Numerical Studies for a Laboratory Astrophysics Experiment of unstable Electron-Positron Beams

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

vorgelegt von

Marvin Beck
aus Karlsruhe

Hamburg, 2023

Gutachter der Dissertation:

Prof. Dr. Günter Sigl

Prof. Dr. Florian Grüner

Gutachter der Disputation:

Prof. Dr. Robi Banerjee

Prof. Dr. Florian Grüner

Prof. Dr. Dieter Horns

Dr. Thomas Konstandin

Prof. Dr. Günter Sigl

Vorsitzende/r der Prüfungskommission:

Prof. Dr. Dieter Horns

Datum der Disputation: 12.05.2023

Vorsitzender Fach-Promotionsausschuss PHYSIK:

Prof. Dr. Günter H. W. Sigl

Leiter des Fachbereichs PHYSIK:

Prof. Dr. Wolfgang J. Parak

Dekan der Fakultät MIN:

Prof. Dr.-Ing. Norbert Ritter

Eidesstattliche Versicherung / Declaration on oath

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

I hereby declare upon oath that I have written the present dissertation independently and have not used further resources and aids than those stated.

Hamburg, 16.03.2023

Ort, Datum | *Place, Date*

Unterschrift | *Signature*

Zusammenfassung

Die An- oder Abwesenheit schwacher primordialer Magnetfelder in den Leerräumen zwischen Galaxienhaufen ist ein empfindlicher Test für die Physik des frühen Universums. Bisher konnten durch Beobachtungen nur Grenzwerte für die Stärke und Kohärenzlänge dieser Felder festgelegt werden. Eine indirekte Methode zur Festlegung unterer Grenzwerte für die Stärke dieser Felder und damit zum Nachweis ihrer Anwesenheit ist die Unterdrückung des sekundären Gammastrahlenflusses von weit entfernten VHE-Gammastrahlenquellen. TeV-Gammastrahlen induzieren elektromagnetische Kaskaden aufgrund von Paarbildungswechselwirkungen mit Photonen, die das extragalaktische Hintergrundlicht bilden und den gemessene Fluss bei hohen Energien reduzieren. Die Kaskaden-Elektronen und -Positronen erzeugen einen sekundären Gammastrahlenfluss mit niedrigerer Energie. Die experimentelle Nichtbeobachtung dieser sekundären Flusskomponente kann durch die Ablenkung der geladenen Elektronen und Positronen in den Magnetfeldern erklärt werden. Eine alternative Erklärung, die in der Literatur kontrovers diskutiert wurde, ist der Einfluss von Plasmainstabilitäten auf den gesamten neutralen Strahl, der aus Elektronen und Positronen besteht, die sich durch ein Hintergrundplasma, das intergalaktische Medium, ausbreiten. Diese Instabilitäten könnten dem Strahl Energie entziehen, ohne eine sekundäre Gammastrahlenkomponente zu erzeugen, oder Teilchen ablenken, ohne dass ein äußeres Magnetfeld vorhanden ist.
Moderne Beschleuniger können Strahlen erzeugen, die die Untersuchung der relevanten Instabilitätsprozesse in einer Laborumgebung ermöglichen. Trotz des großen Unterschieds der Skalen kann es möglich sein, mithilfe von Skalierungsbeziehungen, die abgeleitet werden müssen, zu extrapolieren. In dieser Arbeit untersuchen wir die Instabilitätsmechanismen mithilfe von Particle-in-Cell-Simulationen. Wir stellen fest, dass die Instabilität zum Aufbau eines strukturierten elektrischen Feldes führt, das dem Strahl Energie entziehen kann, indem es das Hintergrundmedium erhitzt, oder die Strahlteilchen durch einen Impulsdiffusionsprozess ablenkt, der mit einer Fokker-Planck-Gleichung modelliert werden kann. Die Veränderung der Impulsverteilung des Strahls könnte in einem Laborexperiment als leicht zugängliche experimentelle Nachweißgröße dienen. Wir stellen fest, dass der Instabilitätsmechanismus sowohl in neutralen Paarstrahlen als auch in experimentell leichter handhabbaren reinen Elektronenstrahlen auftritt.

Abstract

The presence or non-presence of weak primordial magnetic fields in the voids between galaxy clusters is a sensitive probe for the physics of the early universe. Currently observations have only been able to set limits for the strength and coherence length of these fields. An indirect method for setting lower limits on the strength of these fields, and thus prove their presence, is the suppression of the secondary gamma ray flux from far away VHE gamma ray sources. TeV gamma rays induce electromagnetic cascades due to pair production interactions with photons that constitute the extragalactic background light reducing the received flux at high energies. The cascade electrons and positrons produce a secondary flux of gamma rays at lower energies. The experimental non-observation of this secondary flux component can be explained by the deflection of the charged electrons and positrons in the magnetic fields. An alternative explanation that has been controversially discussed in the literature is the influence of plasma instabilities on the overall neutral beam consisting of electrons and positrons propagating through a background plasma, the intergalactic medium. These instabilities could drain energy from the beam without producing a secondary gamma ray component or deflect particles without the presence of an external magnetic field.

Modern accelerators can produce beams that allow the study of the relevant instability processes in a laboratory environment. Despite the vast difference of scales it can be possible to extrapolate using scaling relationships that have to be derived. In this work we study the instability mechanisms using particle-in-cell simulations. We find that the instability leads to the build up of a structured electric field that can drain energy from the beam by heating the background medium or deflect the beam particles by a momentum diffusion process that can be modeled using a Fokker-Planck equation. The changing of the beam momentum distribution could serve as an easily accessible experimental probe in a laboratory experiment. Lastly we find that the instability mechanism takes place in both neutral pair beams and experimentally more easily handleable pure electron beams.

Contents

List of Abbreviations

AGN	Active Galactic Nucleus
CMB	Cosmic Microwave Background
EAS	Extensive Air Shower
EBL	Extragalactic Background Light
EGRB	Extragalactic Gamma-Ray Background
EM	Electromagnetic
GRB	Gamma Ray Burst
HE	High Energy
IACT	Imaging Air-Cherenkov Telescope
ICS	Inverse Compton Scattering
IGM	Intergalactic Medium
IGMF	Intergalactic Magnetic Field
LWFA	Laser Wakefield Accelerator
PIC	Particle-in-Cell
PMF	Primordial Magnetic Field
SA	Surface Array
SED	Spectral Energy Distribution
VHE	Very High Energy

1 Introduction

The empty voids between galaxy clusters are a rich laboratory for physics. These voids are far more than just empty space, they contain a very dilute plasma known as the intergalactic medium (IGM), the entire diffuse electromagnetic radiation called the extragalactic background light (EBL) and possibly very weak, large scale magnetic fields (IGMF) that could be remnants of new physics in the early universe [55, 15]. It is of course not possible to probe these magnetic fields with in situ measurements, rather one has to apply indirect probes. So far these searches have only returned limits on the strength and coherence length of the IGMF. One innovative method to measure the strength of the magnetic fields has been dubbed the *gamma ray window* [15] and uses gamma rays, that is extremely high energy radiation, emitted from some of the most violent and energetic environment in our universe: The jets of blazars and gamma ray bursts. These objects are so bright, that their gamma ray emission can be measured from earth using modern telescopes. On their way to us these gamma rays spent many million years passing through the voids and they are so energetic that they can interact with the photon background to produce cascades of electrons and positrons [23, 58], attenuating the received signal earth at the highest energies [77]. The cascade electrons and positrons can up-scatter low energy photons and produce a secondary flux of gamma rays with lower energies. However the electrons and positrons are charged particles and thus susceptible to being deflected in the IGMF. If the IGMF is strong enough it can suppress the secondary gamma ray flux [82, 85, 118], produce a halo of extended emission around the point-like source [43, 17, 58, 6, 83] or produce a time delay that could take the form of a gamma ray echo for variable gamma ray source [78, 52, 83, 56]. So far no secondary gamma ray flux has been measured establishing a lower bound on the strength of the magnetic fields.

An alternative explanation for these observations has been investigated in the form of plasma instabilities [95, 94, 121, 131, 36, 41, 16, 37, 89]: The electrons and positrons in the cascades form a very dilute and highly relativistic beam of plasma that penetrates another background plasma, the IGM. These kinds of configurations have been known to be notoriously prone to unstable behavior of various kinds [34, 30, 33, 31, 32, 35]. A feedback loop between the beam plasma and the background plasma mediated by electromagnetic fields could drain energy from the pair beam without producing a secondary gamma ray flux or modulate the beam in a way that mimics the deflection in an incoherent magnetic field. If the plasma instabilities are actually the dominant process has been widely discussed in the literature with mixed results.

The unstable behavior of neutral pair beams has been studied with various analytical and numerical approaches, but performing a dedicated laboratory experiment that mimics the astrophysical pair beam at vastly reduced scale could improve our understanding of the mechanisms involved. The production of suitable pair beams with high energy and density by hitting a high density

target with an electron beam from a Laser Wakefield Accelerator (LWFA) [59] has already been demonstrated [98, 127, 100, 99, 101]. Nevertheless we need to stress that most conventional beam optic elements are not suitable to handle beams containing both negatively and positively charged particles. In this thesis we want to outline methods for performing such an experiment. They key challenges will be producing a beam that is suitable for driving the instability sufficiently long to produce a detectable signal and finding methods to quantitatively measure the unstable behavior.
We will start in chapter 2 with an excursion into the field of gamma ray astronomy, where will trace gamma rays from their source across the universe until their observation and explain the relevant physics along the way. In chapter 3 we will take a close look at plasma instabilities, introduce relevant quantities and give analytical description of their behavior. Another key point will be the derivation of an analytical model based on the Fokker-Planck equation that describes the feedback of the plasma instability on the momentum distribution of the particle beam. The following chapter 4 is focused entirely on the Particle-in-Cell (PIC) approach, a very powerful numerical method in plasma physics.
In chapter 5 we will use the PIC simulations to study the unstable behavior of cold pair beams in an idealized scenario of a beam that is infinitely large, neutral and propagating through a perfectly homogeneous background medium without any external fields. In chapter 6 we will use the results of the previous chapter to comment on the influence of plasma instabilities on the astrophysical observations. In chapter 7 we will take a look at effects that go beyond the idealized scenario and may become relevant in either astrophysics or a laboratory experiment. This includes the effects of a finite size particle beam, a non-neutral pair beam or an external magnetic field. In chapter 8 we will draw conclusions on how to set up a laboratory experiment. Lastly in chapter 9 we will present an alternative experimental scenario where the unstable behavior of an electron beam is investigated, that could still probe the relevant unstable behavior.

2 | Gamma Rays

This chapter will give an overview of the physics that are relevant for understanding the gamma ray portal for measuring the intergalactic magnetic field (IGMF). Starting from the production of gamma rays in far distance sources we will trace their path to earth including interactions with photons from the extragalactic background light (EBL), the electromagnetic cascades that are induced by these interactions and the production of secondary photons via interactions with the cosmic microwave background (CMB). Furthermore we will also introduce the IGMF and its possible influence on cascade leptons. At the end we will discuss the extragalactic gamma-ray background (EGRB) and how to measure gamma rays using space- and earth-bound observatories.

2.1 Gamma Rays from Blazars

Gamma Rays are a form of very high energy electromagnetic radiation. They are photons that carry more energy and have a shorter wavelength than even X-rays. Typically photons with energies above 100 keV are called gammma rays. Gamma rays with energies between 100 MeV and 100 GeV are often called high energy (HE) gamma rays, where those with energies above 100 GeV are often called very high energy (VHE) gamma rays [51]. Due to their high energy these photons can only be produced in the most energetic processes in the universe. Typical environments are gamma ray bursts (GRB) or the jets of active galactic nuclei (AGN) for extragalactic sources [28]. They can also be produced by galactic sources like Supernova Remnants (SNR), Pulsar Wind Nebulae (PWN), X-Ray Binaries or Pulsars, but galactic sources are too close for significant attenuation and as a consequence also cascade emission and are not considered here for this reason.
AGN are the compact regions found in the center of a minority of galaxies. The unified AGN model (figure 2.1) explains classes of objects that appear very different to an observer as AGNs with or without a jet that are observed from different angles [22]. In the center of an AGN is a supermassive black hole that accretes material from their surrounding environment and produces intense radiation as a result. Consequently the black hole is surrounded by an accretion disk followed by a dust torus further out. Some AGNs produce a jet, a highly relativistic outflow of plasma. The highly boosted jet environment significantly enhances the observed flux and shortens time variability to an observer [28]. The exact mechanism for the formation of a jet is still not very well understood.
Blazars are a subclass of radio-loud AGN with a jet that is directly pointing at the observer. They are know to be sources of electromagnetic radiation across the entire electromagnetic spectrum from radio waves all the way to VHE gamma rays. In the VHE band the spectrum is generally

assumed to be a power-law with a cutoff above some maximum energy. The spectral energy density of the blazar Markarian 421, one of the closest blazars to earth at a redshift of $z = 0.03$, is shown in figure 2.2 to span at least 18 orders of magnitude with two distinct peaks. The double hump structure is very typical for the emission of a blazars. The lower energy peak of the SED, in the radio to optical range, is caused by synchrotron emission from relativistic electrons and positrons in the jet. The exact origin of the high energy peak, in the X-ray to VHE gamma ray range, is not understood yet, but two mechanisms are generally proposed. In the leptonic model the photons from the low energy population are up-scattered by inverse-Compton interactions with the same population of relativistic leptons (the total process is then often summarized as synchrotron-self-compton). In the hadronic model the emission of high energy photons is explained by photo-pion processes caused by high energy hadrons interacting with the background environment [64]. A mixture of the leptonic and hadronic models could also be viable. Only the hadronic model would also produce a significant flux of high energy neutrinos from the source. Measuring the neutrino flux would illuminate the exact process that leads to the second hump. A recent measurement by IceCube indicates that blazars do produce high energy neutrinos [1].

Blazars have been observed to show variability of their flux in all wavebands on timescales of from minutes to years [28]. This variability can also go hand in hand in hand with a change of the spectral index of the VHE emission. The short-time variability indicates that the emission region is very small due to reasons of causality.

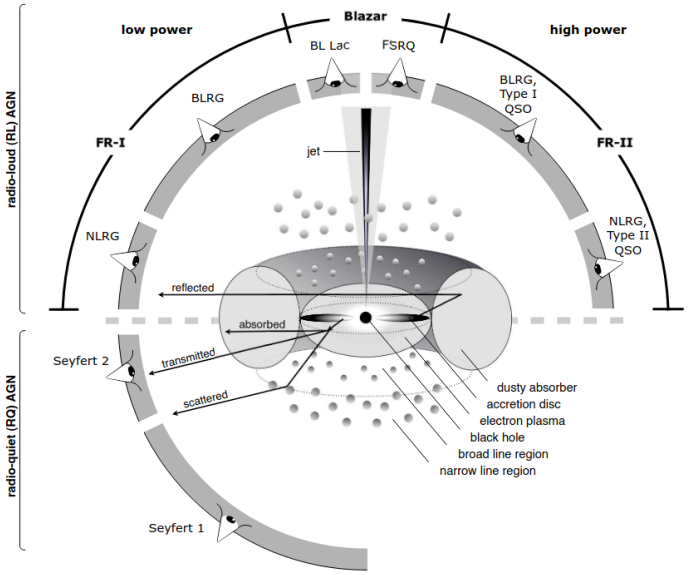

Figure 2.1: Sketch of unified AGN scheme, where the class of object observed depends on the viewing angle, the emission strength and if the AGN produces a jet (Image source: [22]).

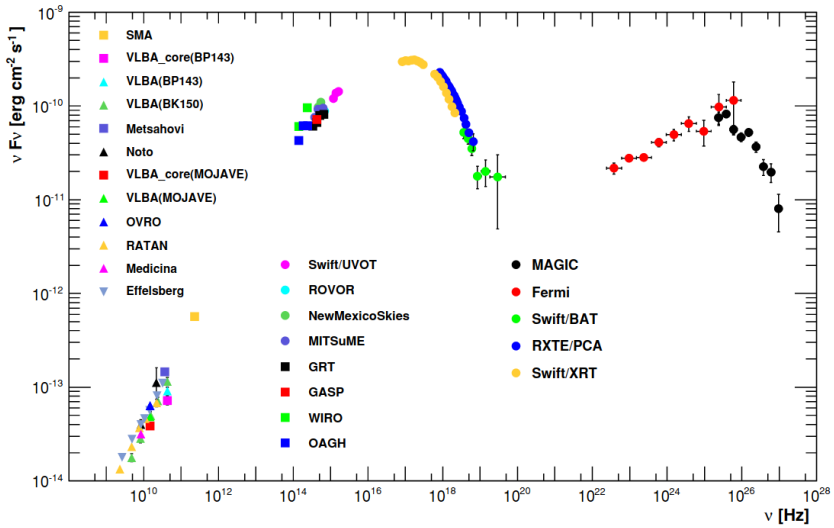

Figure 2.2: Spectral Energy Distribution (SED) of Markarian 421 measured in multiwavelength campaign between MJD 54850 and MJD 54983. Some measurements where corrected for propagation effects and background contributions (Image source: [3]). The typical double hump structure is visible.

2.2 Extragalactic Background Light (EBL)

The extragalactic background light (EBL) covers the entire diffuse radiation emitted by all sources outside of the Milky Way. The EBL refers to the entire electromagnetic spectrum (see figure 2.3) from radio to gamma rays, however sometimes the term is used to only refer to the part of the spectrum from ultraviolet to infrared [45]. The most prominent and also well-understood part of the spectrum is the microwave peak associated with the cosmic microwave background discussed in more detail in section 2.3.
The EBL is of great interest to the field of gamma ray astronomy because highly energetic photons can interact with the low energy photons that make up the EBL to produce e^+e^- pairs. This leads to an attenuation of the gamma ray flux at the highest energies [63, 61, 65]. The dominant range of the EBL for this process is the ultraviolet to infrared [65]. In this range spanning from 10^{-3} eV to 10 eV in energy, or 0.1 μm to 1000 μm in wavelength, the EBL is dominated by starlight. The spectrum follows a double hump structure with one peak at 1 μm due to direct starlight and one at 100 μm due to reemission from dust [61]. Unfortunately direct measurements of the EBL in this regime are difficult due to foreground contributions in our own galaxy [45]. A lower limit on the EBL can be provided by galaxy count surveys [63, 61, 65]. Indirect measurements that use the suppression of the gamma ray flux deliver the best measurements of this part of the EBL, however they suffer from large uncertainties on the intrinsic spectra of gamma ray sources [61, 65].

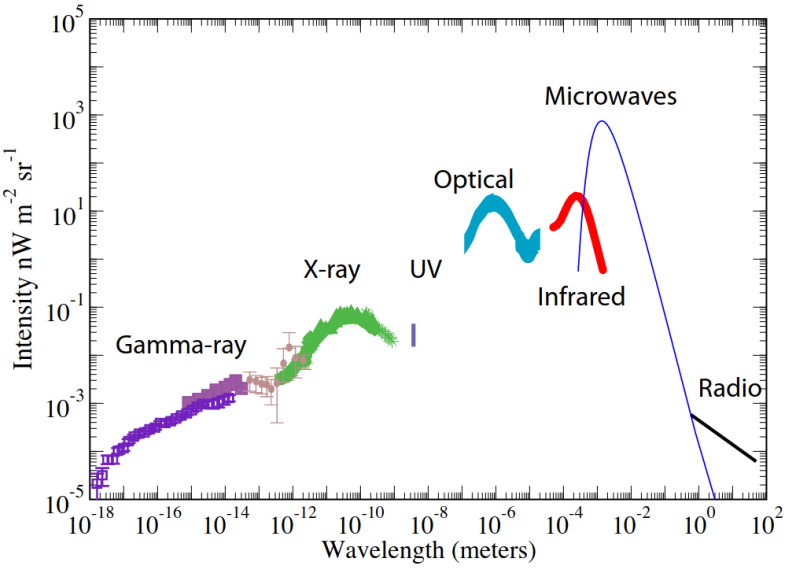

Figure 2.3: The intensity of the EBL as a function of wavelength across the entire electromagnetic radiation spectrum from radio to gamma rays summarizing multiple measurements (Image source: [45]).

2.3 Cosmic Microwave Background (CMB)

The matter in the early, hot universe forms a fully ionized primordial plasma. Photons were tightly coupled to the electrons in this plasma. At temperatures well below the ionization energy of hydrogen (13.6 eV) the electrons and protons recombined. The ionization fraction of the universe rapidly decreased and it became transparent to photons.

$$p^+ + e^- \rightleftharpoons H + \gamma \tag{2.1}$$

The transition is however not instantaneous, rather the backreaction of (2.1) becomes inefficient for temperatures below the ionization energy. The exact redshift of recombination z_{rec} has to be calculated using the Saha-equation for the freeze-out of the ionization fraction yielding $z_{\text{rec}} \sim 1100$ [108].
The formerly tightly coupled photons form a relic called the *Cosmic Microwave Background* (CMB) [39, 54]. They form a perfect blackbody that has been redshifted since the time of recombination and has a temperature of $T = 2.725\,\text{K}$ today. As a perfect blackbody the density and spectrum of the CMB are determined alone by the temperature and are thus precisely known [72]. Since the temperature of the CMB is the temperature of the primordial plasma at the time of decoupling redshifted to today, the CMB traces the density perturbations of the universe at the time of decoupling. Figure 2.4 show the temperature fluctuations of the CMB (without the dipole) measured by the Planck satellite [92]. The fluctuations are of the order of $100\,\mu\text{K}$. The power spectrum (two-point-correlation function) of these temperature fluctuations are a sensitive probe for the evolution of the early universe [39, 54, 92].

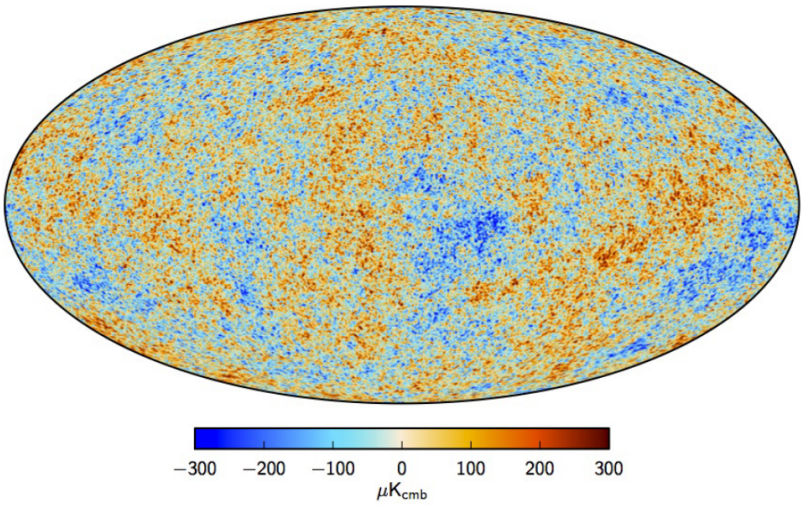

Figure 2.4: Map of the CMB fluctuations (after subtraction of Monopole and Dipole) (Image source: [92]).

2.4 Pair Cascades

Gamma rays propagating through the universe can interact with background photon fields to produce pairs of electrons and positrons (2.2). These electrons and positrons can then interact with the background photon field and up-scatter photons via inverse compton scattering (ICS) to high energies (2.3). If these photons are energetic enough they can again produce an e^+e^- pair.

$$\gamma + \gamma_{\mathrm{bg}} \rightarrow e^+ + e^- \tag{2.2}$$

$$e^\pm + \gamma_{\mathrm{bg}} \rightarrow e^\pm + \gamma \tag{2.3}$$

These processes continue until the photons are no longer energetic enough to produce pairs, leading to an excess of secondary or cascade emission in the HE band. The mean free path of this process depends both on the energy of the gamma ray and the differential energy density of the photon field. The center of mass energy of the two photon system needs to exceed the combined mass of electron and positron, leading to a threshold energy E_{th} for the gamma ray as a function of the background photon energy E_{bg} given by (2.4).

$$E_{\mathrm{th}} = \frac{m_e^2}{E_{\mathrm{bg}}} \tag{2.4}$$

For a background photon with $E_{\mathrm{bg}} = 1\,\mathrm{eV}$ this leads to a threshold of $E_{\mathrm{th}} \sim 260\,\mathrm{GeV}$. Figure 2.5 shows the associated mean free path of a gamma ray as function of energy and labels the energy band of the EBL that is most significant for pair production. Gamma rays in the VHE band typically have mean free paths of the order of 100 Mpc where they dominantly interact with photons in the optical and near infrared band [15]. One of the most dominant uncertainties in the modeling of electromagnetic cascades is the density of the EBL in this energy band. For gamma rays above $\sim 100\,\mathrm{TeV}$ the universe becomes highly opaque because pair production on

the much denser CMB becomes kinetmatically possible. The mean free path of ICS for electrons and positrons of the same energy band is much shorter on the order of 30 kpc [15]. This is in part the case because ICS has no energy threshold and thus is always possible with the much denser CMB.
At the highest energies higher order process like double pair production (2.5) and triplet pair production (2.6) serve as corrections to the evolution of the cascade.

$$\gamma + \gamma_{\mathrm{bg}} \rightarrow e^+ + e^- + e^+ + e^- \tag{2.5}$$
$$e^\pm + \gamma_{\mathrm{bg}} \rightarrow e^\pm + e^+ + e^- \tag{2.6}$$

The pair production process itself produces particle with a small amount of transverse momentum with respect to the orientation of the incident gamma ray, leading to a small increase emission opening angle. More importantly the charged electrons and positrons can be deflected in magnetic fields that could permeate the entire universe. These deflections can alter the observation of a gamma ray source in three ways:

1. They can reduce the secondary gamma ray flux within the point spread function of a detector, reducing the measured flux in the HE band.

