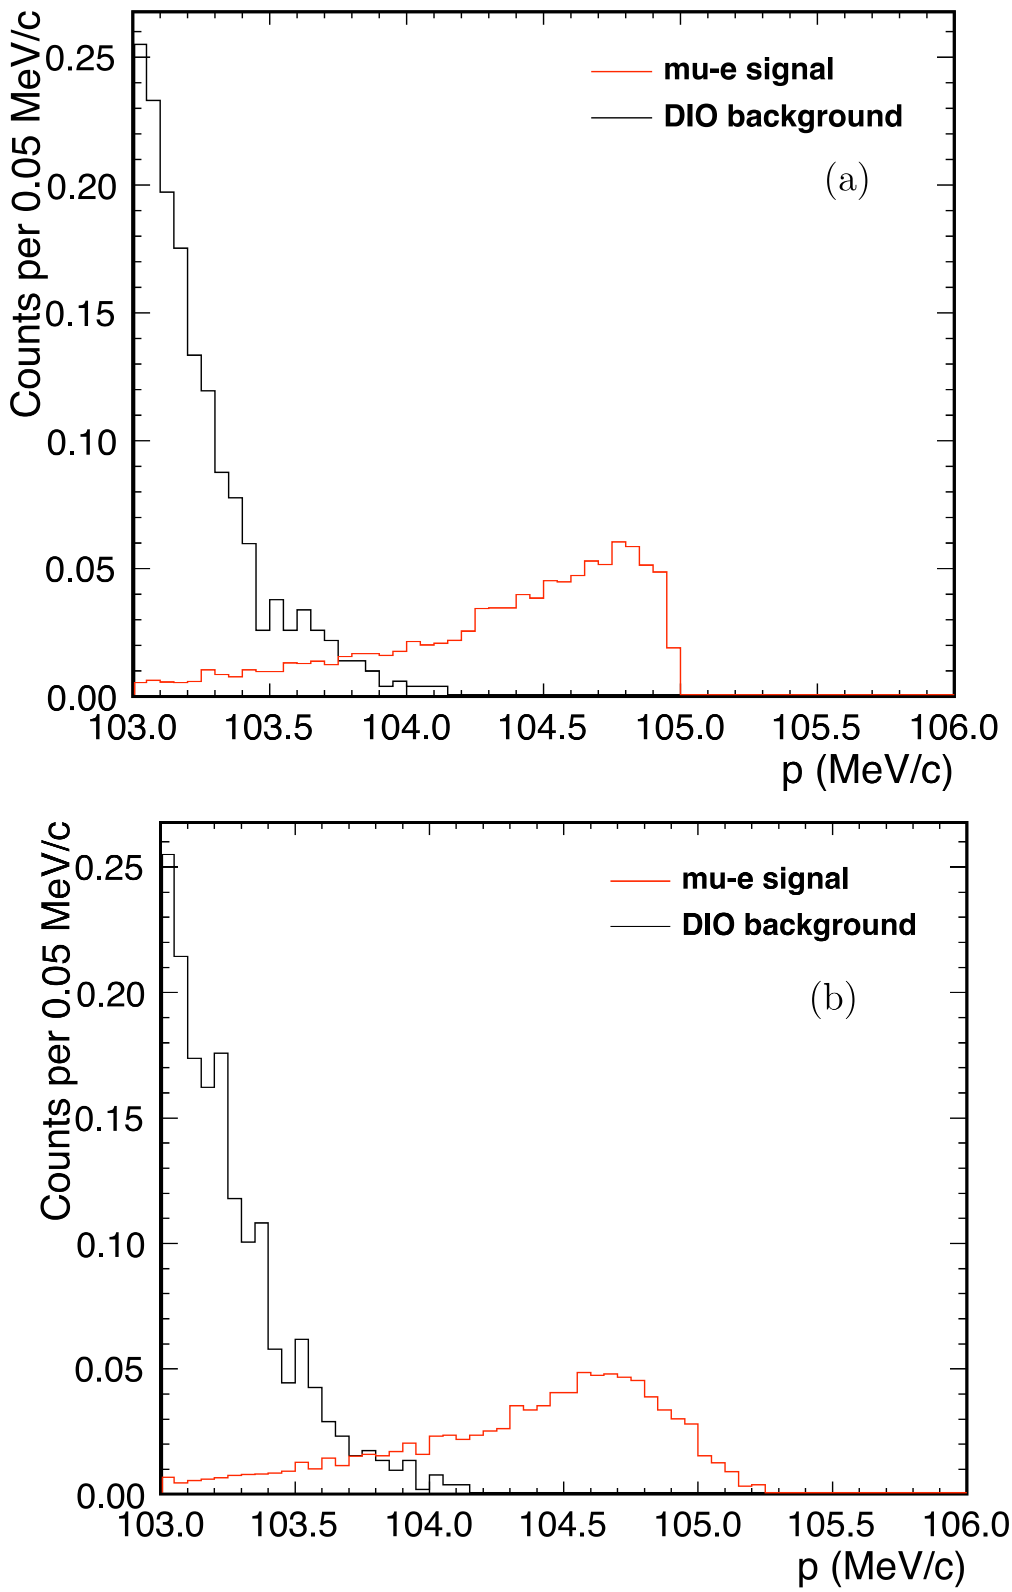

Figure 3.16: The expected $\mu - e$ conversion and DIO background results in the signal region, with the default stopping target configuration, a) as a raw output from the COMET_G4 code, and b) with a Gaussian smear due to a tracker resolution of 150 keV. A BR of 10^{-16} used in both plots.

total number of protons	8.5×10^{20}
beam power	56 kW
proton energy	8 GeV
total running time	2.0×10^{7} s
mean transit time to target	165.8 ns
mean transit time to detector	149.6 ns
timing window	1000 ns
detector resolution	150 keV
$N^{\text{stopped}}_{\mu^-/p}$	0.0024
$N^{\text{stopped}}_{\pi^-/p}$	3.3×10^{-7}
μ^- stopping efficiency	0.46
total number of stopped muons	2.0×10^{18}
signal region	$103.5 < p < 105.2$ MeV/c

Table 3.4: Breakdown of yields, timing cuts, and efficiencies expected for COMET with the default stopping target configuration.

The virtual monitors within the target section are used to quantify the evolution of the beam profile, for both the incident muon beam and resultant $\mu \to e$ electrons. These profiles reveal some pertinent properties of the muon beam: the x-plane remains approximately Gaussian across the entire region, described by almost constant values of σ and μ; the y-plane begins skewed, with an increasingly negative mean; both transverse planes are contained within ± 15 cm. It is this last property that suggests an increase in disk radius from the default 10 cm to one that would allow all muons to be incident on the stopping disks.

Making this increase requires a similar investigation to that previously outlined, calculating the number of stopped muons and the BR[Punzi]. The momentum distribution of incident and stopped muons with a disk radius of 15 cm (analogous to the plot in Fig. 3.15 with $r = 10$ cm) is shown in Fig. 3.17. Making this adjustment to the stopping target geometry decreases BR[Punzi] to 2×10^{-16}. A disk radius of 15 cm is therefore taken as the default from this point onwards.

Increasing the disk radius by 5 cm evidently stops more muons, stopping essentially all muons below 45 MeV/c. This leads to an increase in stopping efficiency of approximately 9%. The high momentum cut-off for stopped muons remains at approximately 60 MeV/c however, suggesting that an increase in the amount of stopping material (whether that be via an increase in disk radius or thickness) won't stop more muons. As this new stopping target configuration is

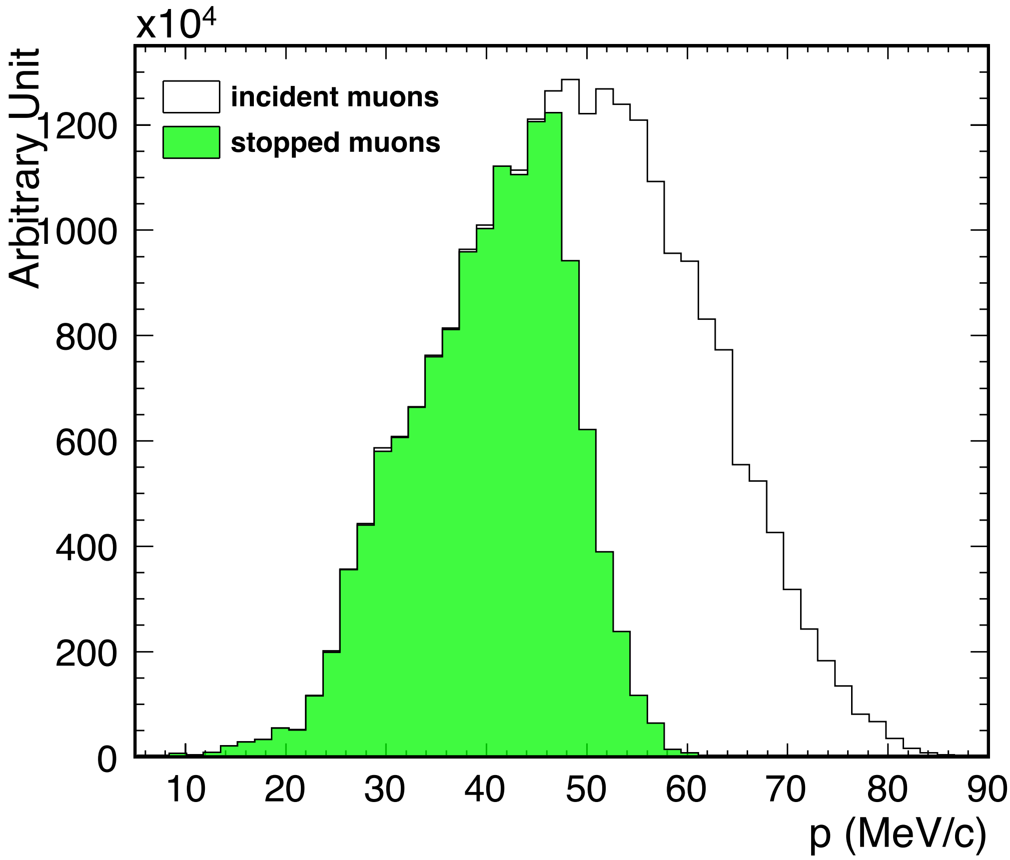

Figure 3.17: Momentum distribution of the muons incident on the stopping target. The momentum spectrum of those stopped is also shown. In this configuration the target disk radius is increased to 15 cm, resulting in the stopping efficiency increasing to 0.55.

well optimised for $p_\mu < 45$ MeV/c, there appears to be little room for geometric improvements without changing the material of the target.

Increasing the disk radius will require the signal electrons to have a larger transverse momentum in order to escape the target section. The distribution of signal electrons is isotropic leading to the possibility of a beam blow-up from more electrons interacting with the next disk. Fig. 3.18 shows the RMS of the transverse size of the incident muon and signal electron beam. The signal electron beam-size demonstrates an initial disparity between the values of the two transverse planes, as well as a net increase over the target region. However, the final (and largest) RMS is less than half the diameter of the disks.

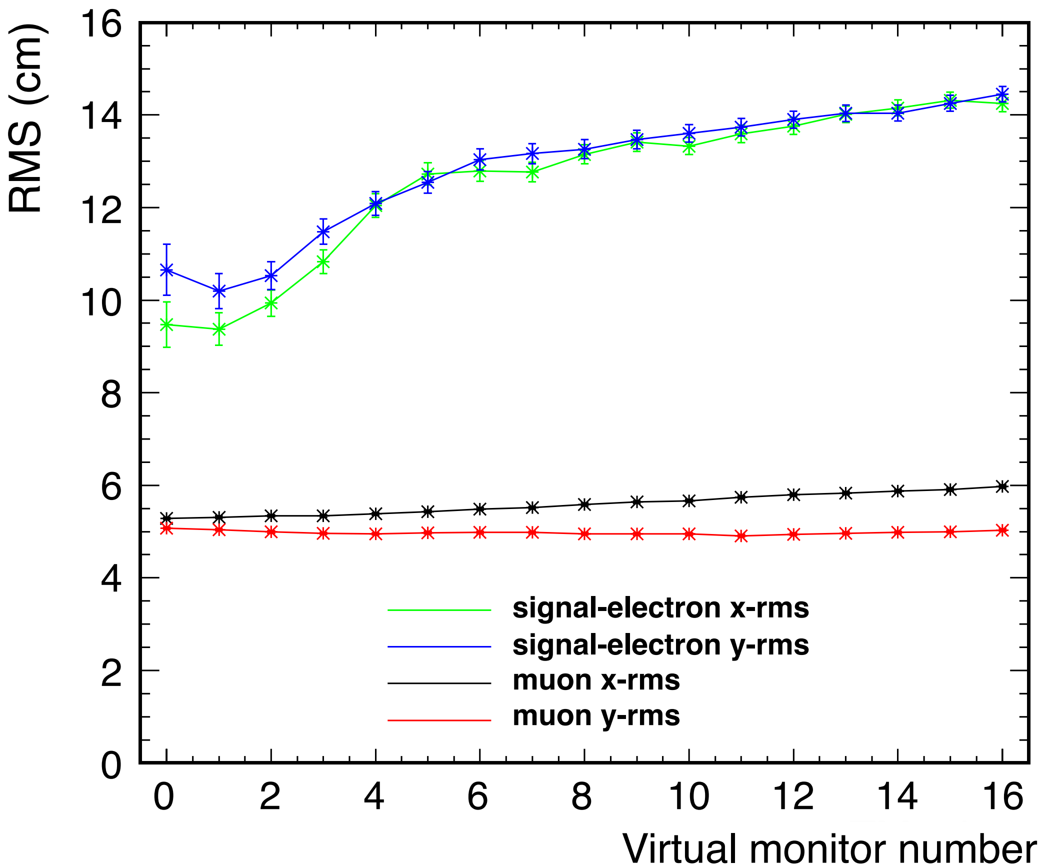

Figure 3.18: The transverse RMS of the incident muon and signal electron beams. The divergence of the muon beam is acceptable across the entire target region. The signal electron beam increases in size but converges and plateaus from the fifth disk onwards.

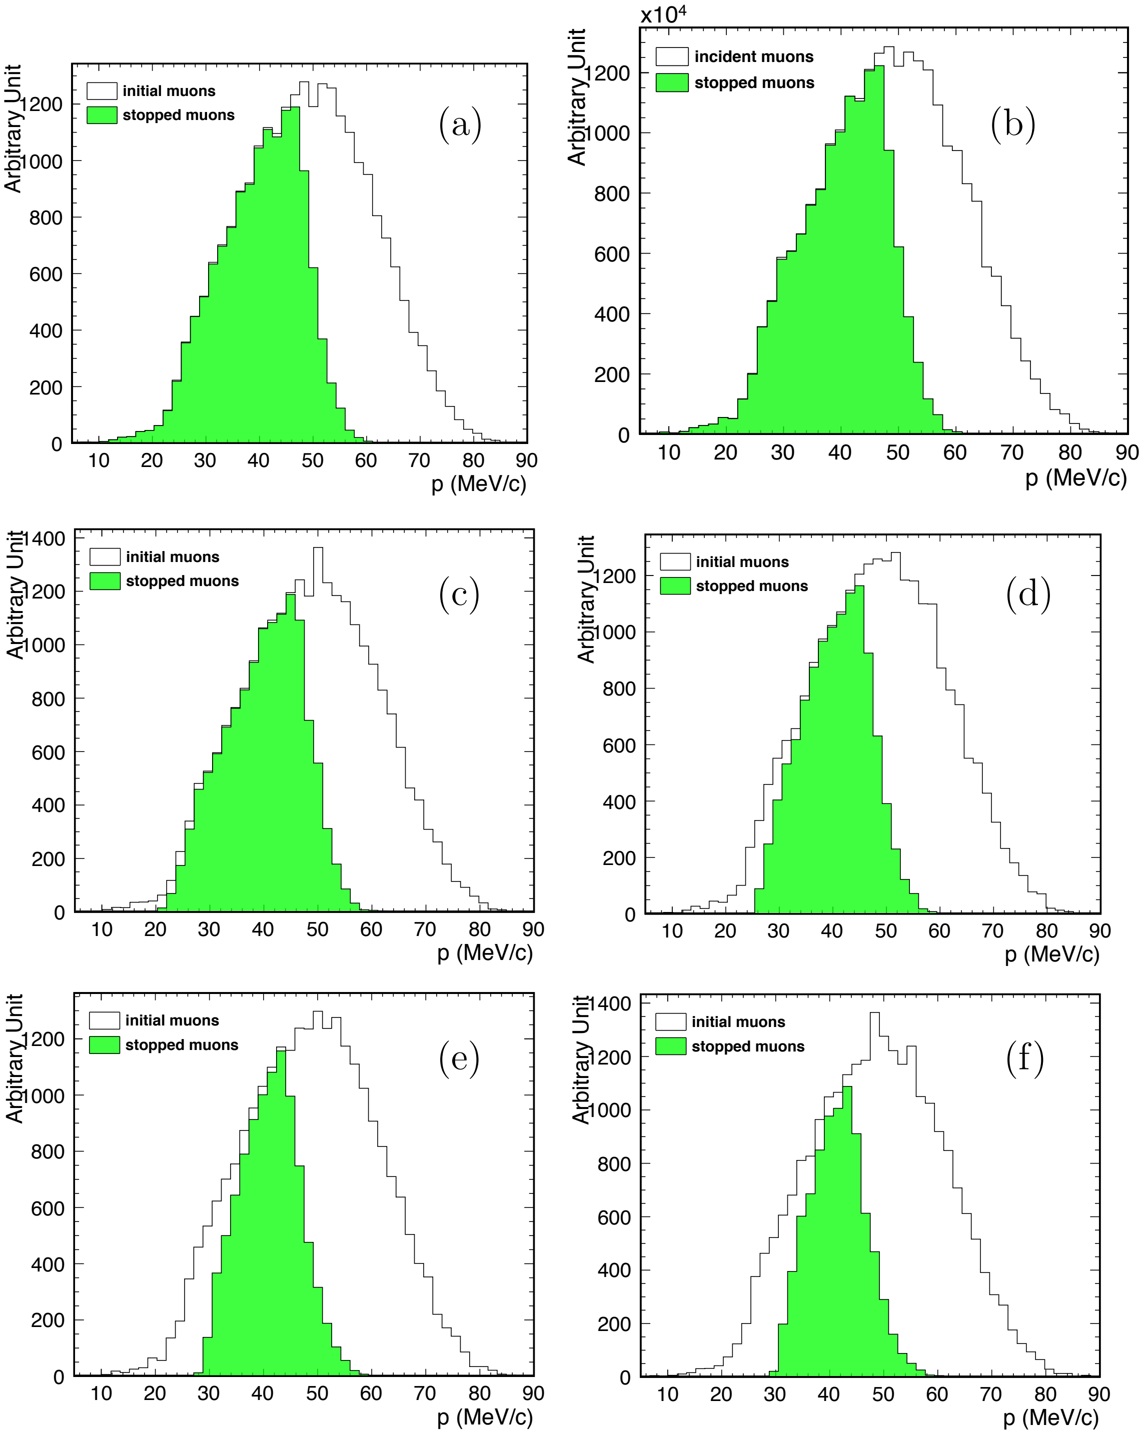

Figure 3.19: Comparison of momentum distributions for stopped muons using the default stopping target configuration but with a) 4 cm, b) 5 cm, c) 6 cm, d) 7 cm, e) 8 cm, and f) 9 cm disk spacing.

It is apparent that increasing the amount of material in the stopping target does not increase the number of stopped muons above $p_\mu = 60$ MeV/c, with the increase in disk radius to 15 cm stopping essentially all muons with $p_\mu < 45$ MeV/c. For this reason changing the disk thickness or increasing the number of disks would be a redundant route of optimisation. The only geometric variable independent of material quantity is the disk spacing, with a default spacing of 5 cm included in simulation thus far. Fig. 3.19 shows the momentum spectra for stopped muons with a spacing ranging from $4-9$ cm. As can be seen the incremental increase in disk spacing decreases the number of stopped muons in the lower momentum region while retaining the upper limit of 60 MeV for stopped muons. Integrating the momentum spectra indicates that the maximum stopping efficiency is found using the default 5 cm spacing. The equivalent BR[Punzi] for these configurations can be seen in Tab. 3.5, also demonstrating the disk spacing as already optimised.

The previous results suggest that the only remaining variable that may improve the stopping efficiency at higher momenta, without increasing the amount of material in the beamline, is the disk material. Tab. 3.2 gives a selection of potential stopping target materials, detailing their muonic atom properties. Aluminium has been taken as the default material but titanium has a branching ratio 1.7 times that of aluminium, with a muonic lifetime still compatible with the pulsed beam structure at J-PARC. The timing window needs to be reconsidered for titanium

disk spacing [cm]	BR[Punzi] [$\times10^{-16}$]
4.0	2.3
5.0	2.2
6.0	2.4
7.0	2.7
8.0	3.1
9.0	3.6

Table 3.5: BR[Punzi] for a range of disk spacings with the now default 15 cm radius stopping target geometry.

as its muonic lifetime is approximately half that of aluminium (see Tab. 3.2). The timing window is therefore opened 550 ns earlier for titanium, accounting for their approximate difference in muonic lifetime. This will open the window to more prompt background, radiative pion background, and beam splash.

The momentum spectrum of stopped muons in titanium is shown in Fig. 3.20. The efficiency has increased by approximately 9% on that of aluminium. The main difference between the spectra of the two materials is the increase in the high momentum cut-off, moving up to approx. 65 MeV/c. This is expected as titanium has a larger dE/dx than aluminium, leading to an increase in the fraction of higher momentum particles being stopped. There is a slight decrease in efficiency from $30 - 50$ MeV/c but this is more than offset by the increase from $50 - 65$ MeV/c.

Sec. 3.1.1 outlined the process employed to simulate the DIO background for different disk materials. Titanium was taken as the example for extending the DIO model as it is the most likely stopping target candidate after aluminium. The momentum of signal and background is displayed in Fig. 3.21.

The signal region for aluminium is selected as $103.5 < E_e < 105.2$ MeV in order to maximise the number of signal electrons and minimise DIO background. If the region is narrowed to $104.0 < E_e < 105.2$ MeV for example, the DIO background will drop but the loss of signal will decrease by a far greater amount. The titanium endpoint is approximately 0.7 MeV/c lower than that of aluminium, so the signal region is shifted down by this value to $102.8 < E_e < 104.5$ MeV. The number of stopped muons and pions per proton is 0.0033 and 4.59×10^{-7} respectively. As well as the adjusted signal region the timing window must be opened earlier in order to account for the shorter muon lifetime in titanium. The difference in muon lifetimes between aluminium and titanium is approximately 550 ns, so the timing window is initially opened after 450 ns. This value was then optimised, with a BR[Punzi] for titanium of 4.89×10^{-16} found with a timing window opened after 515 ns.

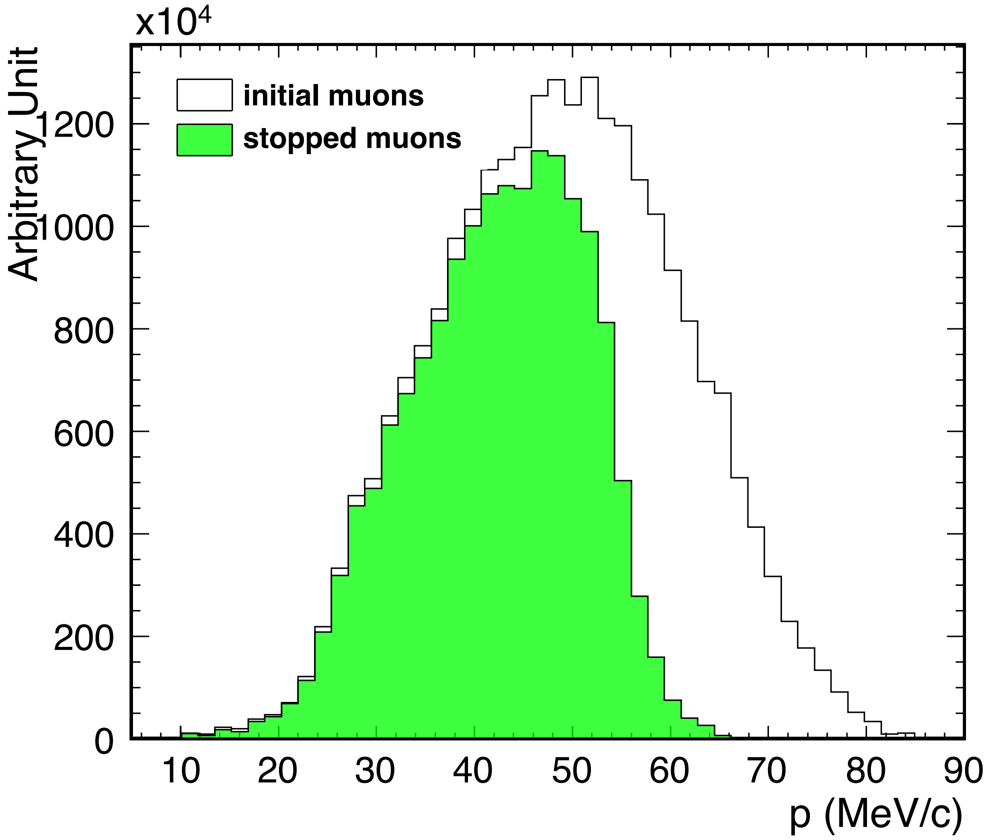

Figure 3.20: Momentum distribution of the muons incident on the stopping target. The momentum spectrum of those stopped is also shown. In this configuration the target disk material is changed to titanium, with an increased stopping efficiency of 0.64.

This new configuration is more efficient at stopping muons and has a BR[Punzi] comparable to aluminium in the adjusted timing window and signal region. The geometry, with parameters listed in Tab. 4.1, is therefore preferential to the default setup currently utilised in preliminary design studies for the COMET experiment.

Tab. 3.7 shows the relevant numbers for the three major stopping target geometries explored in this analysis. As can be seen the BR[Punzi] shows a slight decrease upon a change in target material, however of the same order of magnitude as aluminium. This, coupled with the relative signal strengths of Figs. 3.21 and 3.16, indicate titanium and aluminium will yield comparable results.

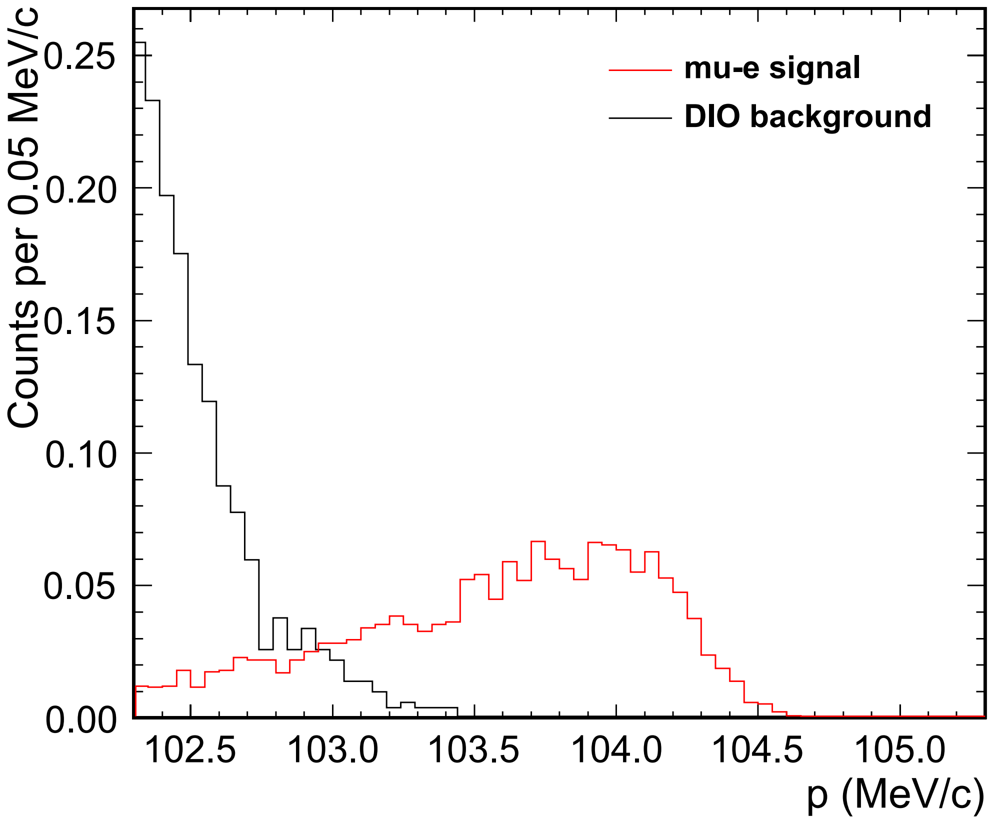

Figure 3.21: Simulated μ-e conversion and DIO background results, with the stopping target disk material changed to titanium.

parameter	value
disk radius	150 mm
disk thickness	200 μm
number of disks	17
spacing	50 mm
material	titanium

Table 3.6: New parameters for the optimised muon stopping target.

material	radius [cm]	stopping efficiency	BR[Punzi]
Al	10.0	0.46	6×10^{-16}
Al	15.0	0.55	2×10^{-16}
Ti	15.0	0.64	5×10^{-16}

Table 3.7: Stopping efficiency and BR[Punzi] for a range of COMET stopping target geometries.

Another way to utilise the additional stopping target monitors is to assess the number of muons stopped by each target disk. This information for the default configuration is displayed in Fig. 3.22. The distribution is not uniform across the disks. This suggests that changing the shape of the target region to a cone is worthy of consideration, but offers two conflicting arguments as to the nature of its geometry: the disk radius should start small and increase downstream in order to stop more muons; the majority of signal electrons are produced in the first few disks so the radius should decrease in order to minimise scattering with further disks.

The SINDRUM-II experiment, searching for the decay $\mu \to eee$, adapted to this non-uniformity of stopped muons by optimising their design to maximise the number of 28 MeV/c muons stopped. The baseline design employed a thin

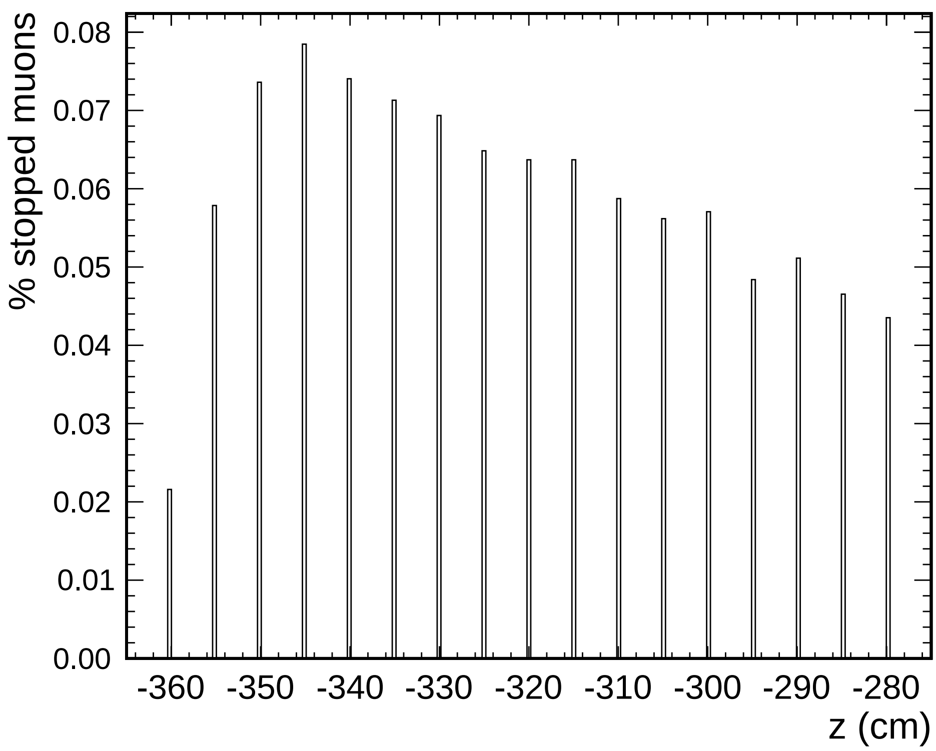

Figure 3.22: The number of muons stopped by each target disk, with the first disk at $z = -360$ cm.

(approx. 60 μm) hollow aluminium double cone, of 100 mm length and 20 mm diameter [44]. This design will be discussed further in Sec. 4.3.

3.1.3 Conclusions

The experimental search for $\mu \to e$ conversion has been proposed as a probe into BSM physics, along with new experimental methods to improve the current sensitivity by $\mathcal{O}(10,000)$.

Benchmarking these experiments is an essential element to achieving this goal and, as such, a sophisticated and robust software framework is required. Both beamline mechanics and physics processes have been incorporated into this framework in the form of magnetic field implementation and DIO calculations respectively.

These successful amendments provided the necessary tools to simulate the experiment in its entirety, with the content of this analysis focussing specifically on the optimisation of the muon stopping target. The necessity to capture low momentum muons in order to minimise stopping target material has been shown, validating the need for a novel type of superconducting pion capture system and momentum selection in the two 90° bend sections.

The stopping target was optimised to maximise the stopping efficiency of the default geometry (see Tab. 3.2) whilst minimising the BR[Punzi]. This was initially achieved by increasing the disk radius from 10 to 15 cm after consideration of the transverse beam size incident on the stopping target region. After an unsuccessful exploration of other variables (disk spacing, thickness, etc.) the `COMET_G4` code was developed to model other disk materials. Titanium was the next choice due to its higher density but also as it has a muonic lifetime compatible with the pulsed beam structure at J-PARC. The pertinent statistics garnered from this analysis are shown in Tab. 3.7. As can be seen the stopping efficiency improves at the expense of a slight deterioration of BR[Punzi]. This suggests that the titanium geometry should be used in preliminary runs to probe as far into the current limit with its higher stopping efficiency and larger signal peak. If this is unsuccessful then the aluminium target should then be utilised to place a lower exclusion limit on the process.

If only one target setup were to be used continuously for the proposed two year run then prudence would suggest the implementation of the 15 cm Al disk geometry as this represents an acceptable compromise between the stopping efficiency and BR[Punzi].

3.2 PRISM

If a cLFV signal is observed by COMET it will be imperative to determine the nature of the beyond SM physics causing it. This will require running with a higher Z stopping target since the ratio of cLFV rates between low and high Z is strongly dependent on the type of interaction e.g. a vector versus a tensor interaction [45]. It will likely also require a significant improvement in experimental sensitivity so that a large signal sample can be collected. If no cLFV signal is observed then an improvement in experimental sensitvity will be mandatory. There are many ideas as to how one could improve the Mu2e/COMET sensitivities by a factor of 100 and one of the most promising, PRISM, would exploit a Fixed-Field Alternating-Gradient (FFAG) accelerator to phase rotate the muon beam. This type of accelerator will be discussed in more detail in Sec. 3.2.3 in the context of the EMMA project.

Phase rotation is a well established technique employed to reduce the momentum range of a beam of particles by accelerating slower moving particles and decelerating faster moving ones. This can then result in an almost monochromatic beam whose momentum can be tuned, through the acceptance, into the FFAG. A low momentum (approximately 20 MeV/c), monochromatic muon beam will improve the sensitivity of muon conversion experiments due to several factors:

- the low momentum means a higher fraction of muons will be stopped,
- the monochromatic nature means that the muons will all stop after a well defined thickness of target and so the origin of signal electrons will be much better localised allowing for greater discrimination against background when combined with the electron trajectory in the tracking detector,
- the low momentum means that a thin stopping target can be used, thus reducing the energy loss of the signal electrons, affording a greater discrimination against the DIO background.

Furthermore, since the muons and pions will cycle the FFAG several times, the additional path length will effectively eliminate the background from pions, as even the most energetic pions will have decayed. The acceptance of the FFAG will also be such that the decay products of both muon and pion decays will not enter the stopping target beamline.

A schematic of the proposed PRISM experiment can be seen in Fig. 3.23, where the FFAG phase rotates the muon beam before it interacts with the stopping target further downstream.

This technique of phase rotation, however, has never been used with such a large emittance beam before (possibly as large as 10000 π mm mrad). An FFAG is the obvious choice in this case since it has a high acceptance for large emittance beams.

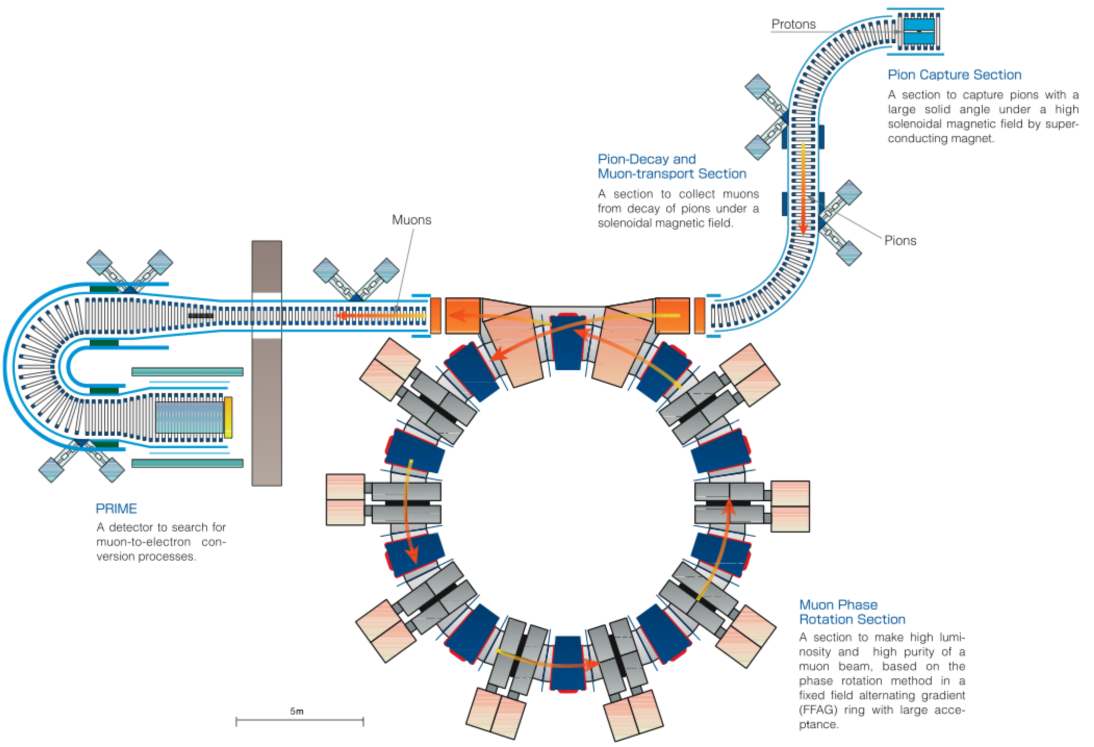

Figure 3.23: A schematic of the proposed PRISM beamline, with the inclusion of an FFAG for phase rotation and additional decay length.

3.2.1 Phase rotation

Phase rotation works on the principle of longitudinal bunching, such that higher energy particles occupy the front of the bunch and lower energy particles sit at the back. The RF cavity has an initial decelerating field, slowing the initial muons, then transitions to a positive field for the slower muons at the back of the bunch. As energy and time are canonical conjugates, Liouville's theorem is preserved.

The effect of phase rotation on a bunch in (t, E) phase space can be seen in Fig. 3.24. The particles initially have a wide energy spread but a narrow temporal spread. They then undergo phase rotation, rotating the phase space ellipse such that the energy spread becomes very narrow. However, as can be seen the time spread of the final beam is increased significantly and so in order for this to be contained within the measurement window between primary proton pulses it is necessary for the original time spread to be very small (20 ns) which requires improvements in the way the protons are bunched. For COMET and PRISM to achieve their expected sensitivities the number of residual protons between pulses needs to be at least 10^9 times smaller than the number of protons in the main pulse. Improvements to the beam extinction, and the pulsed proton beams necessary to achieve this in next generation cLFV experiments, such as COMET, are detailed earlier in this chapter. It is the combined effect of these improvements that will allow PRISM to reach a proposed single-event sensitivity of 3×10^{-19}.

3.2.2 Lattice design

Any accelerator used to perform phase rotation for a muon beam is required to accept a bunch with an emittance upwards of 10^5 π mm mrad. In comparison the electron beam injected into EMMA's non-scaling FFAG ring has a typical emittance of 10 π mm mrad. At present FFAGs are the only accelerator type capable of accommodating these large emittances.