2. They can lead to a halo of extended HE emission around a source.

3. They can increase the time of flight and thus delay the secondary gamma ray emission. This becomes only relevant if the emission of a source it not constant over long time scales.

Neronov & Semikoz [84] derived analytical approximations for the angular size of the secondary halo θ and the time delay Δt under the assumptions that the individual deflections are small and the pairs are produced closer to the source than to Earth. The result strongly depends on the hierarchy of coherence length of the magnetic field λ_B and the ICS mean free path λ_{ICS}. If $\lambda_B \ll \lambda_{\mathrm{ICS}}$ then the electron or positron crosses many independent domains of the magnetic field before up-scattering a secondary photon, whereas if $\lambda_B \gg \lambda_{\mathrm{ICS}}$ it only experiences one domain of the magnetic field.

$$\theta = \begin{cases} 0.07^\circ (1+z)^{-\frac{1}{2}} \left(\frac{\tau_\theta}{10}\right)^{-1} \left(\frac{E}{0.1\,\mathrm{TeV}}\right)^{-\frac{3}{4}} \left(\frac{B}{10^{-14}\,G}\right) \left(\frac{\lambda_B}{1\,\mathrm{kpc}}\right)^{\frac{1}{2}} & \lambda_B \ll \lambda_{\mathrm{ICS}} \\ 0.5^\circ (1+z)^{-2} \left(\frac{\tau_\theta}{10}\right)^{-1} \left(\frac{E}{0.1\,\mathrm{TeV}}\right)^{-1} \left(\frac{B}{10^{-14}\,G}\right) & \lambda_B \gg \lambda_{\mathrm{ICS}} \end{cases} \tag{2.7}$$

Here B is the strength of the magnetic field, E is the energy of the cascade photons when they are observed, z is the redshift of the source and τ_θ is the ratio of the angular diameter distance between observer and source and the pair production mean free path.

$$\Delta t = \begin{cases} 7 \cdot 10^5\,\mathrm{s} \left(1 - \tau_\theta^{-1}\right) (1+z)^{-5} \kappa \left(\frac{E}{0.1\,\mathrm{TeV}}\right)^{-\frac{5}{2}} \left(\frac{B}{10^{-18}\,G}\right)^2 & \lambda_B \ll \lambda_{\mathrm{ICS}} \\ 1 \cdot 10^4\,\mathrm{s} \left(1 - \tau_\theta^{-1}\right) (1+z)^{-2} \kappa \left(\frac{E}{0.1\,\mathrm{TeV}}\right)^{-2} \left(\frac{B}{10^{-14}\,G}\right)^2 \left(\frac{\lambda_B}{1\,\mathrm{kpc}}\right) & \lambda_B \gg \lambda_{\mathrm{ICS}} \end{cases} \tag{2.8}$$

Here κ is a model-specific parameter that normalizes the EBL and is of order 1. To observe the time delay the source needs to show variability on time scales smaller than Δt, either in the form a flaring behavior or a gradual change in the flux state.
Another effect that could severely alter the evolution of the electromagnetic cascades are plasma

instabilities. The electrons and positrons form a neutral plasma beam that propagates through a background plasma, the intergalactic medium (IGM). Beam instabilities are a feedback loop between the beam and the background plasma that build up electromagnetic fields and drain energy from the beam. We will give a more in-depth overview of the physics of plasma instabilities in chapter 3. If the cooling of the beam due the plasma instabilities is more efficient than the inverse compton cooling the secondary gamma ray flux is suppressed. This would render any limits derived for the IGMF strength moot.
The effect of plasma instabilities on the electromagnetic cascade has been a topic of great debate. There is no consensus if they significantly influence their evolution or even which kind of instable behavior would be relevant. In particular the oblique [121, 95, 40, 36, 109, 117, 103] instability as well as non-linear landau damping [42, 120, 80] have been suggested, although there is disagreement on whether these processes would be faster than the inverse compton cooling. Another point of dispute is whether the reactive or kinetic regime of instability growth is applicable [104, 29, 80, 40]. Some authors suggest that the instabilities would saturate before they can drain a significant amount of energy from the beam [89]. Other authors have suggested that even if the instability dominates the system under idealized conditions the inhomogeneity of the background medium [106, 107] or magnetic fields [9] could suppress the instability growth. It has also been suggested that the instability could lead to significant heating of the background plasma [41, 109]. Simulations of cascade evolution including the effect of various models of plasma instabilities have been performed showing that under some circumstances the observed suppression of the cascade contribution in the HE band could be explained by these instabilities [131, 16, 102].
Electromagnetic cascades may also be induced by cosmic rays which could be accelerated by the same sources that emit VHE gamma rays. Further corrections to the cascade process could occur due to gravitational lensing, Lorentz invariance violation or kinetic mixing with axion-like-particles (ALPs) [15].

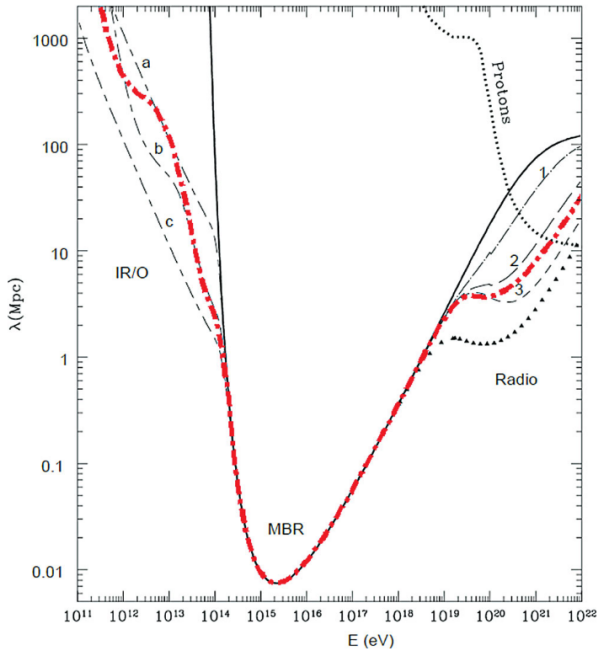

Figure 2.5: The mean free path λ_{PP} of high energy photons for the pair production process $\gamma\gamma \to e^+e^-$ as a function of the gamma ray energy. The universe becomes highly opaque to gamma rays with energies above 100 TeV (Image source: [48]).

2.5 Intergalactic Magnetic Fields (IGMF)

The Intergalactic Magnetic Fields (IGMF) are large scale magnetic fields that permeate the universe. Inside galaxies and galaxy cluster the magnetic field is of the order of Microgauss [55]. These measurements are based on Faraday rotation and synchrotron emission [96]. Galaxies form clusters that are connected via filaments where magnetic fields are of the order of 0.1 nG to 10 nG [15]. Gamma rays from blazars travel over enormous distances and thus pair cascades probe the magnetic field on distances much larger than the size of galaxy clusters or even galaxy filaments. The volume filling factor of the magnetic fields of galaxies as well as filaments are small and the evolution of cascades is entirely governed by the as of today unknown fields in the voids.

Stochastic magnetic fields can be characterized by their strength, more precisely their *root mean square* value, B and their coherence length L_B. The coherence length describes the average size of structures in the magnetic field. The strength and coherence length is intimately linked to the question of the fields origin. While the exact origin is still unknown the possible scenarios can broadly be divided into two groups: Cosmological scenarios were strong seed fields where created in the early universe that later decayed into the magnetic field today. In cosmological scenarios the magnetic field is also called a Primordial Magnetic Field (PMF). And astrophysical scenarios where weak seed fields where produced due to local effects and then later amplified by dynamo or compression mechanisms [15, 55].

Independent of the field strength constraints on the coherence length can be derived from causality (Hubble radius) and the evolution of the magnetic fields (magnetic diffusion) to set an upper

and a lower limit respectively [15]. Further limits on the field strength (as a function of the coherence length) can be set from Zeeman splitting, Faraday rotation or CMB measurements. In the case of cosmological PMFs one can also derive stringent limits on their current strength and coherence length from limits on their generation and calculating their time evolution with freely decaying MHD (early magnetic dissipation) [20]. An overview of these bounds can be found in figure 2.6.

One of the most promising probes of the IGMF in the galactic voids is the deflection of electrons and positrons in cascades. These deflections can lead to multiple phenomena that can be used to measure or constrain the IGMF. A popular choice for studying the effect of cascade emission are extreme blazars with a hard spectrum, where the intrinsic flux in the HE band is smaller relative to the cascade contribution that depends on the VHE flux. Neronov & Vovk [85] used the outright suppression of the secondary emission from four blazar sources to constrain the IGMF to $B \geq 3 \cdot 10^{-16}\,\mathrm{G}$ for coherence lengths above Mpc and improving with $\lambda_B^{-\frac{1}{2}}$ for smaller coherence lengths. In a further work Taylor, Vovk & Neronov [118] used three blazar sources to constrain $B \geq 10^{-15}\,\mathrm{G}$ for $\lambda_B > 1\,\mathrm{Mpc}$ for a constant source and $B \geq 10^{-17}\,\mathrm{G}$ for the same coherence length if the source is variable and the suppression of the flux is due time delay. These constraints are also shown in figure 2.6. The suppression of the GeV cascade flux is shown for one source in figure 2.7. The MAGIC collaboration uses the non-observation of pair echos from a variable blazar source [78] to constrain the magnetic field strength $B > 1.8 \cdot 10^{-17}\,\mathrm{G}$ for long correlation lengths and $B > 10^{-14}\,\mathrm{G}$ for an IGMF with cosmological origin. Current gamma ray space telescopes are not capable of measuring the time delayed signal from GRBs but this could be a feasible method to set similar bounds in the future [56].

The secondary gamma ray emission would take on the form of an extended gamma ray halo in the presence of a magnetic field. Even if this halo is too weak to be detected for individual sources it could be detected by stacking multiple sources, increasing the significance [38, 119]. The angular structure of the halo would give insight on the coherence length and helicity of the IGMF but is also affected by the angle under which the jet is observed [17, 15].

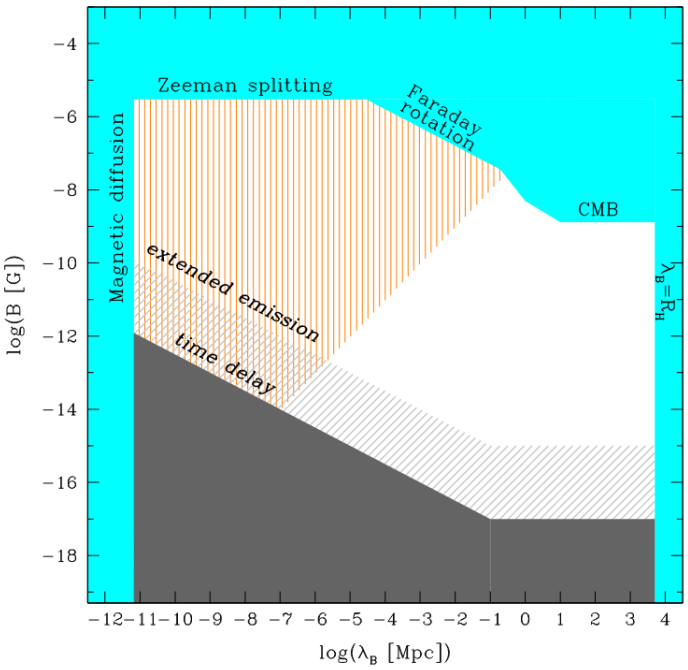

Figure 2.6: Constraints on the IGMF (Image source: [118]). Upper limits due Zeeman effect, Faraday rotation measurements and CMB measurements. Magnetic diffusion and the Hubble radius give a lower and upper limit on the coherence length. The orange shaded area refers to freely decaying MHD in the early universe (only for PMFs). Extended emission & time delay refer to blazar measurements.

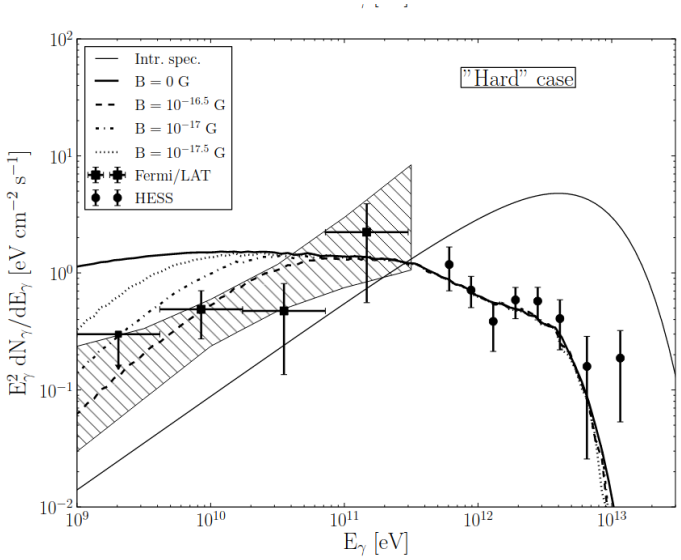

Figure 2.7: SED of 1ES 0229+200 under the assumption of a hard intrinsic spectrum measured by Fermi-LAT and HESS (data points). The different lines give the combined intrinsic spectrum and cascade contribution under the assumption of different field strengths (Image source: [126]).

2.6 Extragalactic Gamma-Ray Background (EGRB)

The gamma ray sky can be divided into resolved point sources and diffuse emission from galactic and extragalactic sources. The diffuse emission from extragalactic sources makes up the extragalactic gamma ray-background (EGRB). The exact origin of the EGRB is as of yet unknown, but unresolved point sources likely contribute significantly. The VHE emission from these sources would also produce pair cascades that contribute to the EGRB at lower energies [124, 73, 71]. The electrons and positrons of these cascades could likewise be deflected in the IGMF or undergo unstable behavior. Deflection in magnetic fields or due to instabilities could increase the flux that is not resolved as point sources increasing the EGRB, while energy loss due to instability growth would reduce the secondary gamma ray flux and thus the EGRB.
Below TeV energies the largest contribution to the gamma ray flux can be attributed to galactic diffuse emission. This emission is mostly confined to the galactic disk where gamma rays are produced by cosmic rays interacting with the interstellar medium and cosmic ray electrons upscatter interstellar radiation fields [49]. An additional source of galactic contribution are the *Fermi Bubbles*, two extended bubble-like structures above and below the galactic center with a significantly harder gamma ray spectrum than the diffuse emission from the galactic disk [115, 116, 53]. These foreground contributions as well as contributions from resolved sources have to be subtracted from the measured flux to determine the EGRB.
Fermi-LAT observations [5] show that the EGRB follows a power law with a spectral index of 2.32 ± 0.02 and a cutoff energy of (279 ± 52) GeV with a total flux above 100 MeV of $(7.2 \pm 0.6) \cdot 10^{-6}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1}$.

2.7 Gamma Ray Detection

For photons with energies above the optical waveband Earth's atmosphere becomes opaque. This motivates the use of space based telescopes to study X-rays and gamma rays. The effective area of these telescopes is given by the size of the instrument, which usually falls in the order of m^2. For the highest energies the gamma ray flux decreases very rapidly and the effective area of space based telescopes is not enough to make meaningful measurements. At these energies ground based telescopes with far larger effective area have to be used. These telescopes do not measure the gamma rays themself, but the extensive air showers (EAS), electromagnetic cascades that are induced by interactions of the photon with air molecules. The energy window of space based and ground based observations roughly follow the energy bands of HE and VHE gamma rays respectively, although there is some overlap in the energy ranges. Figure 2.8 compares the sensitivity, that the is the minimum flux that can be distinguished from the background at a given significance for a certain observation time, for different instruments.

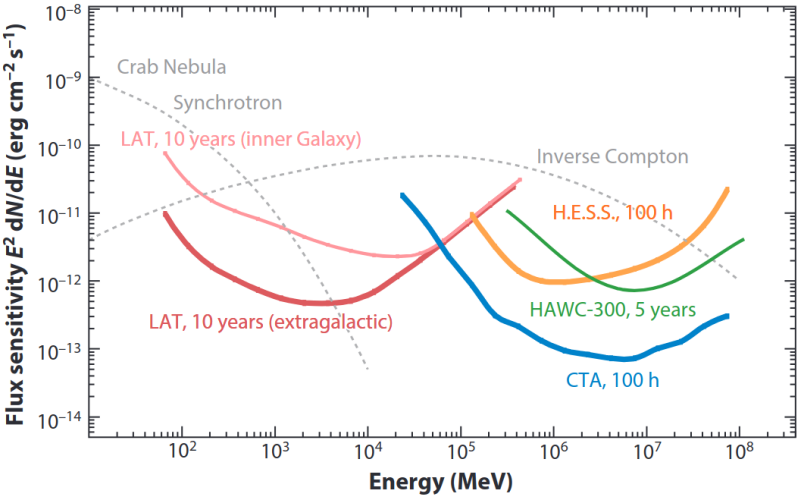

Figure 2.8: Sensitivity to detect a source with 5σ significance as a function of energy for different gamma ray telescopes for specific observation times (Image source: [64]). For Fermi-LAT the reduced background for sources outside the galactic disk increase the sensitivity. In dashed grey lines the flux of the crab nebula is given for comparison.

2.7.1 Space Based Observations

Space based gamma ray telescopes measure gamma rays in the HE energy band. In contrast to lower energy bands of electromagnetic radiation (X-ray, optical, radio, etc.) they can not be focused and have to be detected through their particle-like interactions. Space based gamma ray telescopes are particle detectors that consist of a tracker and calorimeter where gamma rays produce e^+e^- pairs that have their arrival direction reconstructed with the tracker and their energy measured with the calorimeter [49, 51]. Charged cosmic rays are a very large background to the measurement of the gamma ray flux and have to be suppressed using an anti-coincidence detector. Only e^+e^- pairs that are produced well inside the detector volume are considered as gamma ray induced.
The currently most advanced space based gamma ray telescope is Fermi-LAT [4, 2], launched in 2008, with an energy range of 20 MeV to 300 GeV and an effective area of 0.65 m^2 above 1 GeV [64]. The angular resolution, also called point-spread function (PSF), reaches 0.8° above 1 GeV and is better than 0.2° above 10 GeV. The energy resolution is better than 10%. The field of view (FoV) is 2.4 sr at 1 GeV and the instrument achieves an effective live time of 75% [64]. This allows the monitoring of the entire sky for many thousands of hours over the course of the entire mission.

2.7.2 Ground Based Observations

Ground based gamma ray telescopes do not observe the incident gamma ray itself, but the EAS that develops when a high energy photon hits the atmosphere. These air showers can be detected by two complementary techniques: surface arrays and imaging air-cherenkov telescopes. Both techniques suffer systematic uncertainties from the atmospheric conditions and need to employ

sophisticated methods to suppress the overwhelming cosmic ray background.
The first method is using a surface array (SA) of particle detectors, like water-cherenkov tanks, at high altitudes to directly measure the air shower particles. This technique has the advantage of a large field of view and a high duty cycle, continuously monitoring the entire sky above [64]. One of the leading SAs is the HAWC (High-Altitude Water Cherenkov) Observatory [113], located at an altitude of 4100 near the Sierra Negra volcano and it consists of 300 water tanks each equipped with four upward facing PMTs (Photomultiplier Tubes) located at the bottom measuring the cherenkov light of air shower particles that traverse the tank volume. The effective area amounts to 22,000 m^2 [113].
The second method is using a an imaging air-cherenkov telescope (IACT). IACTs observe the Cherenkov light emitted by the high energy particles in the air shower with a camera that is located in the focal point of a large, usually segmented, mirror [64]. A particle shower that reaches its maximum at 10 km above the surface and emits cherenkov light at an angle of 1° illuminates an area with radius of 120 m on the ground [110], leading to an effective area of 10^5 m even though the telescopes themself are much smaller. The camera is able to resolve the air shower over many pixels and with sufficient time resolution to essential take a video of the air shower as it develops in the sky. The total light measured by an IACT is well-correlated with the energy of the incident particle. The shape of the air shower image as well as the timing information is enough to determine the origin of the incident gamma ray on the sky even with a single telescope. Their high sensitivity and angular resolution comes at the cost of a small FoV only allowing to observe a single source at a time. Furthermore they can only be operated during dark, cloudless nights severely reducing their duty cycle. Modern IACTs are operated as an array of at least two telescope that observe the same air shower from different positions to improve the rejection of background, the angular resolution and the energy resolution.
The currently most advanced IACTs are Magic [10, 46] on La Palma, Spain, Veritas in Arizona, USA [70, 129] and Hess in Khomas, Namibia [8, 24, 87]. In the near future CTA (Cherenkov Telescope Array) [67] spread over two sites on La Palma, Spain and near Paranal, Chile will become operational as the next generation of IACTs. It will consist of over 100 telescopes of three different sizes to achieve ten times higher sensitivity than current generation IACTs [81, 81].
The performance of SAs and IACTs is generally dependent on the energy of the incident gamma ray and the zenith angle of the observation. For this reason we can only give a very general overview of the performance of current day instruments. The threshold energy of IACTs is at tens of GeV, whereas for SAs it is at the TeV level. IACTs achieve a better energy resolution of $\sim 15\%$ than SAs of $\sim 40\%$. The biggest difference is the much higher sensitivity of IACTs, reaching roughly 1% of the Crab nebula flux in 25 hours of direct observation time, whereas SAs need 5 years of total observation time to reach similar sensitivity [110].
If the incident particle of an air shower is a photon or a lepton then the air shower will exclusively consist of leptons and photons (electromagnetic shower), whereas a hadron as primary particle will induce a hadronic cascade. These cascades develop very differently and can be distinguished by both SAs and IACTs based on the lateral distribution of particles. Generally the rejection of background improves at higher energies [110]. An additional method of background rejection for a known source is the comparison with the measured flux from an off-source region taken under the same observation circumstances.

3 | Plasma Instabilities

Plasma instabilities refer to an unstable behavior in a plasma. They occur when multiple plasma components, different species of particles and fields, interact. This can lead to self-reinforcing fluctuations of properties like the plasma density, the plasma temperature or electromagnetic fields. In many applications of plasma physics like fusion devices these instabilities can lead to critical failure and have to be suppressed. In this work we are particularly interested in beam instabilities of neutral electron positron pair beams. There are many different classes of beam instabilities. For example the two-stream instability occurs when two different plasma beams with different velocity distributions flow past each other leading to the growth of electromagnetic waves. Another example, the Weibel instability occurs when plasma beam propagates through a magnetic field, inducing fluctuations that lead to the growth of magnetic filaments.
Of course within the scope of this thesis we can not give a full overview of the physics of plasmas, which is why we defer to a number of excellent books written about this topic [90, 88, 86, 74, 44].

3.1 Plasmas

A plasma is a fourth state of matter similar to gases, that is defined by a sizable presence of free electrons or ions, making the plasma electrically conductive. The key property of a plasma are long range, collective effects governed by electromagnetic fields, rather than direct particle-particle interactions. A plasma where particle-particle interactions play no role, typical for low-density plasma like the IGM, is called collisionless. The collective effects govern the response of a plasma to a small perturbation which lead to oscillations of the fields and charge density with a defined frequency that only depends on the density n of the plasma: the plasma frequency ω_p. For this chapter we will use SI units [25] unless otherwise specified.

$$\omega_p = \sqrt{\frac{ne^2}{m_e \epsilon_0}} \tag{3.1}$$

A plasma is opaque to electromagnetic radiation with frequency below its plasma frequency and can only penetrate up to the depth of the plasma skin depth $\frac{c}{\omega_p}$. Since the individual particles that constitute a plasma are free to move around and react to electric fields they can redistribute themselves to screen out any electric fields on length scales larger than the Debye length λ_D.

$$\lambda_D = \frac{v_{\text{th}}}{\omega_p} = \sqrt{\frac{\epsilon_0 k_B T}{n_e e^2}} \tag{3.2}$$

Magnetic fields on the other hand are not screened in a plasma, rather the particles react by gyrating around the magnetic field lines with the gyrofrequency ω_g.

$$\omega_g = \frac{|e|\,B}{m_e} \tag{3.3}$$

The microphysics of a collisionless plasma can be completely described using the four Maxwell equations for the evolution of the fields and the Lorentz force for the evolution of the particles.

$$\nabla \cdot \vec{E} = \frac{\rho}{\epsilon_0} \tag{3.4}$$

$$\nabla \cdot \vec{B} = 0 \tag{3.5}$$

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{3.6}$$

$$\nabla \times \vec{B} = \mu_0 \left(\vec{j} + \epsilon_0 \frac{\partial \vec{E}}{\partial t} \right) \tag{3.7}$$

The Maxwell equations can be used to derive the continuity equation.

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \vec{j} = 0 \tag{3.8}$$

Together with the acceleration due to the Lorentz force the description of the microphysics of a collisionless plasma is complete.

$$\vec{a} = \frac{q}{m} \left(\vec{E} + \vec{v} \times \vec{B} \right) \tag{3.9}$$

For any realistic scenario it is of course unfeasible to follow the individual trajectories of all particles. Instead the Valsov equation uses the conservation of phasespace density $f = f(\mathbf{r}, \mathbf{p}, t)$ to treat the evolution of a collisionless plasma.

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} f(\mathbf{r}, \mathbf{p}, t) &= \frac{\partial f}{\partial t} + \frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t} \frac{\partial f}{\partial \mathbf{r}} + \frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} \frac{\partial f}{\partial \mathbf{p}} \\ &= \frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f - q \left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right) \frac{\partial f}{\partial \mathbf{p}} = 0 \end{aligned} \tag{3.10}$$

The evolution of the field $\mathbf{E}$ and $\mathbf{B}$ is then still described by the Maxwell equations and the quantities ρ and $\mathbf{j}$ can be calculate from the phasespace density.

$$\rho = qn = q \int \mathrm{d}\mathbf{p} f \tag{3.11}$$

$$\mathbf{j} = q \int \mathrm{d}\mathbf{p} \mathbf{v} f \tag{3.12}$$

In the case of a plasma consisting out of multiple species of particles (e.g electrons and positrons or electrons and protons), each species has to be treated with a separate Vlasov equation, while ρ and $\mathbf{j}$ are calculate from the sum of the individual contributions.
To study the unstable behavior of a plasma one can assume an equilibrium state and investigate

the influence of a small perturbation in the form of a plane wave. f_0, $\mathbf{E}_0$ and $\mathbf{B}_0$ denote the equilibrium state, f_1, $\mathbf{E}_1$ and $\mathbf{B}_1$ the amplitude of the perturbation.

$$f(\mathbf{r}, \mathbf{p}, t) = f_0(\mathbf{r}, \mathbf{p}, t) + f_1 \exp\left(i\mathbf{k} \cdot \mathbf{r} - i\omega t\right) \tag{3.13}$$

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{E}_0(\mathbf{r}, t) + \mathbf{E}_1 \exp\left(i\mathbf{k} \cdot \mathbf{r} - i\omega t\right) \tag{3.14}$$

$$\mathbf{B}(\mathbf{r}, t) = \mathbf{B}_0(\mathbf{r}, t) + \mathbf{B}_1 \exp\left(i\mathbf{k} \cdot \mathbf{r} - i\omega t\right) \tag{3.15}$$

This Ansatz can be inserted into the Valsov and Maxwell equations and solved for the dispersion relation $\omega(\mathbf{k})$. If the imaginary part of ω is positive the perturbation with wavevector $\mathbf{k}$ grows exponentially and the mode is unstable. Conversely if the imaginay part is negative then a mode is damped even if there is another source of excitation. For a given equilibrium state the calculation can be quite involved and is often only possible using approximation or numerical methods. For a number of beam plasma configurations the solutions have been calculated in the literature [34, 30, 33, 31, 32, 35]. How quickly an unstable mode grows is determined by the growth rate δ. If multiple unstable modes exists the system is dominated by the mode with the highest growth rate.

$$\delta(\mathbf{k}) = \mathrm{Im}(\omega(\mathbf{k})) \tag{3.16}$$

An unstable mode continues to grow exponentially in amplitude until the linear theory is no longer a valid description, because the phasespace density and/or the fields are too different from their original equilibrium state. At this point non-linear physics become relevant, which usually results in the saturation of the unstable mode at a finite amplitude.

3.2 Plasma Beam Instabilities

We are interested in the unstable behavior of a relativistic particle beam that consists of electrons and positrons with equal momentum distribution propagating through a much denser, neutral background medium that consists of electrons and protons. Furthermore we make the simplifying assumption that the background medium is unmagnetized. This leads to the vanishing field equilibrium where $\mathbf{E}_0 = \mathbf{B}_0 = 0$. The most simple assumption for the equilibrium phasespace density would be a monokinetic, or cold, plasma. There the momentum spread of the beam and background medium is negligibly small.

$$f_{0,\mathrm{b}}(\mathbf{p}) = \alpha\delta(p_x)\delta(p_y)\delta(\beta\gamma m_e - p_z) \tag{3.17}$$

$$f_{0,\mathrm{bg}}(\mathbf{p}) = \delta(p_x)\delta(p_y)\delta(p_z) \tag{3.18}$$

The solution in this scenario are multiple competing classes of instabilities defined by different wavevectors where maximum growth of the instability occurs [35]. We note that often the literature discusses the analog case of counterpropagating electron beams instead of copropagating beams of electrons and positrons, however their instable behavior in the absence of an external magnetic field is the same. In this scenario the entire system can be described by two parameters. The Lorentz boost γ of the pair beam and the density ratio α of the beam plasma to the background medium. Since we are interested in the physics of dilute, relativistic beams we can assume $\alpha \ll 1$ and $\gamma \gg 1$.
The Two-Stream instability is an electrostatic instability, that means only electric field modes

are excited, where the unstable modes are parallel to the beam momentum. The Filamentation instability is an electromagnetic instability, that means both electric and magnetic field modes are excited, where the unstable modes are perpendicular to the beam momentum. The Oblique instability is another electrostatic mode where the unstable modes are oblique to the beam momentum. While the growth rates of all modes influence the total evolution of the system, especially during and after the transition to the non-linear regime, the most important influence is the mode with the highest growth rate. For different modes the growth rate of the dominant mode has been calculated in [35].

$$\delta_{\text{Two-Stream}} = \frac{\sqrt{3}}{2^{\frac{4}{3}}} \frac{\alpha^{\frac{1}{3}}}{\gamma} \omega_p \tag{3.19}$$

$$\delta_{\text{Filamentation}} = \beta \sqrt{\frac{\alpha}{\gamma}} \omega_p \tag{3.20}$$

$$\delta_{\text{Oblique}} = \frac{\sqrt{3}}{2^{\frac{4}{3}}} \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \omega_p \tag{3.21}$$

In figure 3.1 the growth rates of the fastest growing mode of these three instabilities are shown for example parameters. For a dilute and relativistic beam the oblique instability will always dominate. Figure 3.2 shows the maximum growth rate of the oblique instability for a wide range of parameters.
The electrostatic oblique modes are Langmuir waves that are excited by the Cherenkov effect [29]. Langmuir waves are electrostatic waves in an unmagnetized plasma with a dispersion relation $\omega(k)$ that is only mildly dependent on k through the electron thermal velocity v_{th}.

$$\omega^2 = \omega_p^2 + 3k^2 v_{\text{th}}^2 \tag{3.22}$$

The Cherenkov effect gives a direct relationship between the frequency and wavevector of a mode and the velocity of the particles that cause the excitation.

$$\omega - \mathbf{k} \cdot \mathbf{v} = 0 \tag{3.23}$$

For a cold and relativistic beam $v_\parallel \sim c$ and $v_\perp \ll v_\parallel$. In a cold background medium the electron thermal velocity will be negligible and ω is not a function of k. Thus the spectrum of excited modes will be a narrow range around $k_\parallel \sim \frac{\omega_p}{c}$ and a wide range in $k_\perp$.
These results have been derived for a negligible beam velocity spread. This regime is sometimes also called the reactive regime. When the velocity spread becomes too large the results are no longer valid and new calculations for this kinetic regime have to be derived. The condition for reactive regime applies to each individual mode and relates the velocity spread and the growth rate [29]. Thus it is also possible to achieve a hybrid regime, where the velocity spread leads to some modes which are reactive while others are kinetic.

$$|\mathbf{k} \cdot \Delta\mathbf{v}| < |\delta(\mathbf{k})| \tag{3.24}$$

For a highly relativistic beam even a large longitudinal momentum spread will lead to a very small velocity spread in this direction. If the instable modes have a significant transversal component, which is the case for the oblique instability, then the left hand side of 3.24 will be dominated

by transversal velocity spread. Furthermore the wavevector will be of the order of $k \sim \frac{\omega_p}{c}$. The condition simplifies to

$$\frac{\Delta p_\perp}{p} < \frac{\delta}{\omega_p} \tag{3.25}$$

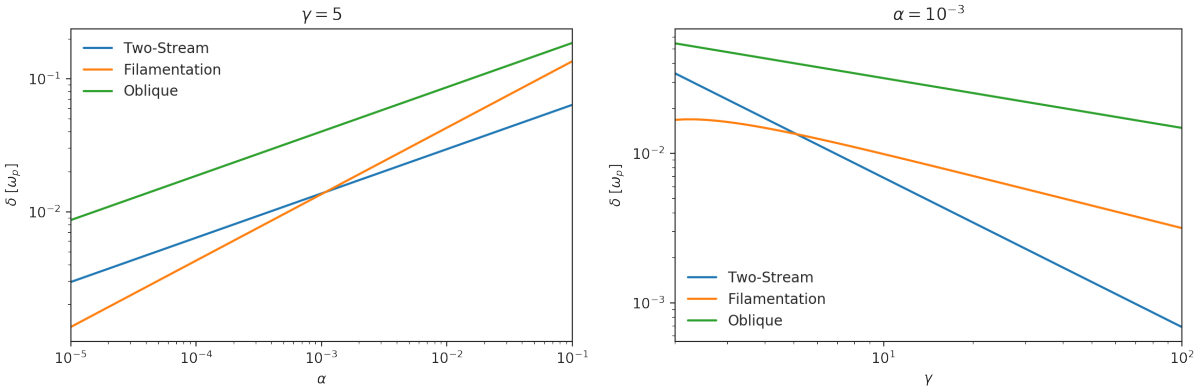

Figure 3.1: Comparison of the growth rate of the respective fastest growing mode as a function of the density ratio α (**left**) and the Lorentz boost γ (**right**) for the Two-Stream, Filamentation and Oblique instability.

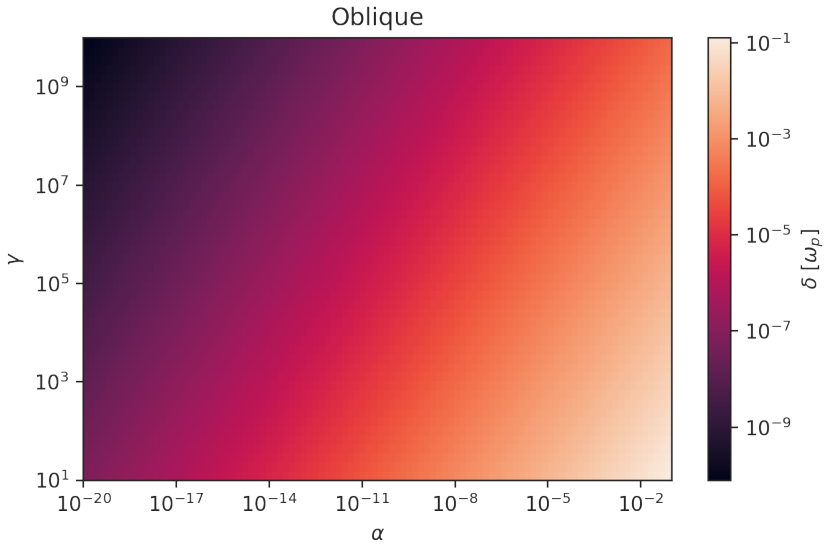

Figure 3.2: Growth rate of the fastest growing mode of the Oblique instability as function of α and γ.

3.3 Analytical Treatments

If the beam does no longer satisfy the conditions for the reactive regime, but nevertheless is highly relativistic and the angular spread as well as the total energy spread are small an analytical

formula for the growth rate can be derived [29].

$$\delta(\mathbf{k}) = -\frac{\alpha}{2k^3}\int_{\mu_-}^{\mu_+} \mathrm{d}\mu \int_1^\infty \mathrm{d}pp \frac{2f(p,\mu) + (\mu - k_\parallel)\frac{\partial f}{\partial \mu}(p,\mu)}{\sqrt{(\mu_+ - \mu)(\mu - \mu_-)}} \tag{3.26}$$

Here μ is the angle relative to the bulk beam momentum and p is the normalized beam momentum (Lorentz factor). The wavevector $\mathbf{k}$ is expressed in units of $\frac{\omega_p}{c}$. The limits of the μ integral are given as a function of the wavevector $\mathbf{k}$.

$$\mu_\pm = \frac{k_\parallel \pm k_\perp\sqrt{k^2 - 1}}{k^2} \tag{3.27}$$

The excited modes are again found in a region around $k_\parallel \sim \frac{\omega_p}{c}$, where the width of the region increases with $k_\perp$. For $k_\parallel \frac{c}{\omega_p} > 1$ the growth rate is positive, whereas for $k_\parallel \frac{c}{\omega_p} < 1$ the modes are damped.
Each excited mode is closely associate to a specific particle velocity via (3.23). The particles of this velocity can gain or lose energy by interacting with this field mode. If the slope of the momentum distribution in this point is positive ($\frac{\partial f}{\partial \mathbf{p}} > 0$) then there are more particles that can lose energy by exciting this mode than there are particles that could gain energy by damping this mode. If the slope of the momentum distribution on the other hand is negative ($\frac{\partial f}{\partial \mathbf{p}} < 0$) then the particles that can damp a specific mode outnumber the particle that can excite it. This gives us the condition that for instability growth, an inversion of the momentum distribution needs to be present. Figure 3.3 sketches how different regions of the momentum distribution lead to growth and damping respectively.
Beyond excitation and damping modes can also interact via scattering with other waves or particles. Wave-particle scattering can alter the momentum of individual particles. The total effect on the ensemble of all beam particles can be described by a momentum diffusion process [29]. Particles can lose and gain energy in all directions. This especially means that a subset of particles can overall gain energy due to the instability mechanism. The diffusion constant is proportional to the mode averaged energy density of the electric field.

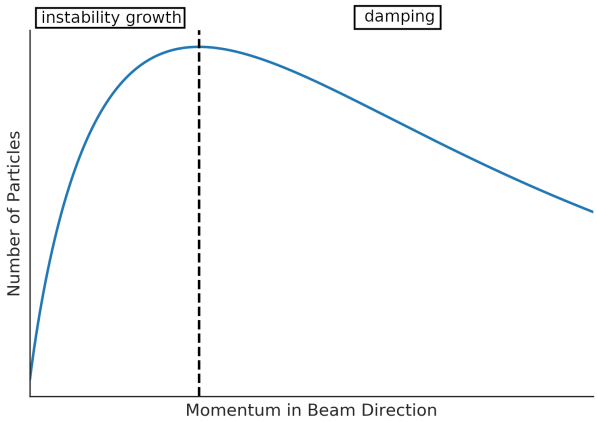

Figure 3.3: Sketch of the effect of the slope of the momentum distribution $\frac{\partial f}{\partial p_\parallel}$ on the instable behavior. A positive slope leads to instability growth, whereas a negative slope leads to damping. For different particle velocities different modes are excited or damped.

3.4 Numerical Treatments

While the fundamental equations that describe the evolution of a plasma are quite simple analytical treatments can only provide useful solutions for a very limited number of problems. To solve these equations for arbitrarily complex problems a number of numerical methods have been developed based on kinetic and/or fluid description of plasmas. In Magnetohydrodynamic (MHD) codes the plasma is described as a fluid by combining electrodynamics with fluid dynamics. MHD is usually applied to large scale physical problems. In kinetic approaches either the Valsov equation is directly discretized and solved numerically (Valsov code) or the phasespace distribution is sampled by numerical particles which are then evolved according to the Maxwell and Lorentz equations. For the particle approach the Particle-In-Cell (PIC) method has become very popular where fields are discretized on a grid to speed up the calculations. Due to the finite number of numerical particles that sample the phasespace distribution particle approaches are usually much more noisy than Vlasov approaches, which typically are numerically much more expensive in multiple dimensions. For some problems it can be useful to take a hybrid approach where some species are treated using a fluid approach while others are treated using a kinetic, usually a particle type, approach. The kinetic methods were originally developed for collisionless plasmas, however there are methods to include the effects of collisions and other particle effects in both the Vlasov as well as the PIC method.
The beam plasma instabilities we investigate in this work are based on perturbations of the momentum distribution of the beam particle, making it necessary to use a kinetic simulation approach, in our case the PIC method, which will be discussed in more depth in the next chapter.

3.5 Momentum Diffusion Model with Fokker-Planck Equation

Previously we have only considered the effect of the beam plasma system on the electromagnetic fields. However a beam instability is characterized by the self-reinforcing back reactions between

fields and beam plasma. Here we will describe two effects of the instability on the beam. First of all due to energy and momentum conservation the beam has to lose energy and momentum when fields are build up. This will lead to a bulk loss of energy and a bulk loss of momentum in beam direction (opposite direction of unstable modes). Secondly the build up of electric fields leads to wave-particle interactions that behave like diffusion in momentum space.
The evolution of the beam momentum distribution can be described by the Fokker-Planck (collisionless Boltzmann) equation.

$$\frac{\partial f(\mathbf{p},t)}{\partial t} = -\frac{\partial}{\partial \mathbf{p}}\left(\vec{V}(t) f(\mathbf{p},t)\right) + \frac{\partial}{\partial \mathbf{p}} D(t) \frac{\partial}{\partial \mathbf{p}} f(\mathbf{p},t) \tag{3.28}$$

Here $\vec{V}(t)$ is called the drift term and describes the bulk momentum loss and $D(t)$ is the diffusion tensor and describes the diffusion of momentum. The diffusion tensor is a function of the spectral field energy density and the drift term is a function of the time derivative of the field energy density. During the linear growth phase both grow exponentially with time.
We will leave the exact determination of the drift term and diffusion tensor aside for now and simply assume that they are known functions of time only and that the diffusion tensor is diagonal. Under these assumptions the Fokker-Planck equation can be solved analytically. For convenience's sake we will only consider a one-dimensional Fokker-Planck equation but our result will generalize to two and three dimensions under assumption that the momentum distribution factorizes.

$$\frac{\partial f}{\partial t} = -V(t)\frac{\partial f}{\partial p} + D(t)\frac{\partial^2 f}{\partial p^2} \tag{3.29}$$

We will use the Ansatz of a Gaussian momentum distribution with a time-dependent mean $\mu(t)$ and standard deviation $\sigma(t)$.

$$f(p,\tau) = \frac{1}{\sqrt{2\pi}\sigma(\tau)} \exp\left(-\frac{(p-\mu(\tau))^2}{2\sigma(\tau)^2}\right) \tag{3.30}$$

Inserting the Gaussian distribution will yield

$$\left[\frac{p-\mu}{\sigma^2}\right]\left(\frac{\partial \mu}{\partial t} - V(t)\right) = \left[\frac{(p-\mu)^2}{\sigma^2} - 1\right]\left(\frac{D(t)}{\sigma^2} - \frac{1}{\sigma}\frac{\partial \sigma}{\partial t}\right) \tag{3.31}$$

Since the solution of this equation has to be independent of p, it can only be solved if the terms in the round brackets both vanish. This means that the equation can be reduced to two rather simple differential equations describing the evolution of $\mu(t)$ and $\sigma(t)$ respectively.

$$\frac{\partial \mu}{\partial t} = V(t) \tag{3.32}$$

$$\sigma\frac{\partial \sigma}{\partial t} = D(t) \tag{3.33}$$

We can give a general solution for these equations.

$$\mu(t) = \mu(t=0) + \int_0^t \mathrm{d}t' V(t') \tag{3.34}$$

$$\sigma(t) = \sqrt{\sigma(t=0)^2 + 2\int_0^t \mathrm{d}t' D(t')} \tag{3.35}$$

We have already remarked that the diffusion term is proportional to the field energy density and that the drift term is proportional to its derivative. During the linear growth phase the field energy density is given by two terms: the initial field (noise) due to for example thermal fluctuation which is constant with time and the exponentially growing field due to the instability.

$$W(t) = W_0 + W_1 \exp(2\delta t) \tag{3.36}$$

W_0 and W_1 are simple (undetermined) constants that both depend on the initial noise of the system and δ is the growth rate of the instability amplitude. The energy density grows with 2δ since it is given by the square of the field amplitude.

$$V(t) \propto \frac{\partial W(t)}{\partial t} \Rightarrow V(t) = 2\delta V_1 \exp(2\delta t) \tag{3.37}$$

$$D(t) \propto W(t) \Rightarrow D(t) = D_0 + D_1 \exp(2\delta t) \tag{3.38}$$

Here we use the undetermined constants V_1, D_0 and D_1. Since the drift term is an energy sink $V_1 < 0$. The explicit evolution of $\mu(t)$ and $\sigma(t)$ during the linear growth phase is thus

$$\mu(t) = \mu(t=0) - V_1 \exp(2\delta t) \tag{3.39}$$

$$\sigma(t) = \sqrt{\sigma(t=0)^2 + 2D_0 t + \frac{D_1}{\delta} \exp(2\delta t)} \tag{3.40}$$

The same equations (with different parameters) can be used to describe further dimensions. We can thus conclude that once the exponential growth of the instability dominates the system the momentum distribution shifts exponentially with rate 2δ and widens exponentially with rate δ. Notably for $\sigma(t)$ there can be an intermediate period where the evolution is governed by the initial noise and thus the width of the momentum distribution grows with the square root of time before eventually the exponential growth overtakes.
However we know that the instability growth does not continue forever and has to stop after some time t_{final} when the field energy saturates at $W(t_{\text{final}})$. When the field saturates the drift term vanishes which implies that $\mu(t)$ is constant and the diffusion term takes on a constant value leading to $\sigma(t) \propto \sqrt{t}$.
In the most general case the saturation will not be instantaneous but field energy density will smoothly transition from exponential growth to a constant value. We model this transition by introducing the parameter κ and replacing δt in the argument of the exponential with an updated expression.

$$\delta t \to -\log_\kappa \left(\kappa^{-\delta t} + \kappa^{-\delta t_{\text{final}}}\right) = -\frac{\log\left(\kappa^{-\delta t} + \kappa^{-\delta t_{\text{final}}}\right)}{\log \kappa} \tag{3.41}$$

This new expression is linear in δt for $t \ll t_{\text{final}}$ and takes on the value δt_{final} for $t \gg t_{\text{final}}$. For $\kappa \to \infty$ the transition is instantaneous. We can use this expression to calculate the evolution of μ and σ into the saturation regime. However this implies that the energy drain from the beam becomes zero in the non-linear regime when the field energy density does not continue to grow. In the non-linear regime it is still possible to transfer energy from the beam to the background plasma with the fields as an intermediate. We can extend this energy transfer by

explicitly assuming a drift term that is proportional to $2\delta W(t)$. In the case of an exponentially growing field energy density this is equivalent to using the derivative.

$$\Delta\mu(t) = -2\delta V_1 \int_0^t \mathrm{d}t' \left(\kappa^{-\delta t'} + \kappa^{-\delta t_{\mathrm{final}}}\right)^{-\frac{2}{\log\kappa}} \tag{3.42}$$

$$\sigma(t) = \sqrt{\sigma(t=0)^2 + 2D_0 t + 2D_1 \int_0^t \mathrm{d}t' \left(\kappa^{-\delta t'} + \kappa^{-\delta t_{\mathrm{final}}}\right)^{-\frac{2}{\log\kappa}}} \tag{3.43}$$

Unfortunately we can not derive the free parameters (V_1, D_0, D_1, t_{final}, κ) except the growth rate from first principle. Nevertheless we can use this model to interpret simulation results. Qualitatively we can deduce that for $t \gg t_{\mathrm{final}}$ $\Delta\mu$ grows linearly with time and σ grows with the square root of time.