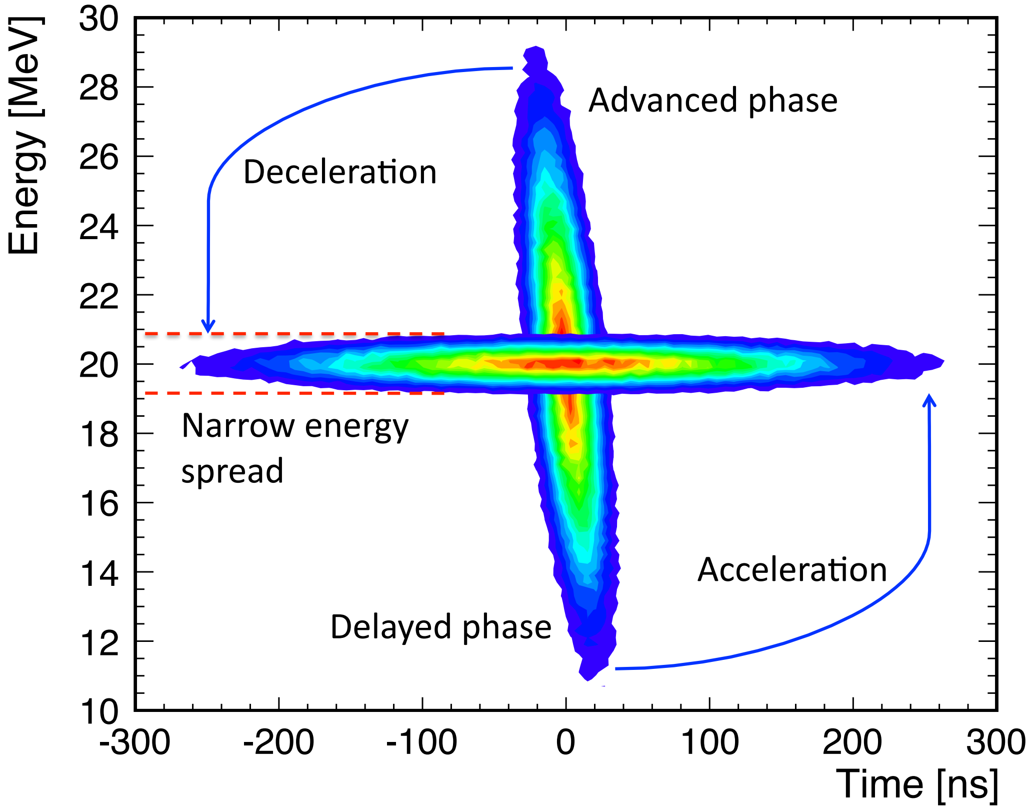

Figure 3.24: A simulation of the phase rotation effect of the PRISM FFAG ring on a beam of muons with a typical energy spread expected in the experiment.

The different types of (already proven) FFAG machines capable of this level of acceptance, as well as the required levels of rapid acceleration, are described in Sec. 6.1.1. The FFAG ring originally suggested for PRISM is a scaling FFAG with 10 identical DFD triplet cells (cells with a combined defocusing-focusing-defocusing series of magnetic elements). A prototype of this accelerator has been constructed at Osaka with six cells, operating with alpha particles [46].

The muons injected into the PRISM ring remain there for approximately five turns, taking about 1.5 μs to traverse this distance. This affords additional time for the pions to completely decay, while a reasonable fraction of the muons remain undecayed.

3.2.3 PRISM@EMMA

When the proposal was made for the phase rotating lattice to use a scaling FFAG the proof-of-principle non-scaling FFAG ring, EMMA, was yet untested. However, since commissioning concluded and acceleration was demonstrated in early 2011 (the results of which are analysed and presented in Ch. 6) non-scaling FFAGs have become a possible avenue of exploration for incorporation into future cLFV experiments.

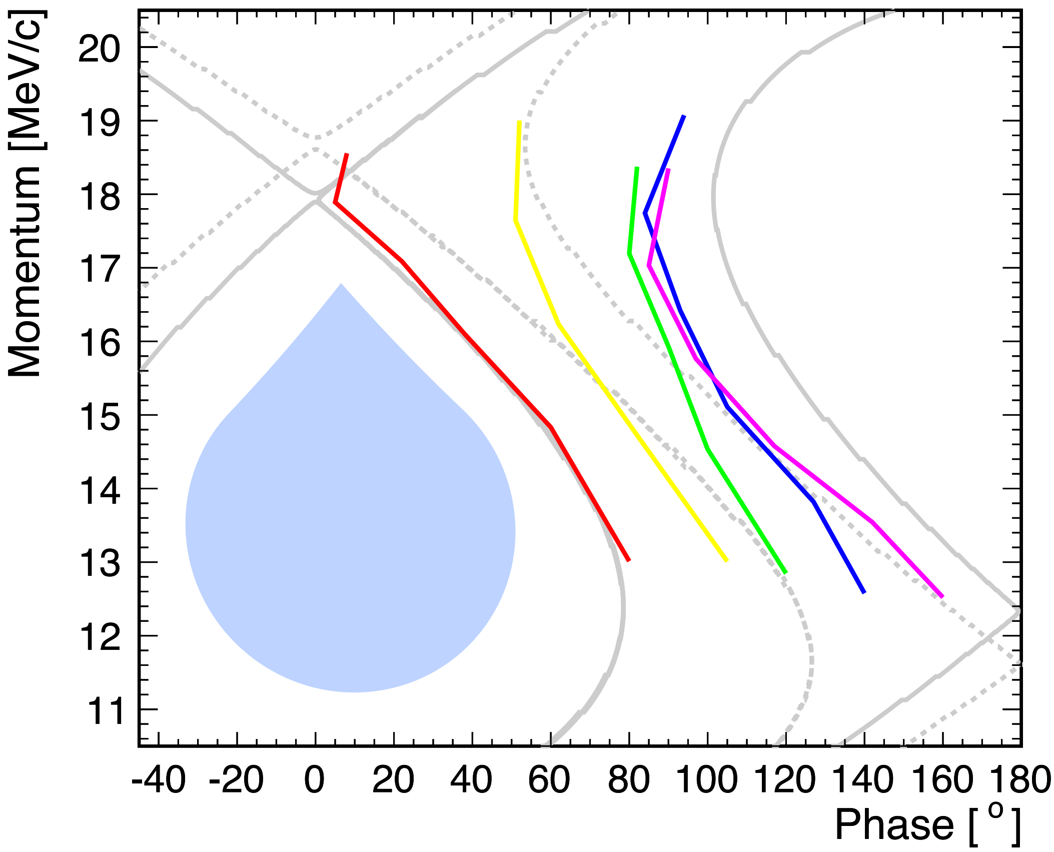

Figure 3.25: Longitudinal phase space trajectories of beams with five different initial phases. Acceleration along the serpentine channel is demonstrated within the separatrix boundary between in-bucket motion denoted by the grey region. This boundary was calculated using the estimated systematic error of ±25 ps in the orbital period measurements of Fig. 6.15. The blue region at low phase and momentum indicates the bucket used for phase rotation on PRISM@EMMA.

The main goal of EMMA commissioning was to observe a parabolic time of flight and serpentine acceleration, indicative of a non-scaling FFAG machine. This novel type of acceleration is described by the particle bunches following an S-shaped path on a plot of energy versus phase, and this has now been observed. Fig. 3.25 illustrates the serpentine acceleration recorded by EMMA while the parabolic time of flight measurements are described in more detail in Ch. 6 and shown in Fig. 6.15.

Having demonstrated these principles EMMA's non-scaling FFAG ring may be used to test phase rotation by accelerating a bunch in the bucket, denoted by the light blue region of Fig. 3.25 (at low momentum and negative phase). This region will then transform into the usual RF bucket seen in EMMA. Achieving this requires setting a high RF gradient to maximise the size of the bucket, measuring the orbit during the period of acceleration to verify the effect.

In accelerator design a buncher is used to shorten the length of the bunch at the expense of an increase in energy spread. Fig. 3.24 indicates phase rotation having the opposite effect of a buncher, lengthening the bunch and reducing the energy spread. An expected muon beam in PRISM is much longer than that available to EMMA from ALICE. In order to mimic the bunch expected in PRISM the bucket is filled with approximately 3 back-to-back 'bunchlets' from ALICE. The bucket denoted in Fig. 3.25 is thus filled either vertically by changing the injection momentum or horizontally by altering the phase. Phase rotation may then be demonstrated by observing a 90° rotation in either plane. PRISM@EMMA proposes to perform both experiments with varied momentum and phase.

After injection into the EMMA ring the bunch is extracted at certain points corresponding to known synchrotron oscillations: 3 turns for 1/4 of an oscillation, 6 turns for 1/2, etc. Extracting the beam at different fractions of its synchrotron oscillation allows the mapping of the longitudinal phase space as a function of the transverse amplitude. In this case the bunch is initially extracted with no

RF after 1/2 a revolution of the machine, then again at 3/2, as it is prudent to perform initial tests with a minimal number of revolutions.

The bunch is monitored using a YAG beam profile monitor in the first dispersive section of the extraction line, immediately after the extraction septum. The transverse dimensions of the beam after 1/2 and 3/2 revolutions, with no RF, can be seen in Fig. 3.27. Both plots demonstrate a correlation in transverse position and size indicating stable propagation after an additional revolution of the ring.

Once the bunch is successfully injected and extracted the RF may be switched on, with the voltage and phase set to produce a stationary bucket. In order to reproduce the results with no RF the magnet strengths of the entire lattice were scaled to reach an effective energy. This was achieved iteratively with a final scaling factor of 1.044 established. The beam was then extracted, again after 3/2 turns, with the result shown in Fig. 3.28. The two profiles after 3/2 turns are not identical but do serve to demonstrate the reproducibility of results with no RF compared to bucket acceleration with RF switched on.

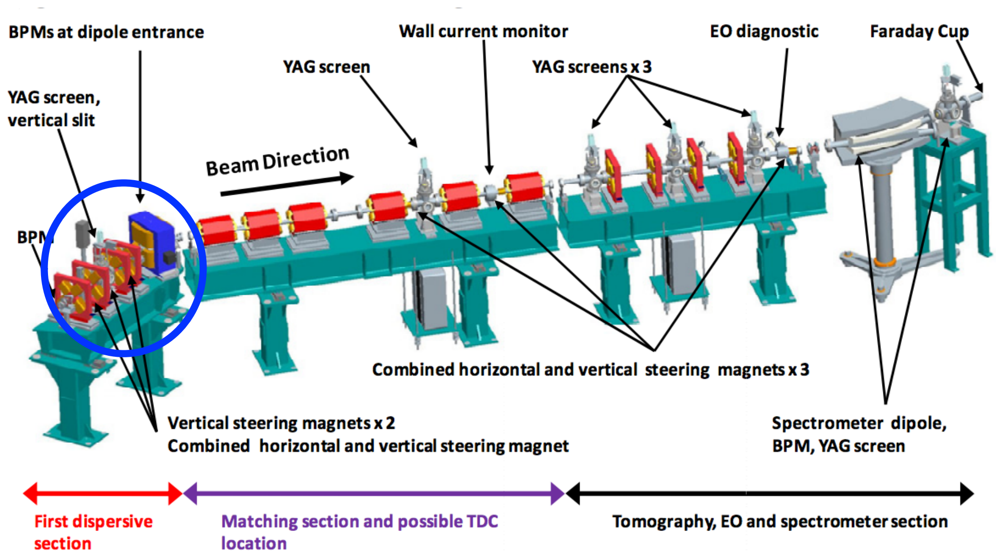

Figure 3.26: Schematic of the EMMA extraction line with the first YAG screen used to output the transverse beam profile highlighted by the blue circle.

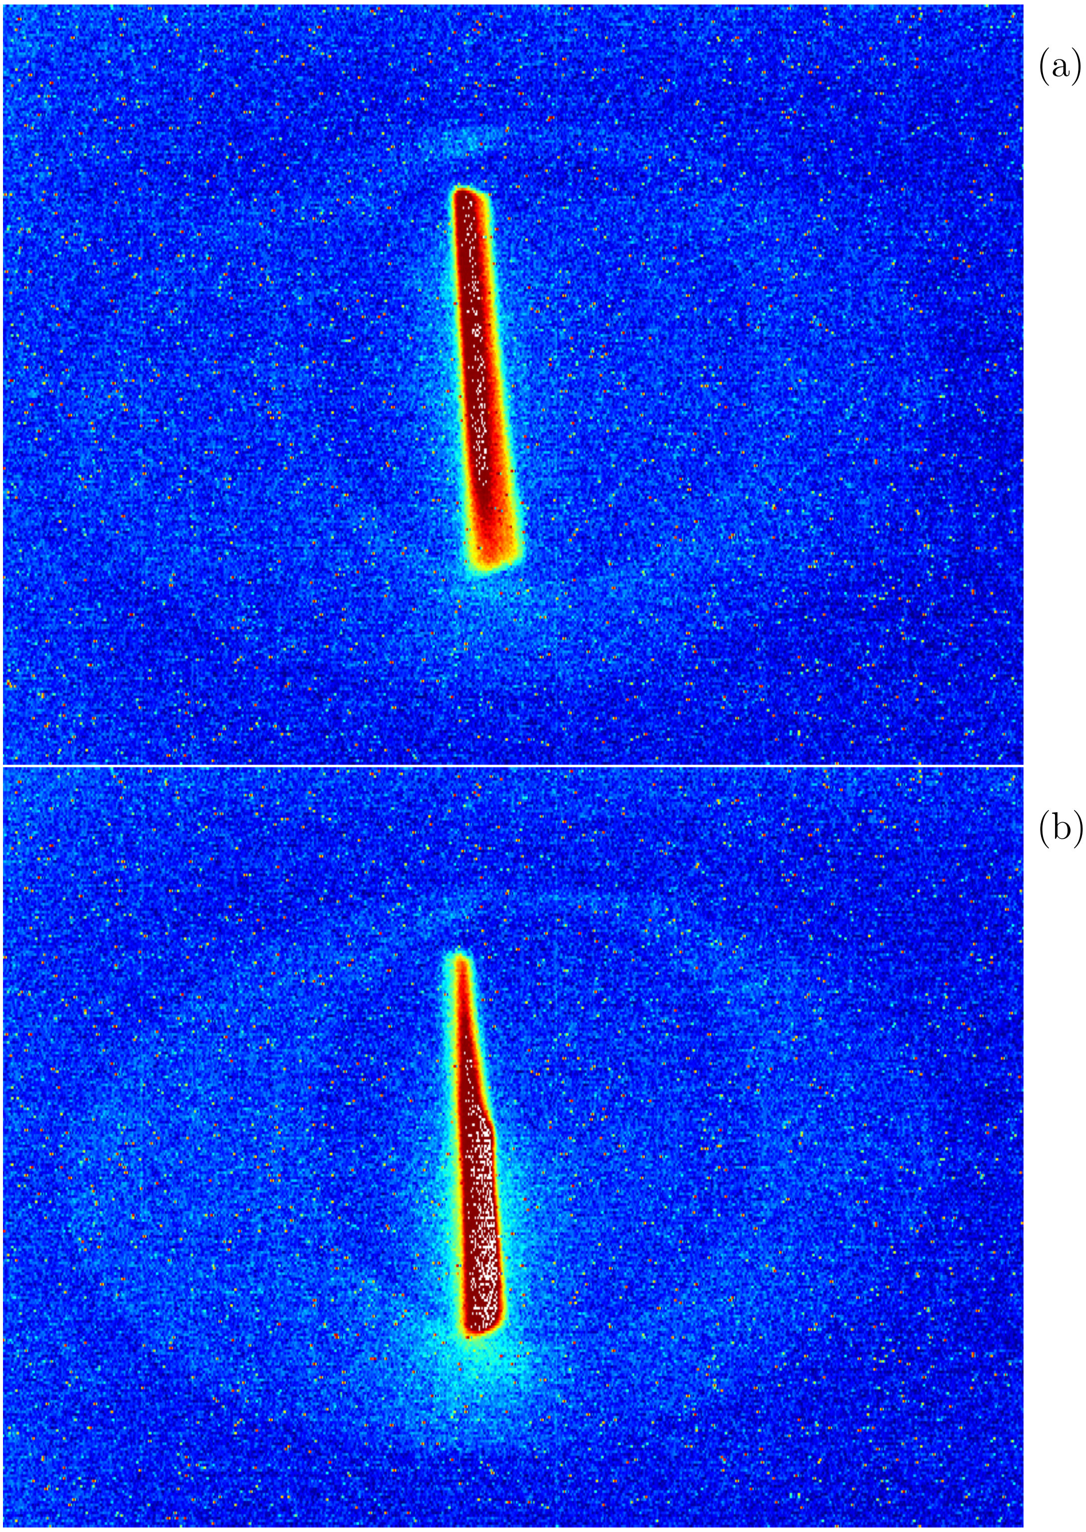

Figure 3.27: Transverse beam profile of the bunch at extraction after a) 1/2, and b) 3/2 turns of the EMMA ring, recorded with a YAG screen. The red region represents a high density beam with the blue region indicating no beam. The CCD camera used to capture these images was placed downstream of the YAG screen with the beam coming out of the page. The y- and x-axis are those of the vertical and horizontal respectively. The same applies to that of Fig. 3.28.

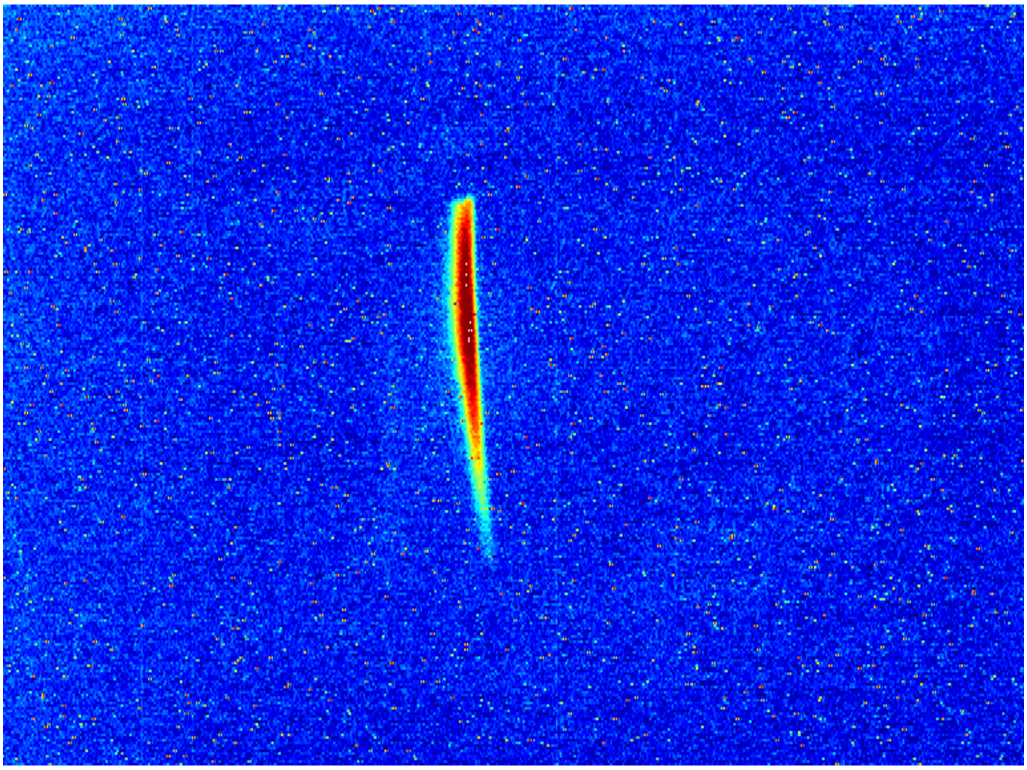

Figure 3.28: Transverse beam profile of the bunch at extraction after 3/2 turns of the EMMA ring with the RF switched on.

This process would then have been repeated for further turns, each time producing the longitudinal phase space in order to map the effects of the synchrotron oscillation. However, due to the the lack of experimental run time this could only be performed up to 3/2 turns of the machine. Similarly better settings may be achieved for 1/2 and 3/2 turns with further tuning of the machine. This work is ongoing but is a successful first step in demonstrating that ns-FFAGs could be used for phase rotation in the next generation of cLFV experimentation.

Chapter 4

COMET Phase-I

The COMET experiment is seeking to improve the sensitivity to the neutrinoless muon-to-electron conversion process by a factor of 100 beyond SINDRUM-II. This will be achieved using a pulsed muon beam produced from a pion capture solenoid using an 8 GeV proton beam. The production of muons of this intensity using this methodology has never been attempted before and as such presents considerable challenges. Extrapolating backgrounds over four orders of magnitude from SINDRUM-II to COMET where theoretical models are only loosely constrained by data has significant uncertainties. The operation and detailed understanding of the detectors, beam, and accelerator in a new regime will also not be without difficulties. In light of these considerations it was decided in January 2012 to stage the construction of COMET so that valuable operational experience could be gained and first physics results obtained on a shorter timescale. The results from this critical initial running with a reduced beamline, dubbed COMET Phase-I, will be used to gather valuable experience and to inform and optimise the design of the final COMET implementation and so reduce the risk that the physics goals of COMET will not be met.

A schematic of the COMET Phase-II beamline is shown Fig. 4.1, with the planned construction of Phase-I highlighted in blue. It comprises the pion capture solenoid and first 90° section of the muon beamline of the final COMET implementation. The primary goals of COMET Phase-I are:

- a direct measurement of the proton beam extinction factors and other potential background sources for the full COMET experiment,
- a search for $\mu \to e$ conversion with a sensitivity at least a factor of 100 better than the SINDRUM-II limit.

However, to achieve this physics goal where the beam power (3 kW) and running time (10^6 s) will be significantly less than the final experiment it will be necessary to improve the acceptance of the experiment. A cylindrical detector in a uniform 1 T field will thus be used in Phase-I as opposed to the lower acceptance (but higher resolution) transverse detector of Phase-II. Fig. 4.2 contains one proposed layout for Phase-I, demonstrating the nature of the cylindrical detector solenoid positioned at the end of the first 90° bend section. The stopping target will thus be at the center of a cylindrical detector in a solenoidal magnetic field which is very different from COMET Phase-II where a graded magnetic field is used upstream of the detector. Initial investigations into the effect of this new magnetic field are performed and detailed in Sec. 4.1, with the implemented field strength shown in Fig. 4.3.

Consequently it has been necessary to determine whether the default stopping target configuration (see Tab. 3.3) is optimal for Phase-I and a double cone target, as used by SINDRUM-II, has been investigated (see Sec. 4.3). The restricted beamline of COMET Phase-I will also result in a large increase in the flux of high momentum muons and pions that, through decay, could result in a significant number of electrons in the signal region. A study of implementing a collimator at the end of the beamline before the detector section has therefore been undertaken (see Sec. 4.2). Both these investigations have required significant changes to the COMET GEANT4 software and these changes are described in the next section.

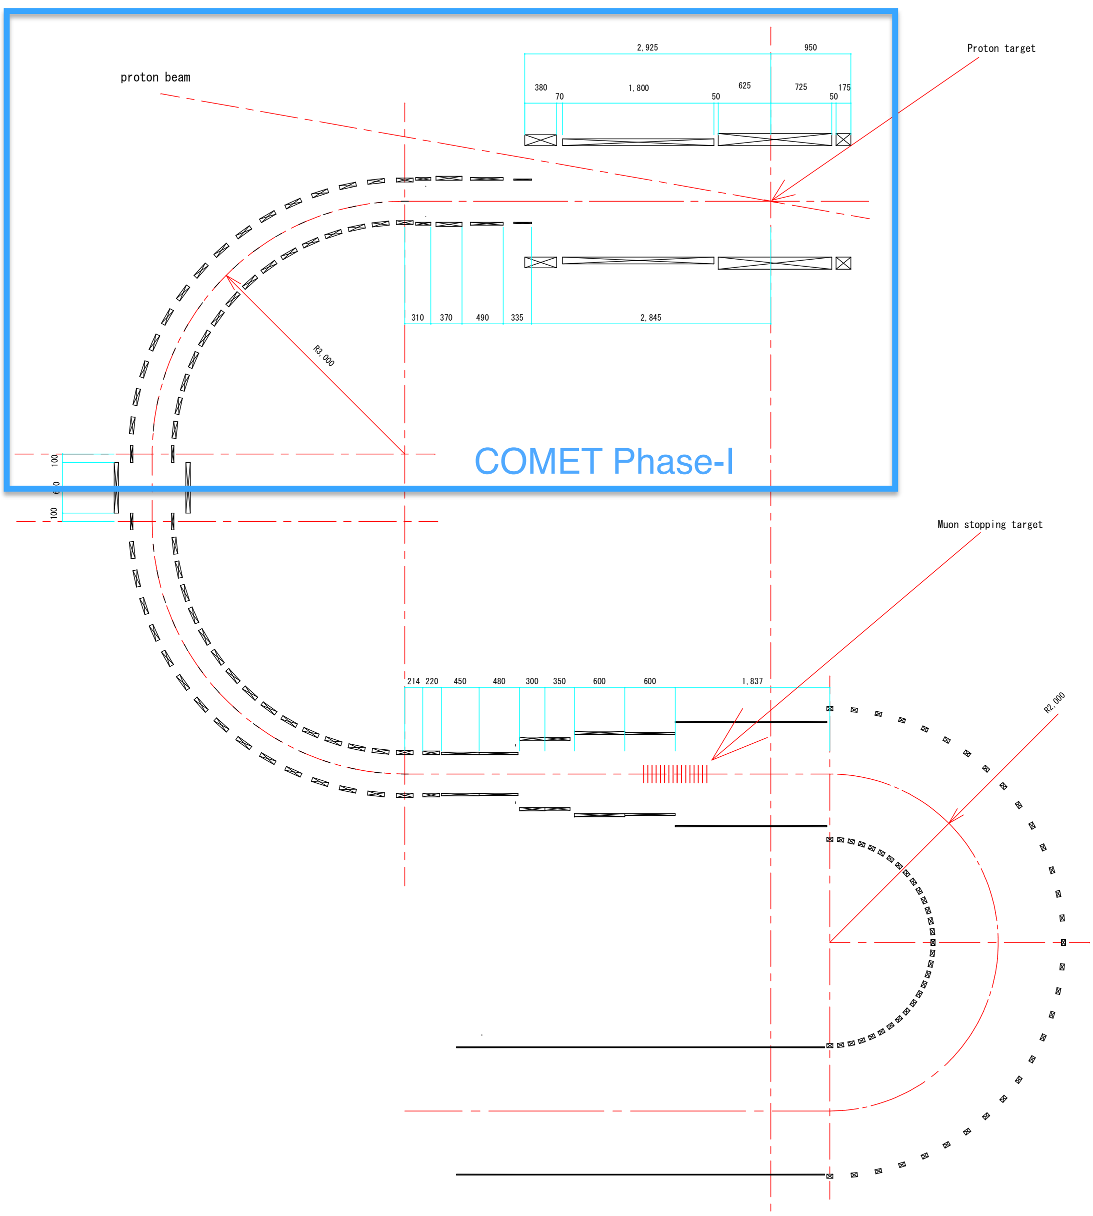

Figure 4.1: A CAD of the COMET beamline, with the Phase-I section of beamline highlighted within the outlined blue region.

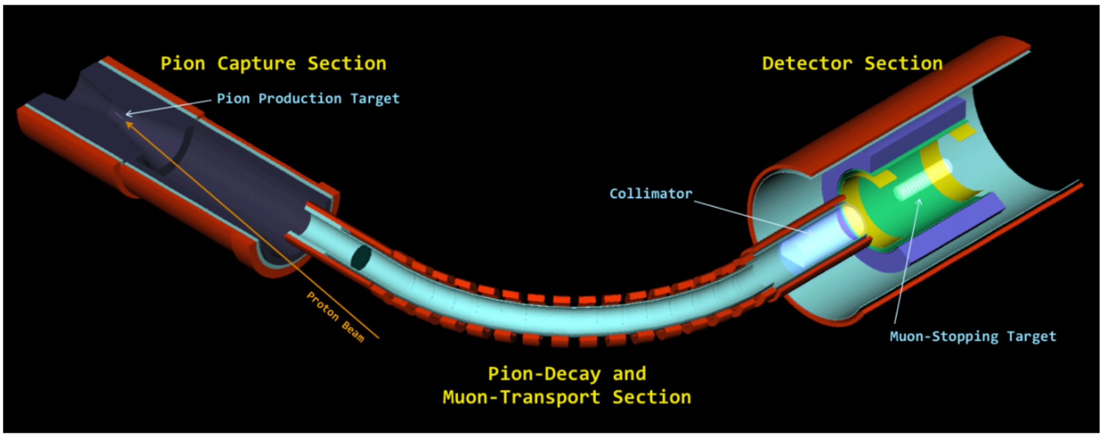

Figure 4.2: A visualisation of the proposed Phase-I beamline, identical to that of Phase-II until the end of the first 90° bend section. The beamline elements under investigation, contained within the cylindrical detector solenoid i.e. the collimator and muon stopping target, are illustrated and highlighted.

4.1 Magnetic field implementation

The Phase-I beamline geometry is identical to that of Phase-II up to the end of the first 90° bend solenoid, as highlighted in Fig. 4.2. Field maps provided by Toshiba, and modelled in TOSCA, are utilised for this section of lattice. Beyond this point detailed field maps do not exist for Phase-I and so approximations have been made to model this region.. The straight section immediately after the bent solenoid is designed to transport the beam to the detector solenoid and remove unwanted particles en route. This selection is achieved by situating a tungsten collimator in this straight section, scraping high momentum particles from the beam profile in a graded magnetic field.

The graded magnetic field across the beamline transport solenoid drops from the total field strength at the end of the 90° bend (approximately 3.2 T) to a constant value of 1 T. The field then remains constant (at 1 T) across the whole detector solenoid. This gradient is also present in the full Phase-II model in order to accommodate the dispersive nature of the beam upon exit of the bent solenoid, attenuating beam blow-up and maximising the functionality of the collimator. This is achieved via the process of adiabatic transition, as outlined in Sec. 3.1.1.

The nature of the magnetic field strength for Phase-I after the 90° can be seen in Fig. 4.3, with the z-axis positions of the beam collimator and stopping target highlighted. This gradient is beneficial since it mitigates the dispersion created by the momentum selection in the bend, which has a dipole field of 0.018 T.

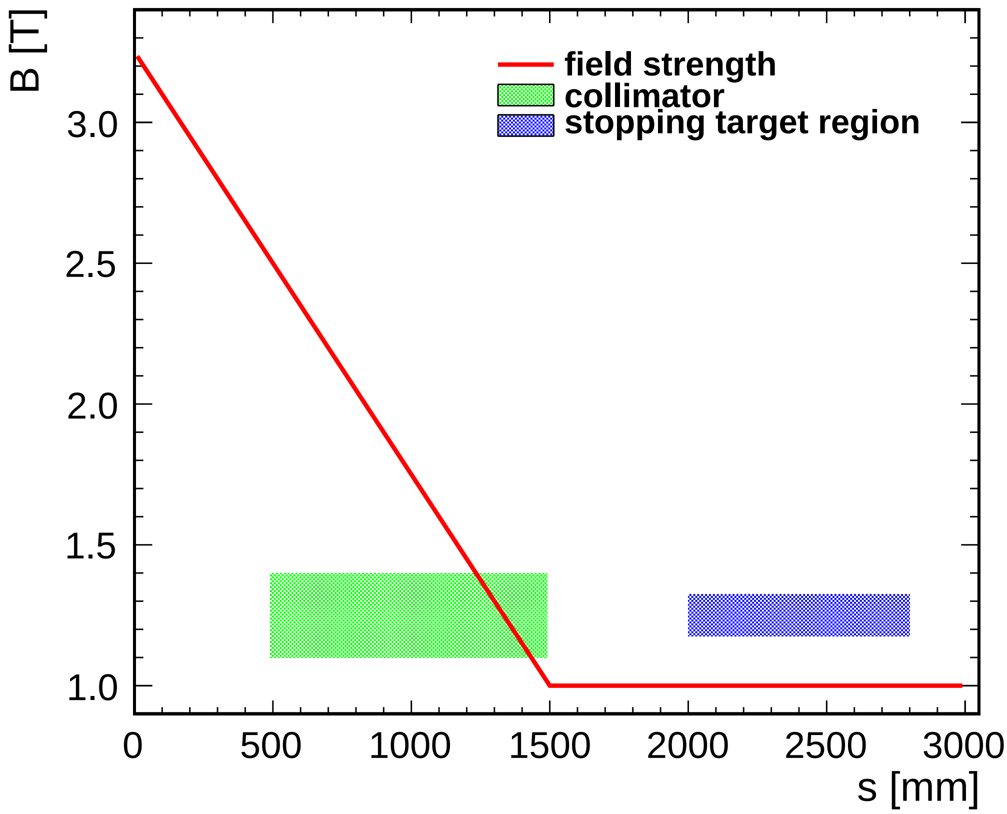

Figure 4.3: The magnetic field strength implemented in the simulation after the truncation of the Toshiba field maps at the end of the first 90° bend section. The field is given a linear negative gradient across the straight solenoid section connecting the end of the bend to the detector/stopping target region.

4.2 Collimator optimisation

As previously mentioned the beam is dispersive at the end of the 90° solenoid. This momentum dispersion is a useful property of the beam for isolating and eliminating high momentum muons and pions ($p > 70$ MeV/c) which could otherwise contribute to background events through decay in flight and radiative pion capture. In order to compensate for this background an asymmetrical cylinder of tungsten is placed after the bend section and before the detector solenoid with the purpose of scraping unwanted particles from the beam. Consideration of the transverse emittance is required in order to optimise the collimator dimension, with the effective emittance in both transverse planes for μ^- and π^- shown in Figs. 4.4 and 4.5 respectively.

Both the desired μ^- and unwanted π^- exhibit a skewed relation between momentum and position in the y-plane only, with lower momentum particles lying closer to the beam axis. In contrast the symmetry in the x-plane means collimation in that plane is redundant since removal of high momentum particles would produce a similar reduction in the number of those with lower momentum. For this reason the tungsten collimator is given a cylindrical shape with an inner and outer radius of 200 mm and 250 mm respectively, and a length of 1 m. The most important physical property of the collimator is the solid section with a negative y-offset. A transverse cross-section of the collimator can be seen in Fig. 4.6.

Fig. 4.4 b) indicates a mean y-position of -70.08 mm for the negative muons, so it is expected that the negative y-offset removing the optimal number of high energy particles will be close to this value. As an initial indication of the effects of the collimator a simulation was performed with a y-offset of -70 mm. The result of this is shown in Fig. 4.7. The shape of the muon momentum spectrum after the collimator indicates approximately 31% of the desired low momentum muons ($p < 45$ MeV/c) reaching the end of the collimator. However, as the collimator is

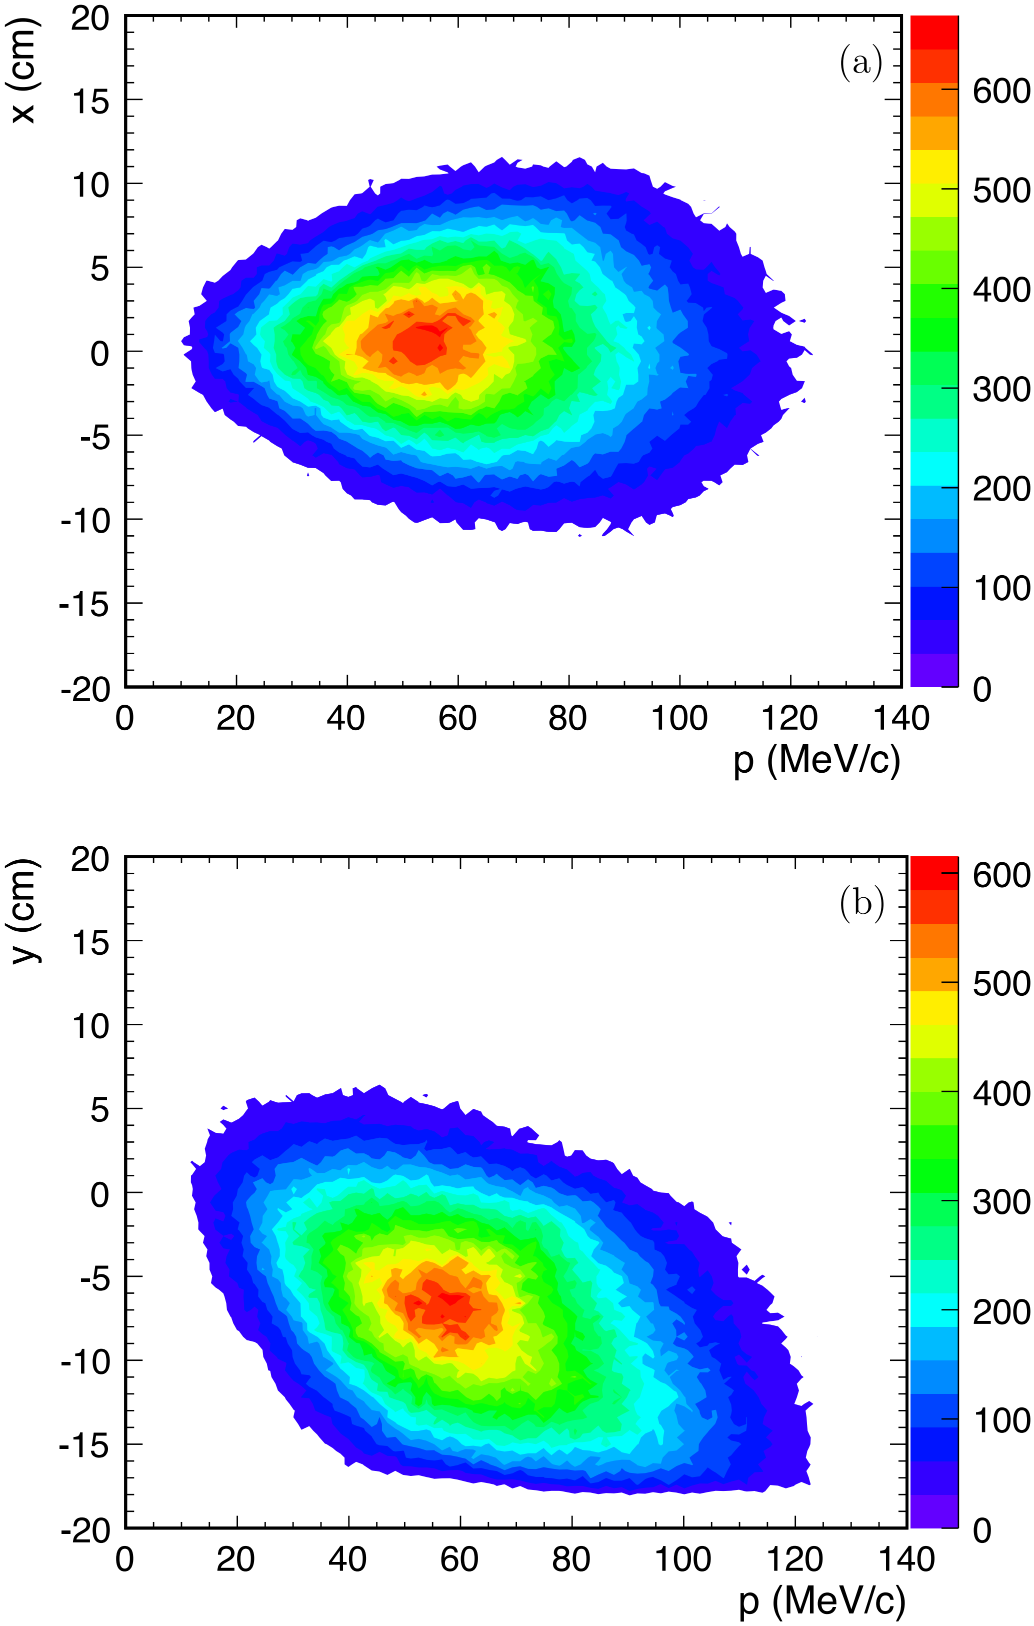

Figure 4.4: The effective emittance of muons in the a) x-, and b) y-plane that are incident on the beamline transport collimator.