4 | Particle-in-Cell (PIC) Method

The Particle-in-Cell method has proven itself as one of the state of the art methods for kinetic simulations of plasmas. The name PIC already implies the two key ideas behind this method:

1. Physical particles are represented by a reduced number of numerical macroparticles that sample the phasespace and are evolved following the equation of motion
2. The electric and magnetic fields are represented on a grid and are evolved following the Maxwell equations

The evolution of both fields and particles are discretized in time. Since the fields and the particles exists in two very different spatial domains (discrete vs. continuous) interpolation schemes have to be applied at every time step to calculate the charge and current densities from the particles and to calculate the Lorentz force acting on each particle from the fields. This additional computational cost for interpolation and calculating the fields comes at the benefit that the total computation time for evolving the particles scales linearly with the number of particles, whereas directly calculating the Coulomb force between all pairs of particles would scale quadratic with time.
In this chapter we give an overview of PIC method and explain how the simulations have to be set up to study the unstable behavior of relativistic beams. First we describe the PIC method in general and then we will focus on the specifics of the PIC code we use in this work: EPOCH.

4.1 PIC Simulations in General

The PIC method can be used to simulate plasma physics in one, two or three spatial dimensions. It is possible to simulate the full three dimensions of momentum and all three directional components of the electric and magnetic fields while still simulating only one or two spatial dimensions. In three dimensions the numerical scheme most closely resembles the actual physics, whereas in one or two dimensions some physical phenomena can not be reproduced. Nevertheless running simulations in one or two dimension can be computationally much cheaper and is sufficient for some problems. For the beam instabilities we want to investigate it is enough to resolve two spatial dimensions, one of which should be parallel to the beam direction.

4.1.1 Simulation Box

The simulation box is the spatial environment the particles and fields exists in. While it is possible to perform the PIC method in an arbitrary geometry we will focus on Cartesian dimensions leading to a simulation box and a grid that are both rectangular. The grid spacing does not

need to be equally spaced in all directions, it could even be unevenly spaced, however we will assume that it is evenly spaced. Then the simulation box and the grid is defined by two numbers per dimension: the length of the simulation box L_i and the number of grid cells N_i. The size of a cell $\frac{L_i}{N_i}$ defines the smallest scale that can be resolved. The plasma wavelength $\frac{c}{\omega_p}$ gives the fundamental length scale of our problem and thus $\frac{L_i}{N_i} \ll \frac{c}{\omega_p}$ is necessary to capture the full physics.
Furthermore the simulation box size and grid spacing also defines the modes that can be resolved.

$$\max(k_i) = N_i \frac{\pi}{L_i} \tag{4.1}$$

$$\Delta k_i = \frac{\pi}{L_i} \tag{4.2}$$

The simulation needs to capture the maximum mode as well as resolve the structure sufficiently. For the oblique instability the relevant modes are found in a narrow region around $k_\parallel \sim \frac{\omega_p}{c}$ and wide range in $k_\perp$. The condition to resolve the plasma wavelength with many grid cells thus also already ensures that the maximum wavevector we can resolve is large enough. We gain the additional requirement that the simulation box has to be long enough to resolve the resonance region.
The simulation box has two boundaries per dimension. The boundary conditions describe how fields and particles behave at these boundaries and have to be tailored to the physical problem. For each boundary we can choose a separate boundary condition from a wide variety of options:

- **periodic:** Fields and particles reaching the edge of the simulation box are wrapped around to the opposing boundary
- **open:** Fields and particles that propagate through the boundary simply leave the simulation
- **reflect:** Particles are reflected at the boundary and fields are clamped to zero (Dirichlet boundary conditions)
- **conduct:** Particles are reflected at the boundary and fields are treated with perfectly conducting boundary conditions (tangential component of electric field is zero)

Of particular interest for simulating large scale astrophysical phenomena are periodic boundary conditions, that essential replicate an infinitely extended system. When simulating relativistic plasma beams without periodic boundary conditions in beam direction special considerations have to be taken. Particles traveling with almost the speed of light will traverse the entire simulation domain in time $t = \frac{L}{c}$. For simulations that want to track the beam over a long period of time this means that also the simulation box has to be increasingly long leading to a high computational cost. A useful alternative is to use a moving simulation box that tracks the beam. by continuously shifting the grid as well as feeding new background particles into the simulation domain at one boundary while dropping particles that exit at the opposing boundary.
The Maxwell equations relate the time derivative of the electric field to the rotation of the magnetic field and vice versa. To make the computations of the rotation easier often a staggered Yee lattice [133] is employed where the charge density ρ is defined on the grid points, the electric field $\mathbf{E}$ and current density $\mathbf{j}$ are defined on the center of the edges of a grid cell and the magnetic field $\mathbf{B}$ is defined in the center of the grid cell surfaces (see figure 4.2 for a sketch of the Yee lattice)

4.1.2 Particle Loading

At the beginning of the simulation the macroparticles that sample the phasespace distribution of the physical particles have to be initialized. A plasma can contain multiple species of particles. For each species we have to define the charge and mass of the physical particles, the phasespace distribution and the number of macroparticles used to sample the phasespace distribution. Each macroparticle is then randomly placed in the simulation box leading to approximately uniform density of macroparticles. For each macroparticle the momentum is determined from the phasespace distribution either by a simple acceptance-rejection method or in the case of a few special momentum distributions (like a Maxwellian or a Maxwell-Jüttner) by transformation from uniform random variables. Then each particle is given a weight that states how many physical particles are represented by a macroparticle.
It can be useful to declare a species of particles immobile, excepting them from the particle push and reducing the computational effort for a time step. This can make sense when one species of particle is much heavier than the others leading to dynamics on longer time scales and smaller length scales that might not be of interest or not resolved by the simulation to begin with.

4.1.3 Particle Shape / Interpolation Schemes

The computational particles exists on continuous space while the derived quantities ρ and $\mathbf{j}$ exist on a discretized grid. It is necessary to give the macroparticles a shape to deposite the charge and current onto the grid. The shape of the particle is determined by a shape function $S(\mathbf{r}-\mathbf{r}')$, where $\mathbf{r}$ is the position of a grid point and $\mathbf{r}'$ is the position of the particle. Some authors distinguish between the particle shape and the weighting function, where the weighting function $S(\mathbf{r}-\mathbf{r}')$ is the result of convolving the particle shape with a zeroth order b-spline. The charge and current densities can then be calculated by summing over all particles N_p of all species.

$$\rho(\mathbf{r}) = \sum_i^{N_p} q_i S(\mathbf{r}-\mathbf{r}_i') \tag{4.3}$$

$$\mathbf{j}(\mathbf{r}) = \sum_i^{N_p} q_i \mathbf{v}_i S(\mathbf{r}-\mathbf{r}_i') \tag{4.4}$$

q_i and $\mathbf{v}_i$ are the charge and velocity of particle i respectively. However directly calculating both ρ and $\mathbf{j}$ will only approximately satisfy the continuity equation (3.8). To satisfy the continuity equation exactly more elaborate density decomposition schemes have to be used that spread the current density of a single particle over all grid cells traversed during a time step [125, 60].
We will only consider shape functions that factorize for higher dimensions.

$$S(\mathbf{r}-\mathbf{r}') = \prod_j^{N_{\text{dim}}} S(r_j - r_j') \tag{4.5}$$

The simplest shape function one could imagine is depositing the entire charge and current density of each particle onto the nearest grid point. This form is called a *tophat*. This shape function is also numerically cheap, however more extended shape functions that spread the charge and current density of each particle on multiple grid points produce smoother charge and current

densities and thus reduce the numerical noise due to a finite number of particles. Usually higher order b-splines are used as shape functions. The *tophat* is a b-spline of order 0. The b-spline b_l of order l can be expressed as a convolution of lower order b-splines or using piecewise polynomial functions [112]. B-splines are a good choice for the particle shape because they are symmetric and have compact support. Compact support means that their value is different from zero only for a small domain and keeps the computational effort low, because only a small number of grid cells near the particle have to be considered for charge and current deposition.

$$b_l(x) = \int_{-\infty}^{\infty} \mathrm{d}x' b_{l-1}(x') b_0(x - x') \tag{4.6}$$

$$b_0(x) = \begin{cases} 1, & \text{for } |x| < \frac{1}{2} \\ 0, & \text{else} \end{cases} \tag{4.7}$$

$$b_1(x) = \begin{cases} 1 - |x|, & \text{for } |x| < 1 \\ 0, & \text{else} \end{cases} \tag{4.8}$$

$$b_2(x) = \begin{cases} \frac{3}{4} - x^2, & \text{for } |x| < \frac{1}{2} \\ \frac{1}{2}\left(\frac{3}{2} - |x|\right)^2, & \text{for } |x| < \frac{3}{2} \\ 0, & \text{else} \end{cases} \tag{4.9}$$

$$b_3(x) = \begin{cases} \frac{1}{2}|x|^3 - x^2 + \frac{2}{3}, & \text{for } |x| < 1 \\ \frac{4}{3}\left(1 - \frac{1}{2}|x|\right)^3, & \text{for } |x| < 2 \\ 0, & \text{else} \end{cases} \tag{4.10}$$

Here x is the distance between a particle and a grid point normalized to the grid spacing. The b-splines of zeroth to third order are plotted in figure 4.1. It is easy to see that for higher order b-splines the particle's charge and current are spread over more grid points. If one distinguishes the particle shape and the particle weighting function then in the case of b-splines the weighting function for a particle with shape b_l is given by the next higher order b-spline b_{l+1}.
When we want to calculate the Lorentz force that acts upon a particle at a given position on continuous space from the discrete fields we need to again employ an interpolation scheme. For reasons of energy conservation it should be the same scheme employed to deposit the particle charge and current onto the grid [93].

$$\mathbf{E}(\mathbf{r}) = \sum_{i}^{N_{\mathrm{dim}}} \sum_{j}^{N_{\mathrm{cell}}} S(\mathbf{r} - \mathbf{r}_{ij}) \mathbf{E}(\mathbf{r}_{ij}) \tag{4.11}$$

$$\mathbf{B}(\mathbf{r}) = \sum_{i}^{N_{\mathrm{dim}}} \sum_{j}^{N_{\mathrm{cell}}} S(\mathbf{r} - \mathbf{r}_{ij}) \mathbf{B}(\mathbf{r}_{ij}) \tag{4.12}$$

Here $\mathbf{r}_{ij}$ is the position of the j-th grid point in the i-th dimension.

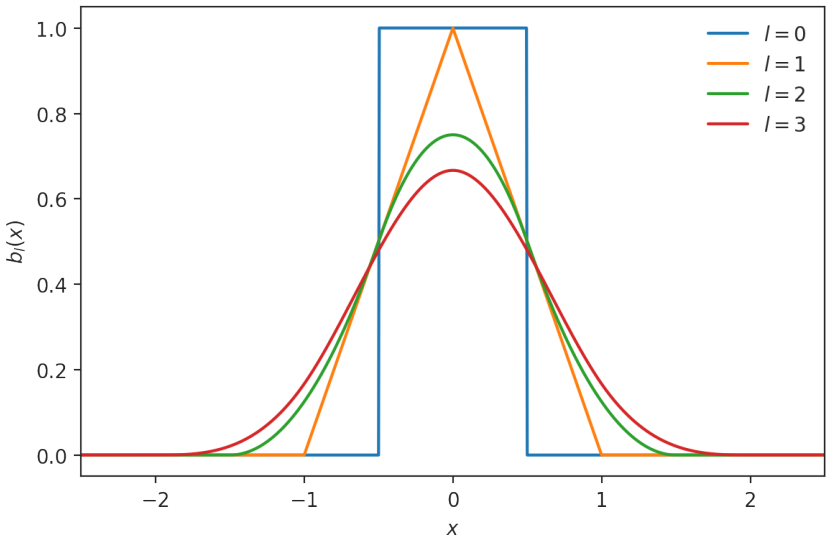

Figure 4.1: B-splines b_l of order l. For higher orders the particles charge and current are deposited on more grid points. x is the distance between a particle and a grid point normalized to the grid spacing.

4.1.4 Field Solver

The field has to be evolved according to Maxwell's equations. Often a finite-difference time-domain (FDTD) method is used to discretize (3.6) and (3.7) in combination with a staggered Yee lattice [18]. The other two Maxwell equations (3.4) and (3.5) are initial conditions. In the case of a neutral beam propagating through a neutral plasma they are fulfilled if both fields initially vanish. Since $\mathbf{E}$ and $\mathbf{B}$ are defined on a staggered Yee lattice calculating the centered second-order accurate derivatives is easy. If we want to calculate the update of $\mathbf{B}$ in position $(i + \frac{1}{2}, j + \frac{1}{2}, k)$ we need to calculate $\nabla \times \mathbf{E}$ at this position for which we need the derivatives of $\mathbf{E}$ in this location. As an example we give the derivative $\frac{\partial E_y}{\partial x}$.

$$\left(\frac{\partial E_y}{\partial x}\right)_{i+\frac{1}{2},j+\frac{1}{2},k} = \frac{(E_y)_{i+1,j+\frac{1}{2},k} - (E_y)_{i,j+\frac{1}{2},k}}{\Delta x} \tag{4.13}$$

Here Δx, Δy, Δz are the grid spacing in the respective direction and Δt is the time step (see next section). It is also possible to include more grid points in the calculation to achieve higher order accurate derivatives [94]. As an example we give the fourth and sixth order accurate derivatives $\frac{\partial E_y}{\partial x}$ which take the field value at four and six positions respectively.

$$\left(\frac{\partial E_y}{\partial x}\right)_{i+\frac{1}{2},j+\frac{1}{2},k} = \frac{1}{24\Delta x}\Bigg(-(E_y)_{i+2,j+\frac{1}{2},k} + 27\,(E_y)_{i+1,j+\frac{1}{2},k} - 27\,(E_y)_{i,j+\frac{1}{2},k} + (E_y)_{i-1,j+\frac{1}{2},k}\Bigg) \tag{4.14}$$

$$\left(\frac{\partial E_y}{\partial x}\right)_{i+\frac{1}{2},j+\frac{1}{2},k} = \frac{1}{64\Delta x}\Bigg(\frac{3}{10}(E_y)_{i+3,j+\frac{1}{2},k} - \frac{25}{6}(E_y)_{i+2,j+\frac{1}{2},k} + 75\,(E_y)_{i+1,j+\frac{1}{2},k} - 75\,(E_y)_{i,j+\frac{1}{2},k} + \frac{25}{6}(E_y)_{i-1,j+\frac{1}{2},k} - \frac{3}{10}(E_y)_{i-2,j+\frac{1}{2},k}\Bigg) \tag{4.15}$$

To update the fields $\mathbf{E}^n$ and $\mathbf{B}^n$ a modified leapfrog scheme is used that calculates the electric and magnetic fields at half time steps $\mathbf{E}^{n+\frac{1}{2}}$ and $\mathbf{B}^{n+\frac{1}{2}}$. The fields at half steps are then used to update the particles position $\mathbf{r}^n$ and momentum $\mathbf{p}^n$ by a full time step and calculate the updated current density $\mathbf{j}^{n+1}$, which is then used to update the fields to the full time step $\mathbf{E}^{n+1}$ and $\mathbf{B}^{n+1}$.

$$\mathbf{E}^{n+\frac{1}{2}} = \mathbf{E}^n + \frac{\Delta t}{2}\left(c^2\nabla\times\mathbf{B}^n - \frac{\mathbf{j}^n}{\epsilon_0}\right) \tag{4.16}$$

$$\mathbf{B}^{n+\frac{1}{2}} = \mathbf{B}^n - \frac{\Delta t}{2}\left(\nabla\times\mathbf{E}^{n+\frac{1}{2}}\right) \tag{4.17}$$

$$\mathbf{B}^{n+1} = \mathbf{B}^{n+\frac{1}{2}} - \frac{\Delta t}{2}\left(\nabla\times\mathbf{E}^{n+\frac{1}{2}}\right) \tag{4.18}$$

$$\mathbf{E}^{n+1} = \mathbf{E}^{n+\frac{1}{2}} + \frac{\Delta t}{2}\left(c^2\nabla\times\mathbf{B}^{n+1} - \frac{\mathbf{j}^{n+1}}{\epsilon_0}\right) \tag{4.19}$$

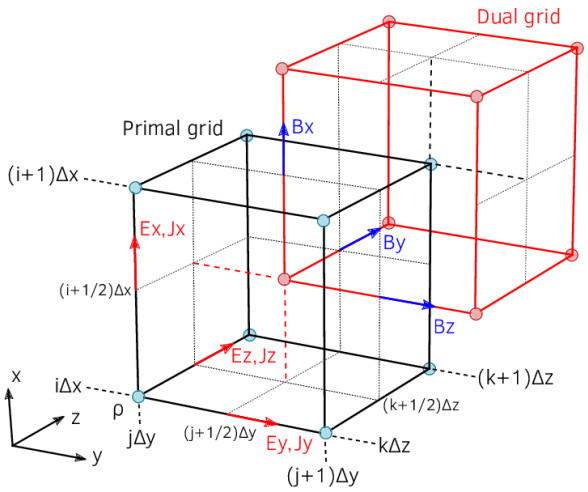

Figure 4.2: Scheme of a Yee lattice where for each simulation cell the charge density ρ is defined on the corners, the current density $\mathbf{j}$ and the electric field $\mathbf{E}$ are defined on the center of the edges and the magnetic field $\mathbf{B}$ is defined in the center of the surfaces [21].

4.1.5 Time Step

The time step Δt is the discretization length in time. The computational effort necessary to run the simulation for a given time is anti-proportional to the length of a time step. Thus it is advisable to choose the time step as large as possible. One upper limit on the time step is given by the time scale of the problem one wants to study. For example the time step should be smaller than the plasma frequency in the case of beam instabilities. However the grid spacing already gives an upper limit on the time step called the Courant-Friedrichs-Lewy (CFL) criterion. When advancing the fields in time the field values from neighboring positions are taken into account. This essentially propagates information through the grid. Information may not travel faster than the speed of light. In one dimension with second order derivatives this leads to the condition

$$c\Delta t \leq \Delta x \tag{4.20}$$

In three dimension with second order derivatives the expression becomes

$$c\Delta t \leq \frac{1}{\sqrt{\Delta x^{-2} + \Delta y^{-2} + \Delta z^{-2}}} \tag{4.21}$$

The CFL criterion directly ensures that any plasma oscillation that is spatially resolved is also temporally resolved. If higher order accurate derivatives are used to calculate the field update then information can travel further than one grid cell in one time step. This leads to a numerical dispersion relation where the phase velocity of modes becomes larger. This increase in phase velocity has to be compensated by the CFL criterion. For fourth and sixth order accurate derivatives the CFL criterion in three dimensions becomes

$$c\Delta t \leq \frac{6}{7} \frac{1}{\sqrt{\Delta x^{-2} + \Delta y^{-2} + \Delta z^{-2}}} \tag{4.22}$$

$$c\Delta t \leq \frac{120}{149} \frac{1}{\sqrt{\Delta x^{-2} + \Delta y^{-2} + \Delta z^{-2}}} \tag{4.23}$$

4.1.6 Particle Pusher

The particle pusher solves the equation of motion to update the position $\mathbf{r}_i^{n+\frac{1}{2}}$ and momentum $\mathbf{p}_i^n$ of the i-th particle using the Lorentz force from the fields at half time to achieve second order accuracy. A leapfrog scheme calculates the position at half time and the momenta at full time.

$$\mathbf{p}_i^{n+1} = \mathbf{p}_i^n + q_i \Delta t \left[\mathbf{E}^{n+\frac{1}{2}} \left(\mathbf{r}_i^{n+\frac{1}{2}} \right) + \mathbf{v}^{n+\frac{1}{2}} \times \mathbf{B}^{n+\frac{1}{2}} \left(\mathbf{r}_i^{n+\frac{1}{2}} \right) \right] \tag{4.24}$$

The velocity $\mathbf{v}_i$ can be calculate from the momentum. The velocity at half time can be calculated by averaging the previous and next velocity.

$$\mathbf{p}_i = \gamma_i m_i \mathbf{v}_i = \left(\sqrt{\left(\frac{\mathbf{p}_i}{m_i c} \right)^2 + 1} \right) m_i \mathbf{v}_i \tag{4.25}$$

The electric field accelerates particle while the magnetic field only reorients the momentum. The most common scheme for particle pushing by Boris [27] separates the two effects by first applying

half of the acceleration due the electric field, computing the rotation due to the magnetic field and then applying the other half of the acceleration. A more advanced scheme by Higuera and Cray [68] uses a different method to average the velocities to achieve higher accuracy for highly relativistic particles.
The leapfrog scheme for calculating the position update can be separated into two steps to also calculate the position at full time $\mathbf{r}_i^{n+1}$ for the calculation of the current density $\mathbf{j}^{n+1}$.

$$\mathbf{r}_i^{n+\frac{3}{2}} = \mathbf{r}_i^{n+1} + \frac{\Delta t}{2}\mathbf{v}_i^{n+1} = \mathbf{r}_i^{n+\frac{1}{2}} + \Delta t \mathbf{v}_i^{n+1} \tag{4.26}$$

4.1.7 Current Smoothing

To reduce numerical noise in the form of high wavenumber modes in the current $\mathbf{j}$ smoothing or filtering schemes can be used. The current at a given position is averaged with the currents at neighboring positions at a distance s [123] to produce the filtered current $\mathbf{j}^f$. The parameter α gives the strength of the damping effect. Applying the same filter multiple times with the same or different strides can damp the noise on different length scales at the same time. Here we give the filter for one, two and three dimensions explicitly.

$$\mathbf{j}_i^f = \alpha \mathbf{j}_i + (1-\alpha)\frac{\mathbf{j}_{i+s} + \mathbf{j}_{i-s}}{2} \tag{4.27}$$

$$\mathbf{j}_{i,j}^f = \alpha \mathbf{j}_{i,j} + (1-\alpha)\frac{\mathbf{j}_{i+s,j} + \mathbf{j}_{i,j+s} + \mathbf{j}_{i-s,j} + \mathbf{j}_{i,j-s}}{4} \tag{4.28}$$

$$\mathbf{j}_{i,j,k}^f = \alpha \mathbf{j}_{i,j,k} + (1-\alpha)\frac{\mathbf{j}_{i+s,j,k} + \mathbf{j}_{i,j+s,k} + \mathbf{j}_{i,j,k+s} + \mathbf{j}_{i-s,j,k} + \mathbf{j}_{i,j-s,k} + \mathbf{j}_{i,j,k-s}}{6} \tag{4.29}$$

The exact choice in the stride and the number of applications has to be carefully tuned to the problem as too much filtering will produce unphysical results, while too little filtering will not reduce the noise.

4.1.8 Output

Storing the position and momentum of all particles as well as the the field values for the entire grid for a single time step can already be quite memory intensive. Thus it is generally not possible to save the entire evolution of the particle population or the fields for every time step and more sophisticated approaches have to be used. We will call a quantity of data that is saved at a single time step or averaged over a range of time steps an observable. An example would be the z-component of the magnetic field at time t_i or the position of all particles of a species at time t_j.
The simplest method to reduce the amount of data that has to be saved is to only write an observable at a select few number of time steps. This could be in regular intervals of $\Delta t_{\text{obs}} \gg \Delta t$ or at a given number of predefined time steps t_1, t_2, et cetera. This can also take the form of only starting to save an observable in regular intervals after a start time t_{start}. However it can be even more efficient to construct observables that take up less memory. The simplest observables take the form of only a single number and can thus be saved as often as one would like. An example would be the total energy density of a field or the total kinetic energy density of a species of particles. This essentially means to average a quantity over the entire simulation domain. In the

case of particle attributes, like the momentum or simply the number, one can average over each grid cell to produce an observable which takes as much storage as saving a field. Another way to average particle quantities is to save their individual data into a histogram. This can also take the form of only averaging over a subset of particles.
A subset of particles is a fraction of the population of a given particle species that can either be chosen randomly or according to some condition. Choosing a random subset can reduce the size of an observable that saves the individual particle attributes at the cost of incurring more noise. It can also be used to only save or only average of particles that fulfill a certain condition, like inhabiting a certain region of interest. A subset that tracks the same particles over multiple time steps is called a persistent subset.

4.2 PIC Simulations with EPOCH

The PIC method has been implemented in a large number of numerical codes. In this thesis we use EPOCH [18], specifically the version 4.17.9, a fully electromagnetic, multi-dimensional, relativistic implementation of the PIC method written in Fortran and parallelized with MPI. The code can be configured using a combination of a configuration file called input deck and compiler flags. A commented example input deck used for simulations discussed in this thesis can be found in appendix C.
EPOCH is available in three separate versions (EPOCH1D, EPOCH2D, EPOCH3D) that allow for the simulation in one, two and three spatial dimensions respectively, while momentum is always resolved in three dimensions. The three versions are only different in this regard and behave the same in all other ways.
EPOCH includes many advanced features like lasers, particle collisions, QED and radiative corrections, ionization and more. These features were not used in this work and will not be discussed here.

4.2.1 Compiler Flags

At compile time a number of flags can be used to configure the code. Many advanced features can be disabled to reduce the computational effort. Two important options that have to be set at compile time are the particle shape and the particle pusher. EPOCH implements both the particle pusher developed by Boris [27] as well as the one developed by Higuera & Cray [68]. Furthermore EPOCH implements tophat, triangle and third order b-spline as particle shape, leading to a particle weighting function that is a b-spline of first, second or fourth order.

4.2.2 Deck variables

The input deck is a text file that allows for simple configuration of the simulation setup. In the input deck it is possible to define constants that allow for simple parameterization of the simulation setup that allows for easy running of an ensemble of simulation. When reading the input deck the code can also parse a limited amount of mathematical expressions.
The configuration options in the input deck include:

- The size of the simulation box as well as the number of grid cells in each dimension.

- The runtime of the simulation can be specified either as number of steps or as a simulation time.
- A moving simulation window can be specified. The simulation window can move at any specified speed and if desired the movement only starts after a specified time.
- The field solver which can be chosen from Yee's method with second, fourth or sixth order accurate derivatives or other methods with second order derivatives that optimize the dispersion relation under some constraint (usually along a specific axis).
- The CFL criterion is automatically calculated for the given field solver and simulation grid. It is possible to designate a multiplier between 0 and 1 that reduces the time step even further.
- Strided current filtering can be applied by specifying the strides and the number of applications after each time step.
- For each particle species the density as a function of position and the momentum distribution can be specified. For thermal particle populations it is also possible to only specify the temperature instead of the entire momentum distribution.
- Outputs can be written in regular intervals, which can be individually chosen for each observable. All observables taken at a time step are written into the same output file in SDF (Self-Describing File) format.

4.3 Units

The EPOCH uses SI units [25] internally. However the fundamental time and length scales of the phenomena we want to investigate are given by the electron plasma frequency. For a laboratory plasma with a density of $n = 10^{16}\,\mathrm{cm}^{-3}$ this plasma frequency is of the order of Terahertz, whereas for an astrophysical plasma with a density of $n = 10^{-7}\,\mathrm{cm}^{-3}$ the plasma frequency is of the order of Hertz. To make our simulations independent of this choice of the background density we normalize all timescales with respect to the plasma frequency and all length scales with respect to the plasma wavelength. We neglect the influence of the dilute pair beam on the total density of the system and only include the density of the background plasma in this transformation.

$$t \to t\omega_p \tag{4.30}$$

$$x \to x\frac{\omega_p}{c} \tag{4.31}$$

Furthermore we also choose to express energies, particle momenta and temperatures in units of energy, with the benefit that for electrons with mass $m_e = 511\,\mathrm{keV}$ the Lorentz boost can be approximately taken as twice the energy in MeV. This defines a set of units we call the "plasma units".

$$\omega_p = c = k_B = 1 \tag{4.32}$$

The benefit of being able to directly compare different density regimes in plasma units makes it necessary to transform to a different set of units if one wants to make predictions for a specific environment.

5 | Cold Beam

In this chapter we will take an in depth look at the instabilities of an idealized, cold e^+e^- beam propagating through an idealized background plasma. A cold beam in this scenario means that the momentum spread is negligible at the start of the simulation. We will assume that the beam and background plasma are infinitely extended and completely homogeneous (up to numerical shot noise). Furthermore the beam shall be perfectly neutral (currentless) and propagate absent any external electromagnetic fields.
The unstable behavior of this system will lead to the build up of electromagnetic fields. But beyond that a number of additional effects can be found, namely the feedback of the instability on the momentum distribution of the beam and the heating of the background plasma. We will be able to show that the instability functions as an energy loss mechanism for the beam, but maybe much more important it also "heats" the beam by increasing its momentum spread.
A number of key parameters with explanations can be found in table D.1.

5.1 Overview of the Simulation Setup

We use EPOCH2D [18] to simulate the cold beam propagating through a background plasma in two Cartesian spatial dimensions to reduce the computational load compared to a full 3D simulation and choose that one dimension should be aligned with the bulk momentum of the beam. This way we can capture the full physics of the longitudinal and the transversal direction. Nevertheless we also performed full 3D simulations at lower resolution to crosscheck this claim. The boundary conditions of the simulation box are chosen as periodic to mimic an infinitely large system. The size of the simulation box and the number of grid points have to be chosen sufficiently large to resolve the unstable modes. Here the dimension of the longitudinal direction is $L_{\parallel,\mathrm{box}} = 500\,\frac{c}{\omega_p}$ and transversal direction is $L_{\perp,\mathrm{box}} = 50\,\frac{c}{\omega_p}$ and for both directions each plasma skindepth is resolved with 8 grid cells. This leads to a maximum wavenumber of $k_{\max} = \pi\frac{N_{\mathrm{cell}}}{L_{\mathrm{box}}} = 8\pi\,\frac{\omega_p}{c}$ in each direction and the resolution of modes of $\Delta k_\parallel = \pi L^{-1}_{\parallel,\mathrm{box}} = \frac{\pi}{500}\,\frac{c}{\omega_p}$ and $\Delta k_\perp = \pi L^{-1}_{\perp,\mathrm{box}} = \frac{\pi}{50}\,\frac{c}{\omega_p}$ respectively.
In the simulations of a neutral beam we use four kinds of particles: Beam electrons and beam positrons with a cold momentum distribution, meaning a low momentum spread, namely a Gaussian with a standard deviation $\sigma_\parallel = \sigma_\perp = 0.5\,\mathrm{keV}$ and a mean value μ that is equivalent to a particle with Lorentz boost γ. Additionally background electrons and background protons with a thermal distribution with a temperature of $T = 200\,\mathrm{eV}\cdot(\gamma-1)\,\alpha$. This definition leads to a fixed ratio of beam energy to background energy of $\epsilon = \frac{m_e}{\frac{3}{2}200\,\mathrm{eV}} \approx 1700$. The mass ratio of protons to electrons takes the physical value ($m_p = 1836 m_e$) and protons are immobile (i.e. the

particle pushing is not applied to them) to reduce computational effort. Per grid cell and per species the momentum distribution is sampled with 25 particles, amounting to a total of $4 \cdot 10^7$ particle per species.
The density of the background plasma is nominally set to a value of $n_{\mathrm{bg}} = 10^{16}\,\mathrm{cm}^{-3}$, however since all relevant quantities are expressed in relation to the plasma frequency this choice does not change the physics. As a point of reference we choose the beam to background density ratio $\alpha = 10^{-3}$ and the Lorentz boost $\gamma = 5$ when varying the other quantity.
During the simulation we use a 6th order Yee-Maxwell solver to update the electric and magnetic fields, a third-order b-spline for depositing charge and current density onto the grid and calculating the force acting on individual particles and the particle pusher by Higuera & Cray [68] to evolve the particle position. Additionally we use strided current filtering with smoothing steps of size 1, 2, 3, 4 and apply this smoothing five times before each particle pushing phase.
During the simulation runs we save the energy density of the electromagnetic fields and the kinetic energy densities of the different particle species every plasm period ω_p^{-1}. The more memory intensive observables of the full electric and magnetic fields and the binned momentum distributions of the different particles species are only saved in intervals of $5\omega_p^{-1}$ initially. After the time $500\omega_p^{-1}$ the fields and momentum distributions are sampled more coarsely with an interval of $50\omega_p^{-1}$.
An overview of the simulation parameters can also be found in table 5.1. The simulation parameters have to be carefully tuned to achieve sensible simulation results. Considerations regarding the numerical convergence can be found in appendix A.

5.2 Energy Density

The key feature of a beam instability is the excitation of electromagnetic field modes. The excitation of a mode takes the form of exponential growth of the amplitude and thus also the energy density of a mode. Usually only a small subset of modes experience significant instability growth. In the case of a cold pair beam we expect the purely electrostatic oblique instability meaning that the instability growth happens for the amplitude of the electric field parallel to the wavevector $\mathbf{k}$ for modes at an oblique angle with respect to the beam direction. However since the instable modes can experience a growth of energy by many orders of magnitude they can usually dominate the energy density of the entire system. This means that after undergoing some time needed to overcome the original energy density of the system (e.g. due to thermal noise or shot noise) the entire energy density grows exponentially.
Due to energy conservation we know that the energy gain of the unstable modes must go hand in hand with energy loss of the pair beam. The energy loss of the beam does not directly translate into the net energy gain of the unstable modes because the modes can dissipate energy to larger scales and they can lose energy by heating the background plasma. Usually dissipation only becomes relevant once the exponential growth of the modes has stopped. On the other hand we find that for a dilute beam and a cold background plasma the field energy density and the background plasma energy density are in equipartition. This suggests that the energy loss of the beam and the energy density of the background plasma energy density will both also grow exponentially with the same rate as the instable modes.

Table 5.1: Overview of the default PIC Simulation settings for the simulation runs.

Settings of the PIC Simulation	
Number of Dimensions	2 (x beam direction, y transverse)
Boundary conditions in x	periodic
Boundary conditions in y	periodic
Box Length $L_{x,\text{box}}$ ($L_{\parallel,\text{box}}$)	$500\,\frac{c}{\omega_p}$
Box Width $L_{y,\text{box}}$ ($L_{\perp,\text{box}}$)	$50\,\frac{c}{\omega_p}$
Number of Cells N_x	4000
Number of Cells N_y	400
Number of particles per cell N_p (per species)	25
Timestep Δt	0.95 CFL-Criterion[a] $\approx 0.068\omega_p^{-1}$
Maxwell Solver	Yee
Field Order	6
Particle Pusher	Higuera & Cray
Particle Shape Function	Third Order B-Spline
Current Smoothing	5-fold (1-2-3-4 steps)
Background particles	Electrons & immobile Protons
Initial Background temperature T	$(\gamma - 1)\,\alpha \cdot$ 200 eV
Beam particles	Electrons & Positrons
Distribution Function	$f(\mathbf{p};\mu,\sigma_x,\sigma_y) \propto \exp\left(-\frac{(p_x-\mu)^2}{2\sigma_x^2} + \frac{p_y^2+p_z^2}{2\sigma_y^2}\right)$
μ [MeV]	$0.511\sqrt{\gamma^2-1}$
σ_x [keV] (alias Δp_x)	0.5
σ_y [keV] (alias Δp_y)	0.5
α	0.001 for γ scan
γ	5 for α scan
Total time T	$5000\,\omega_p^{-1}$

[a] The CFL criterion depends on the grid spacing and the field order. For sixth order field interpolation it is $\frac{120}{149}\frac{\Delta x \Delta y}{\sqrt{\Delta x^2+\Delta y^2}}c^{-1}$

5.2.1 Total Energy Density

The most straight-forward observable[1] of the instability is the growth of the total energy density. Figures 5.1 and 5.2 show the evolution of the total energy density for several simulation runs with varying α and γ. The energy density is normalized to the kinetic energy of the beam. In all simulation runs we find that the energy density starts at a constant noise level. The initial noise level is the same in all simulation runs since the ratio of beam energy density to background energy density ϵ is kept constant by adjusting the background medium temperature. After some time an exponential growth phase sets in until a saturation level is reached.
We can already make some remarks on how the instability evolution is influenced by the major parameters α and γ. For smaller α the growth of the instability is slower, initial noise phase lasts longer and the saturation occurs at a later time and at a lower level. Similarly for simulation runs with larger γ we also find slower instability growth that leads to a longer initial noise phase, later saturation and a lower saturation level.
We find that for a background density of $n_{\mathrm{bg}} = 10^{16}\,\mathrm{cm}^{-3}$ the saturation level of the instability is reached after propagating a distance of the order of centimeters. This would be feasible to reproduce in a laboratory environment. We note that the distance scales with $\sqrt{n_{\mathrm{bg}}}$ and performing the experiment with the same dimensionless parameters but a higher background density leads to saturation over even shorter distances.

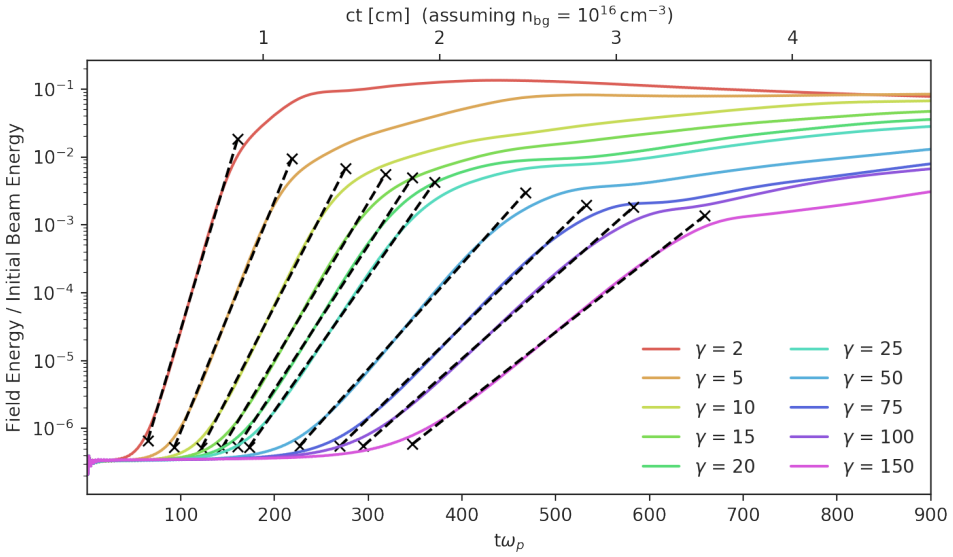

Figure 5.1: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a density contrast of $\alpha = 10^{-3}$ and varying Lorentz boost γ. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

[1] Here we mean a numerical observable that is accessible to us in a simulation run. The energy density of the fields is not a direct observable of a physical experiment.

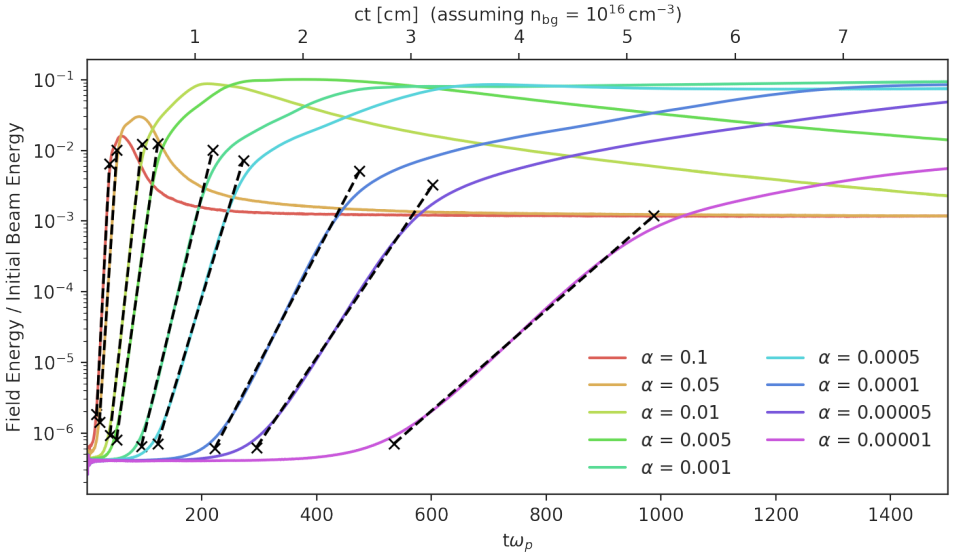

Figure 5.2: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a Lorentz boost of $\gamma = 5$ and varying density contrast α. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

5.2.2 Mode Constrained Energy Density

While we showed that the total energy density can already exhibit significant instability growth we know that this is driven by only a subset of modes that undergo instability growth. On the other hand a large fraction of the initial noise level is due to the initial excitation of modes that do not experience significant growth. We can reduce the initial noise level by only considering the energy density of the modes which have large growth rates. To this end we only consider the energy density of modes around $k_{\parallel} \sim \frac{c}{\omega_p}$ in a range we determined from the Fourier maps discussed in section 5.3. The resulting evolution of the energy density is shown figures 5.3 and 5.4.

We find that this reduces the initial noise by roughly two order of magnitude. For large γ and small α the unstable modes are confined to a smaller region in Fourier space. In these cases the initial noise in figures 5.3 and 5.4 is thus even smaller. Additionally the exponential growth of the instability sets in at earlier times since the energy density of the instability needs to overcome a smaller initial noise level. However we also find that the behavior at late times remains unaffected compared to the total energy density. Most importantly the instability grows with roughly the same rate and saturates at the same time and level.

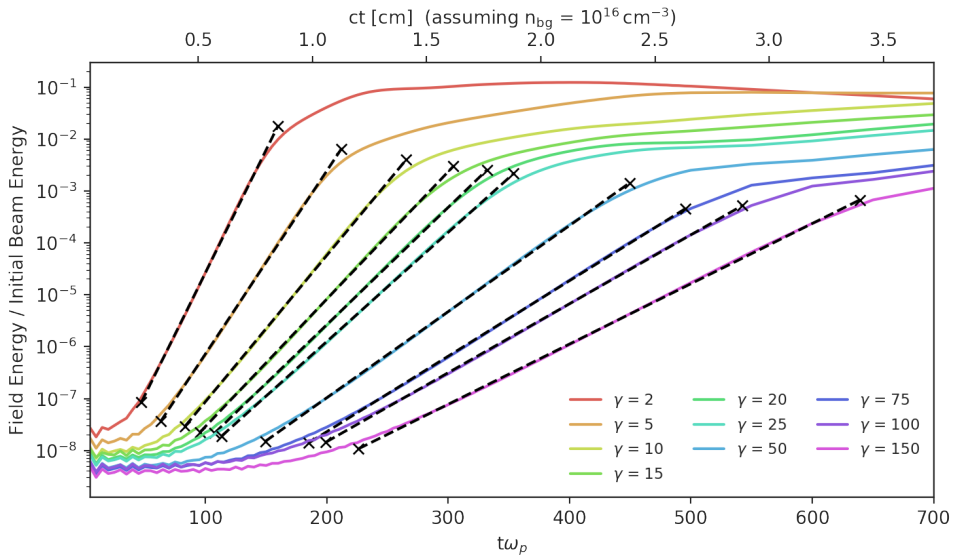

Figure 5.3: The energy density of the electric field modes, constrained to a region around $k_{\parallel} \sim \frac{c}{\omega_p}$ where significant instability growth is expected, as a function of time and normalized to the initial beam energy. Each simulation run has a density contrast of $\alpha = 10^{-3}$ and varying Lorentz boost γ. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

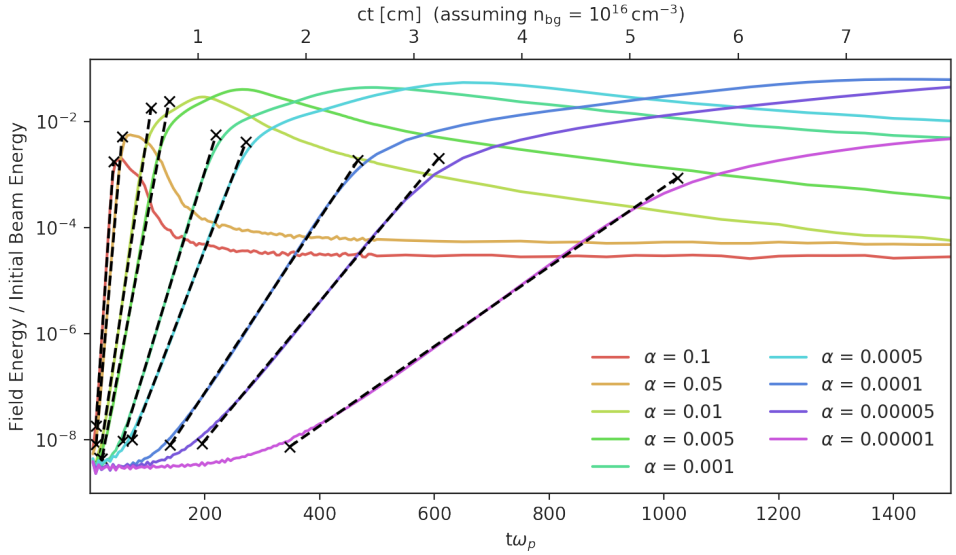

Figure 5.4: The energy density of the electric field modes, constrained to a region around $k_{\parallel} \sim \frac{c}{\omega_p}$ where significant instability growth is expected, as a function of time and normalized to the initial beam energy. Each simulation run has a Lorentz boost of $\gamma = 5$ and varying density contrast α. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

5.2.3 Beam Energy Loss

An observable that is not spoiled at early times by the initial thermal and shot noise is the energy loss of the beam which should undergo exponential growth with the same rate as the field energy density. In figures 5.5 and 5.6 we can see that the exponential growth of the energy loss starts almost immediately at simulation start but otherwise proceeds with roughly the same rate as the growth of the field energy density.
We find that for low γ and high α the beam can lose as much as 10% of it is energy due to the instability, however for larger γ and smaller α the energy loss is reduced. For an astrophysical pair beam this low saturation level could mean that the suppression of the secondary gamma ray flux would be small.

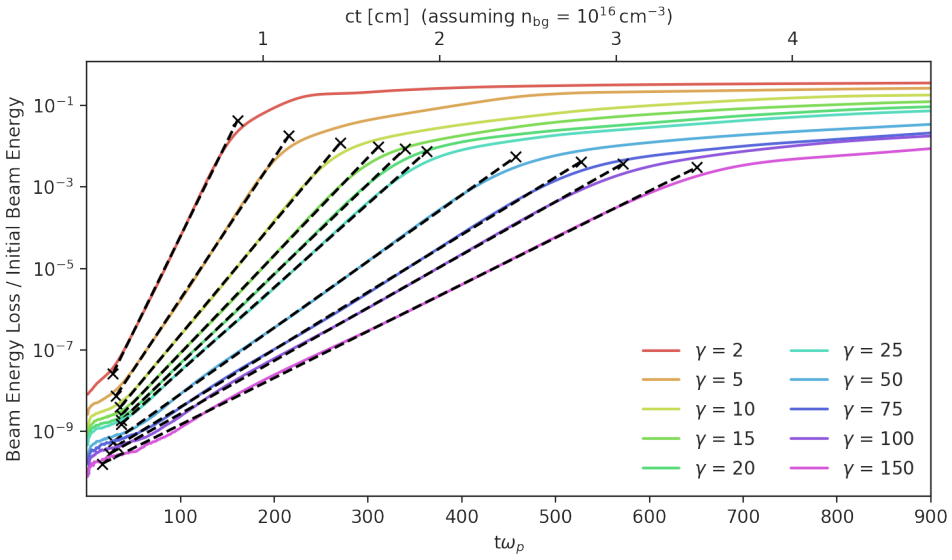

Figure 5.5: The energy loss of the pair beam normalized to the initial beam energy as a function of time. Each simulation run has a density contrast of $\alpha = 10^{-3}$ and varying Lorentz boost γ. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

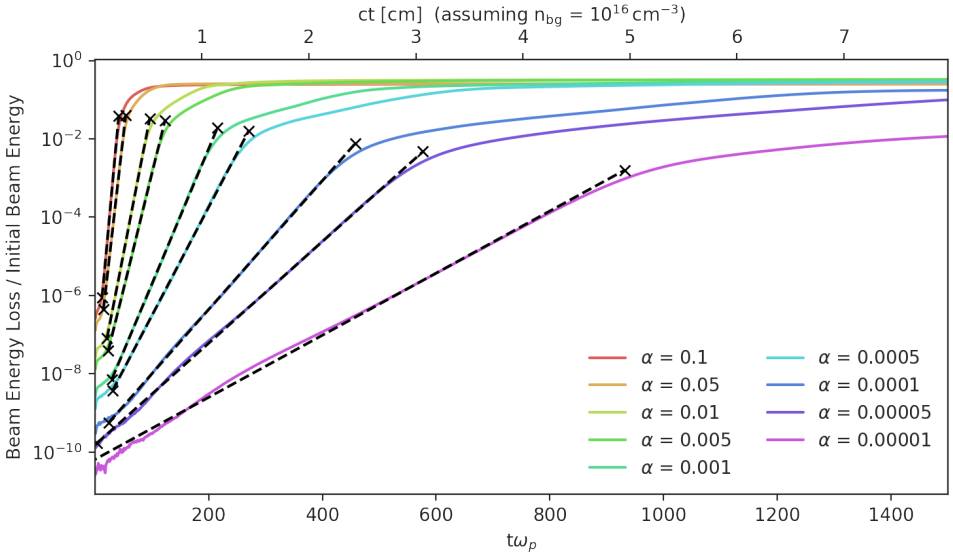

Figure 5.6: The energy loss of the pair beam normalized to the initial beam energy as a function of time. Each simulation run has a Lorentz boost of $\gamma = 5$ and varying density contrast α. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

5.2.4 Energy Budget

For a relativistic beam propagating through a cold background medium almost the total energy of the system is concentrated in the beam. The instability is a mechanism for the beam to lose energy to the electromagnetic fields, which in turn can heat the background plasma. An important question is how this energy allocation changes due to the instability.
Figures 5.7 and 5.8 show the energy budget of different simulation runs. We can see that in the limit of a very dilute beam the energy of the fields and background plasma are roughly in equipartition. For less dilute beams the energy of background plasma can exceed the energy of the fields by multiple orders of magnitude. This means that is not just a matter of the number density ratio α, but also a matter of the energy density ratio. This suggests that a more appropriate condition for a dilute beam would be $\alpha\gamma \ll 1$. This is a stronger condition than $\alpha \ll 1$. Nevertheless even for the combination of $\gamma = 5$ and $\alpha = 10^{-3}$ (see figure 5.7 on the left) the condition $\alpha\gamma \ll$ is satisfied but we still observe that the beam is not dilute enough for the energy of field and background plasma to reach equipartition.
We also find that most of the energy transfer happens after the end of the linear growth phase when the system continues to evolve non-linearly for some time. Critically the beam only looses as sizable fraction of its energy in the case of a low-energy, high density beam. With a larger Lorentz boost γ and a lower density ratio α the fraction of the beam energy that is lost via the instability reduces.