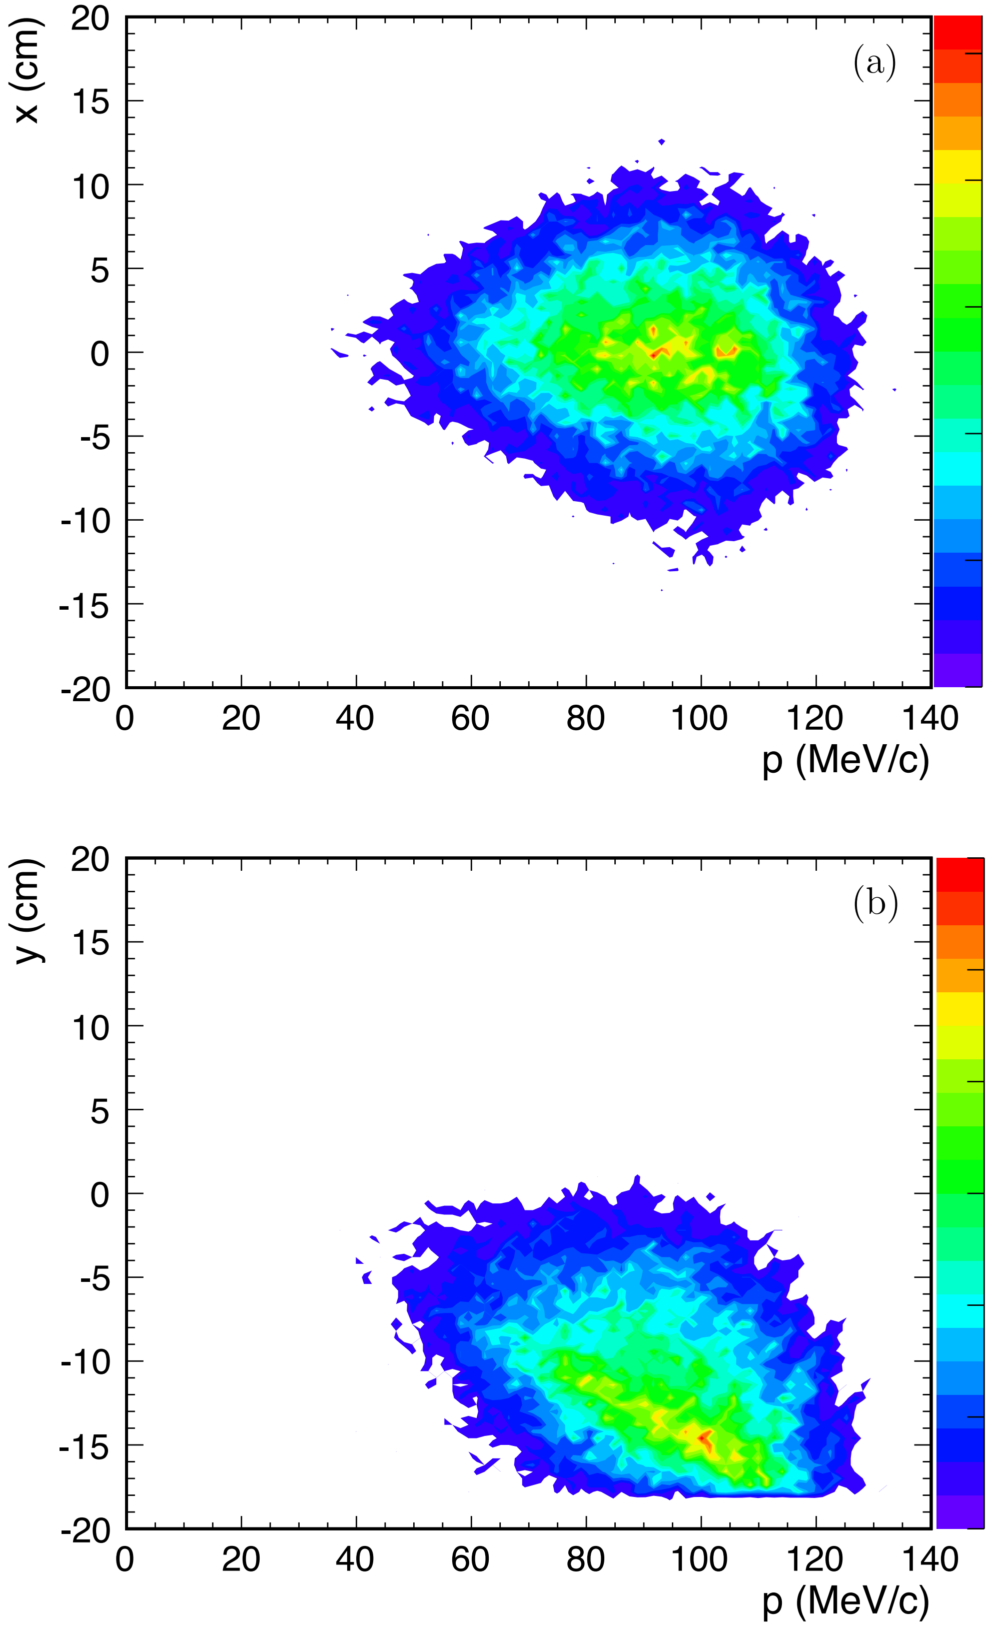

Figure 4.5: The effective emittance of pions in the a) x-, and b) y-plane that are incident on the beamline transport collimator.

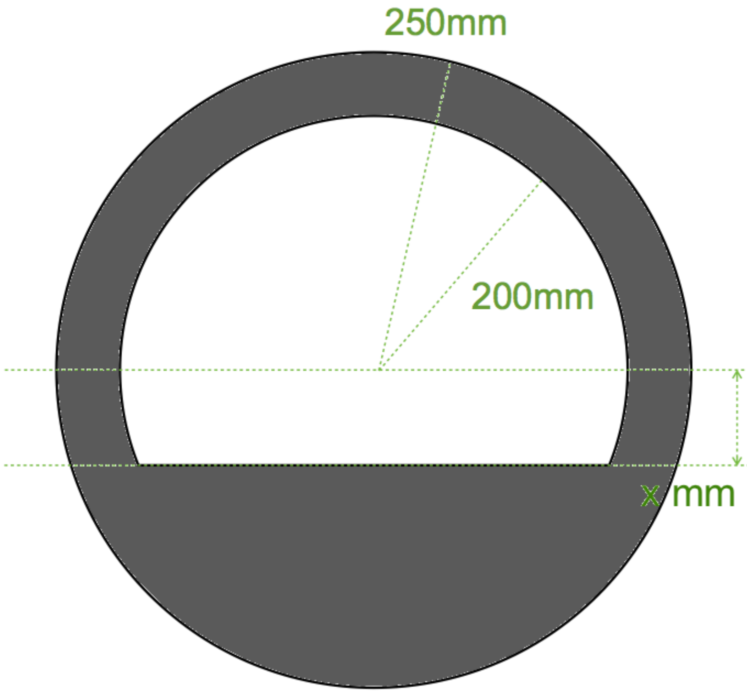

Figure 4.6: A cross-section of the schematic of the proposed beamline transport solenoid collimator, of length 1 m, with the negative y-offset (x) to be optimised by simulation.

designed to maximise the removal of high momentum muons and pions, whilst transmitting a maximum number of low momentum muons, the efficacy of the collimator may be quantified by

$$w = \frac{N_{\mu^-}(p < 45 \ \ \mathrm{MeV}/c)}{N_{\mu^-}(p > 70 \ \ \mathrm{MeV}/c) \times N_{\pi^-}} \quad , \tag{4.1}$$

which is effectively a signal/background ratio since low momentum muons are the most likely to stop in the target and higher momentum muons will not and can decay-in-flight to yield high momenta electrons. Similarly all pions can potentially be radiatively captured by the target nucleus with the emission of high energy photons that can convert, yielding high energy electrons.

Simulations were thus conducted over a range of collimator offsets, centred around -70 mm, using a sample of 7 million μ^- and π^- as generated by the `MARS` Monte Carlo. The muon momentum spectra incident on the stopping target, with

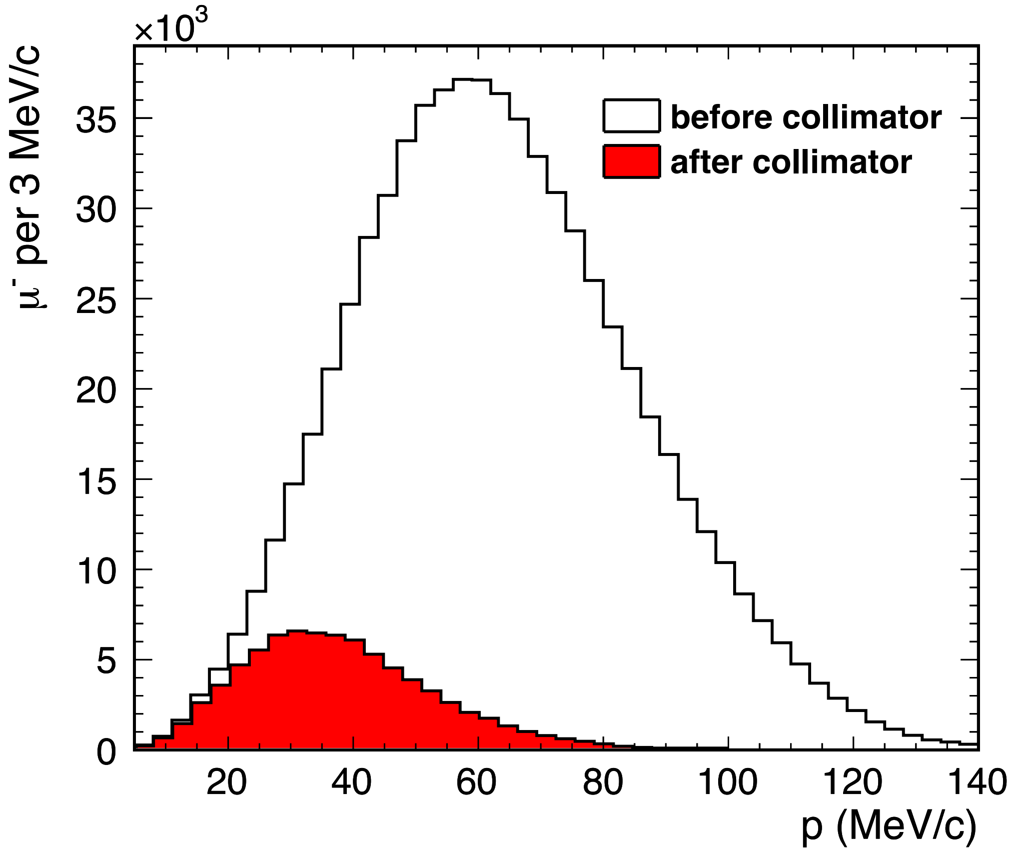

Figure 4.7: Muon momentum spectra for μ^- before and after a collimator with a negative horizontal offset of -70 mm. This transmits 12.5% of the initial μ^- but 31% of μ^- with momenta below 45 MeV/c.

y-offsets of $-65, -70$, and -75 mm, are shown in Fig. 4.8. It is not immediately clear from this plot which configuration is more efficient as the number of muons reaching the stopping target is similar. However, Tab. 4.1 displays the number of muons per proton and the value of w after the collimator. The maximum signal/background is achieved with a horizontal offset of -70 mm, resulting in 0.0027 muons per initial proton incident on the pion production target. This is smaller than the rate of 0.0035 for Phase-II [43], demonstrating the necessity to extend the beamline to include the additional Phase-II elements in order to increase the muon yield. i.e. the second 90^o bend affording additional decay length and further momentum selection.

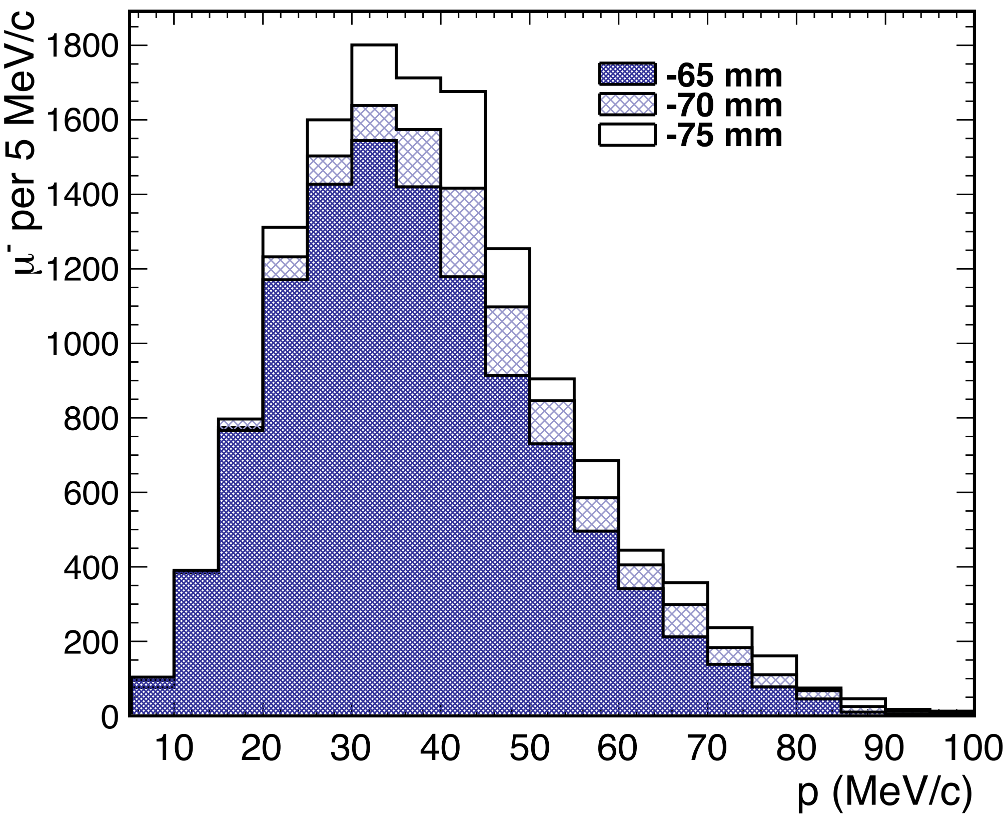

Figure 4.8: Muon momentum spectra incident on the stopping target after passing through a collimator of varying negative horizontal offset.

The collimator offset is therefore taken to be -70 mm as the default in all further simulations. As a verification of its effects, the impact of the collimator on the incident muon and pion beam is shown in Fig. 4.9, which clearly shows that the particles passing through the collimator are predominantly at low momentum: the average muon momentum is reduced from 63 to 39 MeV/c and 99.5% of the pions are eradicated.

4.3 Stopping target optimisation

As mentioned at the beginning of Ch. 4, the significant differences between the detector and magnetic field surrounding the stopping target in Phase-I versus

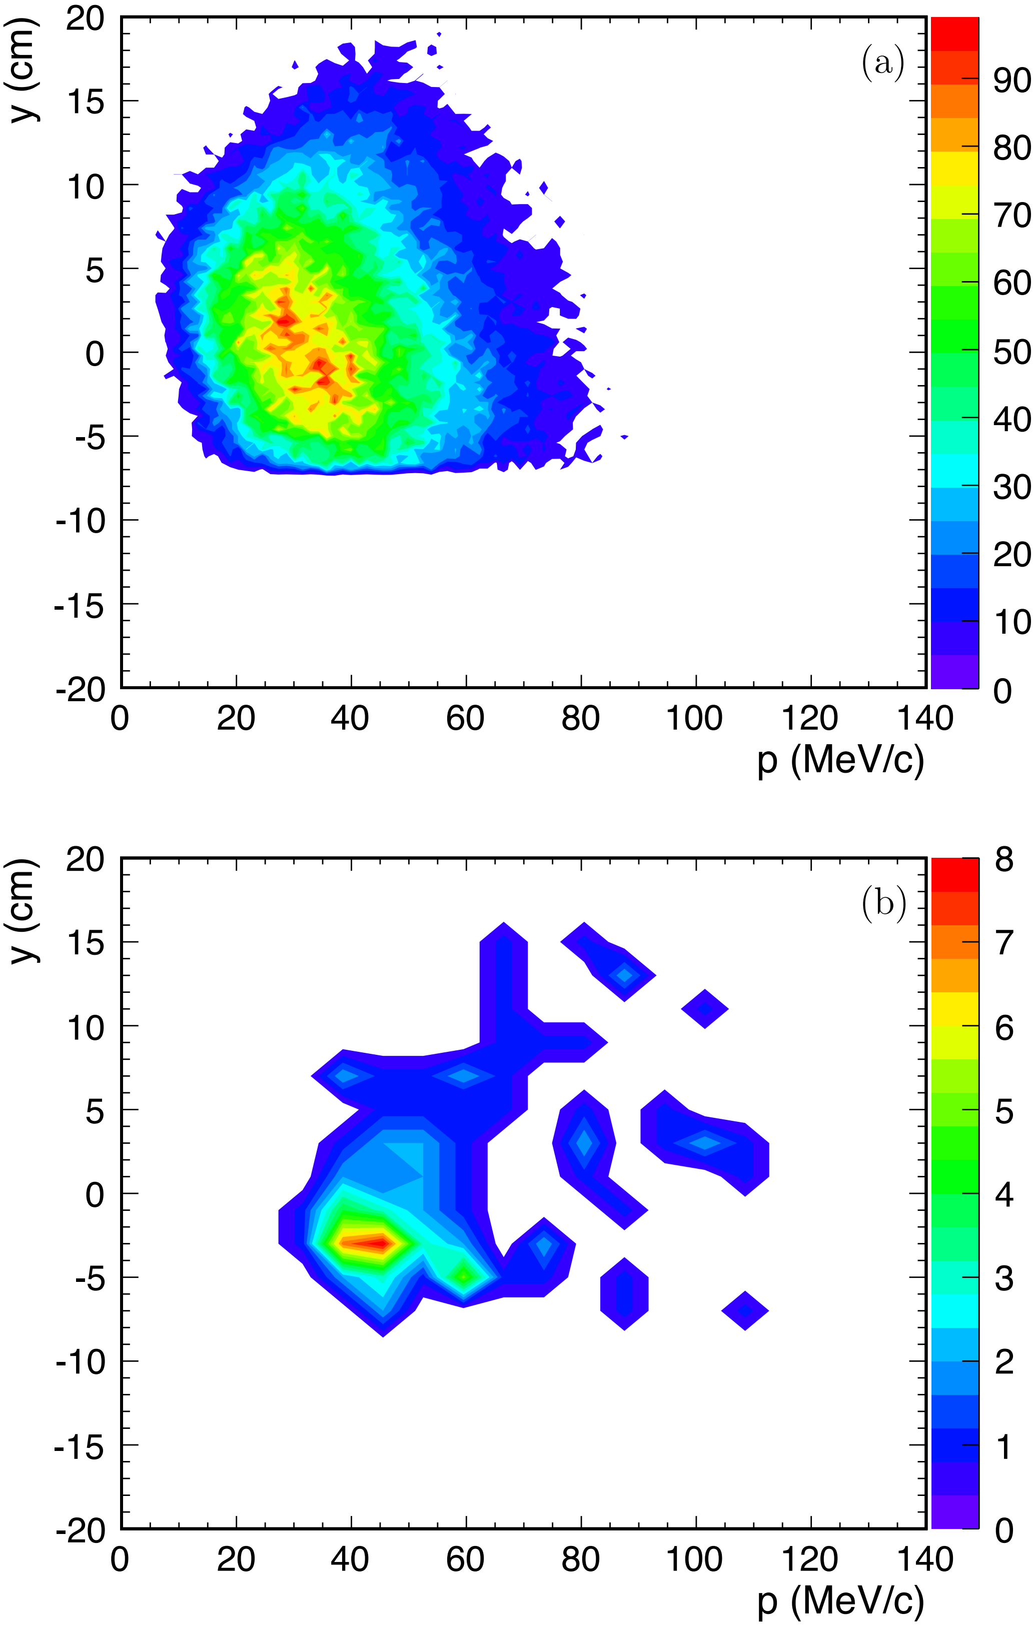

Figure 4.9: The effective emittance of a) muons, and b) pions incident on the stopping target after the effect of the beamline transport collimator with a negative horizontal offset of -70 mm.

y-offset (mm)	$N_{(\mu^-/p^+)}$	w
-55	0.0019 ± 0.0001	0.31 ± 0.01
-65	0.0023 ± 0.0001	0.90 ± 0.01
-70	0.0027 ± 0.0001	0.99 ± 0.01
-75	0.0032 ± 0.0001	0.61 ± 0.01
-85	0.0191 ± 0.0001	0.25 ± 0.01

Table 4.1: The number of muons per proton after the collimator and effectiveness of the collimator for a range of negative y-offsets.

Phase-II warrant a re-evaluation of the stopping target design for Phase-I. The Phase-I geometry has a graded magnetic field before the stopping target, as opposed to a graded field around the target in Phase-II, and the target is placed in the detector volume with a fixed 1 T field.

Two stopping target geometries have been modelled: the disk geometry of Phase-II (see Tab. 3.3) and a double cone geometry that was used by the SINDRUM-II experiment with a cylindrical detector in a solenoidal field. The schematics of both geometries are shown in Tab. 4.2. Each target region is placed at the centre of the detector solenoid, with visualisations of both target regions from the Phase-I simulation (`Phase-I_G4`) shown in Fig. 4.10. In the visualisations an on-axis 30 MeV/c muon is incident on the target region, substantiated upstream in the beamline transport solenoid, scattering off the aluminium target and producing secondary particles.

parameter	17 disks	double cone
radius	100 mm	100 mm
thickness	200 μm	200 μm
region length	800 mm	2 $\times$ 400 mm
volume	106814.2 mm^3	50682.8 mm^3

Table 4.2: Geometric parameters of the disk- and cone-type muon stopping targets.

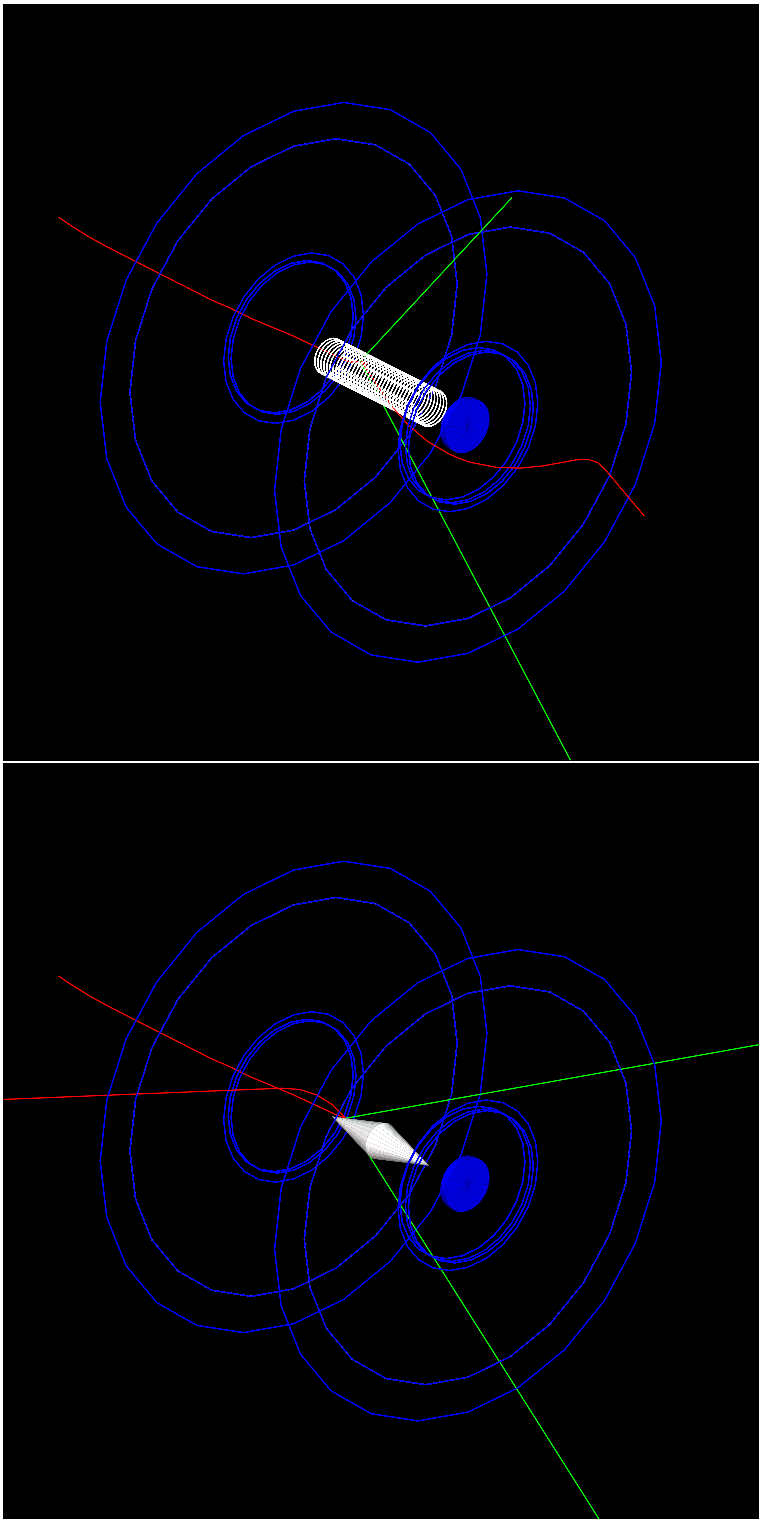

Figure 4.10: `Phase-I_G4` visualisations of a 30 MeV/c μ^- entering the detector solenoid (from the top left of frame) and interacting with the disk (top) and cone (bottom) stopping target geometries.

Using these two different configurations seven million backwards propagating μ^- and π^-, as simulated by MARS, are initiated at the pion capture solenoid and propagated through the experiment. These seven million μ^- and π^- used in the simulation were created by firing 31×10^6 p^+ at the pion production target. After the bend section the beam interacts with the collimator, yielding 0.0027 μ^- per proton after the collimator. These muons are then incident on the stopping target region, with the momentum spectra of the incident and stopped muons (for both geometry types) shown in Fig. 4.11. Phase-I is expected to run for approximately 10^6 s, with a beam power of 3.2 kW and proton momentum of 8 GeV. Using the equation

$$N_{\mathrm{p}} = \frac{P}{eE_{\mathrm{p}}} \times t \quad , \tag{4.2}$$

the number of protons incident on the proton stopping target will be approximately 2.5×10^{15} for the entire run. The number of incident and stopped muons is therefore scaled by the ratio of protons simulated to expected, giving a total of 6.72×10^{15} muons incident on the target.

The original disk configuration stops 66% of incident muons, more than the 42% stopped by the cone. Tab. 4.3 shows that, even though the spatial dimensions of the target regions are similar, the disk configuration has a volume approximately twice that of the cone. While more material will stop more muons, the different geometry plays a role: the disk geometry stops more muons but the cone geometry stops more muons per unit volume of target.

Shifting from a transverse to cylindrical detector requires a redesign of all the components therein. A cylindrical drift chamber (CDC) is proposed as the baseline detector for Phase-I. This has a number of benefits over the transverse detector in Phase-II:

- a reduction in background events. Only DIO electrons above a certain transverse momentum will reach the inner layer of the CDC,

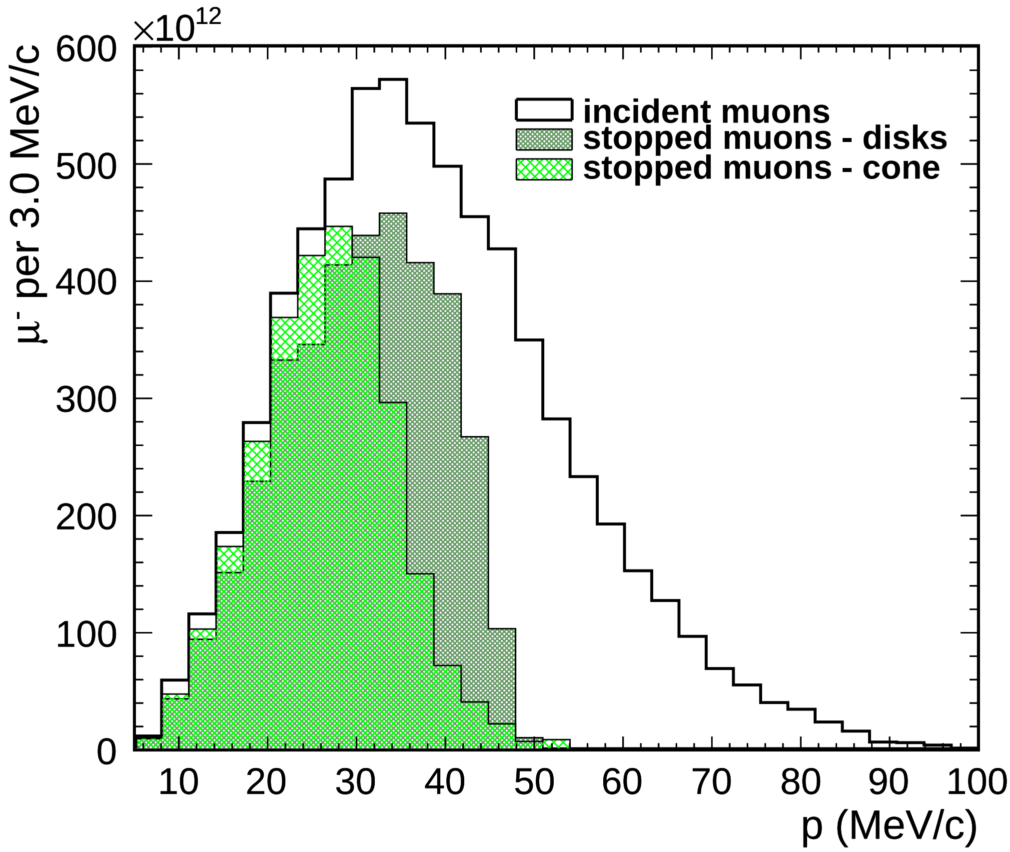

Figure 4.11: Comparison of the momentum spectrum of stopped muons for the Phase-I disk and cone geometries, simulated using the same incident beam. The disk geometry has a stopping efficiency of 0.66 whereas the cone stops 0.42 of the incident muons.

- detector saturation from beam particles that are not stopped in the muon stopping target is avoided as the particles continue downstream and out of the detector solenoid.

The CDC comprises a thin carbon-fibre reinforced plastic inner and outer cylinder, filled with a 50:50 helium-ethane mixture, with aluminium endplates. Hit information generated in the simulation is smeared by the expected position resolutions of the drift chamber, $\Delta x = \Delta y = 100$ μm and $z = 2$ mm, resulting in a core momentum resolution of 600 keV and a tail resolution for 10% of 1.2 MeV. The detector system also includes a scintillator trigger counter at the start and end of the CDC inner cylinder providing trigger timing insensitive to protons. These physical components have been included in the `Phase-I_G4` simulation, with all interactions included in the results shown. The simulated resolutions are included in the Punzi sensitivity calculations. These are key factors for inclusion as stopping target optimisations will be rendered redundant if the gain is washed out by poor resolutions and energy losses.

A conical target presents less material to escaping electrons and so conversion electrons will have a smaller energy loss in comparison to the the disk target and the signal peak should be sharper and so more distinguishable from the DIO background. The DIO and signal momenta (for a conversion BR of 10^{-15}) are shown in Fig. 4.12 where the signal peak is indeed seen to be better defined for the cone target. There are also more signal events for the cone geometry but to quantify whether the cone geometry improves the experimental sensitivity we evaluate the Punzi sensitivity.

As with the analysis in Sec. 3.1.2, in order to establish a firm conclusion a Punzi calculation of the 90% confidence level upper limit on the branching ratio (BR[Punzi]) is required, in addition to the number of events and signal shape. These results can be found in Tab. 4.3. In this case the BR[Punzi] is of the same order of magnitude with similar values for both cases. If only one type of stopping target were to be manufactured then this, coupled with the relative number of events and signal shape, would suggest that the cone should be chosen for construction.

	N_{events}	BR[Punzi]
disks	1.14 ± 0.08	$(1.2 \pm 0.2) \times 10^{-14}$
cone	1.45 ± 0.17	$(1.8 \pm 0.3) \times 10^{-14}$

Table 4.3: Number of events above 104.4 MeV/c and the BR[Punzi] for the disk- and cone-type muon stopping targets for a BR of 10^{-15}.

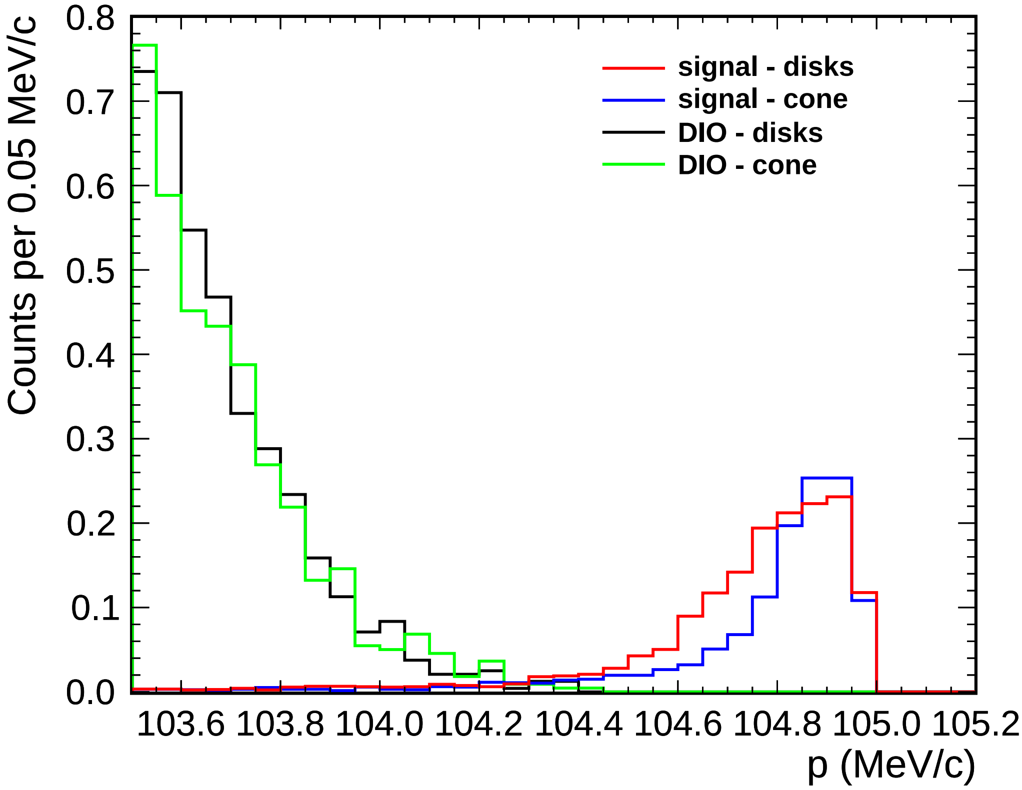

Figure 4.12: Comparison of the signal and DIO electron momentum spectra, as produced from the disk and cone stopping target configurations. The spectra are scaled to the expected number of protons on target for the entire Phase-I run. A muon-to-electron conversion branching ratio of 10^{-15} has been used in this plot.

4.4 Conclusions

A staged approach to the COMET experiment has been proposed in order to expedite the physics results as well as gaining valuable operational experience with these novel experimental components. The first stage of this, Phase-I, will advance the sensitivity to $\mu \to e$ conversion by a factor of 100 on current limits. The components required to achieve this have been described and simulated: implementation of magnetic fields and the inclusion of a cylindrical detector system.

The software framework was used to estimate experimental elements such as the beamline collimator. Transverse beam studies indicated a correlation between negative y and high p, suggesting an asymmetric design with a y-offset. Optimisation of low momentum transmission and high signal/background advocated a -70 mm offset, with 0.0027 μ^- per proton reaching the muon stopping target region with this configuration.

The nature of the muon stopping target geometry was therefore reconsidered due to the shift in detector design and previous optimisation of the collimator. A double cone design, utilised by SINDRUM-II, was tested in conjunction with the disk structure of Phase-II. The stopping efficiency, signal shape, and Punzi sensitivity were calculated for these two geometries, with the cone structure having a slightly higher sensitivity. However, the cone design resulted in more signal events in the required momentum region and had a more distinct peak in Fig. 4.12 which cannot be ignored. Taking these features into account suggests a similar conclusion to that of the stopping target studies of Ch. 3: a cone design would be a logical first step to maximise the likelihood of observing signal, with the disk design then implemented to minimise the exclusion region if no signal is observed.

The current optimised design for COMET Phase-I outlined in this chapter demonstrates a viable first step in achieving the ultimate goal of reaching the proposed sensitivity of COMET: 10,000 times that of current experimental limits. Simulations suggest employing the double cone structure with the cylindrical Phase-I detector will probe $\mu \to e$ conversion to a limit 100 times better than achieved by SINDRUM-II.

Chapter 5

MuSIC

There are many muon related programmes proposed or under discussion, such as searches for muon-to-electron conversion (see Ch. 3), neutrino factories, and muon colliders. They need intense muon beams of $10^{11} - 10^{14}$ muons per second, while the highest muon intensity currently available is approximately 10^{8} muons per second at PSI [16]. In order to achieve a more intense muon beam it is necessary to build a dedicated muon production system using new concepts to improve on current pion/muon production efficiencies. A novel superconducting pion capture system is one such idea. In the system a thick pion production target is located on a proton beamline, with a strong solenoidal magnetic field applied to the target region. According to many simulation studies the pion/muon production efficiency of such a pion capture system produces is more than a 1000-fold increase in muon yield compared to current muon facilities.

The intention of the Muon Science Innovative Commission (MuSIC) experiment in Osaka is to test, and measure, these novel ideas and prototype equipment essential to future research in the field of muon physics. In particular it is seeking to demonstrate (for the first time outside simulation) that pion capture solenoids do indeed result in a dramatically increased muon yield which is a key requirement of the COMET and Mu2e experiments (see Ch. 3). This chapter will detail the studies of position and energy distribution of the MuSIC beam at the end of the existing 36° of beam pipe, highlighted in Fig. 5.1. These measurements will ultimately be used to refine the simulations for COMET since MuSIC is in many ways a prototype of COMET, but with a beam power two orders of magnitude less.

The MuSIC experiment took data in discrete running periods in order to allow the machine to be optimised stepwise and for subsequent runs to have time to assimilate and utilise the data and experience from the previous runs. This chapter will cover the first two of seven runs of the machine in Secs. 5.1 and 5.2 respectively. These sections are then subdivided into simulation and analysis of experimental data. The simulation is intended to aid the analysis of the experimental data, both in terms of defining the expected behaviour, isolating anomalous behaviour, and providing possible explanations for such.

The first phase of MuSIC construction was completed in 2010, incorporating a prototype superconducting pion capture solenoid and a 36° portion of the first pancake solenoidal bend section. A photograph of the current state of the MuSIC

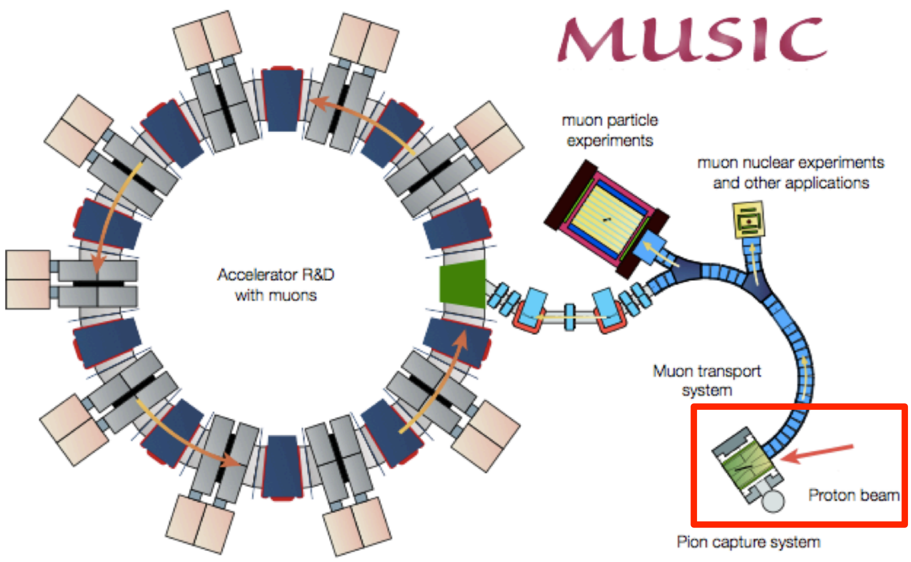

Figure 5.1: A schematic of the proposed MuSIC beamline, with the current 36° of constructed beamline highlighted within the red box.

beamline can be seen in Fig. 3.4. The next phase of construction is due to begin in 2013 so all studies in this analysis are completed with the equipment present in the first phase.

5.1 Run-I

In order to assess the pion capture and transport efficiency of the constructed beamline a detector system was required to make timing, energy, and hit count measurements. A plastic scintillating cuboid bar was placed in the beamline with a series of photon counters used to measure the amount of scintillation caused by the beam. This information was recorded by a data acquisition (DAQ) system which facilitated the subsequent extensive offline analyses. This section outlines the experimental process of simulation, design, and construction of the online system for Run-I of the MuSIC experiment.

5.1.1 Simulation

The first aspect of the simulation developed was the interaction of the expected MuSIC beam at the end of the 36° bend with the plastic scintillating bar. This was a necessary step to benchmark the experiment and define thresholds for the DAQ system. The simulation was also used to develop techniques that could be useful for particle identification, particularly the discrimination of muons from other particles.

Scintillator design

A plastic (polyethylene) scintillating bar is modelled in GEANT4 with dimensions $38.0 \times 3.0 \times 1.0$ cm. The bar is wrapped in a layer of aluminium foil of thickness 0.2 mm, with the surrounding volume given the properties of air.