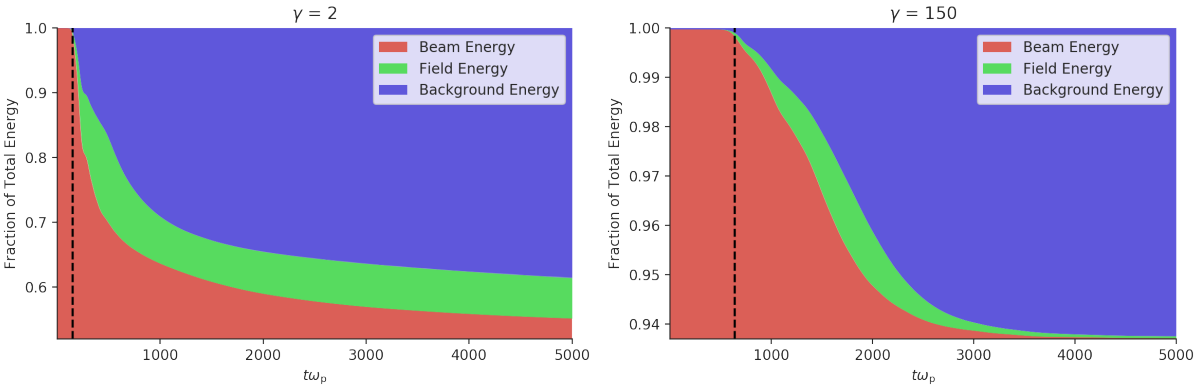

Figure 5.7: The allocation of the (fractional) energy on the beam plasma, the fields and the background plasma as a function of time. On the **left** a simulation run with $\gamma = 2$ is shown, on the **right** a simulation run with $\gamma = 150$. In both cases $\alpha = 10^{-3}$. The end of linear growth phase is indicated by a dashed vertical line.

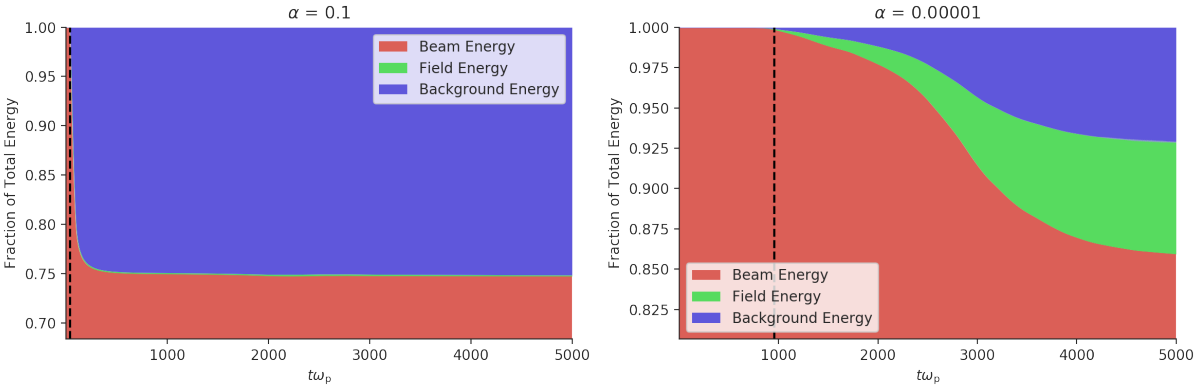

Figure 5.8: The allocation of the (fractional) energy on the beam plasma, the fields and the background plasma as a function of time. On the **left** a simulation run with $\alpha = 10^{-1}$ is shown, on the **right** a simulation run with $\alpha = 10^{-5}$. In both cases $\gamma = 5$. The end of linear growth phase is indicated by a dashed vertical line.

5.3 Fourier Maps

The growth rate is a characteristic of an individual field mode. Some modes experience very fast growth leading to very high amplitude whereas many other modes experience little to no growth or even damping. To determine the growth rate of each individual mode we employ the Fourier transformation of the fields. Let $\tilde{E}_i$ be the Fourier transform of the electric field component E_i. The power of mode is then given by the square of Fourier amplitude.

$$\tilde{E}_i = \mathcal{F}(E_i) \tag{5.1}$$

$$P(E_i) = |\mathcal{F}(E_i)|^2 \tag{5.2}$$

For electrostatic modes the amplitude of the excited field modes is parallel to the wavevector $\mathbf{k}$ Then we can calculate the power in an electrostatic mode $P(\mathbf{E} \parallel \mathbf{k})$ by transforming the

individual field components and adding them weighted by the wavevector.

$$\begin{aligned}\frac{\mathrm{d}P\left(\mathbf{E} \parallel \mathbf{k}\right)}{\mathrm{d}k} = \Bigg[&\cos\left(\arctan\left(\frac{k_\perp}{k_\parallel}\right)\right)\left|\frac{\mathrm{d}\tilde{E}_\parallel\left(k_\parallel, k_\perp\right)}{\mathrm{d}k}\right| \\ &+\sin\left(\arctan\left(\frac{k_\perp}{k_\parallel}\right)\right)\left|\frac{\mathrm{d}\tilde{E}_\perp\left(k_\parallel, k_\perp\right)}{\mathrm{d}k}\right|\Bigg]^2\end{aligned} \tag{5.3}$$

Here electric fields are give in units of $\frac{\omega_p c m_e}{e}$. For each dimension the Fourier transform adds a factor of $\frac{c}{\omega_p}$ and thus the two-dimensional Fourier transform of the electric field is given by $\frac{c^3 m_e}{\omega_p e}$. The power of each mode is then given in units of $\left(\frac{c^3 m_e}{\omega_p e}\right)^2$.
The power of a given mode is proportional to its energy density and thus also grows exponentially during the linear growth phase. This can also be observed in figures 5.9 and 5.10 showing the differential power spectrum. We find that the most dominant modes are located in a narrow region around $k_\parallel \sim \frac{c}{\omega}$ while existing in a wide range of $k_\perp$. Interestingly at the beginning of the instability growth the most dominant modes can be found at an oblique angle with $k_\perp > k_\parallel$, but during the non-linear phase of the instability modes with smaller $k_\perp$ become more dominant. We will see that the momentum distribution widens during the linear growth phase. This will transition the instable behavior of the beam from the reactive into the kinetic regime suppressing the instability growth. According to (3.24) modes with larger $k_\perp$ transition more quickly into the kinetic regime. This indicates that the saturation of the instability we observe is caused by the widening of the momentum distribution.
We can determine the growth rate of each individual mode by fitting a linear slope to the logarithm of the power against time. We use the saturation time as determined by the fit to the total energy density to determine the fit interval for the exponential growth. We only consider time steps between the start of the simulation and the middle of the linear growth phase.
A comparison of the growth rate maps for different α and γ can be found in figures 5.11 and 5.12. Here we find that not only the growth rate of the fastest growing mode is reduced with higher Lorentz boost γ and lower density contrast α, but also the spectral width becomes narrower. Since the phase space of the instable modes is intimately related to the energy they carry we can speculate that this leads to the lower saturation level of these scenarios and ultimately less energy drain from the beam.
We calculate the width of the resonance for the constrained energy density by taking the full width half maximum (FWHM) in $k_\parallel$ of the growth rate $\delta\left(k_\parallel, \hat{k}_\perp\right)$ for a specific $\hat{k}_\perp$, which is the $k_\perp$ of the fastest growing mode. This is the region of Fourier space that we considered for the mode constrained energy density.

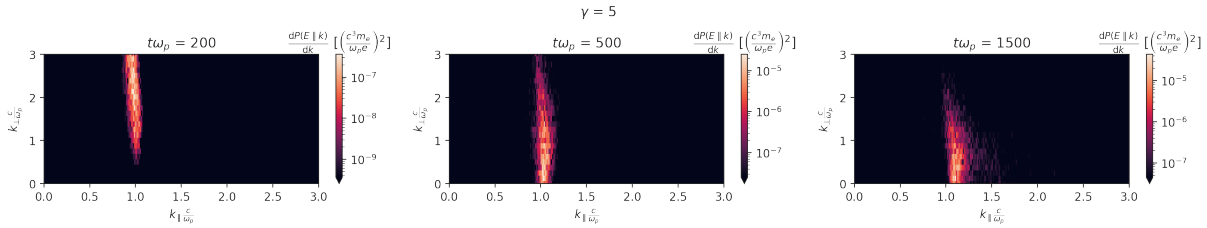

Figure 5.9: The differential power spectrum for a simulation run with $\gamma = 5$ and $\alpha = 10^{-3}$ at three different time steps. The modes with significant growth can be found at $k_\parallel \sim \frac{c}{\omega}$. From the **left** to the **middle** plot the power of the dominant mode grows by two decades. From the **middle** to **right** plot the instability has saturated and the power does not continue to grow.

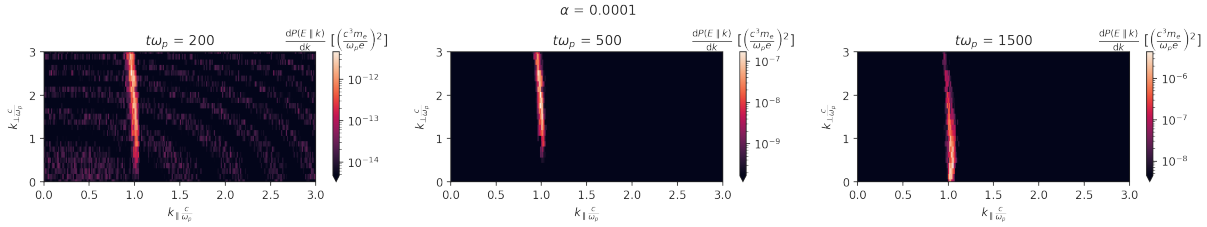

Figure 5.10: The differential power spectrum for a simulation run with $\gamma = 5$ and $\alpha = 10^{-4}$ at three different time steps. The modes with significant growth can be found at $k_\parallel \sim \frac{c}{\omega}$. From the **left** to the **middle** plot the power of the dominant mode grows by two decades. From the **middle** to **right** plot the instability has saturated and the power does not continue to grow.

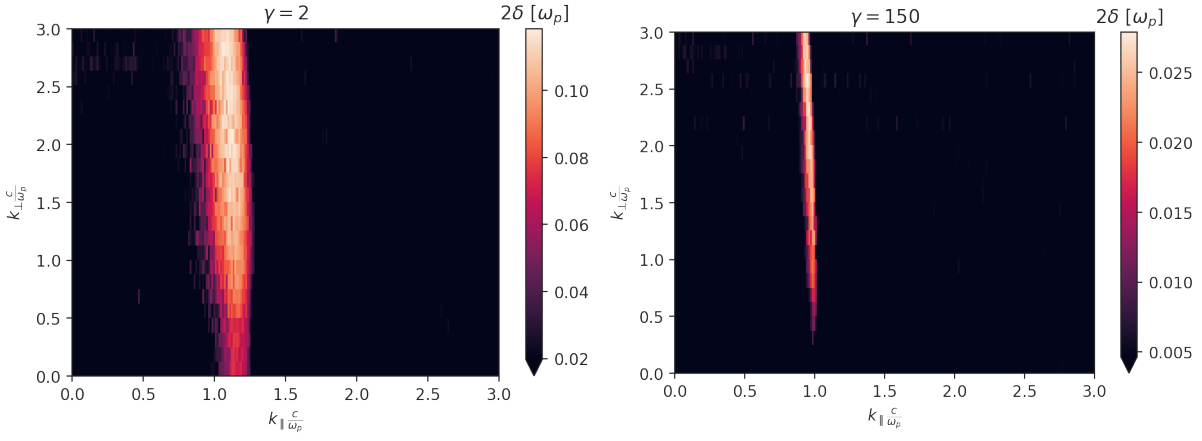

Figure 5.11: Map of the growth rate as a function of wavevector for each mode for two simulation runs. On the **left** the Lorentz boost $\gamma = 2$ is lower than on the **right** where $\gamma = 25$. In both cases the density contrast $\alpha = 10^{-3}$ is the same. For a higher Lorentz boost the maximum growth rate is lower and also the width of the resonance is much smaller.

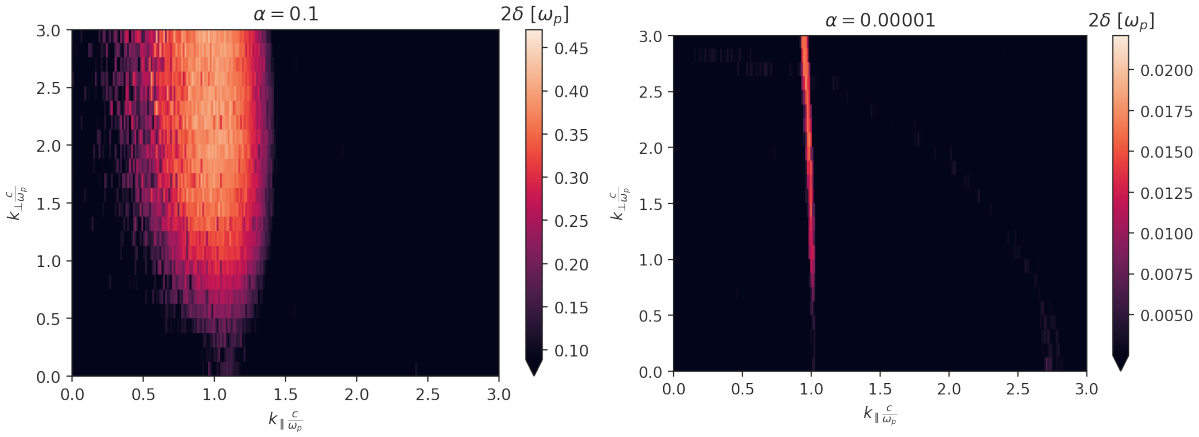

Figure 5.12: Map of the growth rate as a function of wavevector for each mode for two simulation runs. On the **left** the density contrast $\alpha = 10^{-1}$ is lower than on the **right** where $\alpha = 10^{-5}$. In both cases the Lorentz boost $\gamma = 5$ is the same. For a lower density ratio the maximum growth rate is lower and also the width of the resonance is much smaller.

5.4 Growthrate

In the previous section we have described multiple scenarios for quantifying the exponential growth of a beam instability. First of all there is an ambiguity when defining the growth rate between the amplitude of a mode and its energy. The energy of a mode is calculated from the square of the amplitude introducing a factor of 2 in the rate. Here we will always express the rate with respect to the field amplitude. When calculated from the energy density we will divide by 2 to make the rates comparable.
The most obvious definition of the growth rate is the differential growth rate of each individual mode. If one is interested in the evolution of the entire system one can consider the growth rate of the fastest growing, most dominant mode $\delta_{\max}$.
The evolution of the field in its entirety is given by the integrated growth rate δ_{int} which is calculated by mode-averaging the differential growth rate. Thus the integrated growth rate is necessarily smaller than the maximum growth rate. Directly related to the energy gain of the fields is the energy loss of the beam. In the limit of a dilute beam, where field energy and background plasma energy are in equipartition, the loss rate δ_{loss} should be equal to the integrated growth rate.
Between the integrated and the maximum growth rate we can consider many growth rates achieved by mode averaging only in different constrained regions of phase space around the most dominant modes. This gives us the constraint growth rate $\delta_{\text{constraint}}$.
All of these definitions of the growth rate should only deviate slightly from another and follow an approximate hierarchy:

$$\delta_{\text{loss}} \approx \delta_{\text{int}} \lesssim \delta_{\text{constraint}} \lesssim \delta_{\max} \tag{5.4}$$

In all cases the growth rate can be extracted without further assumptions from simulation data by fitting a ramp function (5.5) to the time evolution of the logarithm of the quantity in question.

$$\log u(t;c) = \begin{cases} c_0 & \text{if } t < c_1 \\ c_2 & \text{if } t > c_3 \\ \dfrac{c_2 - c_0}{c_3 - c_1}(t - c_1) + c_0 & \text{otherwise} \end{cases} \tag{5.5}$$

Here c_0 is the initial noise level (if applicable), c_1 is the time the instability emerges from the background noise, c_2 is the saturation level and c_3 is the saturation time of the instability growth. The growth rate is given by $\delta = \frac{c_2 - c_0}{c_3 - c_1}$. This ramp function-fit can be applied to the total energy density, constraint energy density, the amplitude of individual modes and other quantities that undergo a phase of exponential growth between two more or less constant phases like the mean and width of the beam momentum distribution.
We performed this fit for the different definitions of the growth rate and show the results as a function of α and γ in figure 5.13. We find that the definitions of the growth rate only deviate slightly from each other and that our prediction of the hierarchy is fulfilled with $\delta_{\max}$ always being the highest rate. In the figures 5.1 to 5.6 we indicated the fitted linear growth phase by a dashed line, as well as the time of emergence and the time of saturation by a cross.

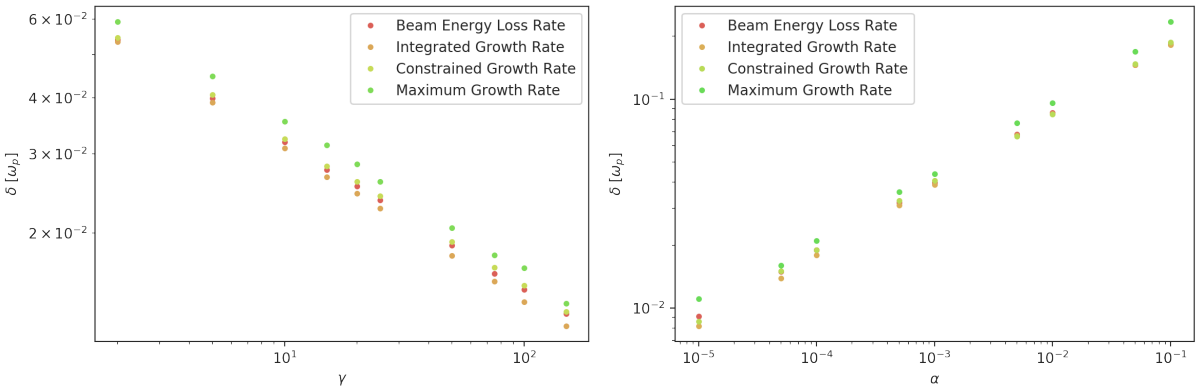

Figure 5.13: The different definitions of the growth rates compared for simulation runs with various α and γ choices. For all runs on the **left** the density contrast is kept at a constant value of $\alpha = 10^{-3}$. For all runs on the **right** the Lorentz boost is kept at a constant value of $\gamma = 5$.

5.4.1 Scaling Relationships

An important factor of laboratory astrophysics is scaling relationships to connect the laboratory system to the much different astrophysical environment. There the energy of the beam particles will be much higher and their density ratio with respect to the background medium much smaller than in the laboratory system. Furthermore the distances and time scales will be many orders of magnitude larger than in the laboratory.
The scaling of time and distances is trivial since all quantities are expressed with respect to the plasma frequency and the plasma skin depth. For α and γ we find that the growth rate follows

a power-law and confirms to theory calculations for the oblique growth rate.

$$\delta \propto \frac{\alpha^{0.34}}{\gamma^{0.34}} \approx \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \tag{5.6}$$

The scaling with α and γ as well as the fitted power laws are shown in figure 5.14 for the integrated growth rate. The scaling of other definitions of the growth rate follow the same behavior.

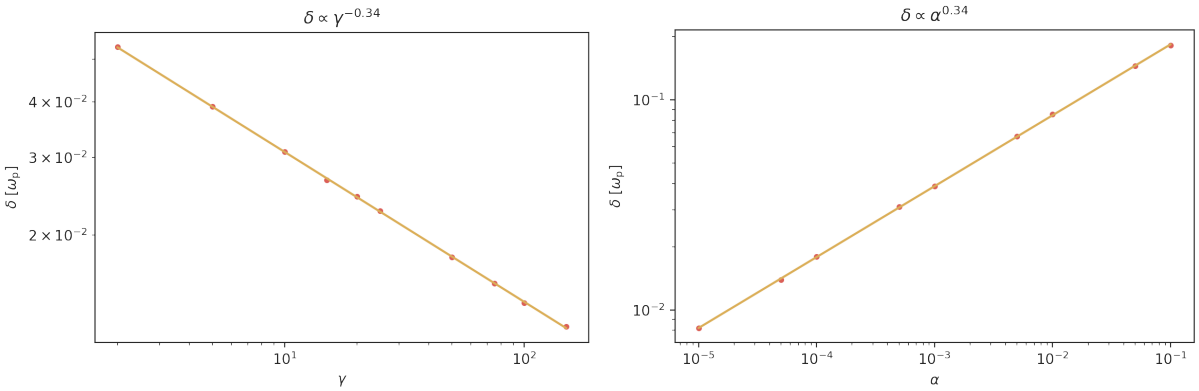

Figure 5.14: The integrated growth rate as a function of Lorentz boost (**left**) and the density contrast (**right**). Both follow a power-law. The spectral index has been fitted to be 0.34 for the α scaling and -0.34 for the γ scaling.