Four Multi-Pixel Photon Counters (MPPCs) are placed at the end of the scintillating bar, both positioned 0.75 cm from the central x-axis of the scintillator. Four MPPCs were chosen to optimise the detection rate within the spatial limitations of the scintillator bar. The face of the MPPCs have a circular shape with an area of 1 mm^2, each with 400 pixels per MPPC. Fig. 5.2 shows the end of the scintillating bar with two attached MPPCs.

The origin of the experimental volume is taken to be at the centre of the scintillating bar (in all three dimensions) and a particle beam is initiated at -20.0 cm along the z-axis, unless otherwise stated.

Radioactive sources (e.g. ^{60}Co, ^{90}Sr, etc.) are used to calibrate the detector by analysing the analogue-to-digital converter (ADC) output corresponding to known emissions of that isotope. A calibration curve was then determined yielding a relation between channel number and energy. In simulation this process is unnecessary as particles with a continuous energy range may be fired into a material, recording the energy deposited by individual scintillations. An example of such an event can be seen in Fig. 5.3, highlighting the detected photons for a 1 MeV incident electron. In simulation the number of photons incident on every MPPC is recorded, along with the energy deposited by particle interactions, but digitisation in the TDC and ADC has been omitted from the simulation. Experimental threshold values are incorporated, however.

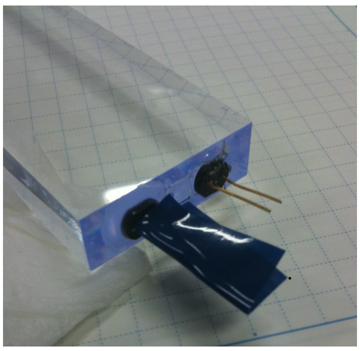

Figure 5.2: A photograph of the polyethylene scintillating bar with MPPCs mounted on one end. One is exposed, the other wrapped in insulating black tape.

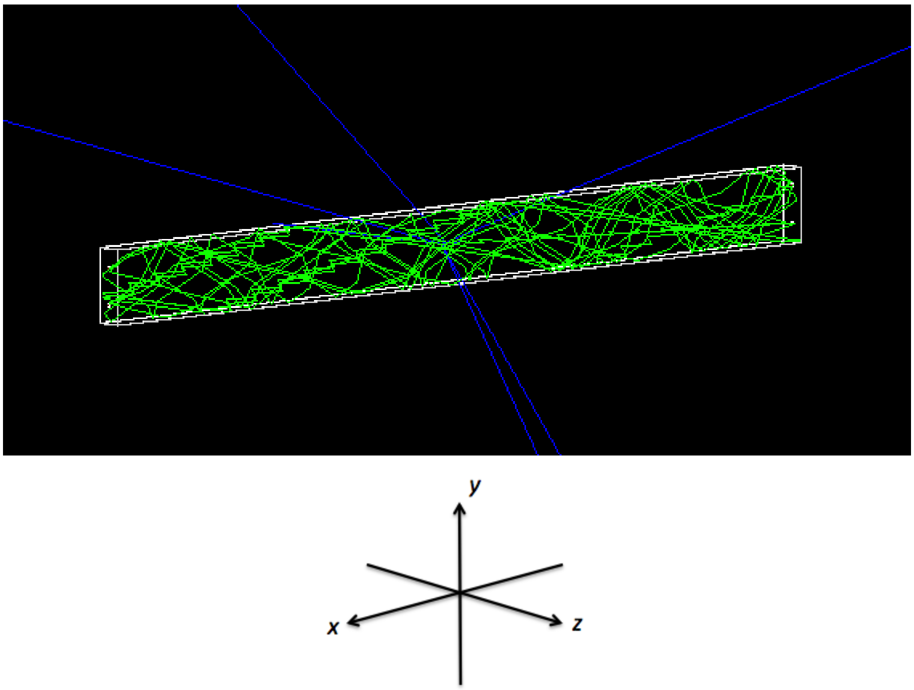

Figure 5.3: An example output of the plastic scintillating bar with a 1 MeV incident electron (arriving along the negative z-axis) producing optical photons (green) detected by the MPPCs and secondary electrons (blue).

Timing (TDC)

Time differences between the left- and right-hand MPPCs are recorded with a time-to-digital converter (TDC). Once calibrated, TDC readings can then be used to estimate the hit-location of the incident particle.

The first method for calculating the hit-location in the simulation was to take a ratio of the number of detected photons from each side (i.e. the sum of the two left MPPCs divided by the two right MPPCs). For a particle incident on the centre of the scintillator the ratio between the two sides should be equal to 1.0, with any offset obeying the relation

$$x = \frac{N_R - N_L}{N_R + N_L} \frac{L}{2} \,, \tag{5.1}$$

where x is the position of the hit (relative to the centre at $x = 0$), N_L and N_R are the total number of hits of the left- and right-hand side respectively, and L is the length of the scintillating bar. In the geometry of the experimental setup the right-hand side is the positive end of the scintillator. It is essential to insulate the scintillating bar with a layer of aluminium foil (to prevent the photons from escaping the bar before detection) and black tape (to isolate external photons). This insulation however rendered this method of timing measurement redundant due to the high reflectivity (≈ 0.99) of the aluminium foil. This meant that the photons, that did not travel directly to an MPPC from scintillation, were liable to scatter within the scintillator until either being re-absorbed by the scintillator or absorbed by another MPPC.

It was therefore necessary to account for the effects of scattering by developing a portion reconstruction method based on time distributions rather than hit count ratios. This was achieved by recreating the results of the DAQ setup by recording the time at which each scintillation photon hits any one of the four MPPCs. The time of each hit is recorded then modelled as a pulse-shape using the equation

$$F(t) = \frac{\tau_1}{\tau_1 - \tau_2}\left(e^{-\frac{t}{\tau_1}} - e^{-\frac{t}{\tau_2}}\right) , \tag{5.2}$$

where τ_1 and τ_2 are arbitrary time constants. A single-photon signal for one MPPC circuit was then isolated using the dark current readout from the DAQ oscilloscope. An example of this can be seen in Fig. 5.18. This shows a 50 ns pulse width from -5 to 45 ns. The time constants of Eq. 5.2 were optimised to recreate this pulse width and shape. The result of this optimisation can be seen in Fig. 5.4.

The model in Fig. 5.4 represents a hit time of 0.0 ns. Each simulated pulse-shape is shifted so the peak is centred on the hit time. This hit time is calculated by subtracting the time a photon interacts with the MPPC from the time of scintillation creating the photon. The pulse-shapes of all the hits for each primary event may then be summed, with an example of the result shown in Fig. 5.5.

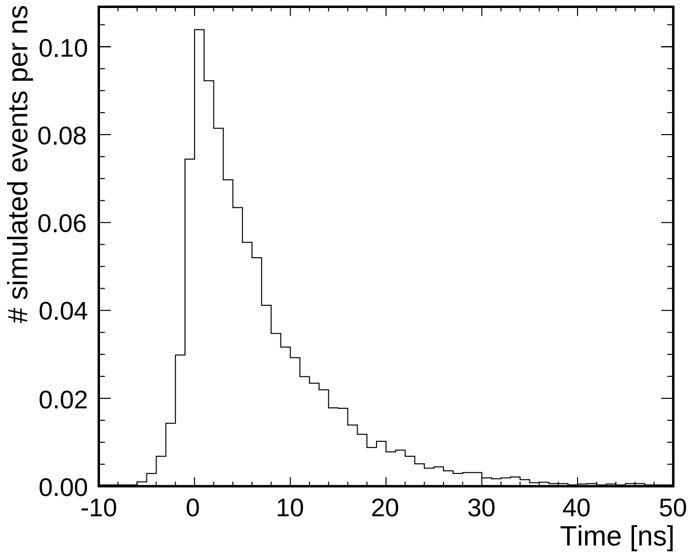

Figure 5.4: A recreation of the one-photon signal pulse-shape produced by the DAQ oscilloscope, with $\tau_1 = 1.0$ and $\tau_2 = 0.15$ for a hit-time of $t = 0$.

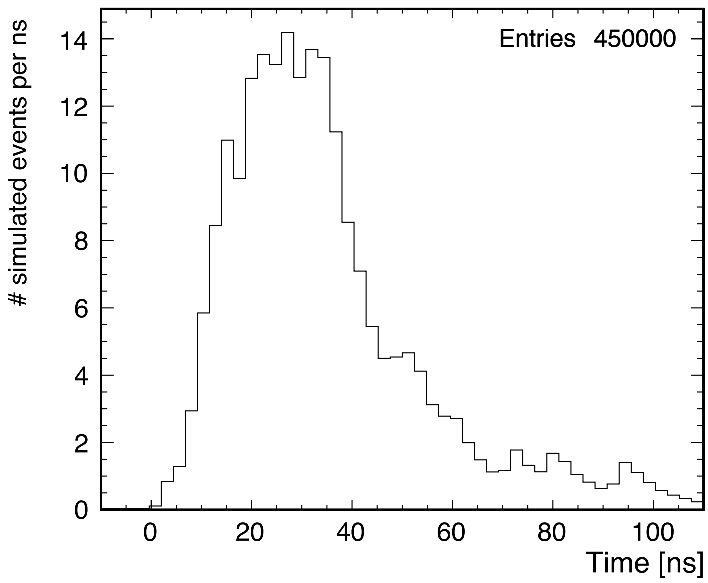

Figure 5.5: A summation of pulse-shapes (as in Fig. 5.4) from one initial event, at one MPPC. This resulting pulse shape is an approximation of what is observed by the DAQ.

This example is for one of four MPPCs with only one initial event. It is a representation of what is seen by the TDC DAQ readout and is used to discriminate between the dark-current and real data. A trigger value is chosen to minimise dark-current noise whilst maximising signal. A trigger value of 7 photons was chosen after analysis of the dark-current by the DAQ setup (see Sec. 5.1.2). This trigger value was then used in the simulation to select the time of the first bin with a content greater than this value. This process is then repeated for each MPPC across 1000 events, with examples of the resulting histograms shown in Fig. 5.6.

A Gaussian curve may then be fitted to each MPPC histogram in Fig. 5.6, with an average of the RMS and mean taken across the four MPPCs. The timing

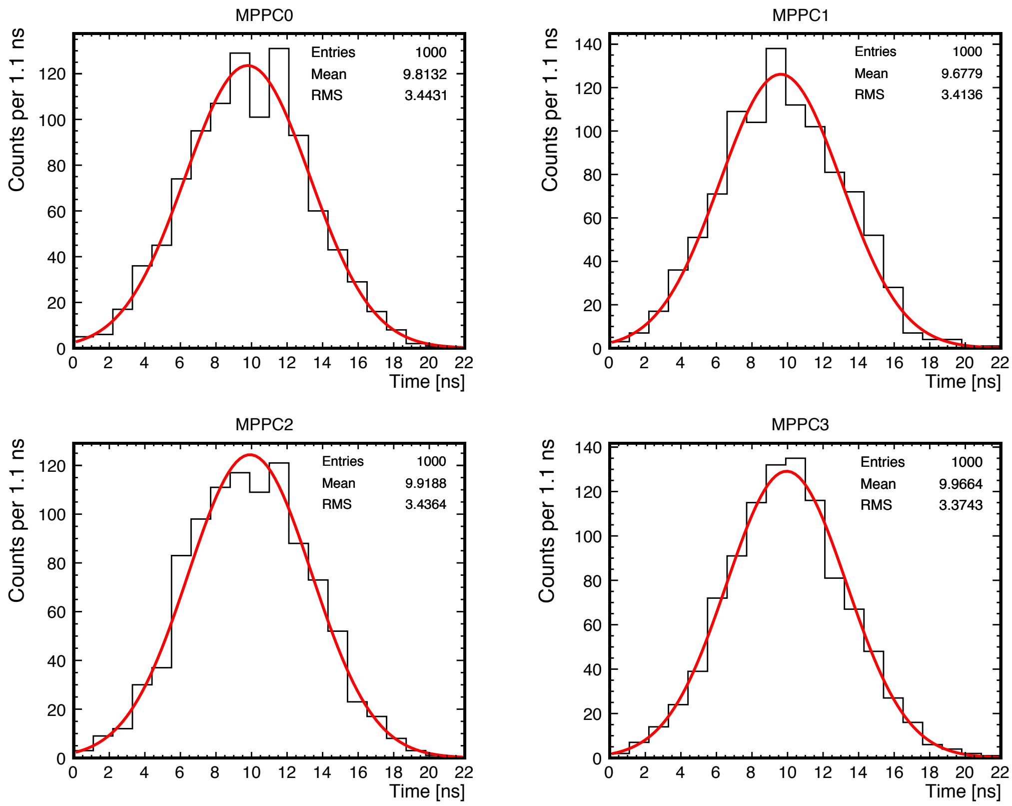

Figure 5.6: Histograms for each MPPC of the first-hit time for 1000 initial 0.511 MeV electrons, incident on the centre of the scintillating bar. A Gaussian distribution is fitted to each histogram.

resolution of the MPPCs (for 0.511 MeV electrons incident on the centre of the scintillator) is then given by

$$R = \frac{\Delta t}{t} = 2\sqrt{2\ln 2} \cdot \frac{\sigma}{\mu} = 0.83 \pm 0.21 \ , \tag{5.3}$$

where σ and μ are the RMS and mean, respectively, derived from the fit.

Using the simulated data of MPPC hit-times it is then possible to create a calibration curve relating the hit-position of the incident particle with the difference between the average hit-time of the left-hand MPPCs and the right. This yields a linear relationship that can be seen in Fig. 5.7.

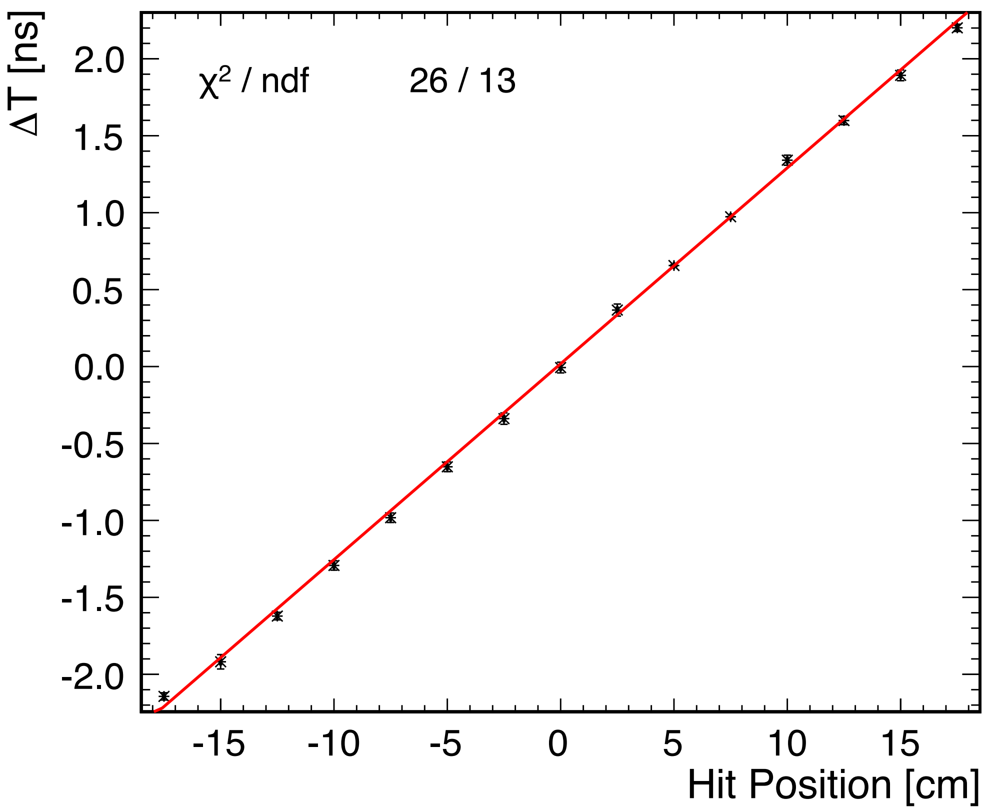

Figure 5.7: A calibration curve (from simulation) for the hit-position of the initial event, related to the difference in average hit-time of the left- and right-hand side MPPCs. The errors are from Monte Carlo statistics.

The calibration curve is calculated by firing particles into the scintillator at different positions along the x-axis (only in the negative x-direction however as the relation is symmetric about the origin). This plot can then be used to analyse the real data from the TDC, estimating the hit position of individual particles from the difference in hit-times of the left and right side of the scintillator. In this case the relation is described by

$$\Delta T = (0.120 \pm 0.002) \cdot x + (-0.094 \pm 0.021)\ . \tag{5.4}$$

Energy (ADC)

The ADC is a device used to convert a continuous signal to a discrete digital number. In the case of this experiment the continuous signal is that of optical photons incident on the MPPC and the discrete digital number is the total count within the timing window.

Unlike in experiment the simulation is able to initialise a string of particles with the same initial energy, incident upon the same point of the scintillator. It may then be used to calculate the exact energy deposited ($E_{\text{dep.}}$) in the detector, via scintillation, per event. The number of photons detected by all four MPPCs in one event is plotted against the energy deposition in Fig. 5.8. It is described well by a linear relation:

$$N_\gamma = (157.4 \pm 8.4) \cdot E_{\text{dep.}} + (-2.1 \pm 15.7)\ , \tag{5.5}$$

where N_γ is the number of photons detected. This relation may then be applied to further results, from the simulation or experimental data, in order to estimate the energy deposited per event.

The simulation was then modified to output ADC data analogous to that of the experimental DAQ. The ADC extracts the total number of MPPC hits per event, and then stores that number for however many events are incident in the

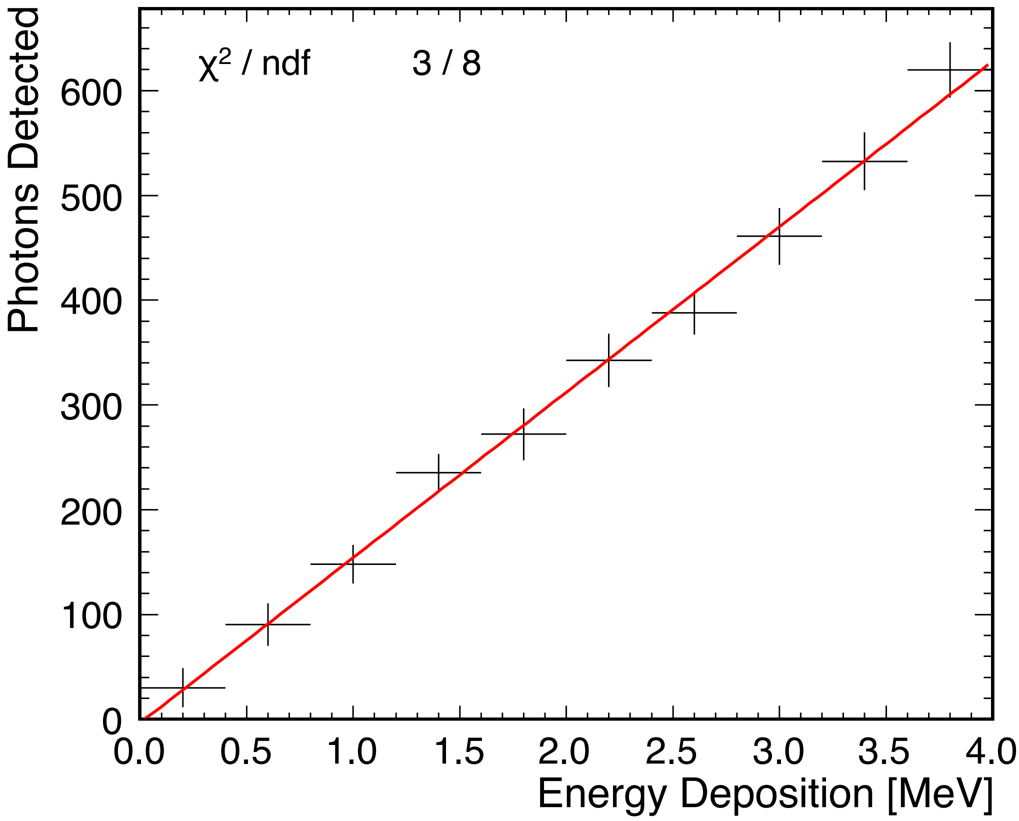

Figure 5.8: The total number of optical photons detected for a series of electrons, with differing initial energies ranging from 0.511 to 7.0 MeV. The energy deposition is plotted (rather than the initial energy of the particle) so as to fit a linear relation between photons detected and $E_{\mathrm{dep.}}$. Each point has error-bars corresponding to $\sqrt{N_\gamma}$.

timing window. The plot in Fig. 5.9 displays the simulation data for 1000 events of three particles observed in the MuSIC experiment; negatively charged electrons, muons, and pions. Each particle type is given an initial energy of 10.0 MeV in this example. The lower x-axis is the expected output of the ADC. However, using the relation in Eq. 5.5 (derived from the fit in Fig. 5.8) it is possible to convert the number of photons detected to an approximation of the energy deposited by the incident charged particle. This calibration is based on electrons.

The significance of the plot in Fig. 5.9 is that the simulation software has the ability to make cuts on particle type, whereas the DAQ does not; the DAQ can

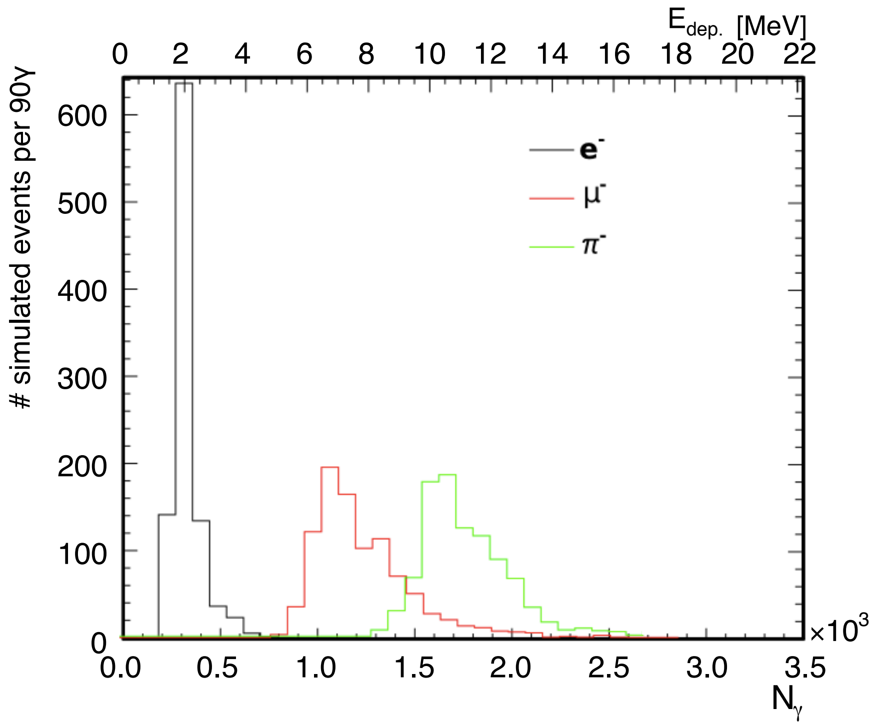

Figure 5.9: The number of hits for 1000 initial events of negatively charged 10.0 MeV e^-, μ^-, and π^-. The calibration curve used to calculate the energy deposited is based on e^-.

only report times and number of ADC hits which need to be interpreted in terms of particle type/flux, position, and energy.

Calibration

The simulation package has also been used to model the spectra of the calibration sources, in this case ^{22}Na and ^{60}Co. The isotopes used in calibration, including the decay channels modelled, can be seen in Tab. 5.1.

The resulting spectra from this simulation can be seen in Fig. 5.10. In this case each incident particle has been simulated 5000 times. These plots have then been compared with the calibration data taken during the experiment.

Isotope	Primary decay channel	Energy [MeV]
^{22}Na	$^{22}\text{Na} \rightarrow {}^{22}\text{Ne} + e^{+} + \gamma$	0.54 (1.27)
^{60}Co	$^{60}\text{Co} \rightarrow {}^{60}\text{Ni} + e^{+} + \gamma$	1.17 (1.33)

Table 5.1: Isotopes used to calibrate the TDC and ADC, and therefore used in simulation. The numbers in brackets are energies of secondary photo-peaks. ^{22}Na's second photo-peak is due to the detection of both photons produced in the annihilation of the positron. ^{60}Co has two photo peaks from the gamma-rays emitted by the de-excitation of the ^{60}Ni.

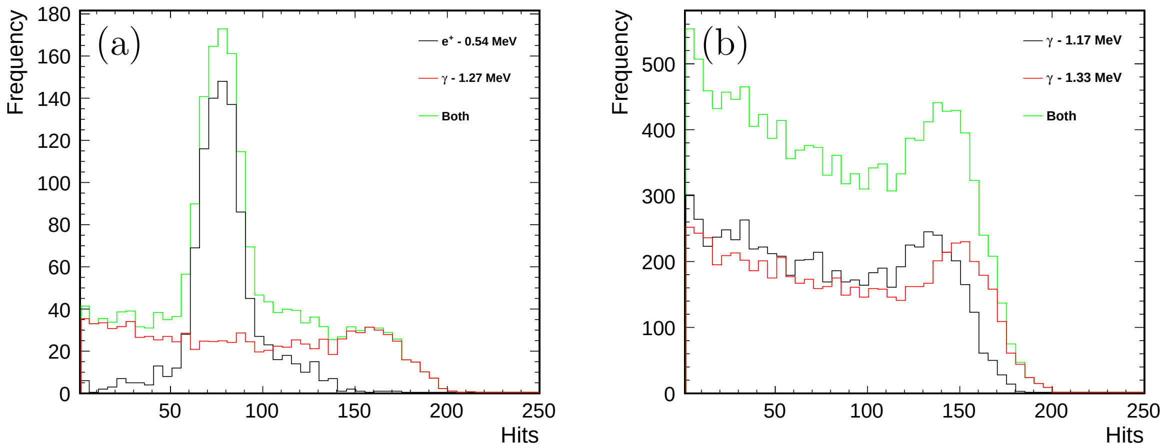

Figure 5.10: Simulation of the calibration spectra for a) ^{22}Na, and b) ^{60}Co, firing incident particles of energy corresponding to the different radioactive-decay channels of that isotope.

Hit count

In order to simulate the hit rate it is necessary to simulate a realistic particle distribution at the position of the scintillator. This is done using the particle tracking code `G4beamline` [34]. An initial beam of protons with energy 392 MeV is simulated to strike a graphite target. The resulting particle cascade is then tracked through a 36° solenoidal bend to the position of the scintillator - 5.0 cm away from the end of the beampipe. The spatial distribution of the beam at the scintillator (in the transverse plane) can be seen in Fig. 5.11.

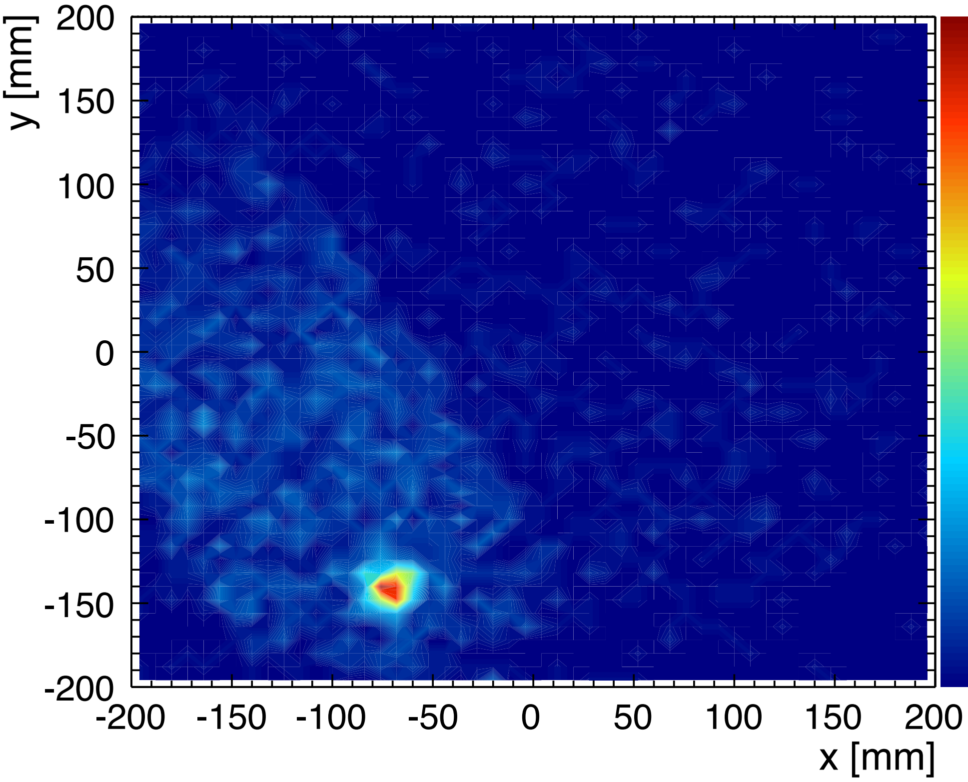

Figure 5.11: Spatial distribution of the beam (with a cut removing neutral particles that are not tracked through the solenoidal magnetic field) incident on the scintillating bar.

Fig. 5.11 demonstrates that the beam of charged particles is not centred around the origin of the beamline axis but skewed in both the negative x- and y-direction. This is due to the nature of the solenoidal magnetic field resulting in a helical trajectory to the beam. Once this beam exits the end of the beam pipe, assuming a hard-edged magnetic model with $B = 0$ beyond the beam-pipe, it will continue on an off-axis trajectory due to the transverse momentum of the beam. A breakdown of the constituent particles in the beam at this point can be seen in Tab. 5.2.

Tab. 5.2 demonstrates that the majority of the beam is composed of electrons, whereas the positively charged counterparts are approximately 50% e^+. The magnetic field in the MuSIC experiment is tuned to select positively charged particles as, due to charge conservation, the majority of pions produced in the

Particle	e^-	e^+	μ^-	μ^+	π^-	π^+	p
number	2146	288	11	121	6	69	85
percentage	78.7	10.6	0.4	4.4	0.2	2.5	3.1

Table 5.2: A breakdown of the number of particles, and relative abundance, in the beam as simulated at the position of the scintillator.

proton-graphite cascade will be positively charged. This is confirmed by simulation with the ratios $7\pi^+$:$1\pi^-$ and $5\mu^+$:$1\mu^-$ emitted from the target. One explanation for the abundance of electrons is the interaction of the beam with the magnetic elements/beam chamber, as well as the decay of in-flight particles. This is a raw particle output however and doesn't account for the threshold of the ADC. If a momentum threshold is applied e.g. $p > 1$ MeV/c then the number of e^- halves.

The particle beam is then tracked through the entirety of the MuSIC beamline, lending a distribution to fire into the scintillator simulation. An example of the output distribution produced by the scintillation software can be seen in Fig. 5.12. There is one definitive peak at approximately 2 MeV that corresponds to the abundance of low momentum electrons and positrons. The other less well defined peaks at higher energies correspond to the muon and pion content of the beam (both positive and negative) which are present in Fig. 5.9.

Fig. 5.12 also illustrates the number of particles that have produced scintillation at that geometry (in this case the origin). Dividing this number by the initial number of protons incident on the graphite target (1×10^6) gives the expected hit rate for the scintillator. The scintillator is then repositioned at discrete vertical intervals of 5 cm, corresponding to the positions recorded for the real data. The plot of hit counts for the simulation, using the `G4beamline` MuSIC beam, can be seen in Fig. 5.13.

This distribution takes into account a flip in the sign of the y-axis, such that the distribution in Fig. 5.11 is reflected in the line $y = 0$. This arises from a

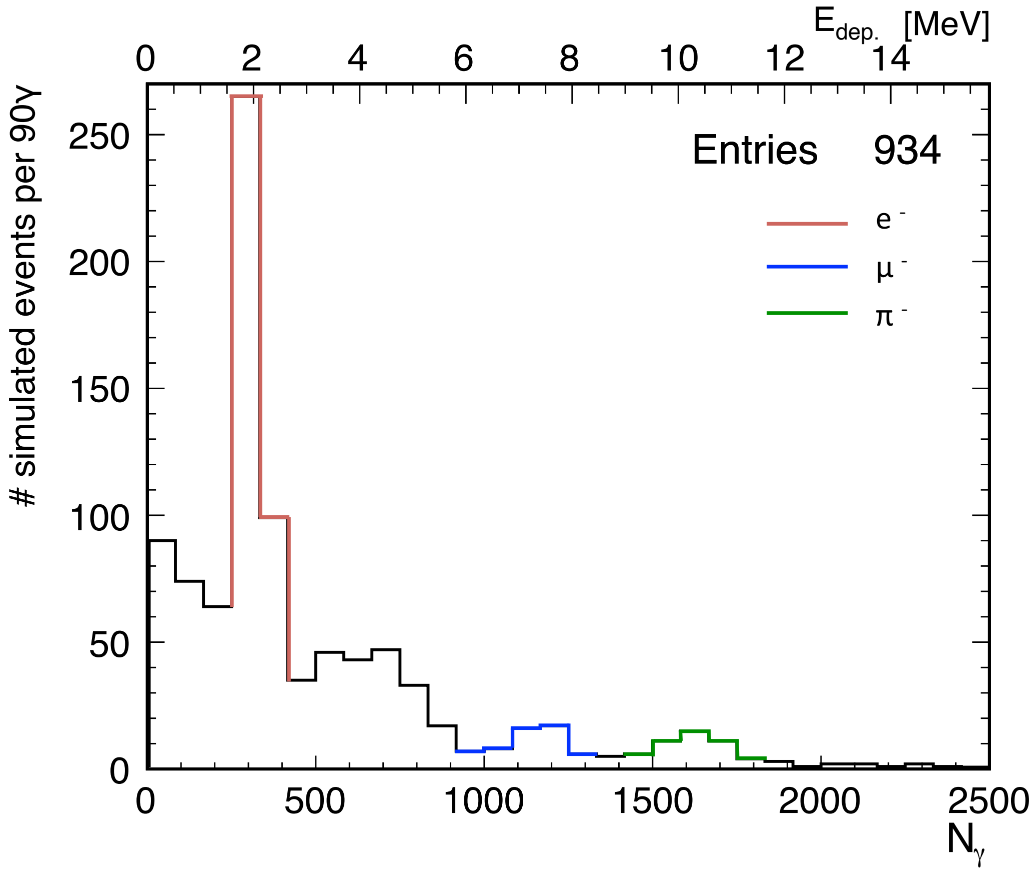

Figure 5.12: A simulation of the spectrum of photons produced - corresponding to the energy deposited by the scintillation process - analogous to the output of the DAQ ADC system.

discrepancy between the particle tracking and scintillation codes, such that the results are analogous to within a co-ordinate flip in the y-axis.

5.1.2 Experiment

After benchmarking the experiment through simulation, ascertaining expectation values for particle counts and corresponding energy depositions, the detector system was then constructed. Three systems were employed for readout: a TDC, ADC, and a scalar for hit counting. The TDC and ADC can be used in tandem but the scalar has to be used separately to calculate the trigger rate. The TDC is

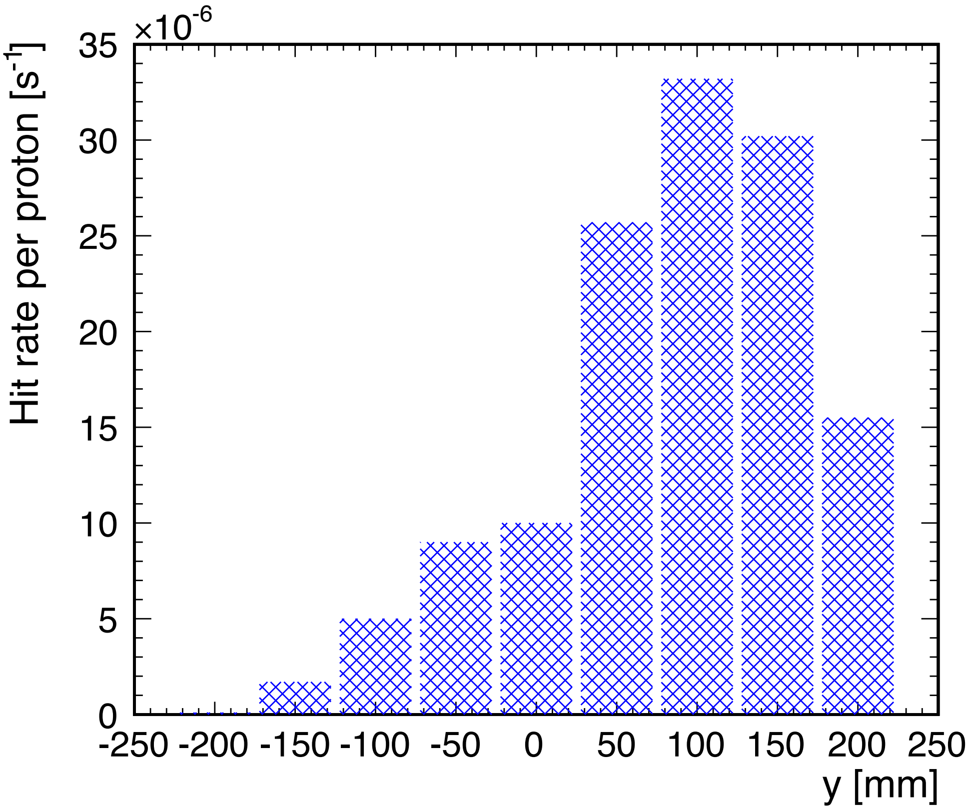

Figure 5.13: The number of particles causing scintillation, for 1×10^6 initial protons incident on the graphite target, at discrete positions along the y-axis of the beampipe.

used to provide positional data on the events by analysing the difference between the time of hits on one side of the scintillator to the other. The ADC meanwhile is used to measure the size of the analogue signal in order to ascertain the strength of the signal; a function of the energy deposited.

DAQ design

During the experiment a scintillator with dimensions identical to those described in Sec. 5.1.1 was employed. Two MPPCs are attached to each end (see Fig. 5.2), the MPPCs used (Hamamatsu S10362-11-050c) having 400 pixels in a 1 mm^2 area. The MPPCs output a charge proportional to the number of photons detected, these signals forming clear bands corresponding to the number of photons detected (see Fig. 5.14) these peaks are then a measure of the energy deposited within the scintillator. The simulations imply that no event should produce more than 200 photons per MPPC so with 400 pixels no significant pileup is expected.

Each MPPC has an operational voltage (V_{op}) in the range 70.92 V to 70.99 V, supplied using a standard power source with an output error of ± 0.05 V. The other pertinent specifications of the four MPPCs can be found in Tab. 5.3. Of these specifications the one of particular interest is the dark current. Dark current is a measure of the average number of non-photon 'signals' produced by the MPPC. These signals are primarily caused by thermal excitation of the MPPC and are random in nature. While they cannot be distinguished from genuine signals they are very rarely of a strength greater than 1 photon.

In order to ensure the highest fidelity of signal, each MPPC was attached to a CR-RC circuit (Fig. 5.15) to reduce noise. This combination of a low- and high-pass filter allows the passage of only frequencies within a desired range i.e. the signal region. The resistor and capacitor values, displayed in Fig. 5.15, are calculated to isolate this region.

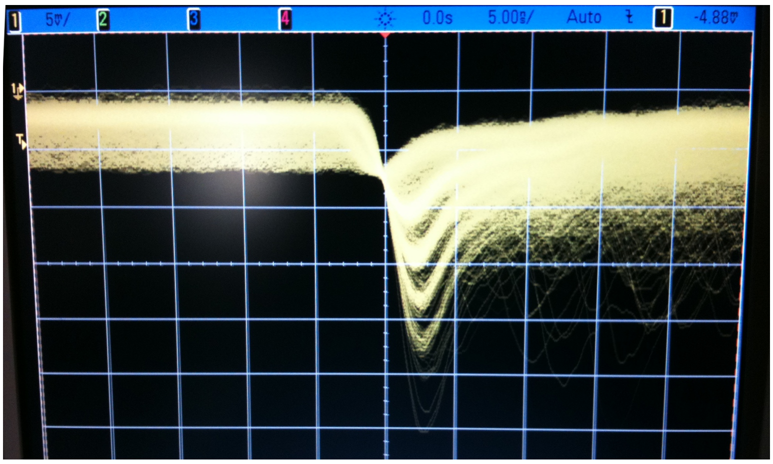

Figure 5.14: The MPPC response, with each integer photon signal corresponding to a curve of higher magnitude.

MPPC	V_{op} [V]	Dark current [$\times 10^3$ count s^{-1}]	Gain [$\times 10^5$]
1	70.99	308	7.47
2	70.92	331	7.52
3	70.92	355	7.50
4	70.92	332	7.49

Table 5.3: Specifications of the four Hamamatsu S10362-11-050c MPPCs used, as supplied by the manufacturer.