5.5 Beam Momentum Distribution

The beam instability is a self-reinforcing effect of small perturbations of the beam's homogeneity, where small changes in charge density or momentum distribution lead to a net charge or current that produces electromagnetic fields which in turn act on the beam in a way that reinforces these perturbations. So far we have mostly considered the effect of the beam on the fields, but hand in hand with field growth we should find a back reaction of the fields on the beam. This can for example come in the form of the bulk energy loss or a heating of the beam.
In the figures 5.15 to 5.18 we have plotted the evolution of the momentum distribution for a number of simulation configurations. For easier visibility we only show the momentum distribution after some broadening has occurred. Initially the momentum distribution is cold, almost equivalent to a Delta-function, with longitudinal and transversal width of $\sigma_\parallel = \sigma_\perp = 0.5\,\mathrm{keV}$. In all cases we find that instability leads to significant broadening of the momentum distribution in both directions. Furthermore also the peak of the $p_\parallel$ distribution shifts to lower momentum. The mean of the $p_\perp$ remains at 0. This effect does not continue indefinitely but at late times the distribution function settles into a final state. We can broadly say that the momentum distribution function remains Gaussian during most its evolution in the linear growth phase. However especially at late times large deviations from a Gaussian distribution are possible. Comparing the figures we also find that a larger Lorentz boost γ and a larger density ratio α lead to more pronounced shift and broadening of the momentum distribution.
A Gaussian distribution is well defined with only two numbers: the mean μ and the standard deviation σ. Since the momentum distribution of the beam particles closely resembles a Gaussian

we will use the first and second momentum of these distributions to describe them. When running the simulation we write an output file with the binned momentum distribution every $5\omega_p^{-1}$ until a total simulation time of $500\omega_p^{-1}$, then we continue to write with an interval of $50\omega_p^{-1}$. Using the binned momentum distribution, we calculate the mean and standard deviation. Due to the binning this can lead to a minor inaccuracy especially when considering the shifting of the mean value which does not start from a finite value to begin with. For the quantities calculated from the full distribution of p_i we will use the names $\mathrm{Mean}(p_i)$ for the mean and $\mathrm{Std}(p_i)$ for the standard deviation to indicate that they are values extracted from the simulation. In section 3.5 we already derived a model for the evolution of a Gaussian momentum distribution under exponential growth of the field energy density with the Fokker-Planck equation. We showed that the evolution can be fully described by the evolution of the mean and standard deviation.
For symmetry reasons the mean of $p_\perp$ should not shift from 0 beyond minor fluctuations. We are left with three quantities that evolve due to the instability growth: $\mathrm{Mean}(p_\parallel)$, $\mathrm{Std}(p_\parallel)$ and $\mathrm{Std}(p_\perp)$. However instead of the mean we will consider the shift of the mean $\Delta\mathrm{Mean}\big(p_\parallel\big)$ which is constructed to grow exponentially like the energy density. For each quantity the evolution depends on α and γ.

$$\Delta\mathrm{Mean}\left(p_\parallel, t\right) = \mathrm{Mean}\left(p_\parallel, t=0\right) - \mathrm{Mean}\left(p_\parallel, t\right) \tag{5.7}$$

The evolution of $\Delta\mathrm{Mean}(p_\parallel)$ is shown in figures 5.19 and 5.20 for varying γ and α respectively. We find the typical exponential growth associated with the instability. As expected from earlier analytic considerations the slope of the exponential growth is the same as for the energy density of electromagnetic fields, namely 2δ. At very early times the quantity is dominated by numeric inaccuracy because the mean is calculated from a binned momentum distribution. The evolution of the mean momentum shift was further analyzed by fitting the model derived from the Fokker-Planck equation, which describes the evolution during the linear regime and the early non-linear regime very well. The very late evolution after twice the saturation time is excluded. The fit parameters for each simulation run can be found in table **??**.
Similarly the evolution of $\mathrm{Std}(p_\parallel)$ and $\mathrm{Std}(p_\perp)$ is shown in figure 5.21 to 5.24. Here we find a more complex behavior starting from the initial value (here 500 eV) and eventually reaching an exponential growth regime with the slope of the field amplitude. Again the model derived from the Fokker-Planck equation was fitted with success. The fit parameters for each simulation run have been listed in tables 5.3 and 5.4. The smoothness parameter κ defines how quickly the transition from the linear regime with exponential growth to the saturation regime takes place. For $\kappa \to \infty$ the saturation is instantaneous and consequently for smaller κ the transition takes longer. For $\Delta\mathrm{Mean}\big(p_\parallel\big)$ and $\mathrm{Std}\big(p_\parallel\big)$ we find no clear dependency of κ with respect to α and γ and κ takes on a value roughly between 1 and 50. In the case of $\mathrm{Std}(p_\perp)$, where the fit shows the best agreement with the simulation data, we find that the fit prefers larger κ and thus a faster transition for larger γ and larger α. This is perhaps remarkable since a larger γ leads to a longer saturation time, but a larger α on the other hands leads to a shorter saturation time. Instead of the time scale of the instability, the smoothness of the transition seems to follow the behavior of the saturation level, with a higher saturation being associated with a smoother transition.
The model derived from the Fokker-Planck equation also describe the evolution of the momentum distribution into the non-linear regime and we have shown that for some time after saturation that is at least as long as the duration of the linear growth phase this description holds true. After t_{final} the mean continues to shift linearly with time and the standard deviation grows with the

square root of time. At saturation time the drift term $V(t)$ and diffusion constant $D(t)$ take on their maximum value and the momentum distribution can still change very strongly. At very late times the model does not hold true anymore since the electric field evolves independently, mostly by shifting energy to more long-ranged modes that do not contribute to the diffusion constant. At this point $V(t)$ and $D(t)$ decrease again and the evolution of the momentum distribution essential halts.

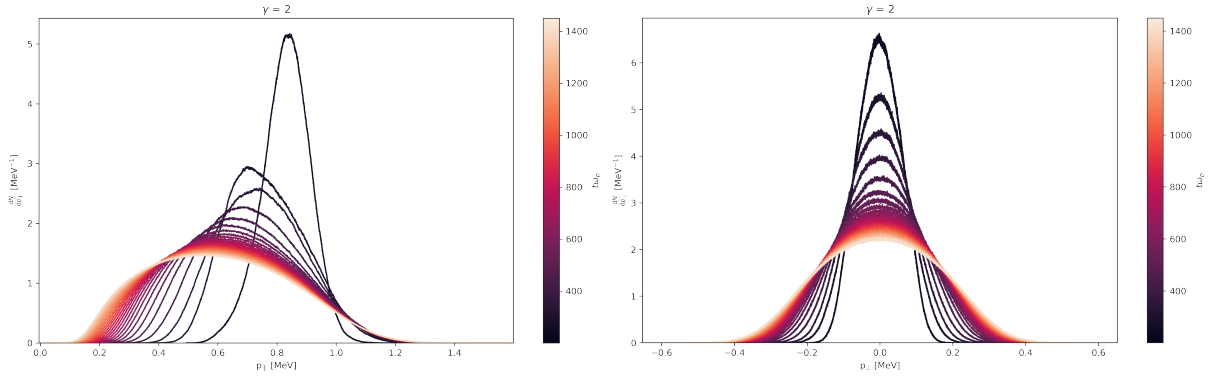

Figure 5.15: The beam momentum distribution for a simulation run with $\gamma = 2$ and $\alpha = 10^{-3}$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **right** plot shows the transversal momentum distribution. In both cases all other momentum directions were integrated over.

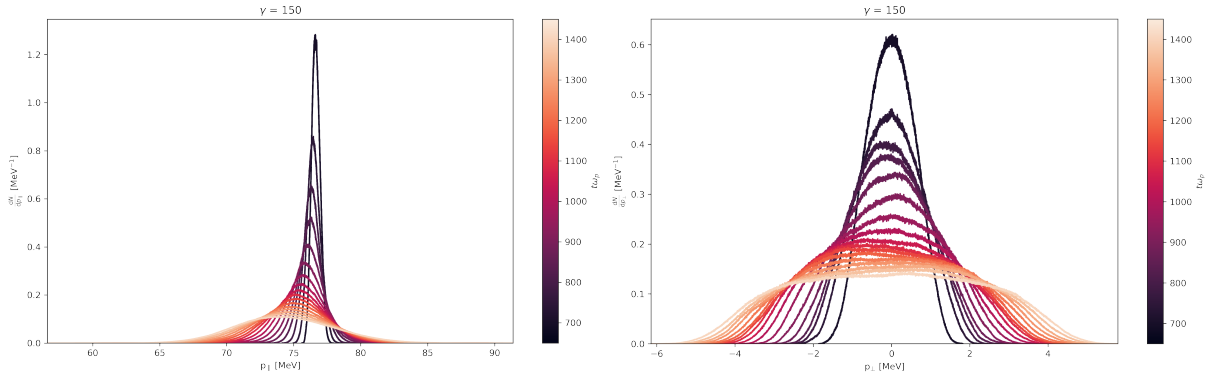

Figure 5.16: The beam momentum distribution for a simulation run with $\gamma = 150$ and $\alpha = 10^{-3}$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **right** plot shows the transversal momentum distribution. In both cases all other momentum directions were integrated over.

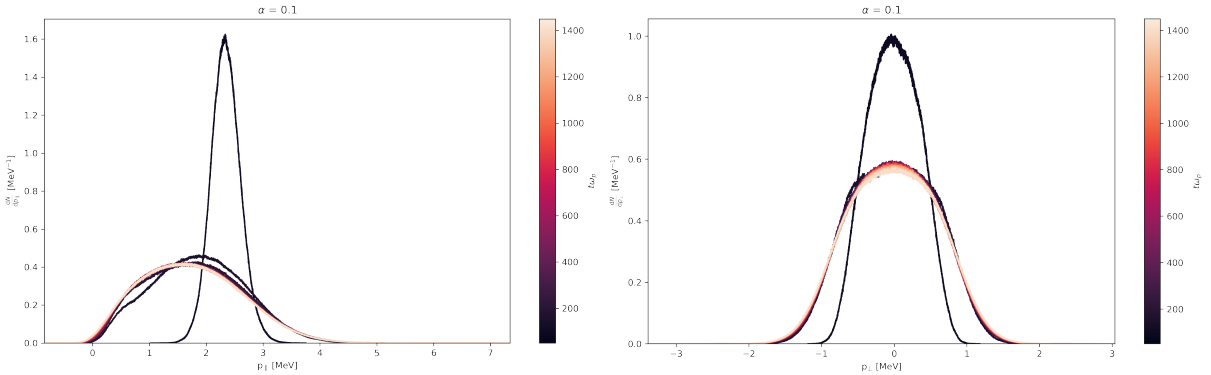

Figure 5.17: The beam momentum distribution for a simulation run with $\gamma = 5$ and $\alpha = 10^{-1}$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **right** plot shows the transversal momentum distribution. In both cases all other momentum directions were integrated over.

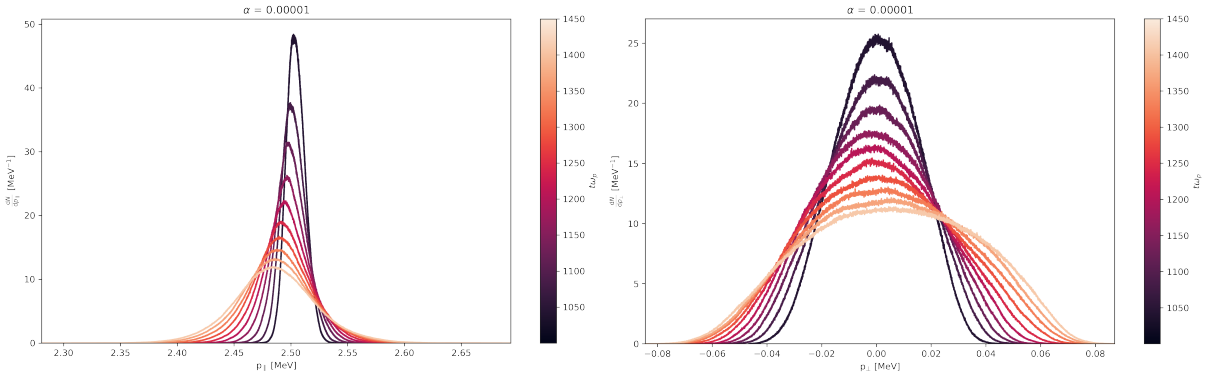

Figure 5.18: The beam momentum distribution for a simulation run with $\gamma = 5$ and $\alpha = 10^{-5}$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **right** plot shows the transversal momentum distribution. In both cases all other momentum directions were integrated over.

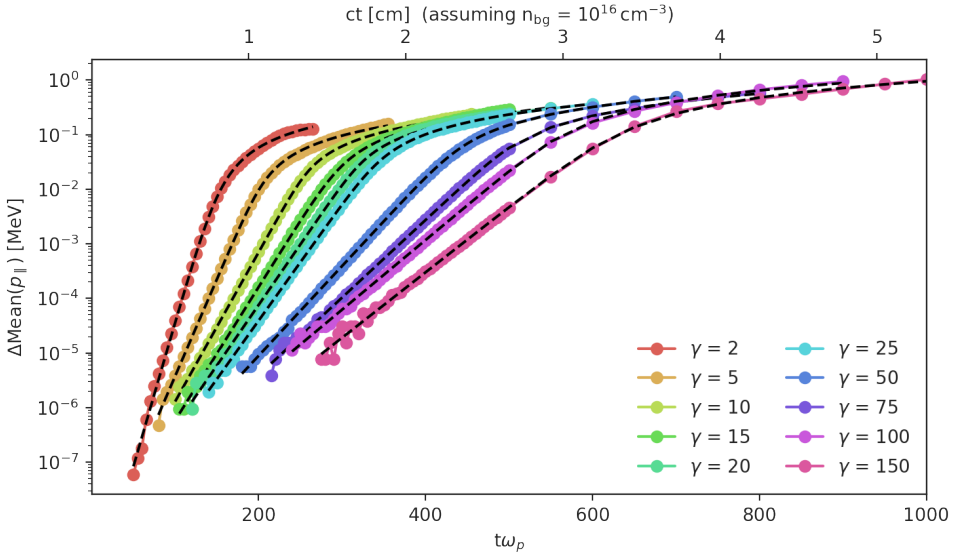

Figure 5.19: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ is shown for runs with varying Lorentz boost γ and $\alpha = 10^{-3}$. The mean value is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for early times the shift of the mean momentum is dominated by numeric effects (finite binning) and not shown here. Furthermore for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown either since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.42) to the data is shown. The fit parameters are listed in table 5.2.

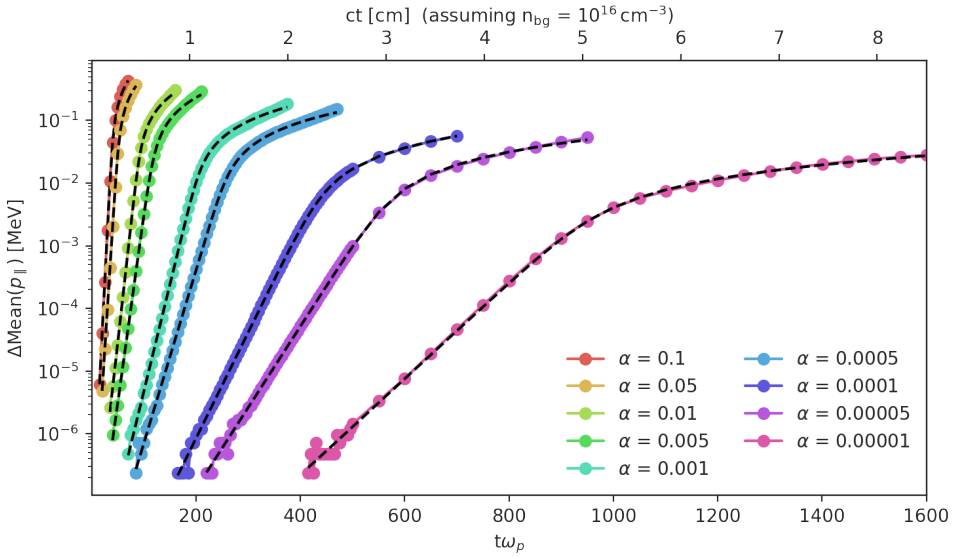

Figure 5.20: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ is shown for runs with varying density ratio α and $\gamma = 5$. The mean value is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for early times the shift of the mean momentum is dominated by numeric effects (finite binning) and not shown here. Furthermore for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown either since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.42) to the data is shown. The fit parameters are listed in table 5.2.

Table 5.2: The fit parameters (δ, V_1, t_{final}, κ) of (3.42) for $\Delta\text{Mean}(p_\parallel)$ from simulation runs with various values for α and γ. The derived quantity Λ is the value of $\Delta\text{Mean}(p_\parallel)$ at time t_{final}. The quantity $(\gamma m_e)^{-1}\,\Lambda$ estimates the relative energy loss of the pair beam.

γ	α	δ $[\omega_p]$	V_1 [MeV]	t_{final} $[\omega_p^{-1}]$	κ	Λ [MeV]	$(\gamma m_e)^{-1}\,\Lambda$
2	0.001	$6.03 \cdot 10^{-2}$	$2.07 \cdot 10^{-10}$	146	4.04	$5.61 \cdot 10^{-3}$	$6.33 \cdot 10^{-3}$
5	0.001	$4.09 \cdot 10^{-2}$	$1.06 \cdot 10^{-9}$	197	10.9	$8.41 \cdot 10^{-3}$	$3.35 \cdot 10^{-3}$
10	0.001	$3.15 \cdot 10^{-2}$	$2.38 \cdot 10^{-9}$	249	33.7	$1.39 \cdot 10^{-2}$	$2.74 \cdot 10^{-3}$
15	0.001	$2.81 \cdot 10^{-2}$	$2.11 \cdot 10^{-9}$	287	11.5	$1.66 \cdot 10^{-2}$	$2.16 \cdot 10^{-3}$
20	0.001	$2.55 \cdot 10^{-2}$	$2.79 \cdot 10^{-9}$	313	33.8	$2.13 \cdot 10^{-2}$	$2.08 \cdot 10^{-3}$
25	0.001	$2.46 \cdot 10^{-2}$	$2.06 \cdot 10^{-9}$	332	15.8	$2.16 \cdot 10^{-2}$	$1.69 \cdot 10^{-3}$
50	0.001	$1.89 \cdot 10^{-2}$	$4.58 \cdot 10^{-9}$	425	19.3	$3.70 \cdot 10^{-2}$	$1.44 \cdot 10^{-3}$
75	0.001	$1.64 \cdot 10^{-2}$	$5.63 \cdot 10^{-9}$	490	33.3	$4.51 \cdot 10^{-2}$	$1.17 \cdot 10^{-3}$
100	0.001	$1.47 \cdot 10^{-2}$	$9.25 \cdot 10^{-9}$	544	17.7	$6.55 \cdot 10^{-2}$	$1.28 \cdot 10^{-3}$
150	0.001	$1.38 \cdot 10^{-2}$	$4.82 \cdot 10^{-9}$	604	9.39	$6.31 \cdot 10^{-2}$	$8.22 \cdot 10^{-4}$
5	0.1	0.187	$2.26 \cdot 10^{-8}$	38.2	13.0	$2.76 \cdot 10^{-2}$	$1.10 \cdot 10^{-2}$
5	0.05	0.150	$1.18 \cdot 10^{-8}$	49.3	14.6	$2.53 \cdot 10^{-2}$	$1.01 \cdot 10^{-2}$
5	0.01	$8.55 \cdot 10^{-2}$	$5.63 \cdot 10^{-9}$	88.8	26.6	$1.85 \cdot 10^{-2}$	$7.36 \cdot 10^{-3}$
5	0.005	$6.90 \cdot 10^{-2}$	$3.16 \cdot 10^{-9}$	112	14.7	$1.48 \cdot 10^{-2}$	$5.89 \cdot 10^{-3}$
5	0.001	$3.93 \cdot 10^{-2}$	$1.85 \cdot 10^{-9}$	198	37.2	$9.65 \cdot 10^{-3}$	$3.85 \cdot 10^{-3}$
5	0.0005	$3.17 \cdot 10^{-2}$	$1.28 \cdot 10^{-9}$	249	29.9	$7.89 \cdot 10^{-3}$	$3.15 \cdot 10^{-3}$
5	0.0001	$1.95 \cdot 10^{-2}$	$3.52 \cdot 10^{-10}$	423	11.7	$3.90 \cdot 10^{-3}$	$1.56 \cdot 10^{-3}$
5	0.00005	$1.49 \cdot 10^{-2}$	$3.41 \cdot 10^{-10}$	545	13.5	$3.09 \cdot 10^{-3}$	$1.23 \cdot 10^{-3}$
5	0.00001	$8.83 \cdot 10^{-3}$	$1.89 \cdot 10^{-10}$	921	14.1	$1.74 \cdot 10^{-3}$	$6.94 \cdot 10^{-4}$

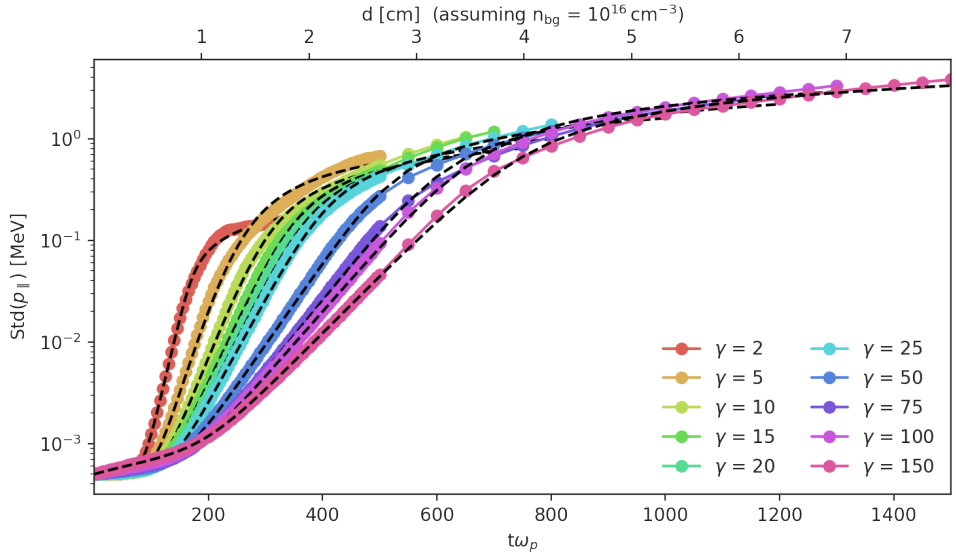

Figure 5.21: The evolution of the width of the longitudinal momentum Std $(p_\parallel)$ is shown for runs with varying Lorentz boost γ and $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 5.3.

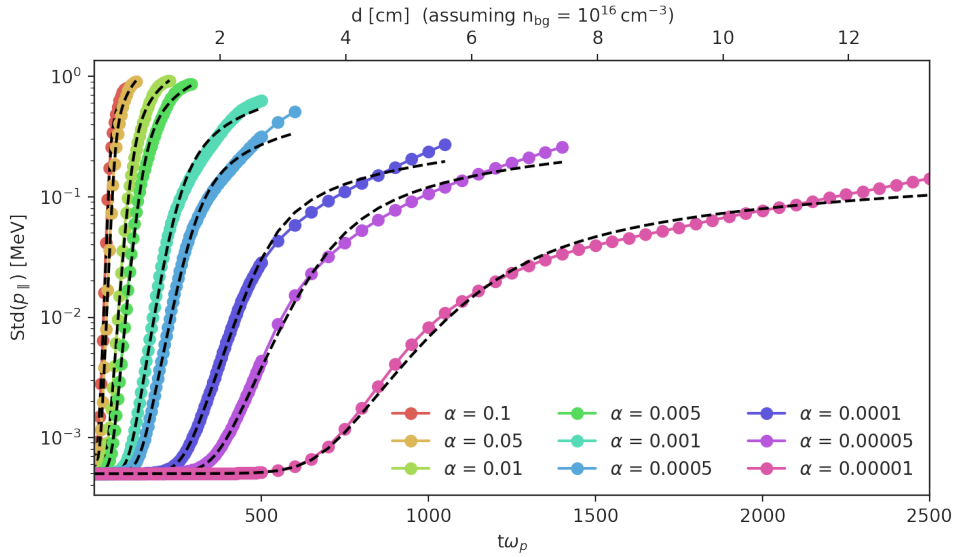

Figure 5.22: The evolution of the width of the longitudinal momentum Std $(p_\parallel)$ is shown for runs with varying density ratio α and $\gamma = 5$.. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 5.3.

Table 5.3: The fit parameters (δ, D_0, D_1, t_{final} and κ) of (3.43) for $\text{Std}(p_\parallel)$ from simulation runs with various values for α and γ. The derived quantity Λ is the value of $\text{Std}(p_\parallel)$ at time t_{final}.