As well as correct implementation of appropriate resistors and capacitors, one of the biggest problems in reducing noise is achieving satisfactory grounding. To minimise this noise all the wires were shielded and copper tape was used as grounding strips. To ensure easy repositioning of the scintillator each MPPC was attached to its respective circuit using a short length of LEMO cable (a prototype without the cable extension can be seen in Fig. 5.16).

After the signal shaping and noise reduction of the CR-RC circuit the largest source of background is the dark current of the MPPCs. As was noted in Sec. 5.1.2, the dark current signals are rarely larger than that of a single photon, so by using a voltage discriminator with a threshold equivalent to that of several photons it is possible to remove most of the dark current.

Using the discriminator removes the majority of the dark current but there are a small number of dark signals that will exceed the threshold. These signals can be removed by imposing a coincidence condition prior to triggering; by requiring

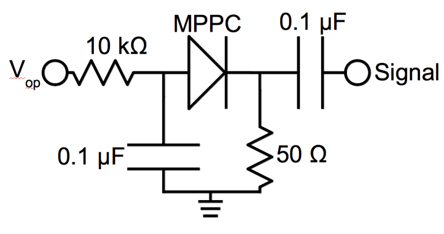

Figure 5.15: Diagram of the CR-RC circuit used to reduce noise. The 0.1 μF capacitor connected to ground is ceramic due to its need to accommodate high potential differences.

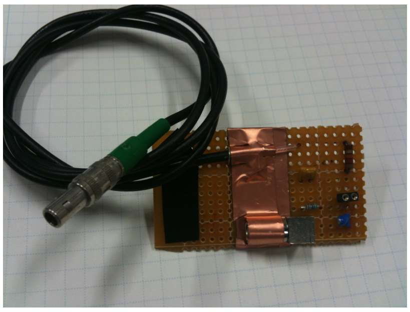

Figure 5.16: Prototype CR-RC circuit without MPPC mounting extension.

that all MPPCs have produced a suitable signal it is possible to vastly reduce the background. The dark current signals that are allowed through the discriminator are very rare and the likelihood of one occurring on each MPPC is very small, of the order 10^{-5}.

The DAQ is split over two systems: the trigger and the readout. The trigger is primarily implemented using Nuclear Instrumentation Module (NIM) other than the veto release. In contrast the readout is performed by Computer Automated Measurement and Control (CAMAC) modules. A block diagram of the logic used can be seen in Fig. 5.17.

There are four sets of delays needed in the logic, the first being the ADC cable delays. The delay gate generator ensures that the ADC has had the necessary 500 ns to process the signals. ADCs operate by charging a capacitor, which must fully discharge before the ADC can be read. The delay on the TDC start signal ensures the 4-fold coincidence detector's signal is '1' first, i.e. it ensures that MPPC1 is the start signal. The delays on the stop signals guarantee that they always arrive after the start signal, and that the start-stop interval is greater than the TDC's minimum time.

Calibration

Calibration of the experiment falls into two distinct sections: calibration of the trigger and calibration of the readout system. The main functions of the trigger are

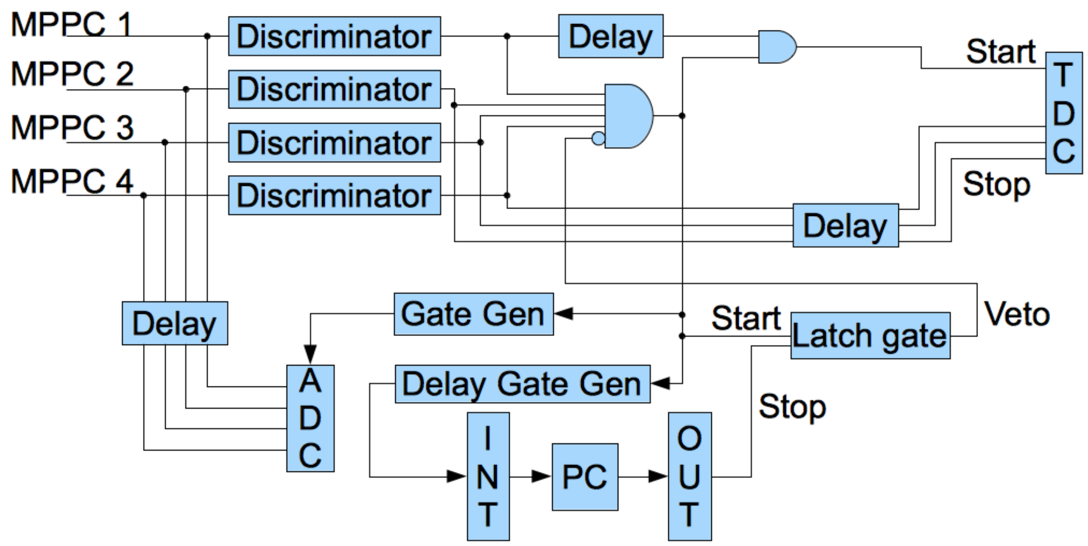

Figure 5.17: A block diagram of the DAQ system. Vertical boxes correspond to CAMAC modules while horizontal and logic symbols represent NIM modules. Signals move left to right unless otherwise indicated. The only analogue signals are between the MPPCs, the discriminator, and the ADC.

to reduce the number of background events and maintain synchronicity between the various signals. The calibration of the read-out focuses on converting the digital results into physical values.

Determining the timing of the delays and the signal duration are the primary goals of calibration. Each signal has to arrive at the right time and stay in the correct state for the trigger to properly respond.

The prescription for ensuring the correct timing between sections is a general solution: the relevant signals are read out to an oscilloscope and then adjusted until they produce the desired result, normally the triggering of the next signal. The exception to this is setting the delay for the analogue signals in the ADC. The ADC signals have to be delayed long enough such that they are wholly contained within the ‘gate’ signal originating from the coincidence detector. An output of this can be seen in Fig. 5.18. The delay on the gate generator attached to the interrupt register ensures that the ADC is ready to be read when the PC starts requesting data. The TDC delays ensure that the start signal is due to MPPC1 and that the stop signals arrive after the minimum interval measurable by the TDC. The values for the delays can be seen in Tab. 5.4.

As well as setting the delays to ensure proper execution of the DAQ, the signal widths need to be carefully set. This is achieved via a similar method to the delay calibration: each input signal is inspected via the oscilloscope and then the timing of the output signal is checked. The simplest widths are those that only need to meet a minimum requirement to trigger their module: these are the start/stop signals of the TDC (where the separation of these signals is what matters) and the input to the interrupt register. This only needs to be strong enough to alert the PC that there is data to collect. As well as these simple signal widths there are several signal widths that directly impact the quality of the data recorded. The first of these is the width of the ADC gate: it needs to be long enough to receive the entire signal without picking up extra signals or dark current. The other critical width is that of the discriminator output. This width is salient as the co-incidence detector will only trigger when all four signals are present. The requirement that all four signals are high means that the longer the discriminator signal is, the more likely it is that background signals will cause false positives. Conversely, too short a signal and the peak of another signal may not have reached

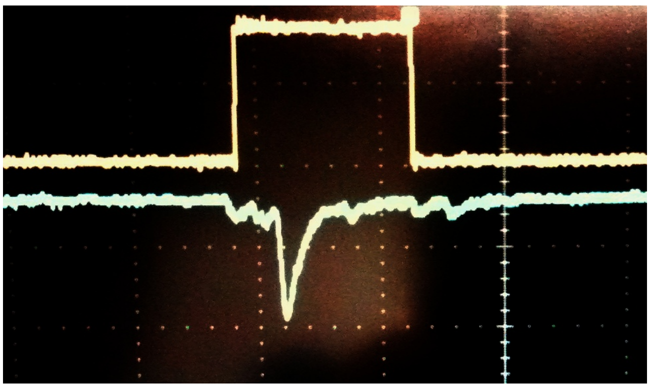

Figure 5.18: An example of the oscilloscope display showing the ADC gate and the signal.

Line	Delay [ns]
MPPCs to ADC (cable)	100
Start (pre ‘AND’)	50
Stop	200
Int gate delay	500

Table 5.4: The delays used to synchronise signals.

sufficient strength. The widths optimised to mitigate these issues can be seen in Tab. 5.5.

Lastly the threshold used by the discriminators requires calibration. This process is vital in removing dark current signals that otherwise could swamp the trigger. The majority of dark current signals only have the strength of a single photon, whereas even a low energy particle will produce hundreds of photons. Because of the comparatively high energy of the particles a photon limit of 7 photons is set, equivalent to $E_{\text{dep.}} = 0.057$ MeV according to the relation in Eq. 5.5. This corresponded to a threshold of 30 mV: the minimum threshold that could be set, thus removing the sensitivity to very low-energy particles. The effectiveness of the discriminator and co-incidence detection can be quantified in the ratio of trigger rates between beam-on and beam-off.

Collimated radioactive isotopes are used to calibrate the ADC and TDC: in this case ^{90}Sr, ^{22}Na, and ^{60}Co. The ^{22}Na and ^{60}Co isotopes have well defined photo-peaks so are used to calibrate the ADC, generating an energy calibration curve. The ^{90}Sr is used to calibrate the TDC as it is the most radioactive of the

Line	Width [ns]
Start/Stop and Int	100
Discriminator	75
ADC Gate	100

Table 5.5: The widths of the various signals within the DAQ.

three sources. It also decays via beta emission so is not suitable for precise energy calibration. The isotopes, decay processes, and peak energies used for calibration are the same as in Tab. 5.1.

Calibration of the ADC is performed in two stages: measuring the pedestal of the ADC and calibrating it for energy measurements. The pedestal of the voltage is the offset gained from the background and the modules. It is easily accounted for by taking measurements using the ADC but without triggering. This is done by manually triggering the interrupt and ADC gate generator using the OUT register, then disabling the TDC (as it would time out with no start/stop signals). To account for it the mean value of this is subtracted from all further measurements.

To calibrate the ADC the sources were placed on the scintillator for extended amounts of time. The peak(s) found were used to associate a bin with a certain energy, thus establishing a calibration curve. Unfortunately, due to time constraints, not enough events were gathered. The photo-peaks couldn't be distinguished in order to generate a calibration curve.

Calibration of the TDC is carried out with a collimated ^{90}Sr source placed at various points on the centre line of the scintillator. The difference between the TDC data for the left and right side gives a ratio of separation of the event from the MPPC. Repeating these readings in different lateral positions generates a calibration curve. The result of these measurements is shown in Fig. 5.19. Initially the average of the differences between the two left and one right MPPC were used but this introduced a bias to the results. Additionally two of the MPPCs broke during data taking, making direct comparison difficult.

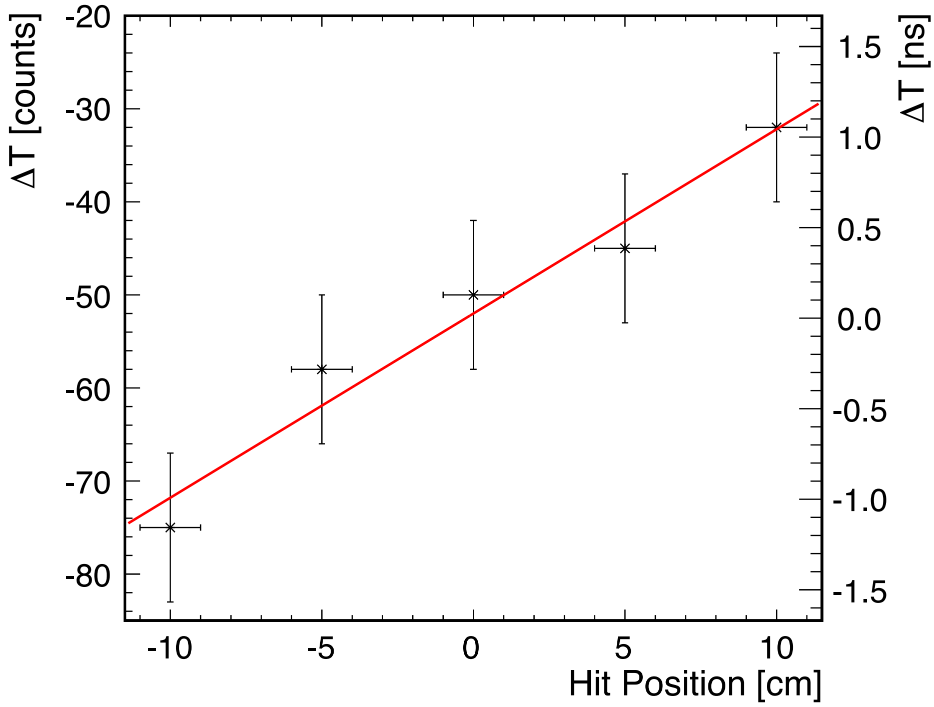

Figure 5.19: The TDC calibration curve. The left-right time difference was taken for different positions, in each case exposed to a collimated ^{22}Na source. This source was attached flush to the scintillating bar. The left-right count difference is converted to time using the relation in Eq. 5.4 from simulated calibration data.

5.1.3 Results and analysis

Hit count

The hit rate was calculated using scalars to record the beam intensity and trigger count in a 50 s period. Taking these values when the beam was on, then off, gives the bias of the results. The hit rate per initial proton can then be calculated using these values and the relation

$$\frac{dH}{dt} = \frac{T_{\text{on}} - T_{\text{off}}}{I_{\text{on}} - I_{\text{off}}} \frac{1}{q} \,, \tag{5.6}$$

where H is the hit count, T refers to the trigger rate, I is the proton beam intensity, and q is the proton charge.

The hit rate is calculated for seven vertical positions along the end of the beam pipe. These values are shown in chronological order of recording in Tab. 5.6, and displayed graphically in Fig. 5.20. Between the measurements at -150 mm and 50 mm one of the MPPCs broke and another developed an intermittent fault (the reason for the displaced second reading at 0 mm). Taking the ratio of the two separate measurements at 0 mm gives a scaling factor that needs to be applied to all subsequent measurements made with the non-functioning MPPCs. These rescaled hit rates are also displayed in Fig. 5.20. This demonstrates that there is a slight bias in the beam towards the top of the beam pipe. The data of Fig. 5.21 can then be used in a chi-squared analysis to test the validity of the simulation against experimental results. The error used in the calculation are those associated with both the simulation and experimental, added in quadrature. The number of degrees of freedom is equivalent to the number of data points, in this case five. The relation used to calculate this comparison is therefore,

$$\chi^2/\mathrm{ndf} = \frac{1}{n-1} \sum \frac{(N_{\mathrm{sim.}} - N_{\mathrm{exp.}})^2}{\sqrt{\sigma^2_{\mathrm{sim.}} + \sigma^2_{\mathrm{exp.}}}} \ , \tag{5.7}$$

giving the result:

$$\chi^2/\mathrm{ndf} = 1.6 \ . \tag{5.8}$$

This simulated beam bias stands up well against the experimental data, serving as a validation of the scintillator simulation framework.

ADC

As has been noted in Sec. 5.1.2, calibration of the ADC for energy measurements did not produce any useable data.

y [mm]	Trigger rate [Hz]		Intensity [nA]		Hit rate [Hz]
	beam	no beam	beam	no beam	
0	2151736	119	1.26	0.24	2126
−50	1438685	69	1.29	0.38	1579
−150	446302	30	1.21	0.38	537
50	1663702	83	1.30	0.40	1848
150	1080170	40	1.34	0.38	1125
0	1307015	34	1.40	0.40	1314

Table 5.6: Hit rates at a range of vertical positions.

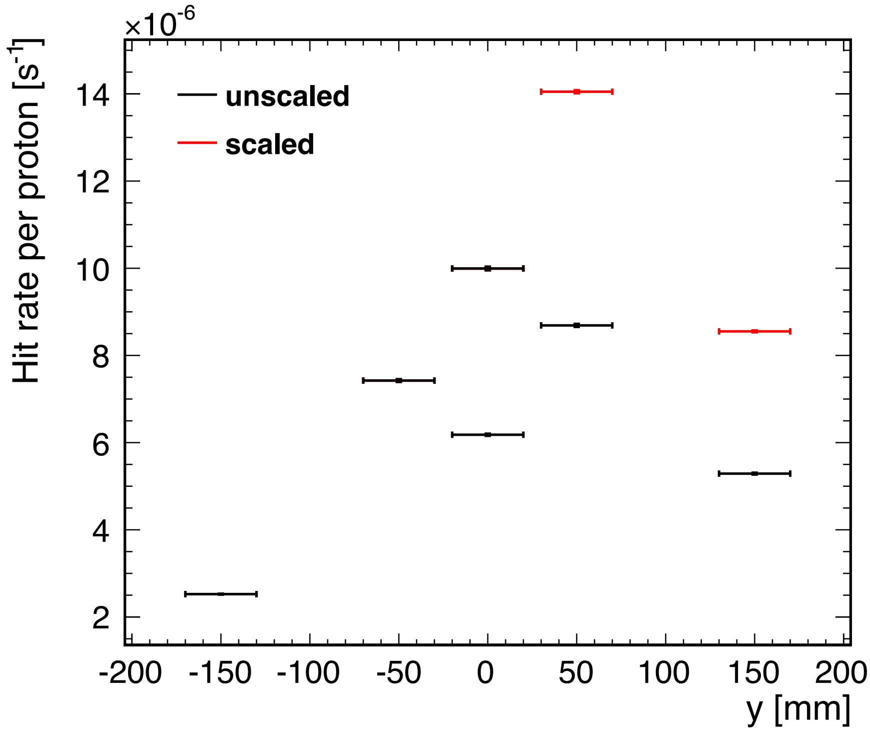

Figure 5.20: Hit rate for a range of vertical positions at the end of the beam pipe. One of the MPPCs broke mid-run (for +50 and 150 mm) so the measurement at 0 mm was retaken, scaling the readings at +50 and 150 mm up by a factor of 1.63.

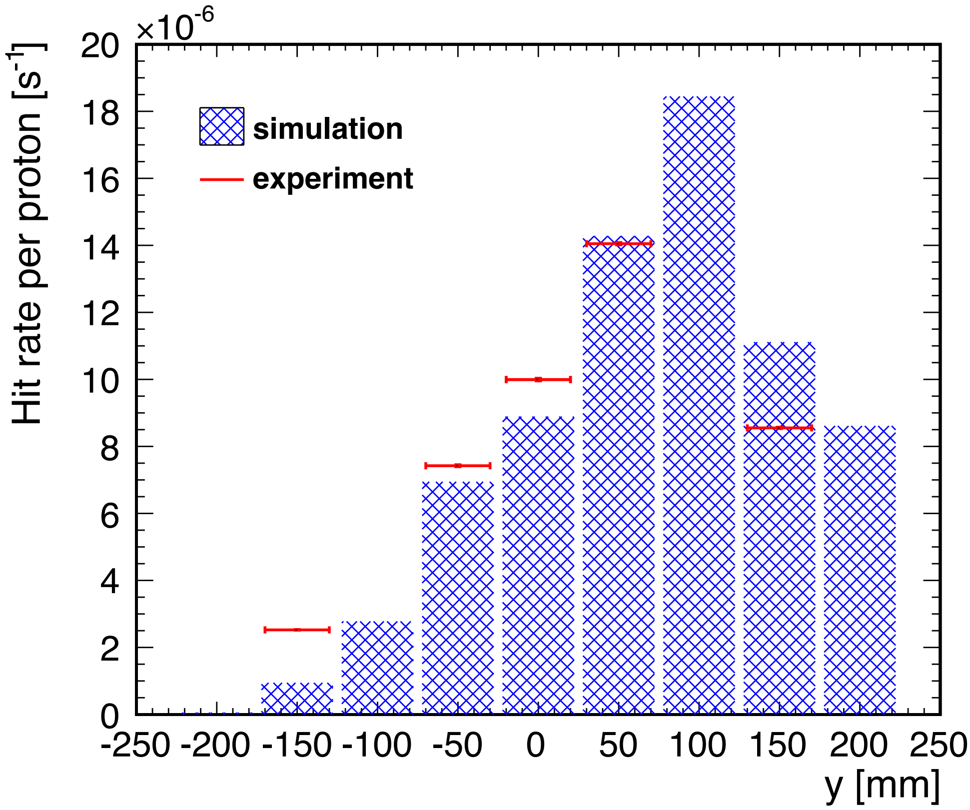

Figure 5.21: Rescaled experimental hit rates of Fig. 5.13 compared with the simulated number of particles causing scintillation per proton of Fig. 5.13. A threshold of 0.057 MeV is added to the simulated data, corresponding to the experimental photon limit of 7 photons.

TDC

As can be seen in Fig. 5.22, even after calibration a large number of events resolve outside the physical bounds of the scintillator, i.e. beyond ±190 mm. This occurs in every data set from each vertical repositioning, shown in Fig. 5.23. The origin of this problem is in the DAQ logic and is outlined in the timing diagram in Fig. 5.24. This diagram shows an earlier signal from MPPC 2 confused with the true stop signal. This flaw in the DAQ means that none of the horizontal position information can be trusted. This flaw does not affect the ADC data, however, as the gate for this is taken from the co-incidence detector. Similarly, the hit rate is unaffected as this is counted from the same signal.

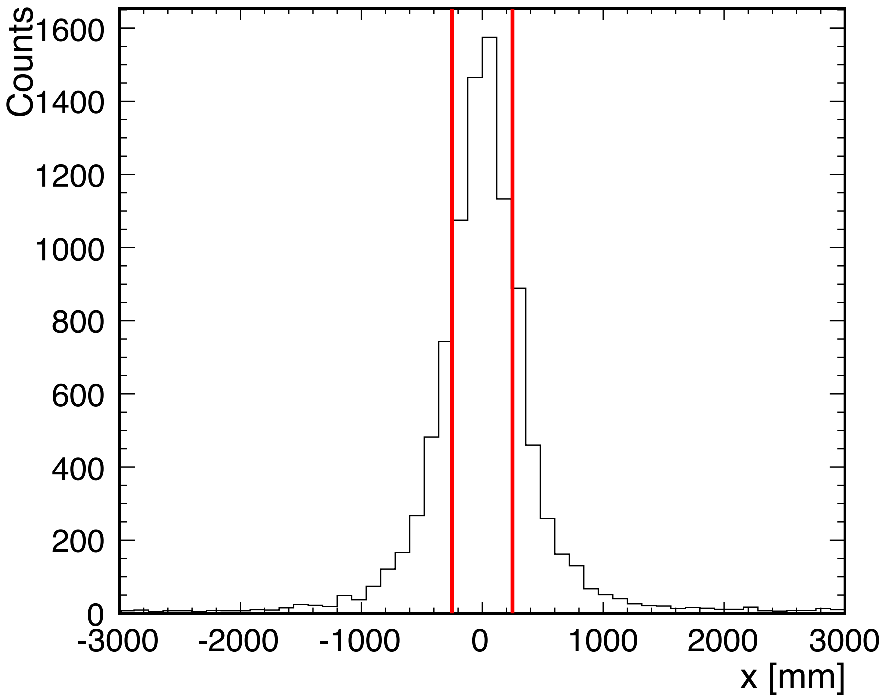

Figure 5.22: Distribution of hits for vertical position of 0 mm. Note the significant number of hits outside of the physical volume of the scintillator (shown by the red lines).

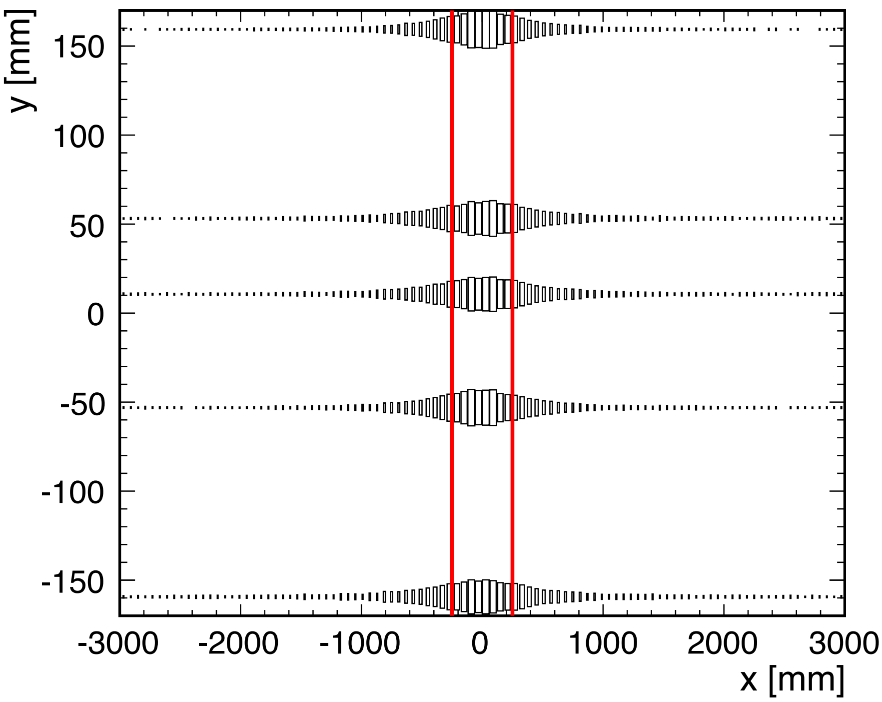

Figure 5.23: 2D plot of particle hits on the scintillator. The size of the box indicates the number of hits. The red lines expound the physical volume of the scintillator.

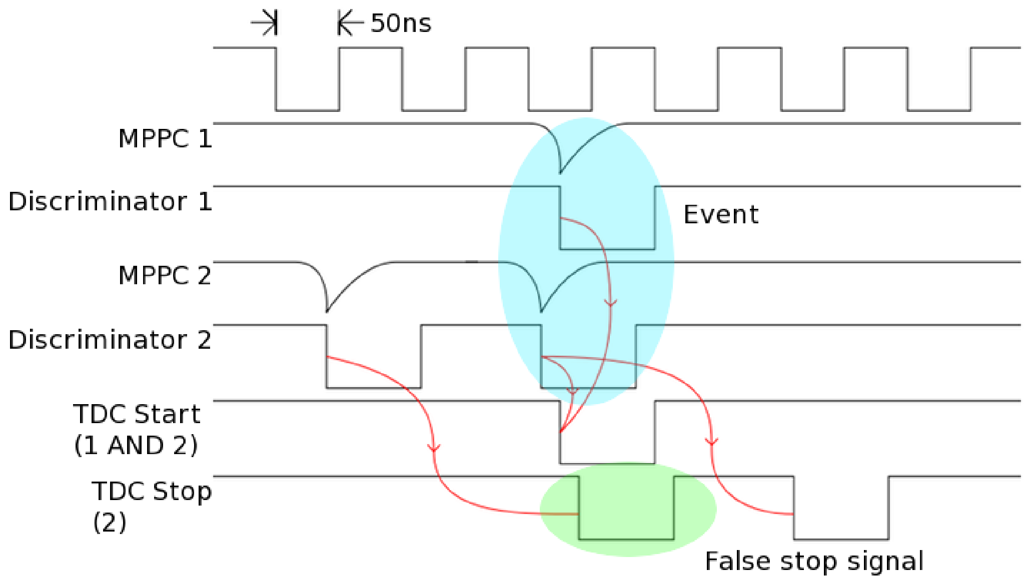

Figure 5.24: Timing diagram showing one possible origin of erroneous TDC measurements. Due to the lack of 'AND' gates on the stop lines earlier signals (either from dark current or other events) can create stop signals that don't correspond to those from the actual event.

5.2 Run-II

Upon analysis of the Run-I data it was decided that the Run-II scintillator design should be optimised to make hit rate measurements only, in order to have comparable results for both periods. For this reason a scintillator with dimensions better suited to such a measurement was chosen: a disk with a radius of 3.5 cm and depth of 2.0 cm. These dimensions allow the scintillator to be moved in both transverse planes for a 2D representation of the beam (limited to 1D in Run-I).

5.2.1 Simulation

Again the Polyethylene scintillating material is modelled in `GEANT4`, wrapped in a layer of aluminium foil of thickness 0.2 mm, with the surrounding volume given the properties of air. In this case the MPPCs are placed on-axis of the $x - y$ plane, at 90° intervals around the scintillator. A visualised output of the disc undergoing scintillation can be seen in Fig. 5.25. In this simulation run the scintillator was placed 85 cm from the end of the beampipe as other detector clamps occupied the intermediate space. In this figure only the photons detected by the MPPCs are displayed, highlighting their positions at 0°, 90°, 180°, and 270°.

A simulated calibration of the scintillator, similar to that of Run-I, was performed. The correlation between energy deposition and number of photons detected is shown in Fig. 5.26, analogous to those in Fig. 5.8 for Run-I.

A ratio between the scintillating volumes of Run-I and -II is

$$R = \frac{V_{\mathrm{bar}}}{V_{\mathrm{disk}}} = 1.48\ . \tag{5.9}$$

Naively this difference in volumes between the two runs should result in 1.48 times more photons in Run-I compared to Run-II. However, by comparing Fig. 5.8 with Fig. 5.26 it is apparent that the simulated ratio is closer to a threefold

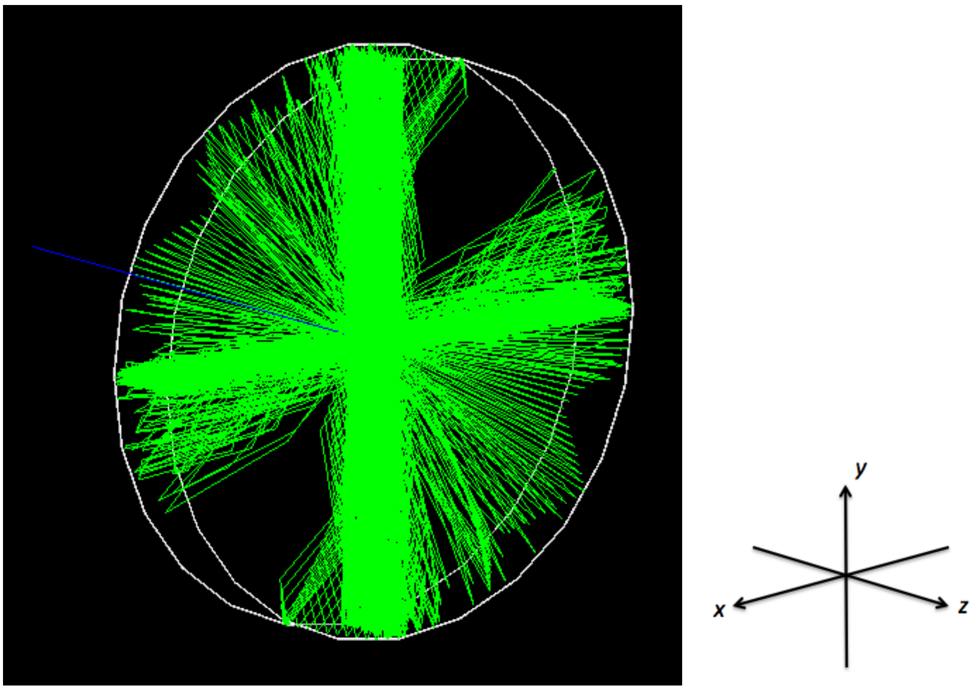

Figure 5.25: An example output of the plastic scintillating disc with a 1 MeV incident electron (arriving along the negative z-axis) producing optical photons (green) detected by the MPPCs.

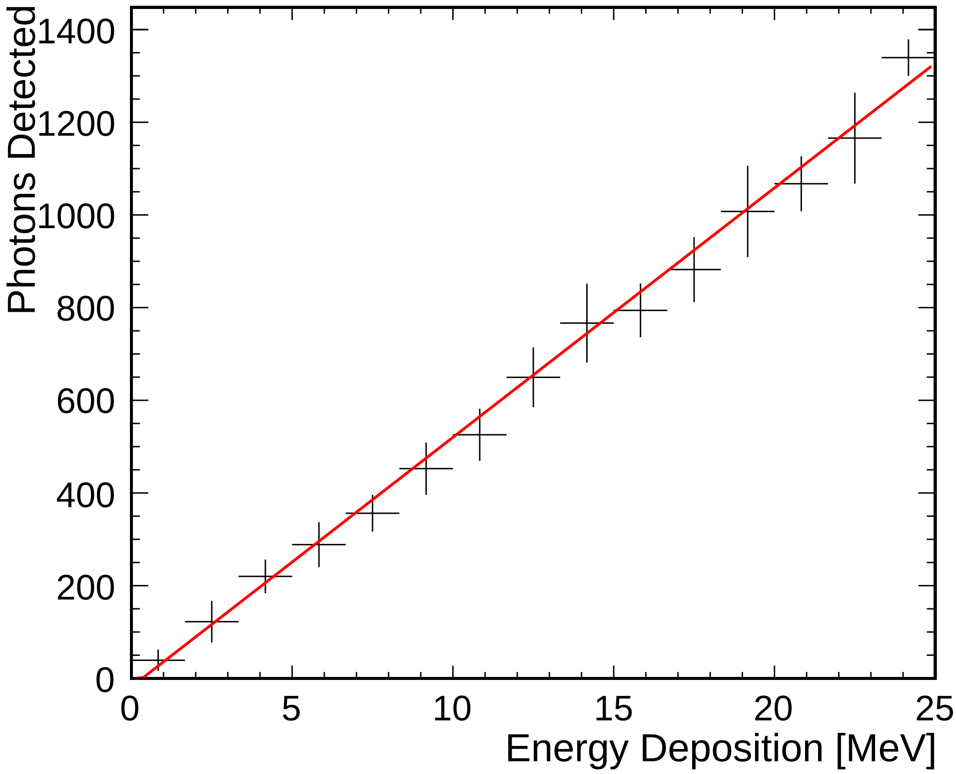

Figure 5.26: The total number of optical photons detected for a series of electrons, with differing initial energies ranging from 0.511 to 7.0 MeV. The energy deposition is plotted (rather than the initial energy of the particle) so as to fit a linear relation between photons detected and $E_{\mathrm{dep.}}$. Each point has error-bars corresponding to $\sqrt{N_\gamma}$.

increase. The discrepancy between these two ratios is due to the geometry of the two scintillators. Owing to the symmetrical nature of the disk the photons have a higher average path length if they do not travel directly to the MPPC from scintillation. This increase in path length in Run-II means the photon is approximately twice as likely to reabsorb than intercept the surface of an MPPC. The number of photons detected from a typical μ^- in MuSIC (with $p_\mu \approx 59$ MeV/c) can therefore be derived from Fig. 5.26 as approximately 800.

As in Run-I the simulated distribution produced from a 392 MeV proton beam incident on a graphite target is tracked through the magnetic lattice. This is then allowed to propagate, in the absence of a magnetic field, for 85 cm until it is outputted from `G4beamline`. The beam profile at the point of interception with the scintillator can be seen in Fig. 5.27.

There are far fewer particles at this scintillator position than in Run-I. Due to the helical trajectory of the beam the particles exit the beampipe off-axis and with a large transverse momentum. If left to propagate unchecked through a drift length (in this case 85 cm) the beam will diverge in the approximate direction it is travelling. This can be seen in Fig. 5.11 with the distribution skewed in the negative x- and y-direction, centred at approximately $(-0.07, -0.14)$ m. In Run-II the scintillator is positioned 80 cm further from the end of the beampipe than Run-I. The beam should therefore reach the scintillating disk with an increase in transverse displacement, shown in Fig. 5.27. To vindicate the simulation a measurement of hit counts across this area was performed.

5.2.2 Experiment

The positioning of the scintillating disc was constrained by the rigging used for its placement. For this reason a hit count was initially taken within the range of the rigging: $|x| < 20$ cm and $|y| < 20$ cm. According to the `G4beamline` simulation the majority of the beam lies outside this region so the rigging was moved to accommodate this. This allows measurements to be made over a larger region in the negative x- and y-direction. The results of these measurements can be seen in Fig. 5.28, with the hits normalised to the intensity of the beam at that point.

Particle	e^-	e^+	μ^-	μ^+	π^-	π^+	p
Number	1066	219	9	122	6	51	82
Percentage	68.6	14.1	0.5	7.8	0.4	3.3	5.3

Table 5.7: A breakdown of the number of particles, and relative abundance, in the beam as simulated at the position of the scintillator (85 cm after the beampipe terminates).

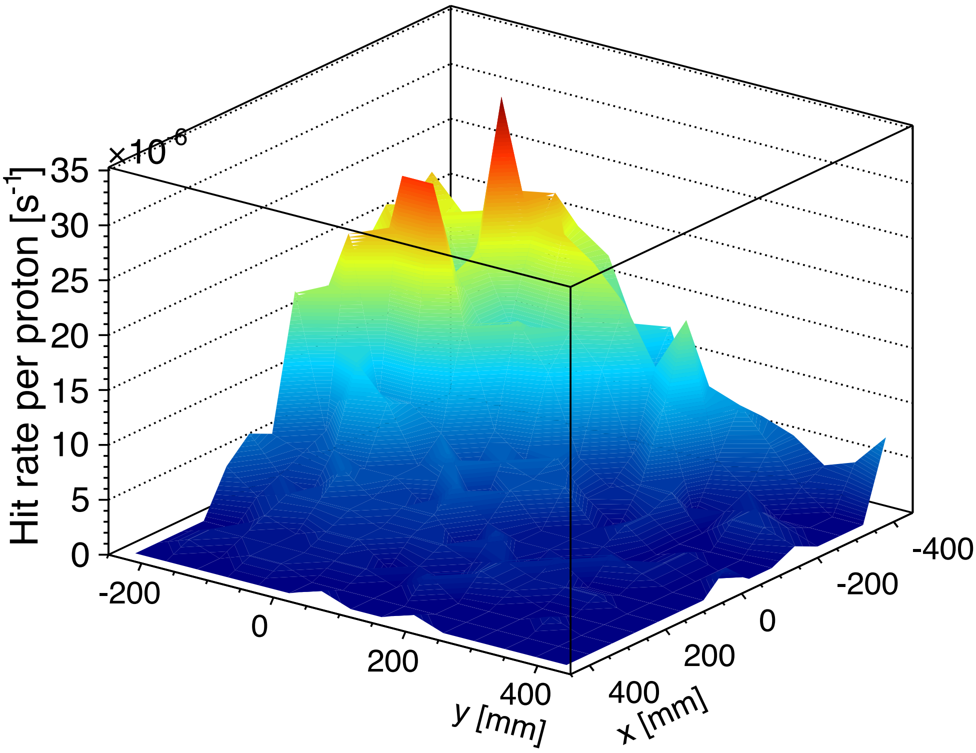

Figure 5.27: Three-dimensional spatial distribution of the beam (with a cut removing neutral particles that are not tracked through the solenoidal magnetic field) incident on the scintillating disc, tracked using `G4beamline`. The profile is taken 85 cm away from the end of the beam pipe, with a 7 photon trigger.

5.2.3 Result and analysis

The hit rates from both experiment and simulation show a strong correlation, with an obvious bias in beam position towards the negative x- and y-directions in both. Upon adding a 7 photon trigger (corresponding to $E_{\text{dep.}} = 0.019$ MeV) to the simulation, analogous to that in the DAQ, the hit rate per proton lies in the region of the experimental data relative to the beam intensity over the measurement window. By selecting the simulated hit rate at the discrete experimental positions, as displayed in Figs. 5.27 and 5.28, a chi-squared analysis may be performed between the two data sets, analogous to that found in Sec. 5.1.3. Again the errors used in the calculation are those associated with both the simulation and

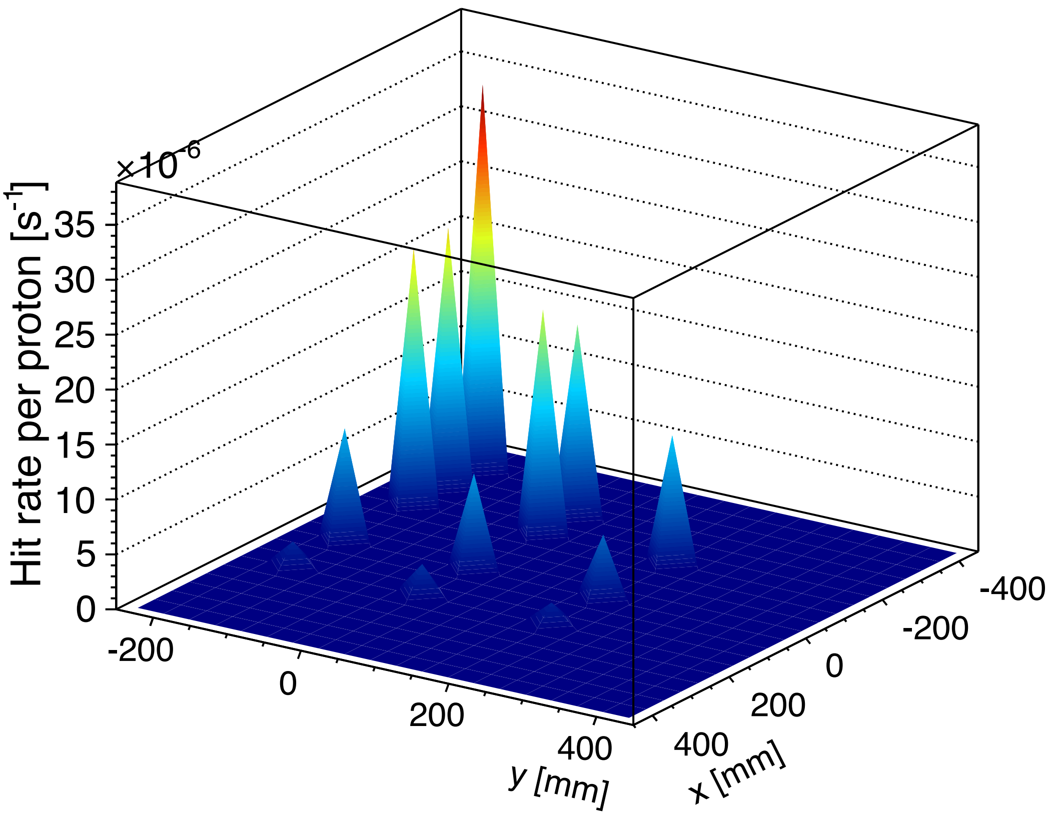

Figure 5.28: A three-dimensional representation of the hit rate on the scintillating disc over a range of discrete transverse positions. The counts are normalised per proton using the beam intensity over each measurement window.

experimental at the discrete experimental positions of the scintillating disk, added in quadrature. The number of degrees of freedom is equivalent to the number of experimental data points, in this case nine. Using the relation in Eq. 5.7 a chi-squared is calculated:

$$\chi^2/\mathrm{ndf} = 1.7\ . \tag{5.10}$$

This again demonstrates the validity of the `G4beamline` and `GEANT4` simulations in modelling the experimental results.

5.3 Conclusions

A novel method of particle capture using superconducting magnets and graded magnetic fields, first proposed by Lobashev in the 1980s, has been demonstrated using a nano-ampere proton beam incident on a graphite target. The increased pion yield from the utilisation of these superconducting solenoid fields has thus been vindicated placing the future of the COMET experiment on an assured footing.

The first two experimental runs of the MuSIC experiment have been analysed, with both data sets modelled using a combination of codes: `G4beamline`, used to simulate the novel superconducting pion capture solenoid and beamline, and `GEANT4`, simulating the detector system and ensuing physics processes.

The integrity of these codes has been verified via direct comparison with experimental results. In the first instance of the experiment 1D hit rates were measured. However, due to a shortcoming of the DAQ software the ADC and TDC results were rendered unusable. The second experimental run focussed solely on 2D hit rate measurements with the simulation again agreeing with experimental data.

Chapter 6

EMMA

6.1 Current status of accelerator science

A significant fraction of the particle physics community is currently focussed on the Large Hadron Collider (LHC) at CERN in Switzerland. This type of accelerator is at the forefront of the energy frontier and aims to push this upper limit by colliding extremely small bunches of particles with small emittances. The desire to probe at ever higher energies has always furthered the technological development in other areas of accelerator science.

In most accelerating lattices, whether they be linear or circular, the particles are guided by magnetic fields. When a particle is accelerated on a circular path its radius of trajectory increases, requiring an increase in the magnetic field to keep it on a fixed orbit. Modulating magnetic fields are difficult to maintain and as such are often limited to a few kHz. Particles therefore spiral outward through acceleration and a larger aperture is needed. For this reason high power machines tend to use fixed magnetic fields, allowing for this radial increase over a few turns of rapid acceleration. These accelerators are known as cyclotrons and are composed of simple magnets to bend the beam.

Synchrotron accelerators (such as the LHC), where the magnetic field strength scales with beam energy, tend to incorporate focusing quadrupole magnets, which focus the beam in one transverse plane but defocus in the other. It is therefore necessary to implement a minimum of two quadrupoles in one cell to focus in both the horizontal and vertical plane (this can be shown through the thin lens approximation). It is a natural progression in accelerator design that a Fixed-Field would be combined with these focusing cells of Alternating-Gradient (FFAG). A description of this type of multi-quadrupole magnetic cell can be found in Sec. 2.2.1.

FFAG accelerators were proposed half a century ago [47, 48, 49], when acceleration of electrons was first demonstrated. Technical issues, such as the modelling and manufacture of complex magnets, stunted the growth of this field of research. FFAGs have become the focus of renewed attention in recent years. A proof-of-principle FFAG has recently demonstrated proton acceleration at KEK [50] and a number of applications for this type of acceleration have emerged.

6.1.1 Scaling and non-scaling FFAGs

To avoid the slow crossing of betatron resonances associated with a typical low energy-gain per turn, the first FFAGs designed and constructed so far have been based on the 'scaling' principle. The scaling term refers to the optics of the machine, which implies that the radial orbits are geometrically similar and scale as the beam is accelerated [51]. This means that the number of betatron oscillations per turn (known as the tune) remains constant and the magnetic field follows the relation $B \propto r^k$ (where r is the radius and $k \gg 1$). Due to this optical property the radial momentum compaction factor is not much greater than that of a normal synchrotron at relativistic velocities. The magnet apertures are therefore large, requiring machining on a large scale which is generally difficult and expensive, leading to recent interest in 'non-scaling' FFAGs.

In non-scaling FFAGs the bending and focusing is provided simultaneously by focusing and defocusing quadrupole magnets repeating in an alternating sequence. The radial orbits and tunes do not scale but have different shapes over a momentum range. This results in a large compaction of orbits through the bending elements of the lattice and therefore a parabolic relation between the mean radius and particle energy. This momentum compaction is the reason these machines are a prime candidate for muon acceleration as they can accept much larger transverse/longitudinal beam emittances than conventional synchrotrons and, conveniently, the magnetic field necessary to accomplish this compaction is linear.

In circular accelerators, such as synchrotrons and FFAGs, the beam repeatedly encounters the same magnetic structure. Under certain conditions the magnetic lattice may cause the circulating beam to resonate, causing the oscillation amplitude of the beam to grow. This results in beam blow-up and possibly complete beam loss; a phenomenon known as optical resonance. The beta function is periodic for a repeating lattice, with the tune defining the number of betatron oscillations per revolution.

In the case of the FFAGs the tunes are not held constant during acceleration, resulting in the beam crossing many transverse resonance conditions, again potentially leading to beam blow-up and loss. If rapid acceleration can be achieved then resonance crossing can be effectively ignored as the beam occupies the ring for only a small number of turns. Cyclic accelerators are generally designed to avoid resonance crossing which provides freedom to choose machine parameters which are beneficial to acceleration: minimising the circumference to control intensity loss and maximising magnet aperture size by using only linear elements [52].

6.1.2 Possible future applications of FFAGs

FFAG devices are an attractive alternative to other accelerating means due to their unique mode of acceleration. The rapid acceleration of particles practically eliminates beam-degrading properties which hamper the design of other accelerating machines. As a consequence of this, specific parameters can be chosen which are conducive to acceleration. Specifically:

- smaller ring diameters,
- larger transverse/longitudinal acceptance,
- simpler, cheaper magnets.

As previously mentioned, these properties are attractive to future particle physics experiments which require a cheaper method of rapid acceleration to high intensities, whilst also demanding the unique properties of alternating gradients applicable to applications outside of particle physics.

Neutrino factories

Neutrino Factories are designed to determine the neutrino mass hierarchy and search for evidence of CP violation in the neutrino sector thus providing an insight, lacking in the Standard Model, into the dominance of matter over antimatter in the universe. Current experiments, such as T2K [53], are designed on a similar premise but have limited sensitivity to CP violation owing to the number of neutrinos produced. Neutrino factories will provide a far higher flux of neutrinos (approximately 10^{21} per year) from muon decays [54]. The muons are produced by irradiating a liquid mercury target with high energy protons, producing pions. These pions quickly decay to muons, producing a muon beam with a high angular dispersion which has to be cooled before it can be accelerated. The schematic in Fig. 6.1 utilises a muon ionisation cooling section which 'cools' the spatial dispersion of the beam. Even after cooling the beam has an extremely large

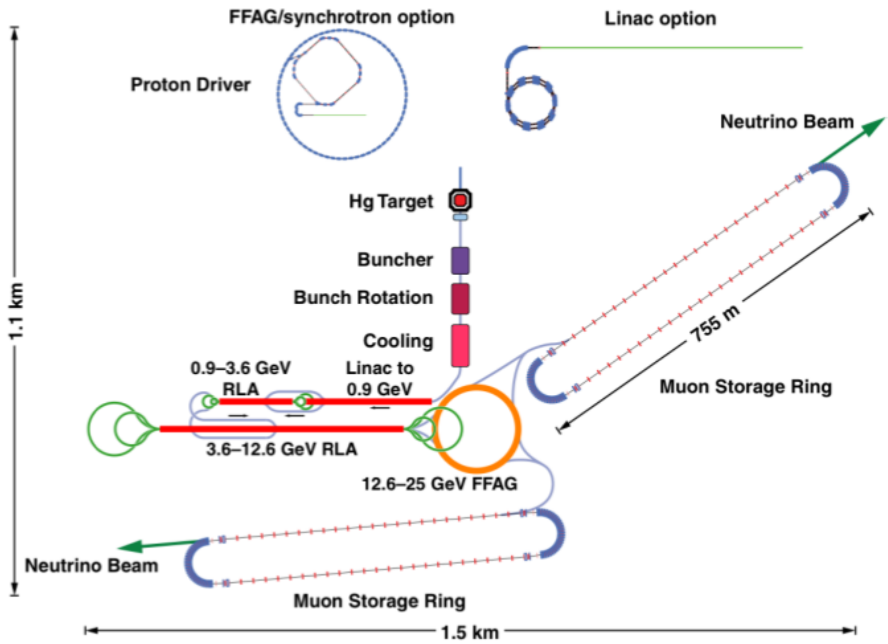

Figure 6.1: A proposed schematic for a Neutrino Factory, incorporating a non-scaling FFAG for muon acceleration [54].

emittance (3×10^4 π mm mrad, compared to e.g. 3×10^{-3} π mm mrad for the LHC) which cannot be accommodated by current synchrotrons, but can by FFAGs due to their large acceptance.

Furthermore, the lifetime of the muon in its rest frame is 2.2 μs so rapid acceleration is also a necessity as the muons need to be accelerated to relativistic velocities before they decay. For these two reasons, which are cheaper to implement than the alternatives, the most attractive option for muon acceleration in Neutrino Factories is to use an FFAG accelerator. The same properties also provide an appealing argument for the use of FFAGs as future high-energy muon colliders.

Proton therapy

In developed countries one in 25 people are diagnosed with some form of cancer per year, with a large fraction of those being treated with some sort of high energy beam [55]. Up until the early 1990s X-rays were used to treat tumours with

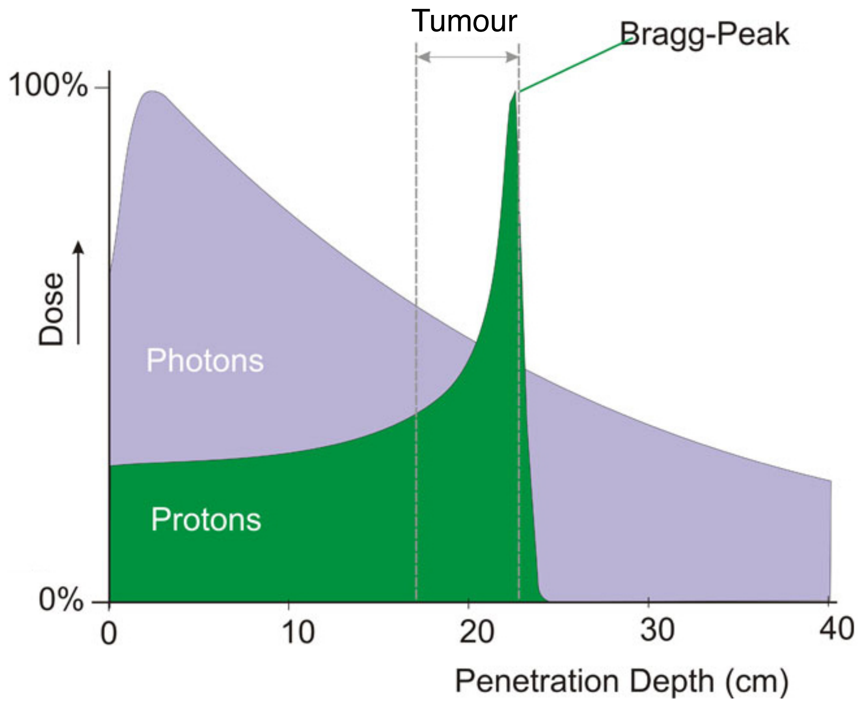

Figure 6.2: The variation in radiation dose against depth for photons (X-rays) and protons. The narrow peak at the largest depths for protons is known as the Bragg peak.

reasonable effect, yet they were inefficient as quite a large volume of surrounding tissue (including the tissue traversed to arrive at the tumour) was irradiated and/or killed. There is another problem, in that the intensity of the beam tailed off as it travelled further into the tissue, meaning the skin received the highest dose of radiation whereas the tumour received a smaller dose depending on its depth within the tissue. Both are problems for the body in general, but particularly for sensitive body parts such as the eye, brain, spinal chord etc. and for children where organs are still developing.

The concept and potential benefits of proton therapy were first proposed by Robert Wilson in 1946 [56] but, due to restrictions in accelerator technology, were not broadly realised until the beginning of the 1990s. Due to the narrow Bragg peak at the end of the depth range, illustrated in Fig. 6.2, it is possible to irradiate a well defined area whilst reducing the irradiation of the traversed tissue. The disadvantage of proton therapy at the time was the huge apparatus (even by medical standards) required for the treatments, often carried out in nuclear physics laboratories with linear accelerators and simple magnet designs. Despite these setbacks 10,000 people had been treated by proton therapy by 1993, indicating the necessity for growth in this area.

For more people to benefit from hadron therapy it is necessary to make the accelerators as easily accessible as possible, namely by making them cheaper and smaller. There are two proposed machines for cheaper hospital-based dual centres, a KEK-like FFAG or a cyclinac (a combination of a low-energy high-current cyclotron and a high-frequency linac), both of which may be suitable for proton therapy. It will be a few years before it is known which proposal offers the solution, both economically and technically, but the results coming from the FFAG sector hold a lot of promise.

Nuclear waste transmutation

The majority of nuclear reactors use uranium as a primary fuel source, which naturally undergoes fission. The emission of neutrons, that then interact with other uranium atoms, causes a chain reaction. In this case the reactor is said to be 'critical' and the reaction continues until an expiration of Uranium. This process has to be controlled to avoid nuclear meltdown.

An alternative, safer option is to use a radioactive isotope that undergoes natural fission at a much lower rate, e.g. thorium. In this case the fission requires a catalyst in the form of firing a beam of low energy protons onto a heavy metal target. The protons create a cascade of neutrons, at an approximate ratio of 30n:1p, which drives the fission process. This type of reactor is said to be 'subcritical' since the fission reactions immediately cease upon termination of the incident beam, eliminating the risk of past nuclear catastrophes.

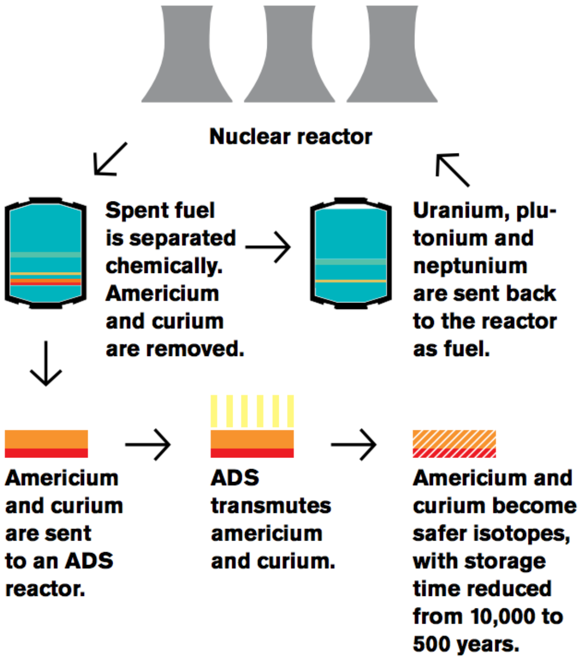

Figure 6.3: Graphic demonstrating the process of accelerator driven nuclear waste transmutation after production in a nuclear reactor [57].

As well as initiating and running a fission process, the accelerator driven system may be used to transmute and reprocess nuclear waste into shorter lived fission products. Americium and curium are two of the most prevalent products of fission, with approximately 10,000 years passing before they decay to the same level of radioactivity as the uranium used to initially fuel the reaction. By transmuting these isotopes through the absorption of slow neutrons, the storage time is dramatically reduced to less than half a century. This process is illustrated in Fig. 6.3.

The accelerator technology required for these processes is currently available, however the capital and running expenses are huge. The next generation of accelerators, specifically FFAG-driven high-power proton sources, will dramatically mitigate costs by utilising a far higher acceleration gradient [58]. The process also requires the down time of the accelerator to be limited to a few seconds per year, otherwise unacceptably large changes in temperature, leading to untenable thermal stresses, occur. Present accelerators are not reliable enough to maintain such a constant level of activity whereas FFAG accelerators potentially are. Both MYRRHA [59], Belgium, and Project X [60], USA, have programmes designed to develop and demonstrate the next generation of ADS technology.

6.2 The EMMA project

The validity of these applications needs to be verified with a similar non-scaling proof-of-principle machine, analogous to the KEK machine outlined in Sec. 6.1. The Electron Model for Many Applications (EMMA) project utilises a $10-20$ MeV linear non-scaling FFAG (ns-FFAG), constructed at the Daresbury Laboratory, UK. The EMMA ring, as well as injection and extractions lines, can be seen in Fig. 6.4. The experiment is part of the Construction Of a Non-scaling FFAG For Oncology, Radiation, and Medicine (CONFORM) project [61], which is also seeking to develop a ns-FFAG accelerator (PAMELA) for hadron therapy. PAMELA is still at the design stage so positive results from EMMA will be of the utmost importance.

The EMMA machine began commissioning in Daresbury in July 2010. It is the first of its kind so extensive simulations, using a spectrum of both analytical and particle tracking codes, had to be performed at the design stage. These simulations are also critical for in situ tuning of the machine, where measurements are being made. The experimental results will then be used to benchmark the codes, allowing optimisation for accurate modelling of future FFAG accelerators. It is shown to give a good description of the data giving confidence that it can be used to optimise the design of future FFAG accelerators.

6.3 Modelling in GPT

The General Particle Tracer (GPT) package [62] is a well established simulation tool for the design of accelerators and beamlines. GPT tracks particles based on comprehensive 3D techniques, allowing for non-linear effects of charged particle dynamics in electromagnetic fields to be modelled. The particle tracking is based on solving the relativistic equations of motion of the individual macro particles. The position x and the momentum $p = \gamma m v$ are used as the coordinates of a particle. The equations of motion for particle i are given by:

$$\frac{d\mathbf{p}_i}{dt} = \mathbf{F}_i \, , \tag{6.1}$$

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{v}_i = \frac{\mathbf{p}_i c}{\sqrt{\mathbf{p}_i^2 + m_i^2 c_i^2}} \, . \tag{6.2}$$

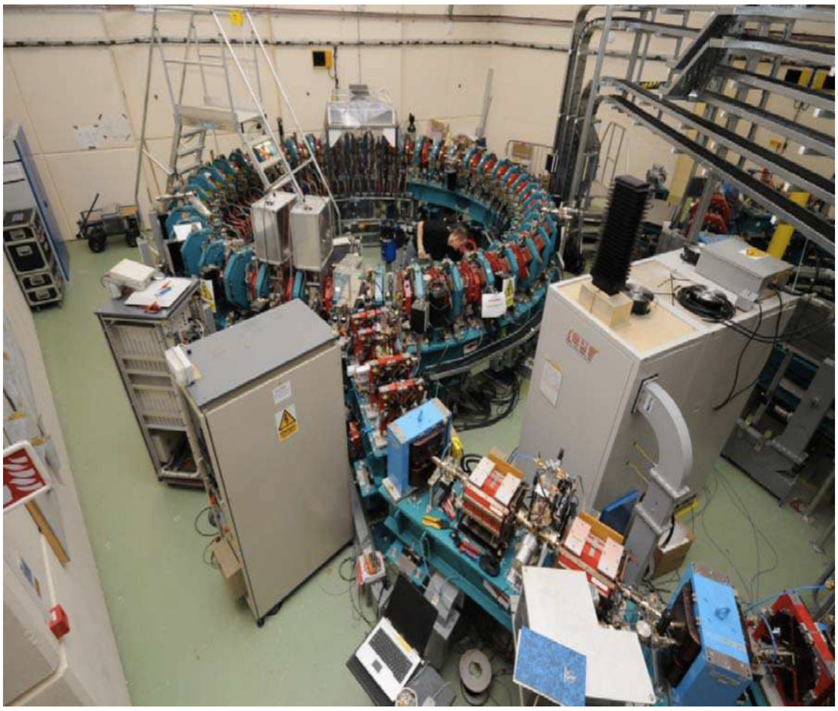

Figure 6.4: The EMMA experiment at the start of commissioning in June 2010. The injection line from ALICE is visible in the foreground, with the ring and extraction lines are shown in the background.

These equations of motion can not be solved for each particle individually because, due to space charge, the force on each particle depends on the position of all other particles. The vector $\mathbf{y}(t)$ is therefore introduced containing the six coordinates of all the particles as a function of time. The equations of motion in Eqs. 6.1 and 6.2 can then be rewritten as

$$\frac{d\mathbf{y}(t)}{dt} = \mathbf{f}(t, \mathbf{y}(t)) \ , \tag{6.3}$$

where $\mathbf{f}(t, \mathbf{y}(t))$ is a combination of Eqs. 6.1 and 6.2 for every particle. The only boundary conditions are the initial particle coordinates at a specified time.

The method utilised by `GPT` to solve Eq. 6.3 is known as fifth-order embedded Runge-Kutta with adaptive stepsize control. The algorithm is efficient and has the capability to specify the accuracy of the calculation. The algorithm advances in steps from $\mathbf{y}(t)$ to $\mathbf{y}(t+h)$, calculating the error between a fifth- and fourth-order formula, keeping the error to a specified degree - not too large so as to indicate an insufficient precision of the result but not too small so as to waste CPU time.

This software was chosen to complement other EMMA tracking codes e.g. `Zgoubi` [63], and analytical transfer matrix packages (e.g. `MAD` (Methodical Accelerator Design)) [28], and because of its inbuilt modelling of space-charge. In this chapter a simulation of EMMA is presented and the results compared to the experimental data. The results of this space-charge modelling will be discussed in Sec. 6.4.

6.3.1 Injection line

The energy recovery linac ALICE [64] serves as an injector for EMMA, providing electron bunches in the energy range of 10 to 20 MeV. The injection line used to transport the bunches from ALICE into EMMA can be seen in Fig. 6.5, where the key magnetic and diagnostic components are shown. The line consists of

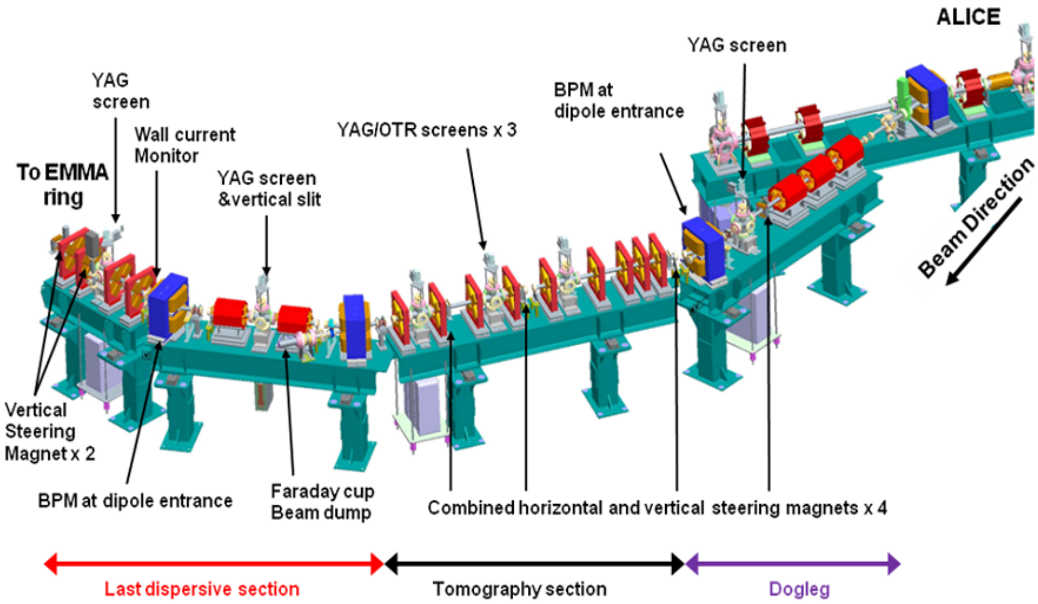

Figure 6.5: A schematic of the ALICE to EMMA injection line.

a symmetric 30° dogleg to extract the beam from ALICE. This is followed by a matching section of quadrupoles with a tomography section for transverse emittance measurements. The line concludes with a transport section to the injection point of the EMMA ring, including further beam diagnostic equipment.

As the injection line consists of well documented magnetic components (dipoles and quadrupoles) its modelling was used to validate the software package. All the magnetic elements were substantiated using built-in GPT elements, thus providing a simulation of all the magnetic fields, including the summed effect of fringe fields. The simulation starts with the two quadrupoles in ALICE upstream of the first dipole of the injection line up to its end (including an approximation of the EMMA injection septum).

The injection line was initially modelled using the matrix transfer method forming the basis of MAD. In this software the lattice is constructed using a hard-edge approximation, calculating the Twiss parameters defining the lattice. MAD differs from GPT in that it analytically tracks the Twiss parameters through the beamline using transfer matrices, whereas GPT numerically tracks the initial particle beam using custom magnetic elements. The MAD software is widely used and deemed highly reliable for modelling simple beam transport systems. As the beam transport is modelled using matrix transformations it does not take into account certain non-linear optics effects, which is why numerical computation is required for comparison. For the EMMA injection line, it is assumed that the lattice is simple enough for the results from the MAD simulation to be reliable and so be used as a robust benchmark against which the performance of GPT can be assessed.

As previously mentioned in Sec. 2.1.2, the beta function is analogous to the beam-size for a system of constant emittance. The beta functions for the injection line calculated by MAD are shown in Fig. 6.6.

The initial Twiss parameters were then extracted from MAD and used to substantiate a Gaussian distribution in GPT. The initial parameters can be found in Tab. 6.1. Using this data, and the theory outlined in Sec. 2.1.2, it is possible to initiate a Gaussian beam.

Once the beam has been initiated the magnetic lattice has to be modelled to observe the effects of propagation. It is a blueprint of Fig. 6.5, consisting of 20 quadrupoles and 5 dipoles, modelled using the elements and co-ordinate changes unique to GPT.

Twiss parameter	Initial value
β_x	13.1 m
β_y	12.7 m
α_x	1.5
α_y	-2.1
$\epsilon_{x,y}$	3.0 πmm mrad

Table 6.1: Initial twiss values for the EMMA injection line, extracted from MAD and inputted into GPT.

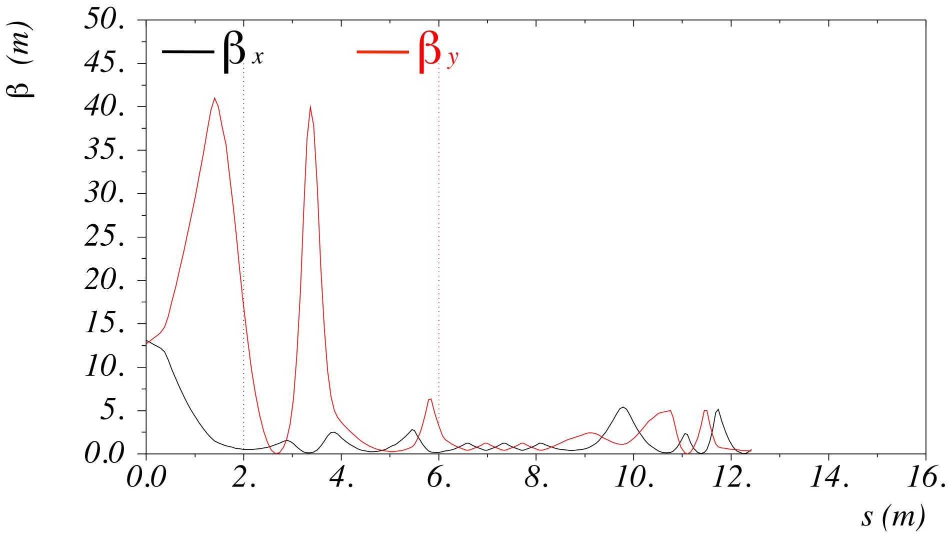

Figure 6.6: Beta functions, $\beta_{x,y}$, for the ALICE to EMMA injection line, as modelled by `MAD`.

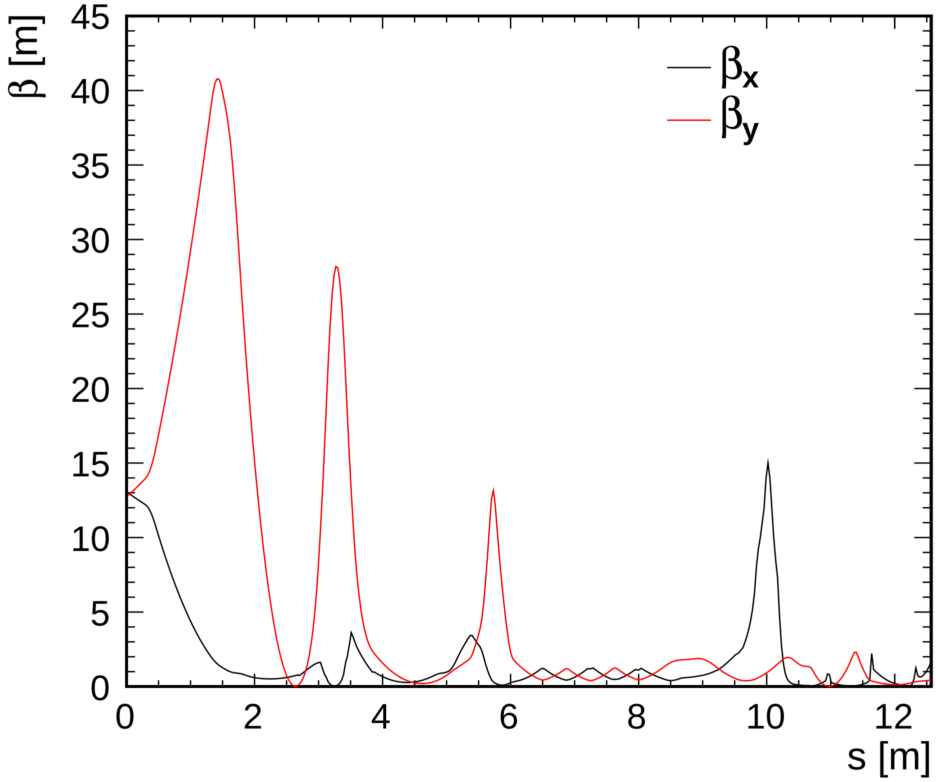

Figure 6.7: Beta functions, $\beta_{x,y}$, for the ALICE to EMMA injection line, as modelled by `GPT`.

The beam is then matched to certain constraints along the line: the tomography section requires the beta values to be equal (and alpha values to be equal and opposite) at the entrance as the quadrupole strengths are fixed and cyclical; the dispersion is required to approach zero at the end of the dogleg section and at the entrance to the ring due to large dipole fields experienced at these locations. As the initial beam parameters are fixed by the `MAD` simulation the quadrupole strengths are the only variables to be optimised. Fig. 6.7 shows the product of this optimisation, analogous to the output from `MAD` seen in Fig. 6.6.

A direct comparison between Figs. 6.6 and 6.7 indicates `GPT` provides an accurate simulation of the injection line. The beta functions differ by a factor of 3-4 before and after the tomography section but the constraints from matching previously outlined are met. These discrepancies between the beta functions are accounted for by the disparate nature of the quadrupole approximations.

6.3.2 EMMA ring

In accelerator design a reference trajectory is chosen, whereby an ideal particle travels at a fixed energy without dissipative forces. Magnets are positioned and aligned such that the particle ends at the exact point it started after one revolution of the machine: a closed orbit. If the orbit is not closed then the bunch will experience a kick every time it sees the same magnetic field, resulting in a first order resonance and leading to beam blow-up.

In synchrotrons the magnet strengths are scaled with beam energy to maintain a constant reference trajectory. By definition a non-scaling FFAG maintains a fixed field so the orbit will be displaced outwards with an increase in energy, meaning the orbit is not closed. It is possible, however, to study the trajectory of a particle in an FFAG without acceleration, allowing for different closed orbits at different energies.

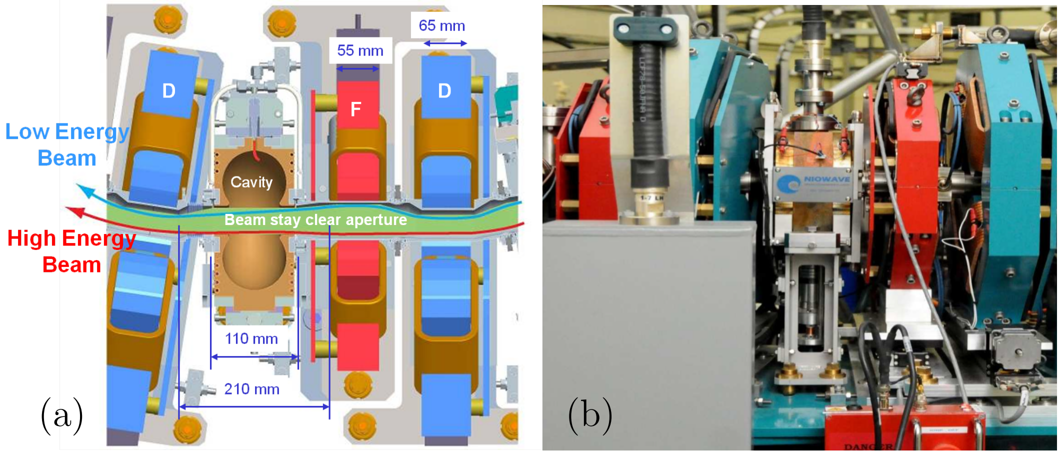

Figure 6.8: A a) schematic, and b) photo of a typical EMMA cell. The defocusing magnet (blue) is the first element in the cell and is larger than the focusing magnet (red) due to its higher magnetic field gradient.

The EMMA ring is composed of 42 identical cells, an example of which is depicted in Fig. 6.8, each containing a focusing and defocusing quadrupole. The lattice has a periodicity of $2\pi/42$ with the intrinsic beam dynamics defined by a single cell. As well as the ability to change the magnetic field (by altering the current) the quadrupoles are on rungs so they can be laterally shifted (see Sec. 2.2.1 for a more detailed explanation). The EMMA cells are therefore characterised by four variable parameters.

Single cell matching

The GPT simulation of the EMMA ring differs from that of the injection line in demanding a closed orbit (both in the cell and the ring) and requiring that this criteria takes precedence over all others. For this reason the quadrupole strengths and offsets were fixed and a field map of am already constructed cell was measured to remove the potential errors arising from quadrupole approximations. The fixed magnet parameters for the field map used in all GPT simulations are detailed in Tab. 6.2. As the field is held constant, the initial Twiss parameters, as well as the radial and angular offsets, were optimised to achieve a closed orbit.

The matched alpha and beta functions for one EMMA cell, as simulated by GPT, are displayed in Fig. 6.9. In order to match the beam, a bespoke GPT element was written requiring a zero difference between the Twiss values at the start and end of each cell, i.e. a closed orbit. As previously mentioned, if periodicity is achieved in one cell then the periodicity of the entire ring is also defined. This matching was achieved to within a self-imposed constrained accuracy of 1×10^{-4}. The matching simulations detailed in this section are for an injection energy of 15 MeV, in the middle of the energy range, and a longitudinal bunch length of 1 mm.

Multiple turns

The electron bunches injected into EMMA are expected to complete hundreds of turns of rapid acceleration before extraction. For this reason the validity of the beam dynamics needs to be reconfirmed at this range, making sure mismatches and resonances are avoided.

The simulation was set up to track 10,000 macro particles around 100 turns of the machine. The Twiss parameters were outputted at the end of each turn and, as for one cell, these values should be constant over each turn if the beam is matched correctly. Fig. 6.10 shows the beta functions and emittance for this range behaving with the properties of a matched beam.

	parameter	value
F magnet	gradient	6.68 T/m
	offset	7.51 mm
D magnet	gradient	-4.70 T/m
	offset	34.05 mm

Table 6.2: The EMMA baseline lattice parameters. The quadrupole magnetic gradients and displacements are shown for both the focusing and defocusing magnets.

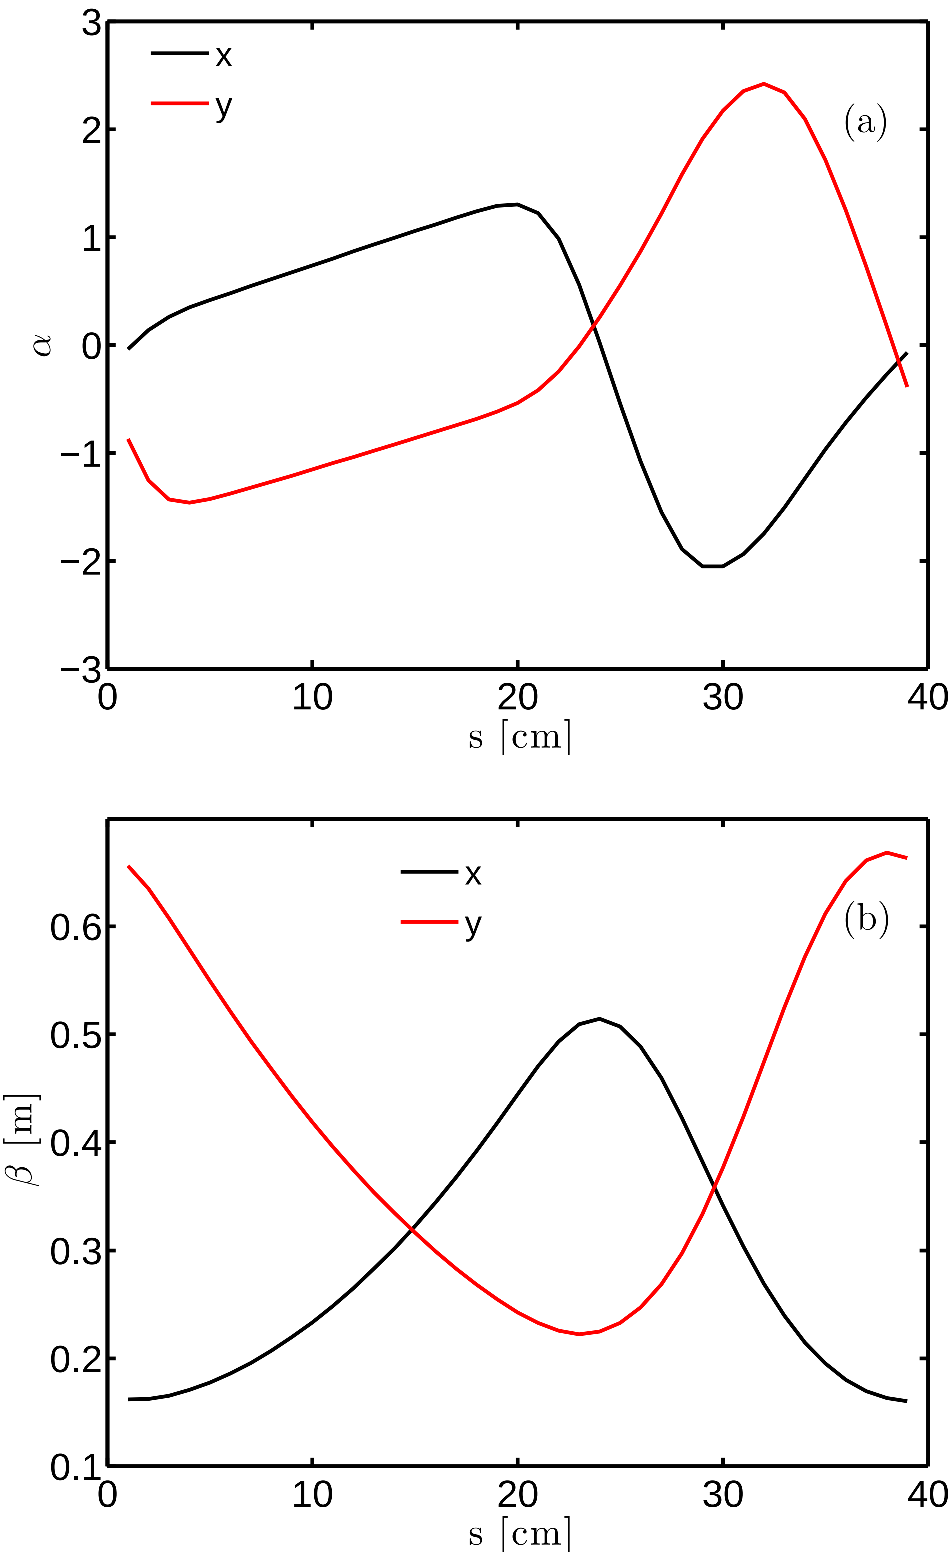

Figure 6.9: a) Alpha functions, $\alpha_{x,y}$, and b) beta functions, $\beta_{x,y}$, for an EMMA cell, as modelled by GPT.

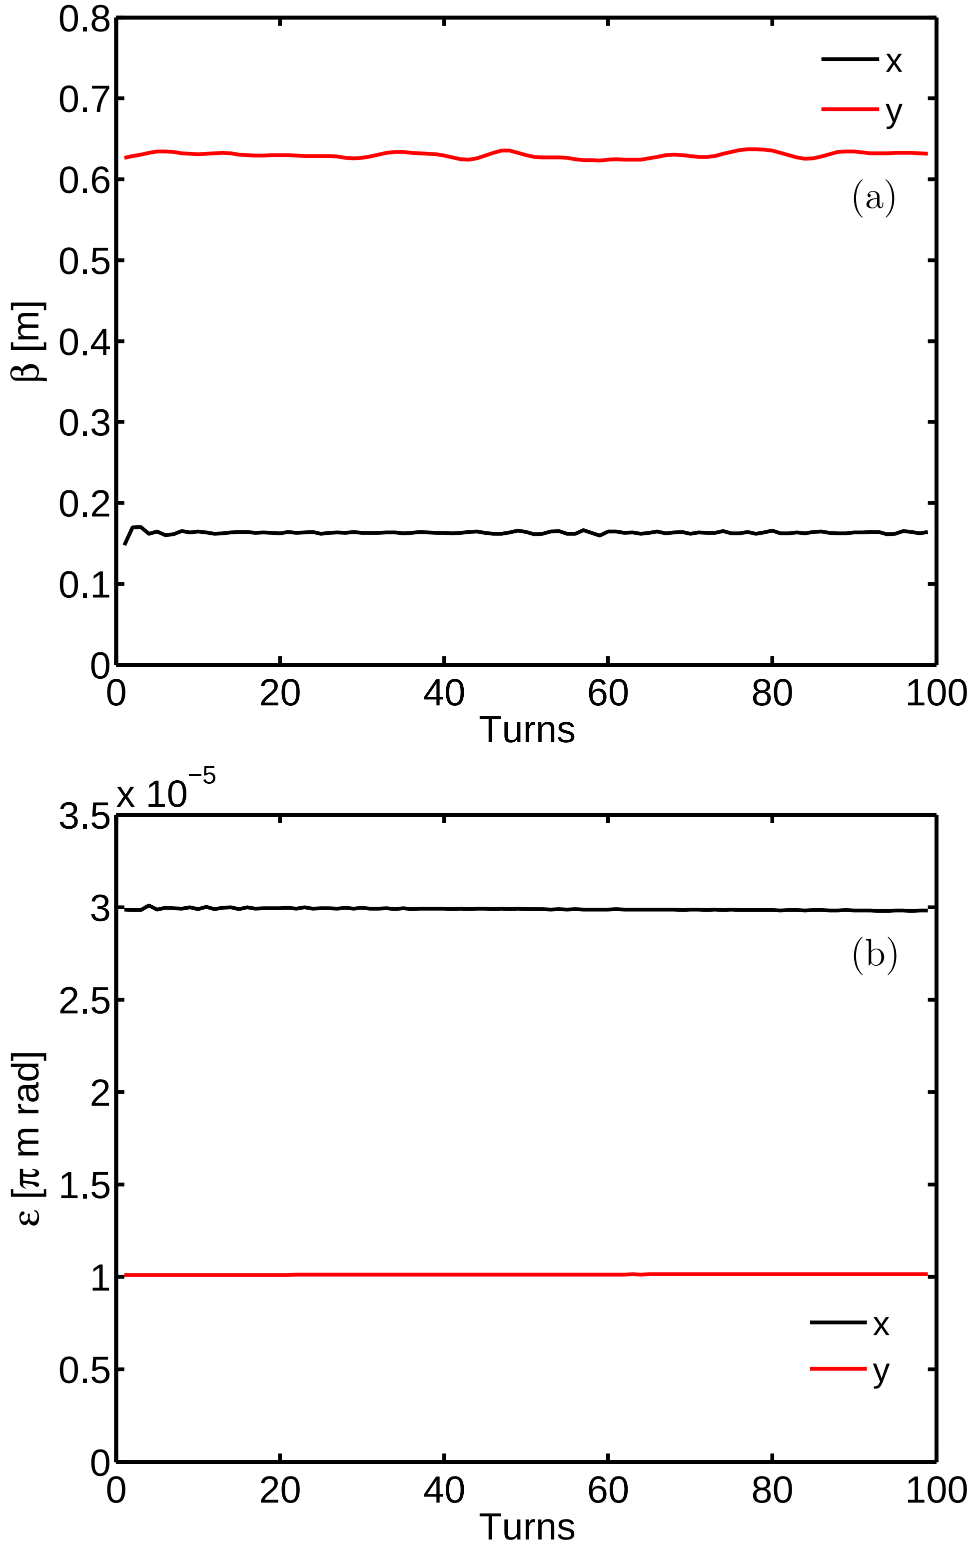

Figure 6.10: a) Beta functions, $\beta_{x,y}$, and b) transverse emittances, $\epsilon_{x,y}$, for 100 turns of the EMMA ring, as modelled by `GPT`.

The small fluctuations around the mean arise from symplectic errors in beam-matching i.e. numerical inaccuracies in integration algorithms that blow-up over many turns. Therefore the fluctuations observed are not dependent on statistics but an intrinsic artefact of the Jacobian matching operator.

Time of flight

In the case of EMMA, the electrons have a kinetic energy in the range of 10−20 MeV and as such experience little increase in velocity with a gain in energy. To overcome orbit radius scaling with energy (i.e. higher energy particles orbit with a larger radius, as seen in traditional FFAG machines like the proof-of-principle machine at KEK), the lattice was designed such that lower energy particles experience stronger magnetic fields than higher energy particles. The effect of this is evident in Fig. 6.11, illustrating that lower energy particles follow a larger bend and higher energy particles travel closer to an on-axis trajectory.

For this reason it is expected that particles with lower energies will take longer to complete one revolution of the machine, therefore having a longer time of flight (TOF). The orbital periods across the EMMA energy range can be seen in Fig. 6.12. The period follows a parabolic relation with a minimum at approximately 17 MeV. Fig. 6.12 also displays the TOF curve from a different code (ZGOUBI) with the same cell geometry, co-ordinate transforms, and baseline fieldmap [65]. These two curves demonstrate a similar trend, with an almost constant offset of 0.01 ns. This displacement is most likely due to nuances between the two codes' numerical solvers used to match the beam.

In EMMA the TOF is not constant, but varies within a small range. By adjusting the cavity frequency to the average TOF of this parabola the entire energy range can experience a positive accelerating voltage in the cavity. The TOF variation with energy depends on the magnet strengths responsible for the bend of the trajectories. It also depends on the transverse positions of the magnets

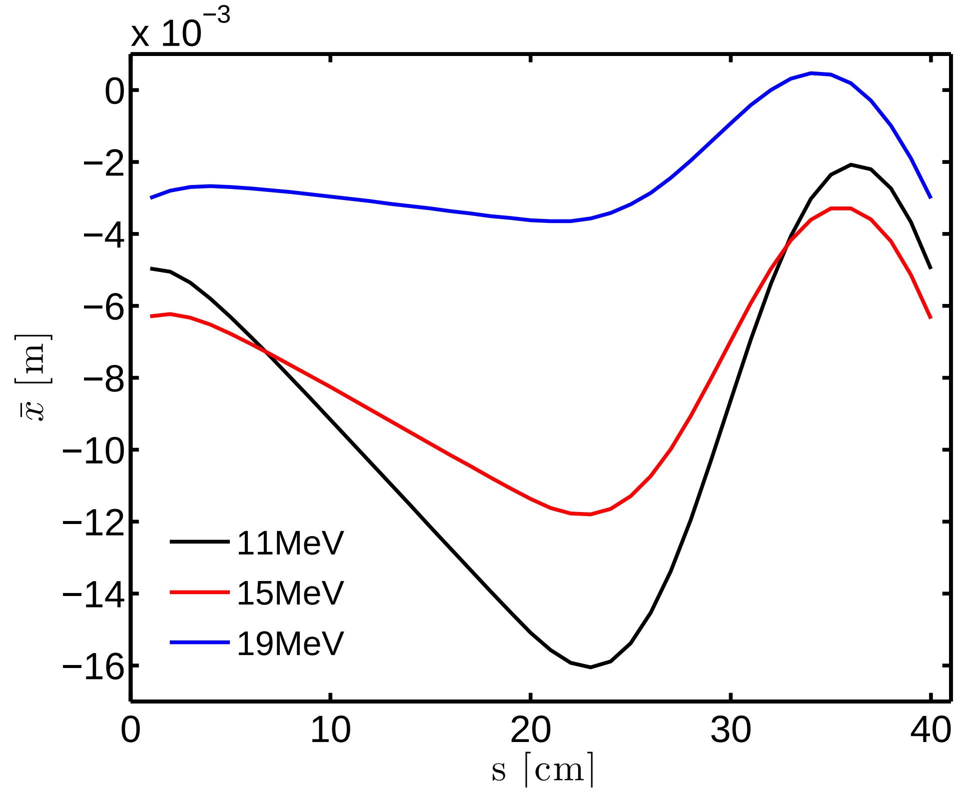

Figure 6.11: Mean horizontal position for three closed orbits of differing energy in one EMMA cell. The transverse co-ordinates at the entrance of the cell are equal to those at the exit after a rotation of $2\pi/42$, not represented in this figure.

since, by moving one magnet with respect to the other, the magnetic field in the cell is changed and the closed orbits are different.

After initial commissioning of the machine was complete, experimental runs were initiated to measure the beam position and orbital period as a function of energy. The TOF is measured using beam position monitors (BPMs), each composed of four electrodes located on the beam pipe. When a bunch passes through that section of the beampipe each electrode measures a change in potential, with a larger gradient indicating a closer beam. The TOF is measured as the time between signals observed by one electrode. To increase the precision the TOF is measured over multiple turns, with 10 turns being the default in this analysis.

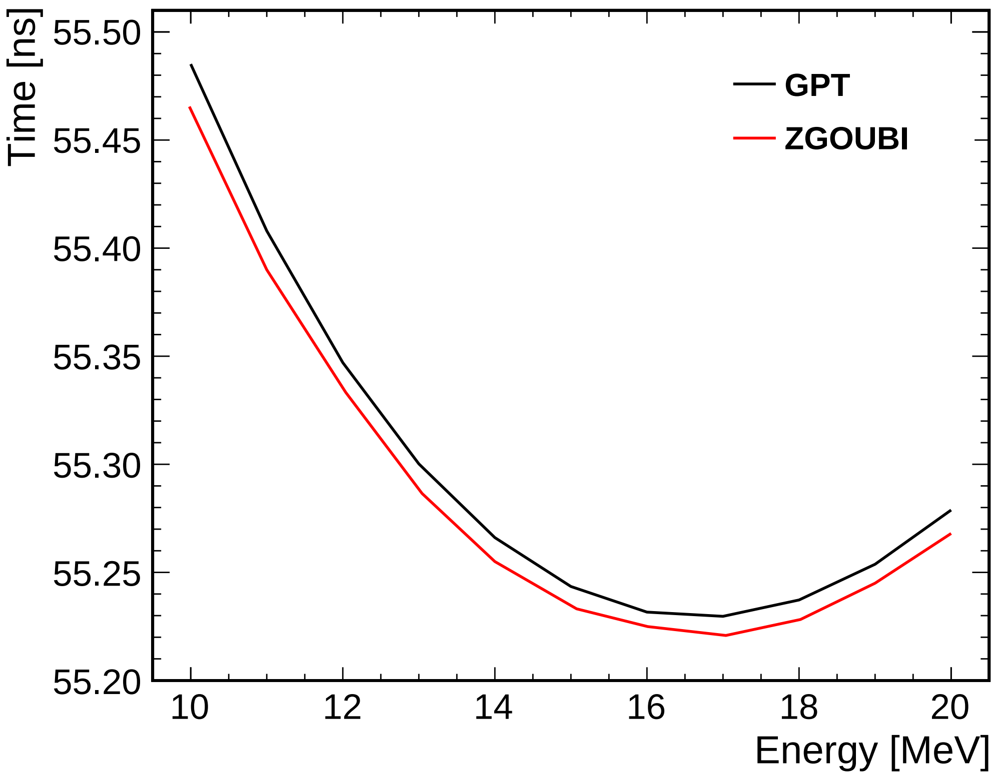

Figure 6.12: Orbital period of the EMMA ring as a function of energy, averaged over 10 turns, output at the end of each revolution.

Each TOF measurement is taken on an oscilloscope and then processed on a computer. Phenomena such as radio frequency jitter or variations in electron gun conditions may lead to a slightly different nominal momentum in each case. To mitigate these potential errors several scope signals are saved for each TOF, with an average and standard deviation calculated for each energy.

The energy recovery linac ALICE used to inject into EMMA is a well-documented, flexible machine with achievable energies up to approximately 27.5 MeV. ALICE could therefore inject into EMMA across the full $10-20$ MeV energy range. However, in terms of technical reproducibility it was decided that ALICE would inject with a fixed energy of 12 MeV, with EMMA creating an equivalent momentum range by changing the quadrupole strengths.

Using the beam dynamics relation that a decrease in magnetic strength is equivalent to an increase in beam momentum, it is possible to characterise the EMMA lattice over an equivalent momentum range. In application this equivalence is only valid if the ratio of the quadrupole strengths is kept constant, with equivalent momenta achieved by scaling the coil currents up or down according to the fixed ratio. For example, if the current coils are reduced by 20% then the equivalent momentum increases by 20%. The baseline lattice configuration for an optimised injection energy of 12 MeV is predicated on nominal magnet currents of 350 A for the defocusing magnet and 320 A for the focusing magnet. All equivalent momenta are therefore scaled according to the ratio of these values, $r = 1.094$.

Fig. 6.13 shows the experimental data for a range of equivalent momenta, where the data has also been fitted with a second order polynomial. The error in the orbital period is the standard deviation of eight measurements at that momenta. It is slightly misleading to plot these two curves concurrently: the GPT simulation is created using a fixed fieldmap with a baseline lattice configuration but a range of injection energies; the experimental data is derived from creating an equivalent momentum range by altering the magnet strengths. It is therefore necessary to recreate the experimental data by implementing an equivalent momentum range in the GPT simulation.

In order to scale the quadrupole strengths the fieldmap has to be removed and GPT magnetic elements reimplemented, analogous to the method described in Sec. 6.3.1. The geometry of the cell remained the same except for an additional 2 mm increase in the quadrupole offsets. This was performed in the experimental runs to make it easier to inject at low equivalent momenta. Using this refined geometry, the equivalent momenta method, and giving the beam a bunch-charge of 40 pC as recorded experimentally, a data set was created. This is shown in Fig. 6.14 along with the previous simulated lattice.

The beam was initiated and matched to the baseline lattice design at each scaling of the quadrupole strengths. The baseline parameters were originally

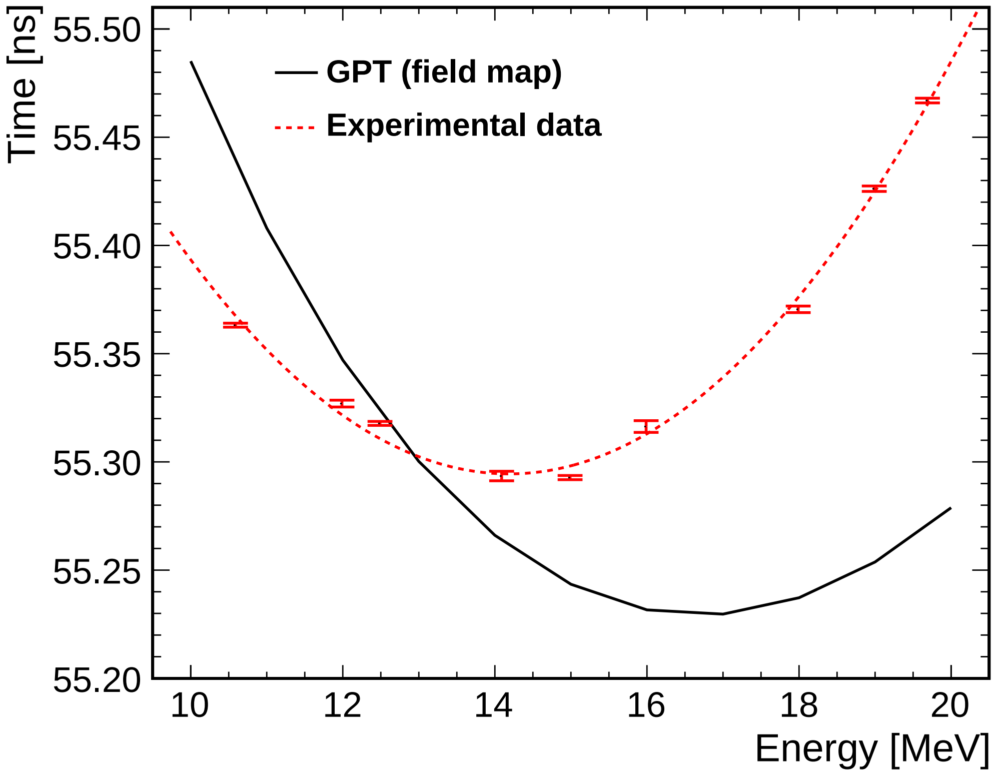

Figure 6.13: Comparison between measurements and the GPT field map simulation of the orbital period variation with energy. A second order polynomial fit is also shown.

optimised using a number of different simulation packages (described in further detail in Sec. 6.3.2) such that the GPT simulated TOF curve was symmetrical around its minimum at 15 MeV. As can be seen in Fig. 6.14 the output demonstrates this property. However, for this reason the parabola shows discrepancies at the upper extreme of the effective momentum range. The experimental lattice has a minimum TOF of approx. 55.30 ns (14 MeV), whereas the equivalent simulated lattice has a minimum TOF of approx. 55.27 ns (15 MeV). Despite the offset at higher energies, the shape of these two curves is similar after their respective minima. It is concluded that this discrepancy arises from errors in the actual lattice that are not accounted for in the simulation: e.g. inconsistencies in individual magnet strengths; misalignment errors across the different cells; etc.

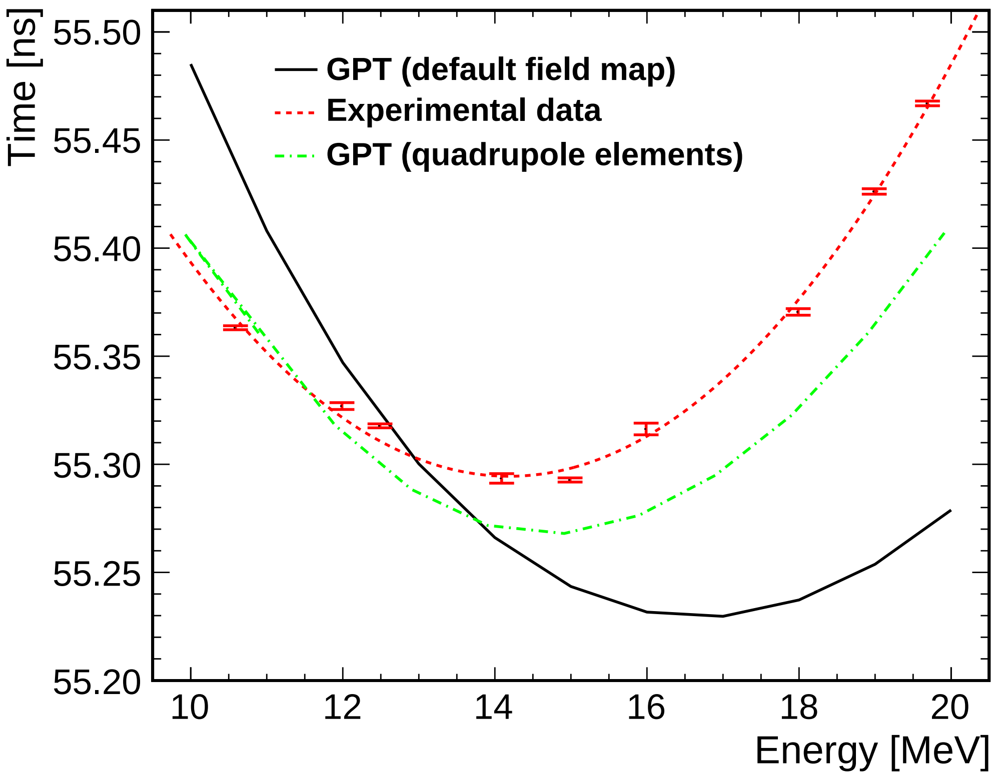

Figure 6.14: Comparison between measurements and simulations of the orbital period variation with energy. As well as the simulation employing the baseline field map a simulation with the `GPT` quadrupole elements, scaled to produce an equivalent momentum range, is also plotted.

The simulation was further optimised by adjusting the four lattice variables (quadrupole strengths and offsets). It remained necessary to maintain a constant ratio between the quadrupole strengths to produce an equivalent momentum range, but the ratio itself was optimised. The baseline offsets were scaled by the same distance as in the experimental lattice i.e. an increase of 2 mm. The optimisation was performed using the `GPT` numerical solver, with the closest approximation of the experimental lattice produced by the ratio $r = 1.108 \pm 0.007$ and 2 ± 0.02 mm offset (back towards the central beam axis), with the error range derived from the variation in results after 100 simulations. The results of this matching are shown in Fig. 6.15.

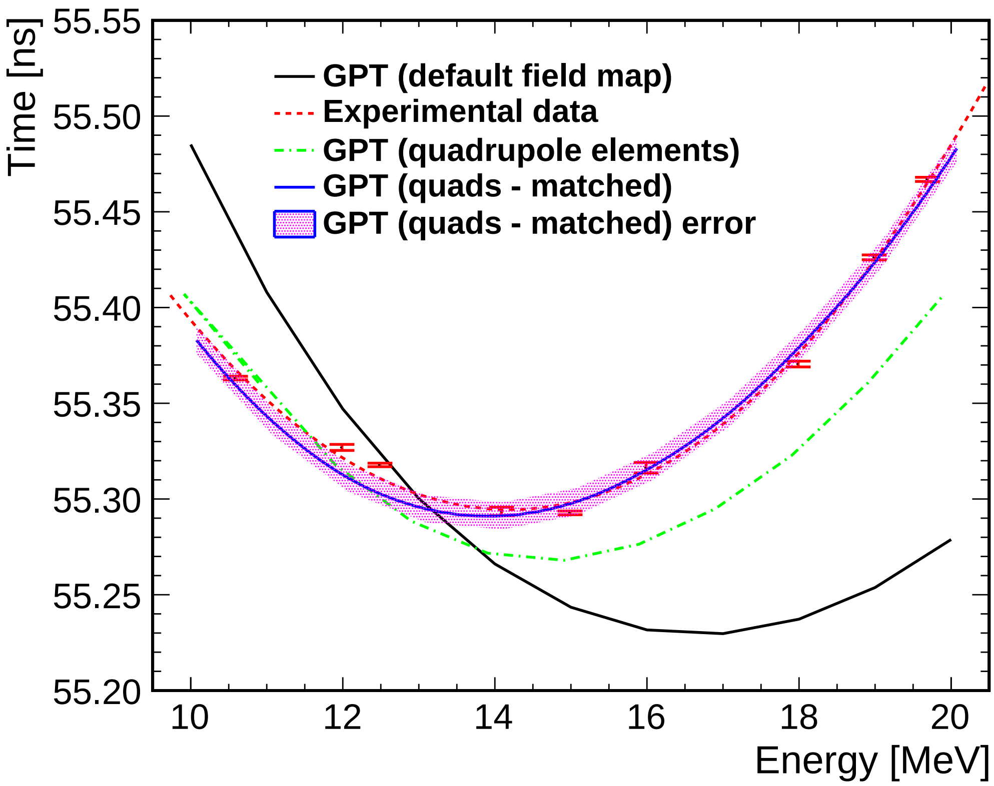

Figure 6.15: Comparison between measurements and simulations of the orbital period variation with energy. The additional curve is produced using the `GPT` quadrupole definitions, with the relative magnetic strength and quadrupole offsets optimised to mimic the experimental data.

The numerical solver in `GPT` requires an accuracy for each variable to be matched. This accuracy affords the solver a tolerance range for faster calculation but gives the values an intrinsic uncertainty. The accuracy required for the four quadrupole variables is one per cent. These errors are non-correlated so, as such, are combined in quadrature. The error range of the matching for the TOF curve is shown by the shaded region in Fig. 6.15. As can be seen, the simulated lattice agrees with the experimental data to within the error range arising from matching.

In addition to the uncertainties in the simulation from the `GPT` solver, the experimental data itself also has a variance from a number of factors with hysteresis in the magnets being the main contribution. There is an associated 0.7% error in the magnet strength due to hysteresis. If this is taken into account then the experimental scaling ratio becomes $r = 1.108 \pm 0.015$. The GPT solution lies just within this range.

Tune calculations

As previously described, the EMMA lattice is a repeating structure of 42 identical magnetic cells and is thus intrinsically periodic; the beam encounters the same magnet structure once every full revolution, therefore experiencing periodically repeating forces. For this reason the focusing function $k(s)$ is periodic with the circumference of the ring, resulting in the same solution to Hill's equation as in Eq. 2.11,

$$x(s) = \sqrt{\epsilon_x \beta_x(s)} \cos\left(\psi_x(s) + \phi_x\right) \ . \tag{6.4}$$

In this solution however the beta function $\beta(s)$ is also periodic, with the same periodicity as $k(s)$. The resonance arising from this repeating lattice is dependent on the betatron phase,

$$\Delta\psi = \psi(s+L) - \psi(s) \ , \tag{6.5}$$

integrated over one complete revolution. The tune is thus defined as

$$\nu_{x,y} = \frac{\Delta\psi_{x,y}}{2\pi} = \frac{1}{2\pi} \oint \frac{ds}{\beta_{x,y}(s)} \ , \tag{6.6}$$

describing the number of betatron oscillations a particle undergoes per revolution [23].

The tune may be approximated using a 'smoothing' method by taking the average beta value, $\bar{\beta}$, and utilising the relation

$$\bar{\beta} = \frac{R}{\nu} \, . \tag{6.7}$$

Eq. 6.7 is very accurate for most accelerating rings but ns-FFAGs have an average off-axis horizontal beam position so it is not a robust enough method for EMMA tune calculations.

If the position of the beam is known during revolution - either experimentally using beam position monitors or, in this case, using virtual screens in the simulation placed at the end of each magnetic cell - then a fast fourier transform (FFT) may be performed. An FFT decomposes a series of values into components of different frequencies. FFTs are a type of discrete fourier transform which reduce the computation time from N^2 down to $N \log N$ by approximating the integral as a summation, either in frequency or time. The algorithm used for the FFT calculation is that included in `MATLAB` [66]. In this case the FFT dissolves the average beam position at the end of each magnetic cell (outputted by `GPT`) into the tune.

Fig. 6.16 shows an example output of the FFT function, demonstrating clear central resonant tunes in both planes. The plots also show two small adjacent resonances. These resonances define the synchrotron tune ν_s and are positioned at $\nu_{x,y} \pm \nu_s$. For a 15 MeV electron beam $\nu_s = 0.0234$. The method used to create the plots in Fig. 6.16 is then repeated across the $10 - 20$ MeV EMMA energy range, with the initial Twiss parameters rematched at each energy. The result of these calculations can be seen in Fig. 6.17. These results were produced using the initial baseline field map model described in Sec. 6.3.2.

Tune calculations have been produced using a number of different codes, all modelling EMMA with the baseline parameters: `MAD` (with displaced quadrupoles), `BERG` [67], and `Zgoubi` [68]. The tune calculation results for these different codes

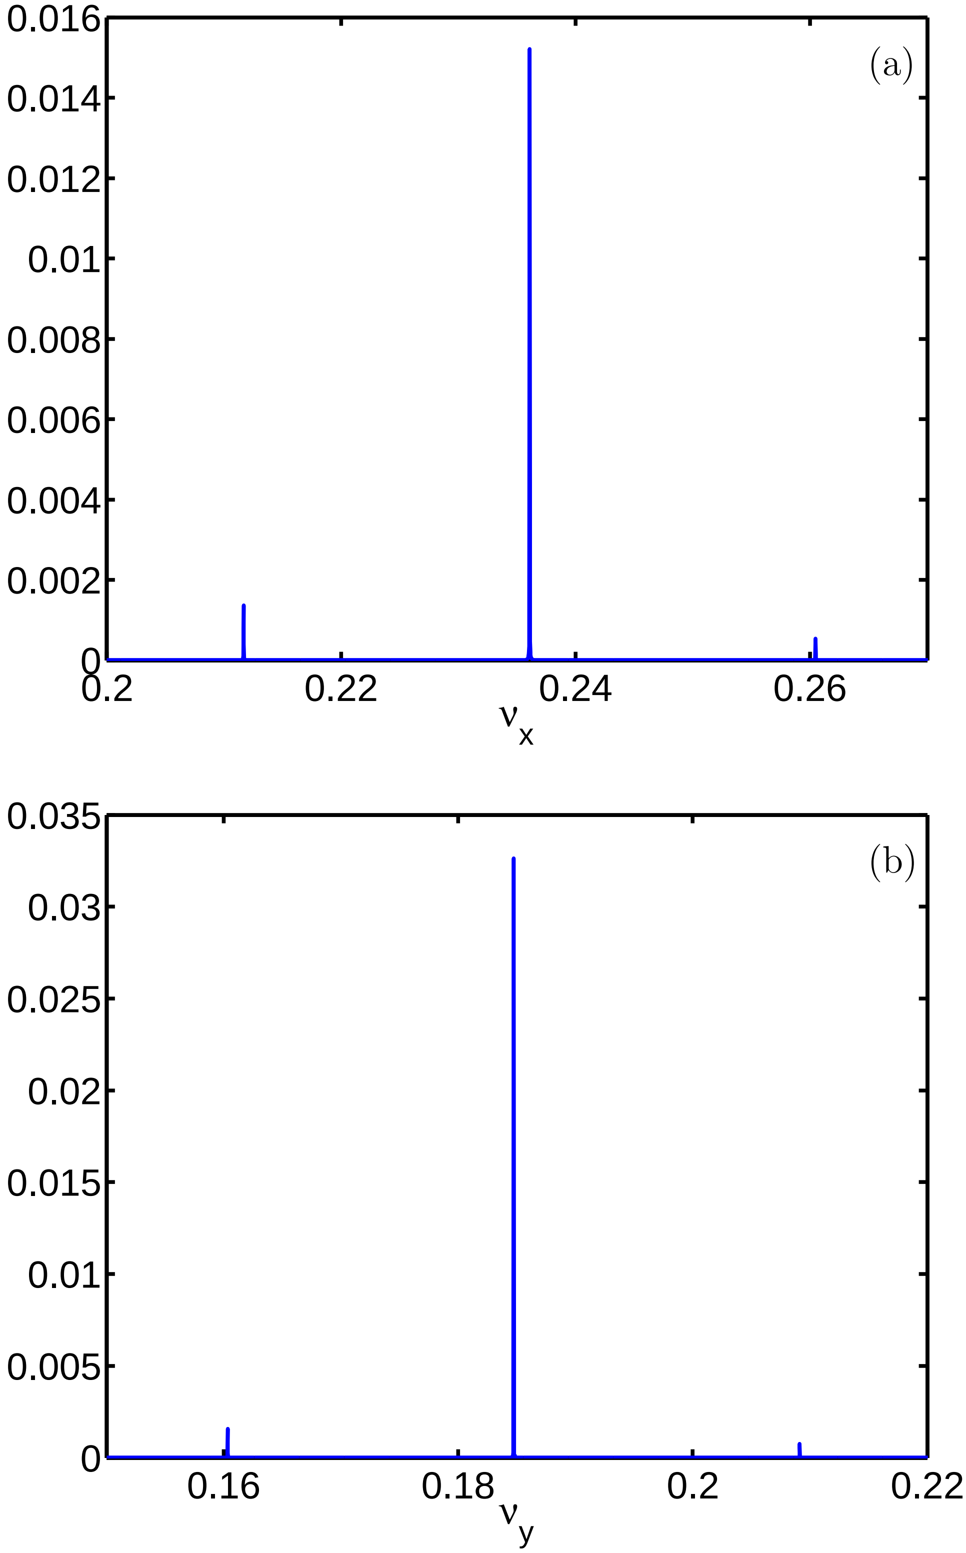

Figure 6.16: Cell tune plots for both the a) horizontal, and b) vertical transverse planes of a single reference particle taken across 100 turns of the ring. The beam is injected with an initial energy of 15 MeV. The y-axis is arbitrary in each case.

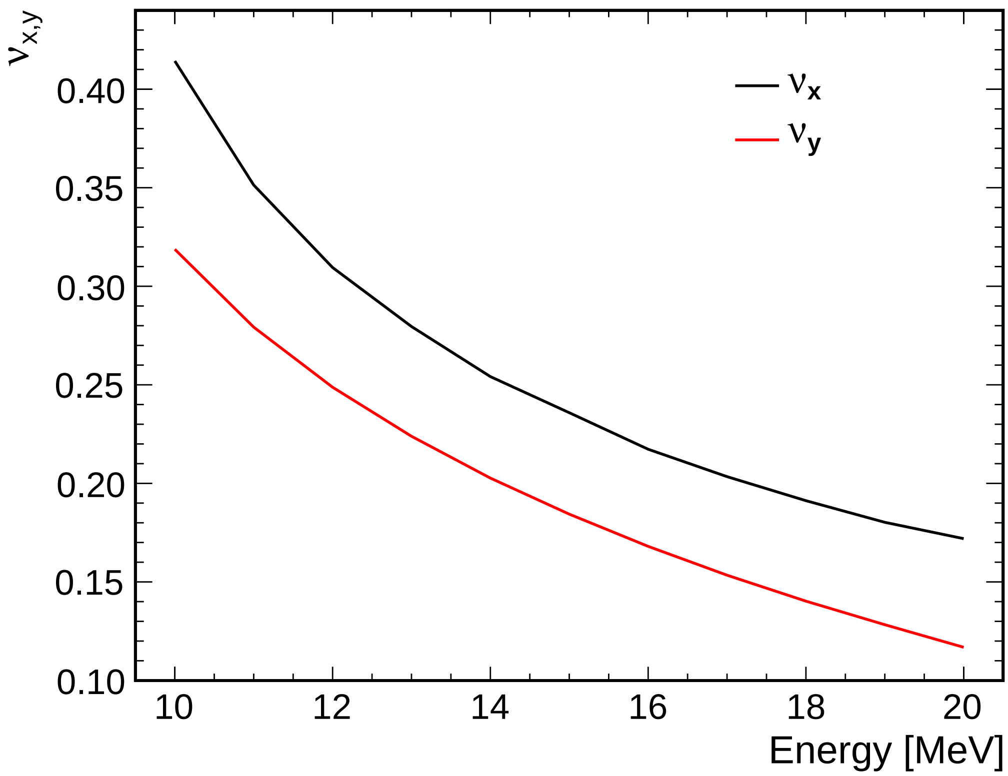

Figure 6.17: Horizontal and vertical cell tunes across the EMMA injection energy range, calculated from 100 turns of the machine.

can be seen in Fig. 6.18. All the models demonstrate good agreement, including the GPT simulation: the horizontal tune sits slightly higher than the other codes but shows the same shape; the vertical tune lies in the middle of the other codes but is slightly skewed at lower energies.

Fig. 6.19 shows the tune equivalent plot of TOF from Fig. 6.13, comparing the simulated and experimental lattice. As with Fig. 6.13 it is misleading to plot them on the same axes as they represent different lattice setups. However, as with the TOF results, they are broadly in agreement.

The GPT simulation was then set up to include the quadrupole elements. The matched quadrupole ratio $r = 1.108$ and 2 mm offsets derived in Sec. 6.3.2 were included in the simulation. The tune across the EMMA energy range was then

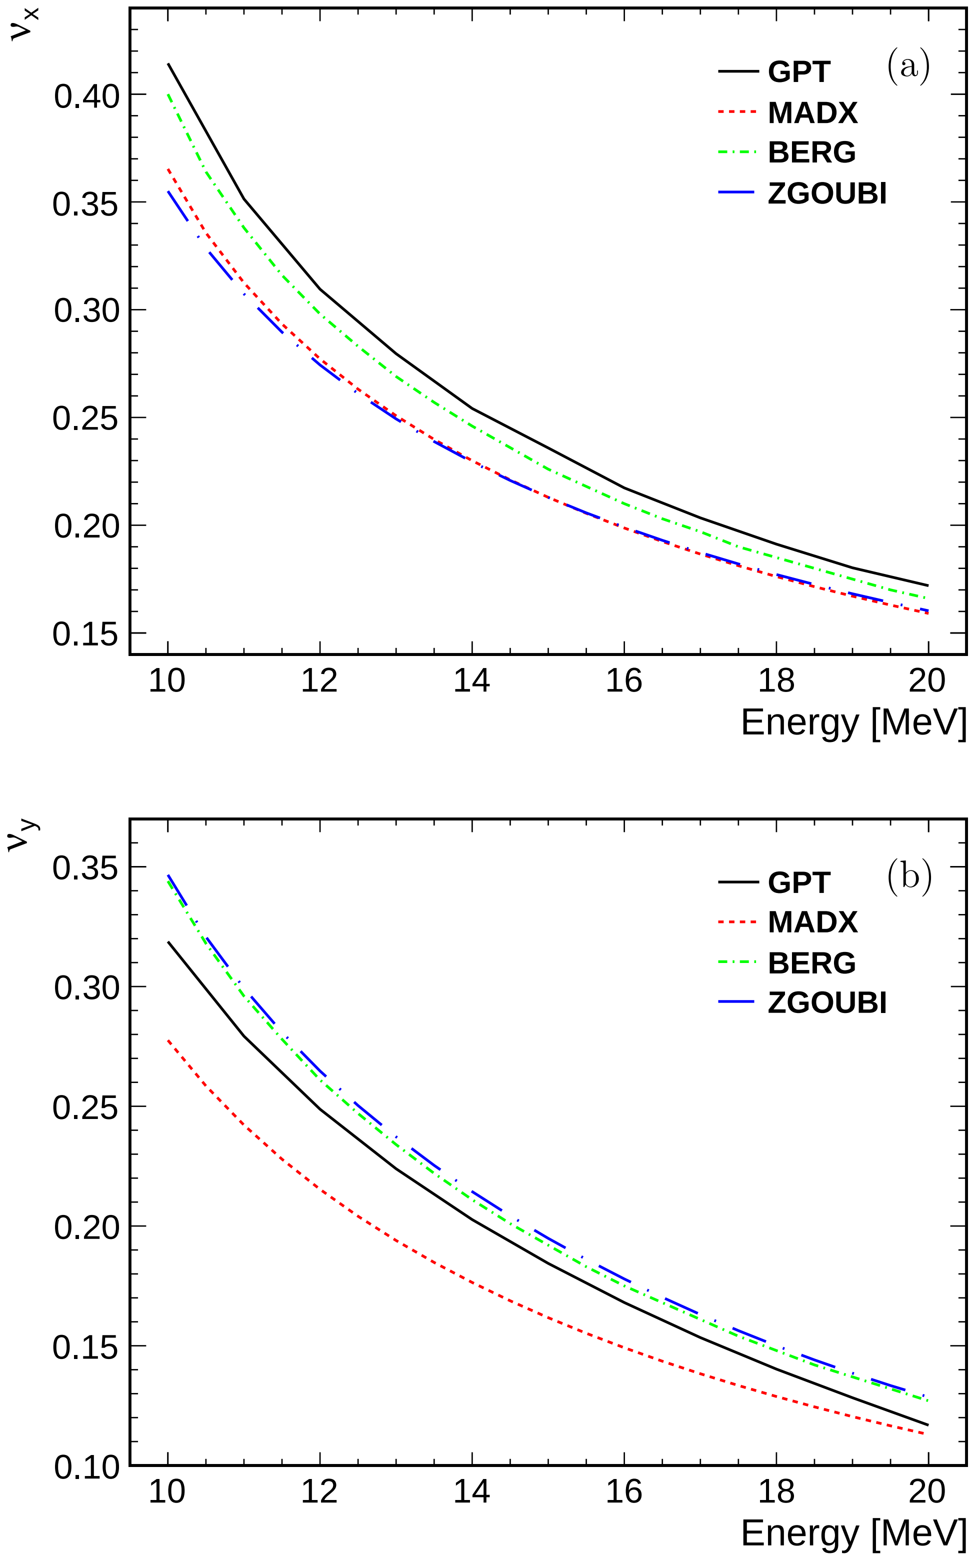

Figure 6.18: Comparison of the a) horizontal, and b) vertical tune produced by different codes modelling the EMMA ring with baseline parameters.

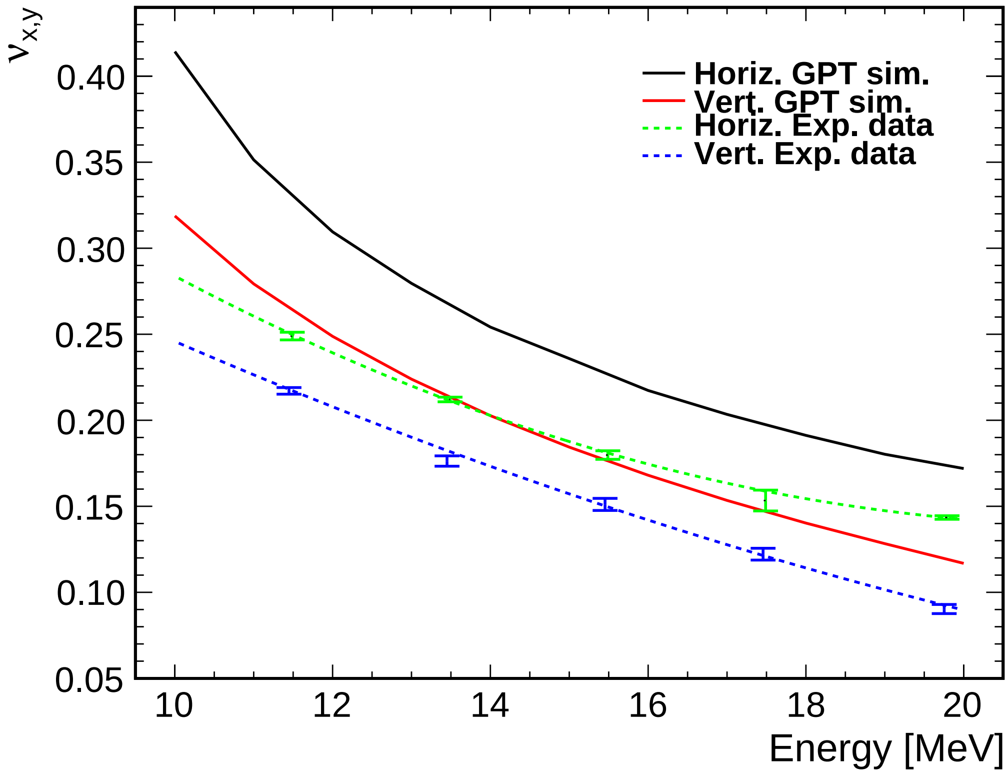

Figure 6.19: Comparison between the tune variation of a cell with energy for the experimental and simulated lattice. The simulation uses the measured field map of the EMMA baseline parameters.

calculated using the method previously described. The results are shown in Fig. 6.20.

Again the simulation with the optimised cell variables demonstrates good general agreement with the experimental data. The simulated lattice shows a slight misalignment at the extremities of the energy range in both transverse planes. However, this misalignment lies within the allowed error range of the simulated lattice.

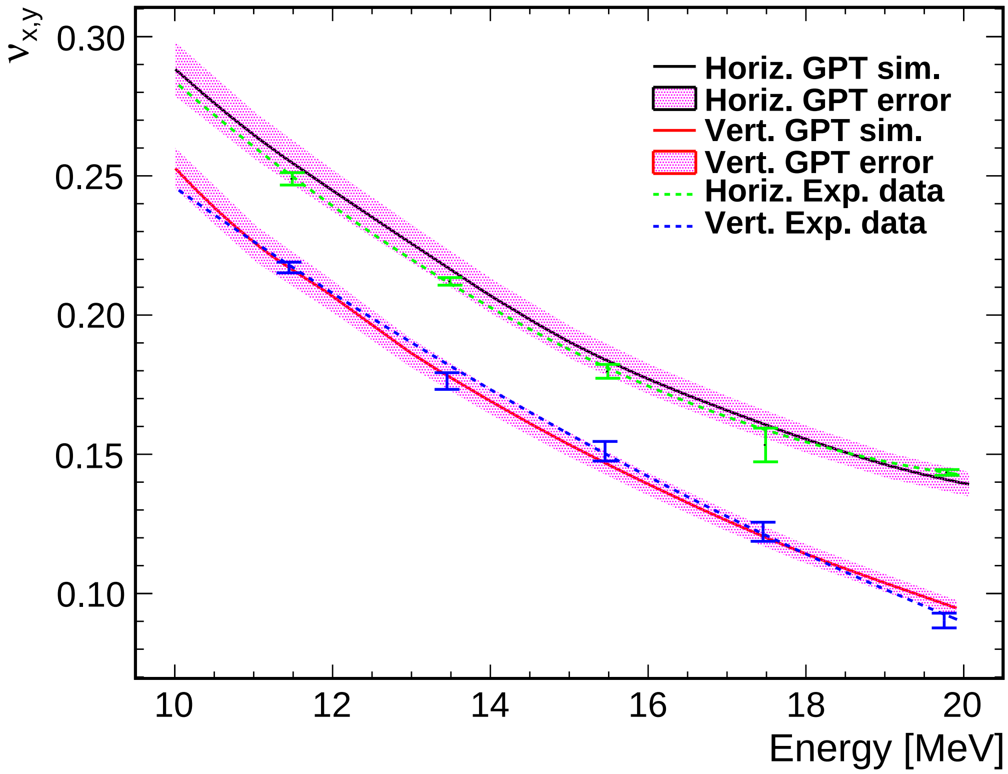

Figure 6.20: Comparison between the cell tune's variation with energy for the experimental and simulated lattice. The simulation employs the GPT quadrupole elements with the matched scaled strengths and offsets optimised in Sec. 6.3.2. The equivalent momentum method is used to create an effective energy range.

6.4 Space-charge calculations

The Coulomb forces between the charged particles of a high-intensity beam in an accelerator create a self-field which acts on the particles inside the beam like a distributed lens, defocusing in both transverse planes. A beam moving at a certain velocity is accompanied by a magnetic field which partially cancels the electrostatic defocusing effect, with complete cancellation at the speed of light. The effect of this direct space-charge alters the number of betatron oscillations per machine turn, ν, by $\Delta\nu$. The space-charge limit of a synchrotron (approximately

$\Delta\nu = 0.5$) may be overcome by increasing its injection energy, thus raising γ to a level where space-charge losses are no longer dominant.

The concept of self fields is most easily represented by two particles of equal charge which, at rest, repel one another due to the Coulomb force. However, when travelling with speed $v = \beta c$, they represent two parallel currents which attract one another by the effect of their magnetic fields. The net effect of these currents is still repulsive but decreases with speed, and thus energy.

Accelerators typically bunch together billions of particles with a simplified circular cross section shown in Fig. 6.21. The Coulomb force serves to push the highlighted reference particle away from the centre of the beam, with the defocusing force increasing towards the edge. This force is also present in a moving beam, however the magnetic force vector focuses the reference particle towards the beam centre.

These radial forces are described by,

$$F_r = e(E_r - v_s B_\phi) \ . \tag{6.8}$$

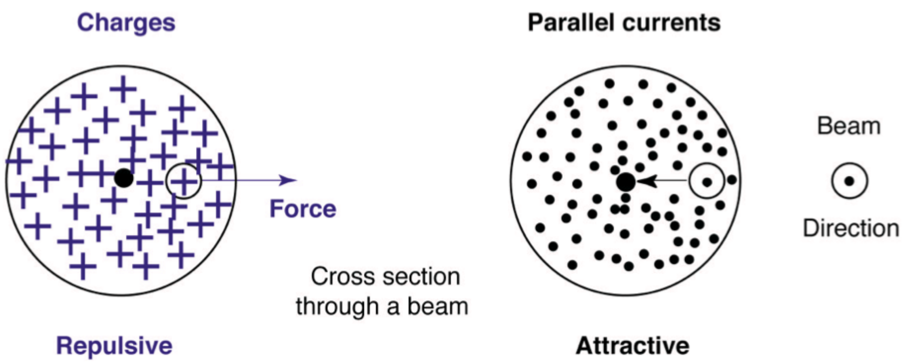

Figure 6.21: The electrostatic and magnetic forces on a reference particle, within a transverse bunched beam.

Inserting the definitions for radial electric field, E_r, and azimuthal magnetic field, B_ϕ, for a uniformly charged cylinder leads to

$$F_r = \frac{eI}{2\pi\epsilon_0\beta c}(1-\beta^2)\frac{r}{a^2} \ , \tag{6.9}$$

where a is the beam envelope size, r the particle radius within the envelope, and I is the beam current. These may be expressed in transverse co-ordinates, such that

$$F_x = \frac{eI}{2\pi\epsilon_0\beta\gamma^2 ca^2}x \ , \quad F_y = \frac{eI}{2\pi\epsilon_0\beta\gamma^2 ca^2}y \ . \tag{6.10}$$

Hill's equation for a FODO cell, defined in Eq. 2.8, may be extended to include the perturbative effects of these forces, denoted by k_{SC},

$$x''(s) + (k(s) + k_{SC}(s))x(s) = 0 \ . \tag{6.11}$$

k_{SC} may be derived by expressing $x''(s)$ in terms of the force F_x:

$$x'' = \frac{1}{\beta^2 c^2}\frac{F_x}{m_0\gamma} = \frac{2r_0 I}{ea^2\beta^3\gamma^3 c}x \ , \tag{6.12}$$

where r_0 is the classical particle radius, 2.82×10^{-15} m for electrons. Hill's equation with space-charge then becomes

$$x''(s) + \left(k(s) - \frac{2r_0 I}{ea^2\beta^3\gamma^3 c}\right)x(s) = 0 \ , \tag{6.13}$$

with a negative space-charge term, reducing the net focusing effect of the magnetic lattice [69]. The incoherent tune-shift may then be calculated by integrating the weighted gradient errors around the circumference of the ring $2\pi R$:

$$\Delta\nu_x = \frac{1}{4\pi}\int_0^{2\pi R} k_{SC}(s)\beta_x(s)\mathrm{d}s \ . \tag{6.14}$$

Taking the definition of k_{SC} from Eq. 6.13 yields

$$\Delta\nu_x = -\frac{1}{4\pi}\int_0^{2\pi R} \frac{2r_0 I}{e\beta^3\gamma^3 c}\frac{\beta_x(s)}{r^2}\mathrm{d}s = \frac{r_0 R I}{e\beta^3\gamma^3 c}\left\langle \frac{\beta_x(s)}{a^2}\right\rangle \ . \tag{6.15}$$

The term $\langle \beta_x(s)/a^2(s)\rangle$ is the inverse of the beam emittance, defined in Eq. 2.13, and thus an invariant. Replacing the beam current I by $Ne\beta c/4\pi R$ (for a Gaussian beam), where N is the number of particles in the bunch, leads to the tune shift in both transverse planes from space-charge induced self-fields,

$$\Delta\nu_{x,y} = -\frac{r_0 N}{4\pi\epsilon_{x,y}\beta^2\gamma^3} \ . \tag{6.16}$$

This form of the tune shift is for a constant, unbroken Gaussian beam filling the entire beamline. In reality the beam is separated into discrete bunches, where space-charge effects are more profound due to an increase in bunch density. The tune is therefore multiplied by a bunching factor, leading to the form

$$\Delta\nu_{x,y} = -\frac{r_0 N}{4\pi\epsilon_{Nx,y}\beta\gamma^2}\frac{C}{\sqrt{2\pi}\sigma_{\mathrm{rms}}} \ , \tag{6.17}$$

where C is the circumference of the ring and σ_{rms} the RMS of the bunch.

`GPT` has a number of built in space-charge routines, differing in rigour and therefore computation time. The `spacecharge3Dmesh` routine is not the most comprehensive calculation of self-fields but it provides robust results in a realistic timescale [70]. In this model stochastic effects are correctly assumed to be negligible so a non-equidistant mesh can be used in conjunction with solutions of Poisson's equation in the co-moving frame.

6.4.1 Injection line

Quadrupole matching of the ALICE to EMMA injection line was demonstrated in Sec. 6.3.1, with the beta functions outputted from GPT shown in Fig. 6.7. A study of the beam-size variation has been performed to quantify the effects of space-charge. Fig. 6.22 shows this equivalent beam-size. It also shows the effect on beam-size by space-charge-generated self fields within the range of $0 - 30$ pC.

Fig. 6.23 zooms into the σ_x plot of the tomography section from Fig. 6.22. This demonstrates the effect of space-charge on beam-size in more detail. As described in Sec. 6.4, space-charge has a defocusing effect on the bunch, which can be seen by the initial increase in beam-size over the first few quadrupoles of the tomography section. The effect on beam-size increases with an increase in bunch-charge. As the bunch initially grows in size it pushes particles further off-axis where they experience a stronger magnetic effect of the focusing quadrupole. This then over-focuses the bunch leading to a mismatch at the end of the section. One solution to this would be to rematch the line for each different injected bunch-charge by making modest quadrupole adjustments.

6.4.2 EMMA ring

The parameters of any circular magnetic lattice are defined by its tune and TOF. The total tune shift arising from space-charge self fields is defined in Eq. 6.17. As this relation contains parameters that may not be inputted directly into GPT i.e. N, β, and γ, it requires a slight modification to accommodate tangible input parameters for simulation. Using the substitution for β and γ in terms of energy,

$$\beta^2\gamma^3 = \gamma\left(\gamma^2 - 1\right) = \frac{E}{m_o c^2}\left(\left(\frac{E}{m_o c^2}\right)^2 - 1\right) \ , \tag{6.18}$$

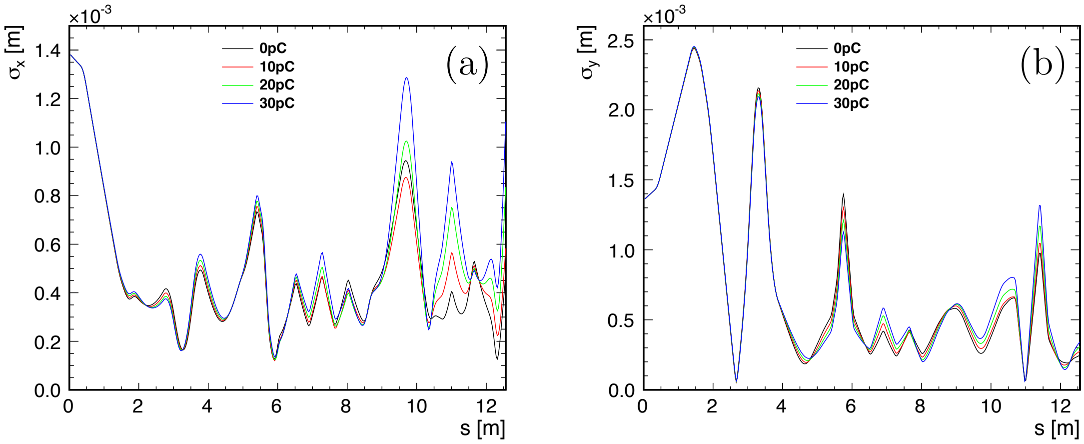

Figure 6.22: Beam-size in the a) x- and b) y-plane for the ALICE to EMMA injection line, detailing a range of bunch-charges with a spread of 0 to 30 pC.

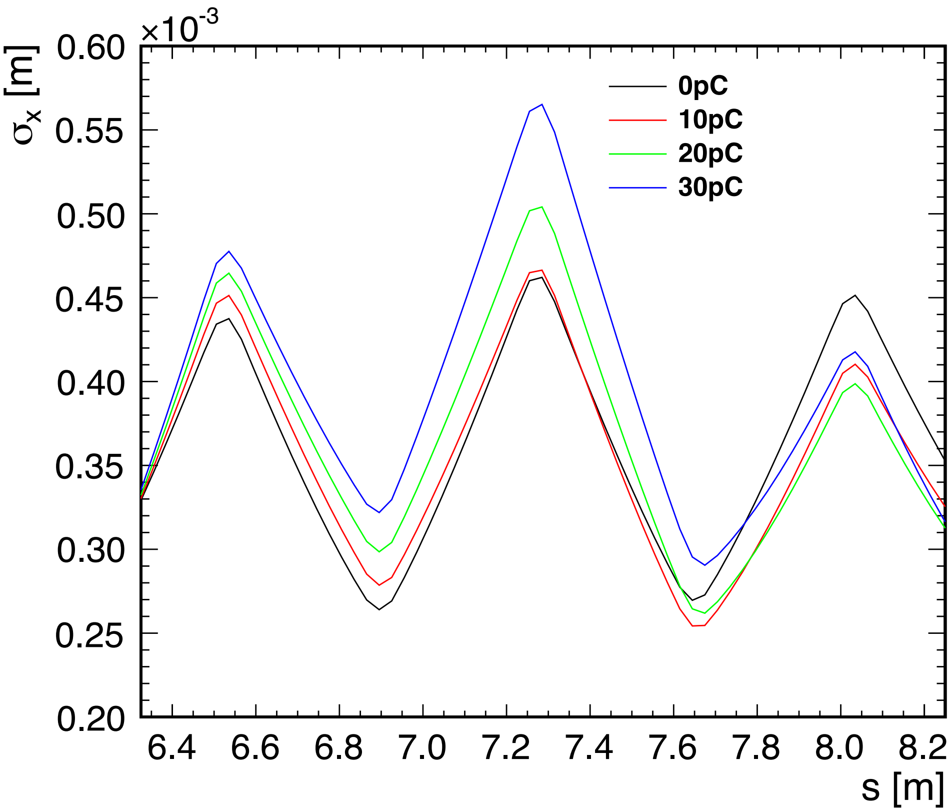

Figure 6.23: Beam-size in the x-plane, demonstrating the effect of a range of bunch-charges on the tomography section of the ALICE to EMMA injection line.

and calculating the number of particles by dividing the total bunch-charge by electron charge,

$$N = \frac{Q_{\text{tot}}}{q_{\text{e}}} \,, \tag{6.19}$$

the tune shift can be calculated using the following equation, where all the variables are well defined injection parameters of the beam,

$$\Delta\nu_{x,y} = -\frac{r_0}{2\pi\epsilon_{x,y}\left[\frac{E}{m_o c^2}\left(\left(\frac{E}{m_o c^2}\right)^2 - 1\right)\right]} \frac{Q_{\text{tot}}}{q_{\text{e}}} \frac{C}{\sqrt{2\pi}\sigma_{\text{rms}}} \,. \tag{6.20}$$

This analytical form can then be used to calculate tune shifts independent of the tune value at that energy. The shift for an 80 pC bunch at 15 MeV represents an approximate 3% shift from the nominal value calculated by GPT at this energy for zero bunch-charge. A 3% shift is more than that which is typically tolerable in a high intensity proton synchrotron or storage ring and, although the beam stays in EMMA for only 10 to 20 turns, resonance crossings excited by space-charge are a concern. Fig. 6.24 shows the relation between the GPT results at 0 pC and an arbitrary space-charge regime of 80 pC, across the full EMMA injection energy region.

Lorentz factor γ	20.5
Number of particles N	5×10^8 (80 pC)
Normalised emittance $\epsilon_{Nx,y}$	10 π mm mrad
Bunch length σ_{rms}	1.0 mm
Circumference C	16 m
Electron radius r_0	2.82×10^{-15} m

Table 6.3: Bunch parameters for a 15 MeV beam, used to calculate an expected tune shift arising from space-charge at the EMMA bunch-charge regime.

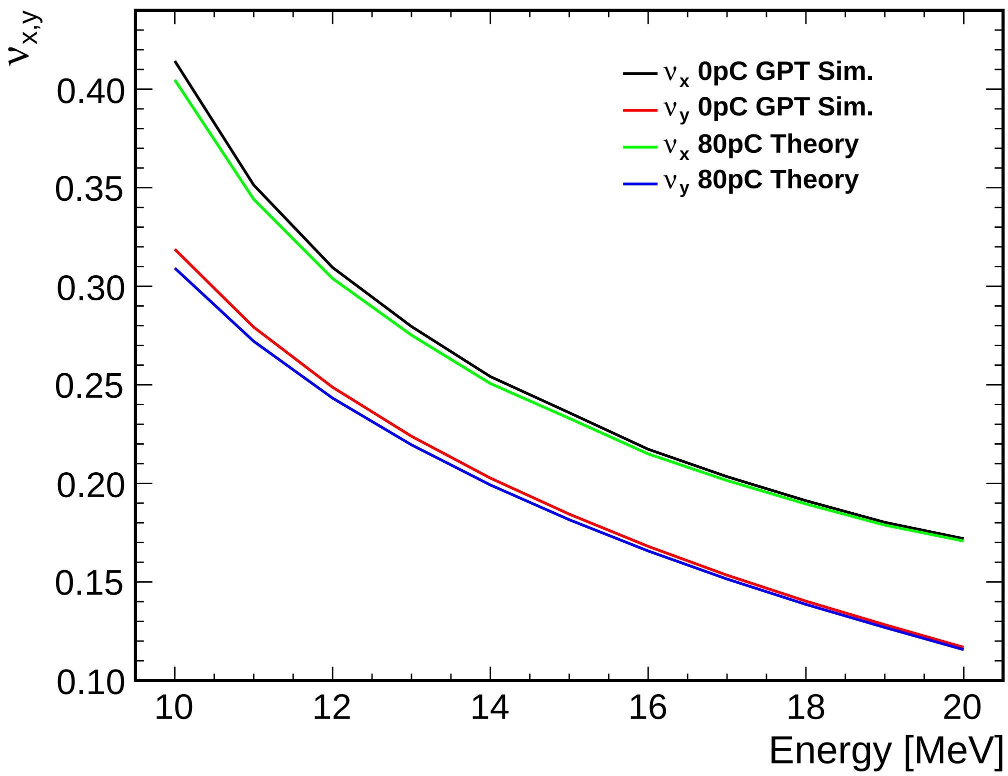

Figure 6.24: Theoretical tune shift from an 80 pC bunch-charge, compared with the tune calculated from GPT simulations with 0 pC.

The FFT routine used to calculate the tune from GPT data sets has a resolution limit inversely proportional to the number of outputs. In the EMMA simulation there is a 'screen' output at the end of each cell, giving 42 outputs per turn. One hundred turns of the machine was simulated for each calculation of the tune, affording 4200 outputs. The beam energy was then fixed to 15 MeV and the bunch-charge scaled up from zero until a tune-shift was observed. Fig. 6.25 shows the expected linear decrease in tune shift with an increase in bunch-charge. The first data point has no associated uncertainty as it is the GPT output for 0 pC, with all theoretical tune shifts stated relative to this value.

The data generated by GPT's space-charge procedure agrees with the theoretical predictions to within the resolution limit of the FFT routine, thus confirming the validity of the `spacecharge3Dmesh` routine.

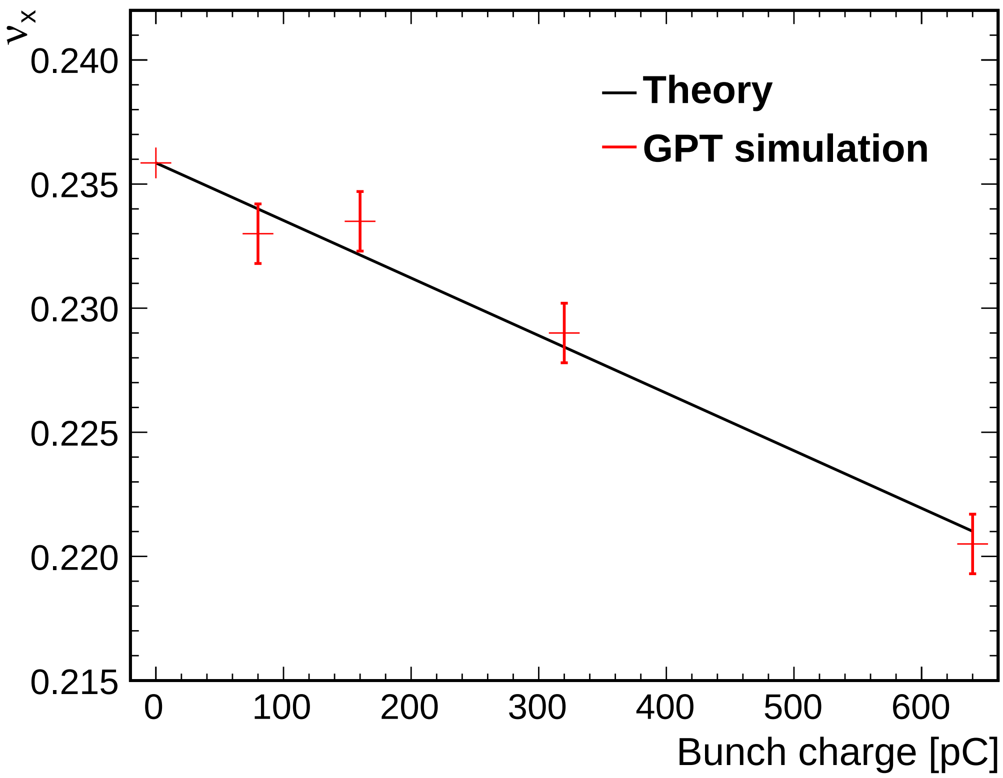

Figure 6.25: A demonstration of tune shift caused by space-charge over a range of bunch-charges (within and beyond the EMMA regime), with a fixed beam energy of 15 MeV. The theoretical data is taken with respected to the 0 pC simulation as performed by GPT, with no uncertainty on this data point.

6.5 Conclusions

This chapter serves to demonstrate the functionality of the next generation non-scaling FFAG accelerator, EMMA. Tune and parabolic TOF measurements have been recorded, with the results as expected of a ns-FFAG accelerator.

As well as experimental data the machine has also been uniquely modelled in the particle tracking code GPT. Due to the novel nature of the machine this was a necessary course of action in order to benchmark the code against other more established simulation packages and ultimately experimental data. The injection line was initially used to validate the code, concurring with the hard edged models of MAD.

The code was successfully adapted to accommodate repeating lattices with results showing positive correlations with the other pre-existing software packages of EMMA. These results required optimisation when held up against experimental data in order to compensate for physical effects unaccounted for by GPT i.e. hysteresis effects in the magnets.

The ultimate goal of the code was to simulate space-charge effects within the EMMA energy and bunch-charge regime; a piece of work not studied prior to this analysis. The validity of the GPT spacecharge3Dmesh routine was ultimately vindicated through demonstration of a negative transverse tune shift, in agreement with theory.

The validity of this proof-of-principle machine suggests a number of wide-ranging future uses both within the particle physics community (muon colliders, neutrino factories, etc.) and beyond (nuclear waste transmutation, proton therapy, etc.).

Chapter 7

Conclusions

Charged lepton flavour violation (cLFV) experiments represent a promising avenue of exploration, complementing those presently being pursued at the LHC. With the inclusion of beyond Standard Model (BSM) physics such as super-symmetry, rates of cLFV processes are dramatically increased to the point where observation may be achieved with near-future experiments. The process of muon-to-electron conversion presents the most promising route to substantially increasing the current experimental limits on cLFV, as well as resolving degeneracy between different BSM models.

The COMET experiment proposes a number of experimental improvements in order to drastically increase current experimental limits on the $\mu^- + N(A, Z) \rightarrow e^- + N(A, Z)$ process. The design of the experiment has been simulated and validated using a combination of the `MARS` and `GEANT4` software packages. Particular emphasis has been placed on optimising the experiment such that the number of stopped muons per proton is maximised, whilst minimising signal energy loss. A number of geometries have been tested, with the optimal disk-based stopping target configuration for the full COMET experiment detailed in Tab. 4.1. The implementation of the 15 cm radius titanium disks gives a 90% confidence level upper limit on the branching ratio (BR[Punzi]) of 5×10^{-16}, calculated using the

Punzi sensitivity routine detailed in Sec. 3.1.2. This is an improvement of $\mathcal{O}(1000)$ on the current results for this process presented by the SINDRUM-II experiment.

As well as collimator optimisation, similar studies have been performed for Phase-I of the COMET experiment, identical to the full design up to the end of the first 90° bend. In this case a shift in the nature of the detector system has resulted in a similar redesign of the muon stopping target. As such a double-cone design, analogous to that employed in SINDRUM-II, has been modelled. The simulation of this conical design has demonstrated an improvement on the disk design optimised for Phase-II, with a cleaner more distinct signal peak and a BR[Punzi] of 1.8×10^{-14} : 100 times better than the current experimental limit of SINDRUM-II. This calculation demonstrated the most rigorous determination of the potential reach of the COMET experiment to date. These figures were thus instrumental in securing funding in early 2013 for the second phase of the COMET experiment.

Initial results from the MuSIC experiment - a staged prototype experiment used to test the preliminary technical and experimental ideas of COMET - have been presented. A simulation framework was developed to aid the design of a DAQ system for hit count measurements. This package demonstrated a strong correlation for both 1D and 2D hit count measurements, with $\chi^2/\mathrm{ndf} = 1.6$ and 1.7 for the two runs respectively. The data sets taken during these two experimental runs, showing strong statistical similarities to simulation, were key initial steps towards proof of the predicted particle capture rates of Lobashev's superconducting solenoid design of the 1980s.

The EMMA experiment demonstrates the functionality of the next generation non-scaling FFAG accelerator, EMMA. Tune and parabolic time-of-flight measurements have been recorded, with the results as expected of a ns-FFAG accelerator. These measurements have been compared to a novel simulation of the ring, modelled using the particle tracking software `GPT`. Minor discrepancies were shown between simulation and data, with the simulation successfully adapted to compensate for physical effects unaccounted for by `GPT`. The code was chosen for its inherent ability to model space-charge effects, with the validity of the `spacecharge3Dmesh` routine upheld.

The vindication of this proof-of-principle machine opens many doors to further applications of the machine. One such application touched on is the proposal to use a non-scaling FFAG for phase rotation of a muon beam for future cLFV experiments. The integration of which, into an extension of the COMET experiment, PRISM, has already been explored. This liaison of these two projects proposes to improve on the current limits for cLFV by $\mathcal{O}(10^6)$.

Bibliography

[1] D. Griffiths, *Introduction to Elementary Particles*, Wiley (1987).

[2] S. Weinberg, *The Quantum Theory of Fields*, Cambridge Univeristy Press (2005).

[3] S. Abdus, *Symmetry concepts in modern physics*, Atomic Energy Centre (1966).

[4] R. Feynman, *Space-time approach to non-relativistic quantum mechanics*, Rev. Mod. Phys. **20**, 367 (1948).

[5] Y. Fukuda *et al.*, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81**, 1562-1567 (1998).

[6] J. Thomas and P. Vahle, *Neutrino Oscillations: Present Status and Future Plans*, World Scientific (2008).

[7] A. Halprin *et al.*, *Double-beta Decay and a Massive Majorana Neutrino*, Phys. Rev. D**13**, 2567 (1976).

[8] G. Aad *et al.*, *Combined search for the Standard Model Higgs boson in pp collisions at $\sqrt{s}$ =7 TeV with the ATLAS detector*, Phys. Rev. D**86**, 032003 (2012).

[9] V. Azabov *et al.*, *Search for neutral Higgs bosons in the multi-b-jet topology in 5.2 fb^{-1} of $p\bar{p}$ collisions at $\sqrt{s}$ =1.96 TeV*, Phys. Lett. B**698**, 97-104 (2011).

[10] T. Aaltonen *et al.*, *Search for a Low-Mass Standard Model Higgs Boson in the $\tau\tau$ Decay Channel in $p\bar{p}$ Collisions at $\sqrt{s}$ =1.96 TeV*, Phys. Rev. Lett. **108**,

181804 (2012)).

[11] B. Campbell and D. Maybury, *Triviality and the (supersymmetric) see-saw*, JHEP **0704**, 077 (2007).

[12] A. de Gouvea *et al.*, *Lepton-Flavour Violation in Supersymmetric Models with Trilinear R-parity Violation*, arXiv:hep-ph/0008085v1 (2000).

[13] B. Dutta *et al.*, *Supersymmetry Signals of Supercritical String Cosmology at the Large Hadron Collider*, Phys. Rev. D**79**, 055002 (2009).

[14] H. Brown *et. al.*, *Improved Measurement of the Positive Muon Anomalous Magnetic Moment*, Phys. Rev. D**62**, 091101 (2000).

[15] H. Hatanaka *et al.*, *The Gauge hierarchy problem and higher dimensional gauge theories*, Mod. Phys. Lett. A**13**, 2601-2612 (1998).

[16] J. Adam et al., *New limit on the lepton-flavour violating decay $\mu^+ \to e^+\gamma$*, Phys. Rev. Lett. **107**, 171801 (2011).

[17] E. Hincks and B. Pontecorvo, *Search for Gamma-Radiation in the 2.2 microsecond Meson Decay Process*, Phys. Rev. **73**, 257258 (1948).

[18] A. van der Schaaf *et al.*, *SINDRUM-II*, J. Phys. G: Nucl. Part. Phys. **29**, 1503 (2003).

[19] F. Deppisch *et al.*, *Heavy Neutrinos and Lepton Flavour Violation in Left-Right Symmetric Models at the LHC*, arXiv:1206.0256 (2012).

[20] A. de Gouvea, *(Charged) Lepton Flavor Violation*, Nuclear Physics B**188** (Proc. Suppl.), 303308 (2009).

[21] Y. Kuno *et al.*, *Background suppression for $\mu \to e\gamma$ with polarized muons*, Phys. Rev. D**55**, 5 (1997).

[22] T. Suzuki, *et al.*, *Total nuclear capture rates for negative muons*, Phys. Rev. C**35**, 26 (1987).

[23] K. Wille, *The Physics of Particle Accelerators*, Oxford University Press (2000).

[24] K. Urwin and F. Arscott, *Theory of the Whittaker-Hill Equation*, Proc. Roy. Soc. Edinburgh **69**, 28-44 (1970).

[25] D. Kerst, *The Acceleration of Electrons by Magnetic Induction*, Phys. Rev. **60**, 4753 (1941).

[26] J. Liouville, *Note on the Theory of the Variation of Arbitrary Constants*, Journ. de Math. **3**, 349 (1838).

[27] K. Brown, *Single Element Optics*, World Scientific (1999).

[28] F. Schmidt *et al.*, *Advances in MAD X using PTC*, Conf. Proc. **C**070625, 3381 (2007)

[29] D. Brown *et al.*, *A coherent muon-to-electron conversion experiment at Fermilab*, AIP Conf. Proc. **1441**, 596-598 (2012).

[30] A. Kurup, *The COherent Muon to Electron Transition (COMET) experiment*, Nucl. Phys. Proc. Suppl. **218**, 38-43 (2011).

[31] R. Djilkibaev and V. Lobashev, Sov. J. Nucl. Phys. **49**, 384 (1989).

[32] K. Abdel-Waged *et al.*, *GEANT4 hadronic cascade models analysis of proton and charged pion transverse momentum spectra from p+Cu and Pb collisions at 3, 8, and 15 GeV/c*, Phys. Rev. **C**84, 014905 (2011).

[33] N. Mokhov and S. Striganov, *MARS15 overview*, AIP Conf. Proc. **896**, 50-60 (2007).

[34] A. Sato, *G4beamline Simulation for the COMET Solenoid Channel.*, Conf. Proc. **C**100523 (2010).

[35] M. Aiba and F. Meot, *Determination of KEK 150 MeV FFAG parameters from ray-tracing in TOSCA field maps*, CERN-NUFACT-NOTE-140, DAPNIA-

04-188 (2004).

[36] O. Shanker, *High energy electrons from bound muon decay*, Phys. Rev. D**25**, 1847 (1982).

[37] A. Grossheim *et al.*, *Decay of negative muons bound in ^{27}Al*, Phys. Rev. D**80**, 052012 (2009).

[38] R. Watanabe *et al.*, *Asymmetry and energy spectrum of electrons in bound-muon decay*, Atom. Data and Nucl. Data Tables **54**, 165 (1993).

[39] A. Czarnecki, *et al.*, *Muon decay in orbit: spectrum of high-energy electrons*, Phys. Rev. D*84*, 013006 (2011).

[40] Y. Semertzidis *et al.*, *The MECO experiment at BNL*, Nucl. Phys. Proc. Suppl. **149**, 372-374 (2005).

[41] K. Nakamura *et al.*, *Particle Data Group*, JPG **37**, 075021 (2010).

[42] G. Punzi, *Sensitivity of searches for new signals and its optimization*, arXiv:physics/0308063v2 (2003).

[43] Y. Kuno et al., *Conceptual design report for experimental search for lepton flavour violating $\mu - e$ conversion at sensitivity of 10^{-16} with a slow-extracted bunched proton beam (COMET)*, KEK-2009-10 (2009).

[44] W. Bertl *et al.*, *A Search for $\mu - e$ conversion in muonic gold*, Eur. Phys. J. **C47**, 337346 (2006).

[45] V. Cirigliano *et al.*, *On the model discriminating power of $\mu \to e$ conversion in nuclei*, Phys. Rev. D**80**, 013002 (2009).

[46] Y. Kuriyama *et al.*, *Development status of PRISM FFAG ring and phase rotation simulation*, Nucl. Phys. Proc. Suppl. **155**, 284-285 (2006).

[47] A. Kolomensky and A. Lebedev, *Theory of cyclic accelerators*, North-Holland Publishing Company (1966).

[48] K. Symon *et al.*, *Fixed-Field Alternating-Gradient Particle Accelerators*, Phys. Rev. **103**, 1837 (1956).

[49] D. Kerst *et al.*, *Electron Model of a Spiral Sector Accelerator*, Rev. of Sci. Instrum. **31**, 1076 (1960).

[50] M. Aiba *et al.*, *Development of a FFAG proton synchrotron*. EPAC Procs. (2000).

[51] S. Machida *et al.*, *Acceleration in the linear non-scaling fixed-field alternating-gradient accelerator EMMA*, Nature Physics **8**, 243247 (2012).

[52] C. Johnstone and S. Koscielniak, *Recent progress on FFAGs for rapid acceleration*, Conf. Proc. **C**0206031, 1261-1263 (2002).

[53] I. Kato *et al.*, *Status of the T2K experiment*, J. Phys. Conf. Ser. **136**, 022018 (2008).

[54] C. Prior, *The Neutrino Factory and Related Accelerator research and development*, J. Phys. Conf. Ser. **110**, 012010 (2008).

[55] U Amaldi and G Kraft, *Radiotherapy with beams of carbon ions*, Rep. Prog. Phys. *68*, 1861 (2005).

[56] R. Wilson, *Radiological Use of Fast Protons*, Radiology **47**, 487-491 (1946).

[57] E. Clements, Symmetry **9**, 1 (2012).

[58] *Accelerator and Target Technology for Accelerator Driven Transmutation and Energy Production*, `http://science.energy.gov/` (2012).

[59] L. Medeiros-Romao and D. Vandeplassche, *Accelerator Driven Systems*, Conf. Proc. **C**1205201, 6-10 (2012).

[60] R. Tschirhart, *Project-X: A new high intensity proton accelerator complext at Fermilab*, arXiv:1109.3500 (2011).

[61] `http://www.conform.ac.uk/` (2012).

[62] J. Botman and C. Thomas, *Simulations on Harmonic Generation with the Particle Tracking Code GPT*, Conf. Proc. *C*0106181, 2763-2765 (2001).

[63] R. Appleby *et al.*, *Status of the PRISM FFAG Design for the Next Generation Muon-to-Electron Conversion Experiment*, Conf. Proc. **C**1205201, 322-324 (2012).

[64] M. Poole and E. Seddon, *4GLS and the Energy Recovery Linac Prototype Project at Daresbury Laboratory*, Conf. Proc. **C**0505161, 431 (2005).

[65] Y. Giboudot, *Study of Beam Dynamics in NS-FFAG EMMA with Dynamical Map*, Brunel PhD Thesis (2011).

[66] A. Terebilo, *Accelerator modeling with MATLAB accelerator toolbox*, Conf. Proc. **C**0106181, 3203-3205 (2001).

[67] J. Berg, *The EMMA main ring lattice*, Nucl. Instrum. Meth. A**596**, 276 (2008).

[68] F. Meot, *The Ray tracing code Zgoubi*, Nucl. Instrum. Meth. *A*427, 353-356 (1999).

[69] H. Wiedemann, *Particle Accelerator Physics II: Nonlinear and Higher-Order beam dynamics*, Springer (1995).

[70] G. Poplau *et al.*, *Multigrid algorithms for the fast calculation of space-charge effects in accelerator design*, IEEE Transactions on Magnetics **40**(2), 714717 (2004).