γ	α	δ $[\omega_p]$	D_0 $[\text{MeV}^2\omega_p]$	D_1 $[\text{MeV}^2\omega_p]$	t_{final} $[\omega_p^{-1}]$	κ	Λ [MeV]
2	0.001	$5.46 \cdot 10^{-2}$	$3.16 \cdot 10^{-11}$	$1.28 \cdot 10^{-12}$	164	8.83	$3.34 \cdot 10^{-2}$
5	0.001	$4.08 \cdot 10^{-2}$	$7.63 \cdot 10^{-11}$	$2.19 \cdot 10^{-12}$	241	2.10	$6.91 \cdot 10^{-2}$
10	0.001	$3.24 \cdot 10^{-2}$	$1.09 \cdot 10^{-10}$	$3.95 \cdot 10^{-12}$	293	2.60	$8.87 \cdot 10^{-2}$
15	0.001	$2.77 \cdot 10^{-2}$	$1.69 \cdot 10^{-10}$	$6.88 \cdot 10^{-12}$	336	3.08	0.118
20	0.001	$2.57 \cdot 10^{-2}$	$1.44 \cdot 10^{-10}$	$7.81 \cdot 10^{-12}$	360	3.57	0.131
25	0.001	$2.38 \cdot 10^{-2}$	$1.63 \cdot 10^{-10}$	$1.07 \cdot 10^{-11}$	390	3.30	0.160
50	0.001	$1.86 \cdot 10^{-2}$	$2.57 \cdot 10^{-10}$	$2.32 \cdot 10^{-11}$	498	4.52	0.285
75	0.001	$1.61 \cdot 10^{-2}$	$4.30 \cdot 10^{-10}$	$2.78 \cdot 10^{-11}$	584	3.68	0.360
100	0.001	$1.48 \cdot 10^{-2}$	$5.35 \cdot 10^{-10}$	$4.18 \cdot 10^{-11}$	637	3.27	0.464
150	0.001	$1.30 \cdot 10^{-2}$	$8.76 \cdot 10^{-10}$	$5.75 \cdot 10^{-11}$	717	3.10	0.510
5	0.1	0.162	$1.48 \cdot 10^{-8}$	$1.85 \cdot 10^{-9}$	46.2	48.5	0.182
5	0.05	0.142	$6.16 \cdot 10^{-9}$	$4.11 \cdot 10^{-10}$	58.5	4.47	0.164
5	0.01	$8.61 \cdot 10^{-2}$	$9.69 \cdot 10^{-10}$	$4.24 \cdot 10^{-11}$	106	2.49	0.129
5	0.005	$6.91 \cdot 10^{-2}$	$4.29 \cdot 10^{-10}$	$1.77 \cdot 10^{-11}$	136	2.33	0.109
5	0.001	$4.07 \cdot 10^{-2}$	$6.97 \cdot 10^{-11}$	$2.44 \cdot 10^{-12}$	238	2.16	$6.64 \cdot 10^{-2}$
5	0.0005	$3.25 \cdot 10^{-2}$	$3.67 \cdot 10^{-11}$	$1.28 \cdot 10^{-12}$	290	2.63	$4.87 \cdot 10^{-2}$
5	0.0001	$1.92 \cdot 10^{-2}$	$5.69 \cdot 10^{-12}$	$1.81 \cdot 10^{-13}$	499	3.26	$3.03 \cdot 10^{-2}$
5	0.00005	$1.51 \cdot 10^{-2}$	$4.73 \cdot 10^{-12}$	$7.28 \cdot 10^{-14}$	654	2.27	$2.38 \cdot 10^{-2}$
5	0.00001	$9.00 \cdot 10^{-3}$	$8.63 \cdot 10^{-13}$	$1.40 \cdot 10^{-14}$	$1.09 \cdot 10^{3}$	2.21	$1.16 \cdot 10^{-2}$

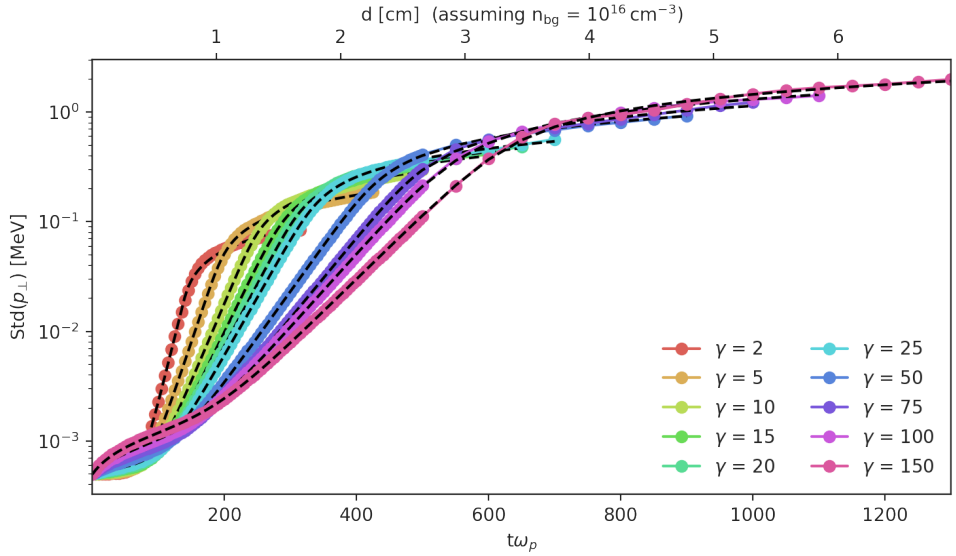

Figure 5.23: The evolution of the width of the transversal momentum Std $(p_\perp)$ is shown for runs with varying Lorentz boost γ and $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 5.4.

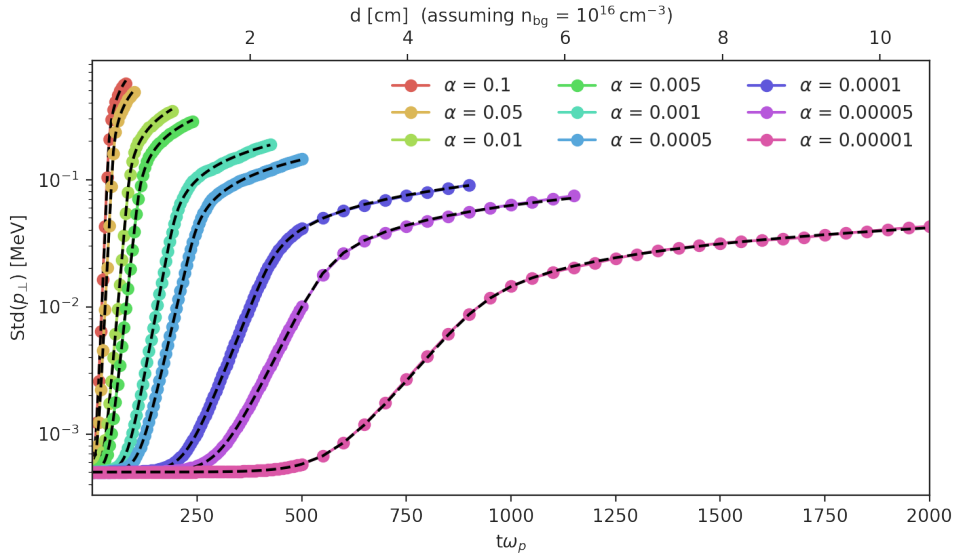

Figure 5.24: The evolution of the width of the transversal momentum Std $(p_\perp)$ is shown for runs with varying density ratio α and $\gamma = 5$.. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 5.4.

Table 5.4: The fit parameters (δ, D_0, D_1, t_{final} and κ) of (3.43) for $\text{Std}(p_\perp)$ from simulation runs with various values for α and γ. The derived quantity Λ is the value of $\text{Std}(p_\perp)$ at time t_{final}.

γ	α	δ [ω_p]	D_0 [MeV$^2\omega_p$]	D_1 [MeV$^2\omega_p$]	t_{final} [ω_p^{-1}]	κ	Λ [MeV]
2	0.001	$5.55 \cdot 10^{-2}$	$3.99 \cdot 10^{-11}$	$4.15 \cdot 10^{-12}$	139	$1.19 \cdot 10^{-24}$	$1.97 \cdot 10^{-2}$
5	0.001	$4.11 \cdot 10^{-2}$	$1.06 \cdot 10^{-10}$	$1.14 \cdot 10^{-11}$	190	$2.26 \cdot 10^{-2}$	$3.86 \cdot 10^{-2}$
10	0.001	$3.27 \cdot 10^{-2}$	$2.16 \cdot 10^{-10}$	$2.16 \cdot 10^{-11}$	241	$2.97 \cdot 10^{-2}$	$6.36 \cdot 10^{-2}$
15	0.001	$2.85 \cdot 10^{-2}$	$3.53 \cdot 10^{-10}$	$3.04 \cdot 10^{-11}$	278	$7.13 \cdot 10^{-2}$	$8.21 \cdot 10^{-2}$
20	0.001	$2.62 \cdot 10^{-2}$	$4.71 \cdot 10^{-10}$	$3.99 \cdot 10^{-11}$	303	$3.96 \cdot 10^{-2}$	$9.94 \cdot 10^{-2}$
25	0.001	$2.44 \cdot 10^{-2}$	$6.26 \cdot 10^{-10}$	$5.10 \cdot 10^{-11}$	324	$4.52 \cdot 10^{-2}$	0.113
50	0.001	$1.94 \cdot 10^{-2}$	$1.35 \cdot 10^{-9}$	$9.21 \cdot 10^{-11}$	413	$8.56 \cdot 10^{-2}$	0.184
75	0.001	$1.69 \cdot 10^{-2}$	$2.10 \cdot 10^{-9}$	$1.16 \cdot 10^{-10}$	478	$9.84 \cdot 10^{-2}$	0.233
100	0.001	$1.55 \cdot 10^{-2}$	$2.96 \cdot 10^{-9}$	$1.68 \cdot 10^{-10}$	522	0.176	0.275
150	0.001	$1.36 \cdot 10^{-2}$	$4.62 \cdot 10^{-9}$	$2.33 \cdot 10^{-10}$	598	0.146	0.363
5	0.1	0.182	$1.13 \cdot 10^{-8}$	$4.97 \cdot 10^{-9}$	37.4	$3.11 \cdot 10^{41}$	0.148
5	0.05	0.148	$7.94 \cdot 10^{-9}$	$1.78 \cdot 10^{-9}$	47.4	$1.50 \cdot 10^{41}$	0.124
5	0.01	$8.84 \cdot 10^{-2}$	$1.77 \cdot 10^{-9}$	$1.95 \cdot 10^{-10}$	84.3	$1.44 \cdot 10^{4}$	$7.95 \cdot 10^{-2}$
5	0.005	$7.03 \cdot 10^{-2}$	$8.51 \cdot 10^{-10}$	$8.64 \cdot 10^{-11}$	107	152	$6.40 \cdot 10^{-2}$
5	0.001	$4.09 \cdot 10^{-2}$	$1.13 \cdot 10^{-10}$	$1.30 \cdot 10^{-11}$	189	86.2	$3.97 \cdot 10^{-2}$
5	0.0005	$3.26 \cdot 10^{-2}$	$5.48 \cdot 10^{-11}$	$6.32 \cdot 10^{-12}$	239	216	$3.26 \cdot 10^{-2}$
5	0.0001	$1.90 \cdot 10^{-2}$	$7.91 \cdot 10^{-12}$	$1.12 \cdot 10^{-12}$	414	51.2	$1.93 \cdot 10^{-2}$
5	0.00005	$1.49 \cdot 10^{-2}$	$4.34 \cdot 10^{-12}$	$5.21 \cdot 10^{-13}$	532	29.7	$1.51 \cdot 10^{-2}$
5	0.00001	$8.75 \cdot 10^{-3}$	$7.37 \cdot 10^{-13}$	$1.14 \cdot 10^{-13}$	897	25.7	$8.59 \cdot 10^{-3}$

5.5.1 Scaling Relationships

In order to apply the simulations to astrophysical environment with very different scales or transfer insight from a laboratory experiment to the astrophysical environment it is important to understand how the investigated quantities can be scaled. The most important parameters here are the density contrast α and the Lorentz boost γ. The astrophysical values can not be replicated in a laboratory environment or simulated with reasonable computational effort.
In figures 5.25, 5.26 and 5.27 the scaling of the growth rate and the saturation level of the linear regime are shown to follow power-laws with respect to α and γ. For the evolution of the mean longitudinal momentum we already accounted for the fact that it grows with twice the growth

rate of the field amplitude. We can summarize the scaling relationships as follows:

$$\delta_{\Delta\mathrm{Mean}(p_\parallel)} \propto \frac{\alpha^{0.33}}{\gamma^{0.34}} \tag{5.8}$$

$$\Lambda_{\Delta\mathrm{Mean}(p_\parallel)} \propto \alpha^{0.31}\gamma^{0.61} \tag{5.9}$$

$$\delta_{\mathrm{Std}(p_\parallel)} \propto \frac{\alpha^{0.32}}{\gamma^{0.34}} \tag{5.10}$$

$$\Lambda_{\mathrm{Std}(p_\parallel)} \propto \alpha^{0.30}\gamma^{0.65} \tag{5.11}$$

$$\delta_{\mathrm{Std}(p_\perp)} \propto \frac{\alpha^{0.33}}{\gamma^{0.33}} \tag{5.12}$$

$$\Lambda_{\mathrm{Std}(p_\perp)} \propto \alpha^{0.31}\gamma^{0.67} \tag{5.13}$$

More generally we find that all growth rates and saturation levels can be approximately described with a single scaling law each.

$$\delta \propto \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \tag{5.14}$$

$$\Lambda \propto \alpha^{\frac{1}{3}}\gamma^{\frac{2}{3}} = \gamma\delta \tag{5.15}$$

In summary the growth of the field amplitude and energy density happens at the same time and with the same rates as the change of the momentum distribution. This confirms our model of the field growth due to the instability driving the change of the momentum distribution. Thus measuring the changed momentum distribution is in a sense equal to measuring the electric field. Furthermore the saturation of the instability seems to be intimately linked to heating of the beam momentum distribution. As a consequence it is sufficient to measure the saturation level of the momentum distribution change to infer the growth rate of field amplitude.

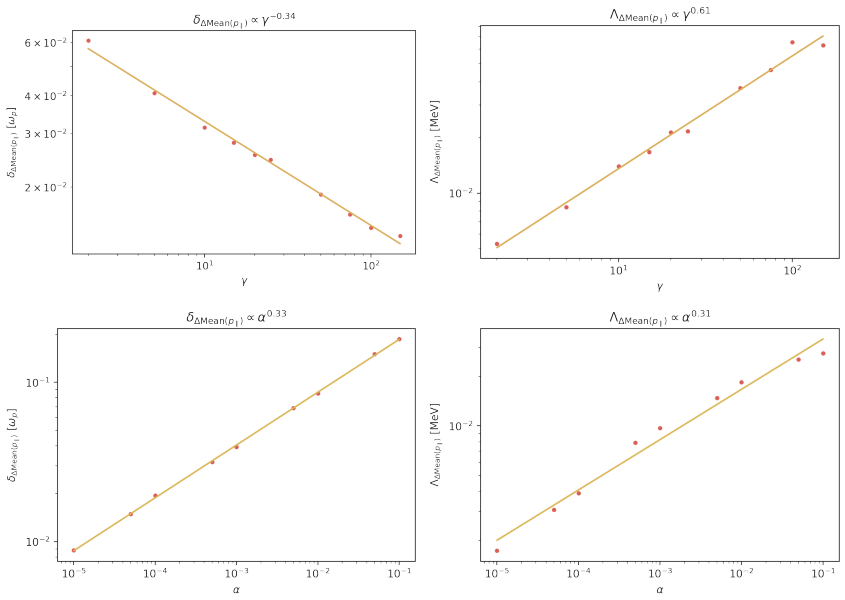

Figure 5.25: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the shift of mean parallel momentum $\Delta\text{Mean}(p_\parallel)$ with γ (**top**) and α (**bottom**). In all cases the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.42).

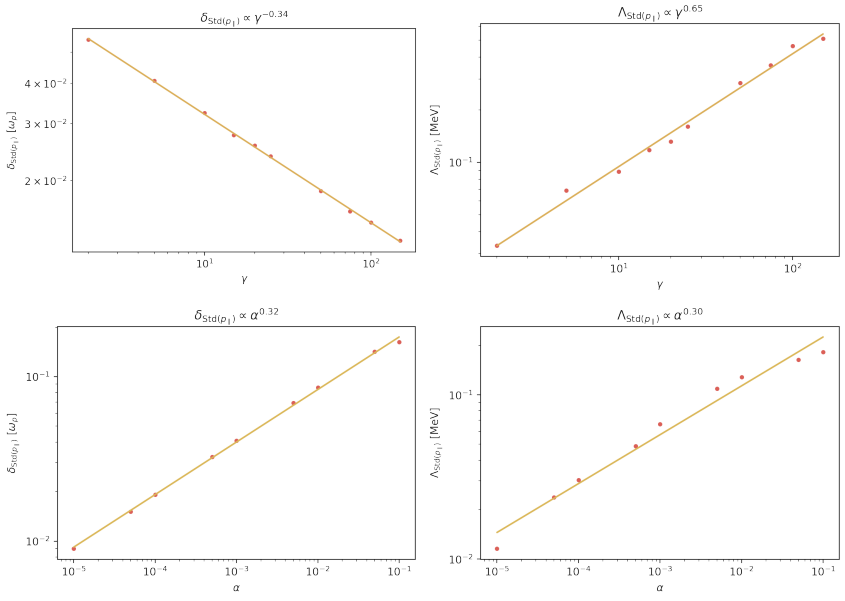

Figure 5.26: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the width of parallel momentum $\text{Std}(p_\parallel)$ with γ (**top**) and α (**bottom**). In all cases the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.43).

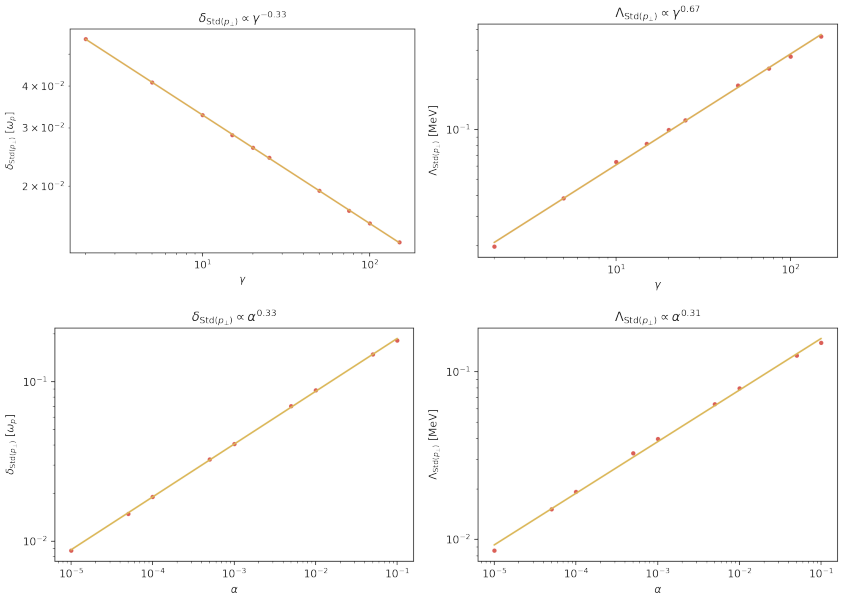

Figure 5.27: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the width of transversal momentum $\mathrm{Std}(p_\perp)$ with γ (**top**) and α (**bottom**). In all cases the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.43).

5.6 2D Correlations

In the previous sections we considered the change of the momentum distribution as separate drift and diffusion processes that lead to a shifting and broadening Gaussian distribution. Most notable the behavior of different momentum dimensions did not affect each other leading to no correlation between $p_\parallel$ and $p_\perp$. But we already saw that even in a single dimension there can be deviation from a purely Gaussian shape from around the time of saturation onwards.
In figures 5.28 and 5.29 the two dimensional momentum distribution is shown for various simulation runs at different points in time. One needs to pay attention that because of the different growth rates the linear growth phase ends at different times for different parameters. However we can see that at the end of the linear growth phase, when non-linear effects start to play a role, significant correlations between longitudinal and transversal momentum become apparent. Namely particles with high transversal momentum are more likely to have lower longitudinal momentum leading to a curved crescent or boomerang shape. This shape does not persist indefinitely though as at very late times we find a significantly more broadening but also a more symmetrical shape again.

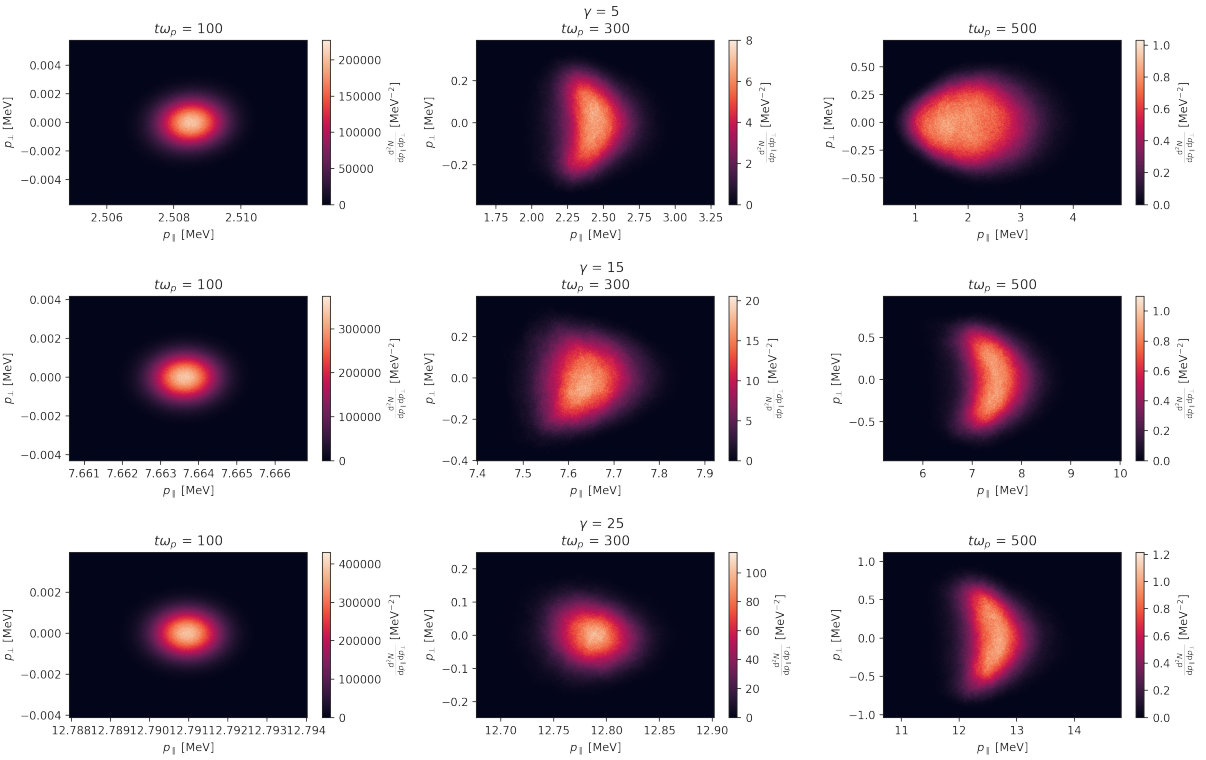

Figure 5.28: The two-dimensional momentum distribution at different simulation times ($100\omega_p^{-1}$ on the **left**, $300\omega_p^{-1}$ in the **middle** and $500\omega_p^{-1}$ on the **right**) for simulation runs with varying γ and $\alpha = 10^{-3}$. Beyond the broadening in longitudinal and transversal direction we can observe a correlation between the two momentum directions at intermediate times (around saturation time), where particles with high transversal momentum have lower longitudinal momentum. Overall this leads to a curved shape (like a boomerang).

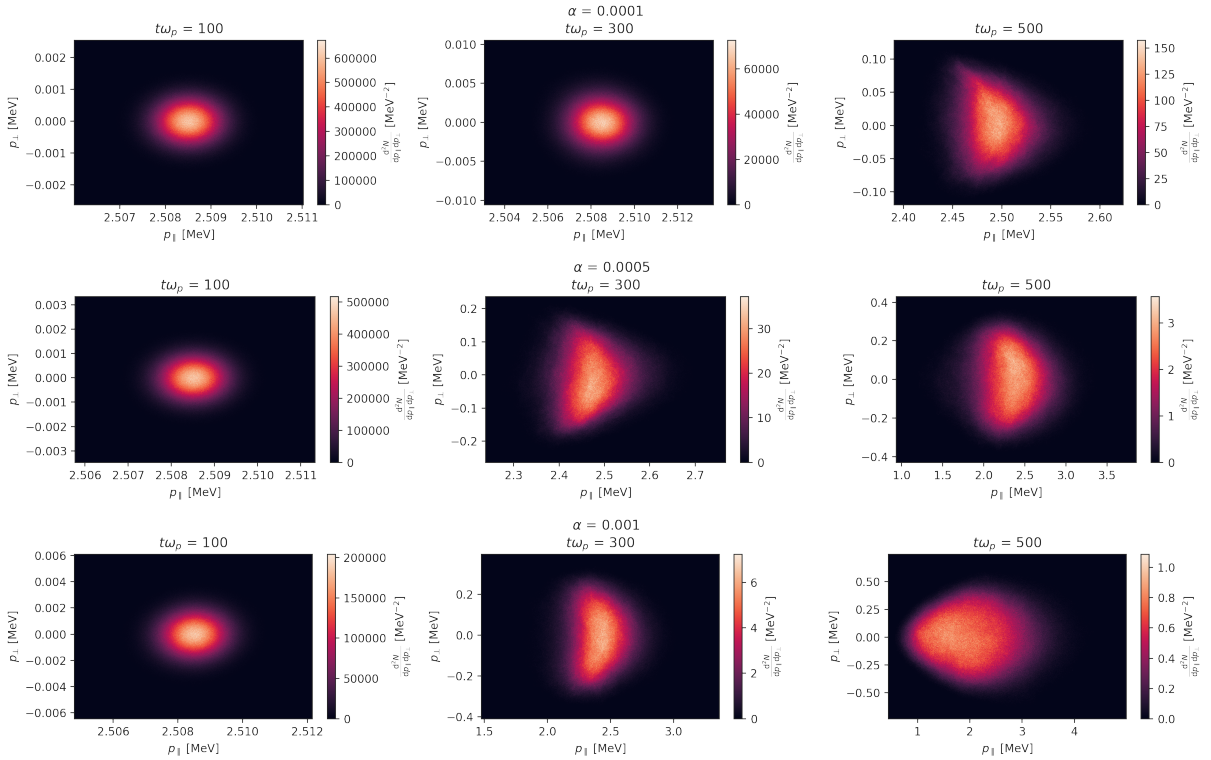

Figure 5.29: The two-dimensional momentum distribution at different simulation times ($100\omega_p^{-1}$ on the **left**, $300\omega_p^{-1}$ in the **middle** and $500\omega_p^{-1}$ on the **right**) for simulation runs with varying α and $\gamma = 5$. Beyond the broadening in longitudinal and transversal direction we can observe a correlation between the two momentum directions at intermediate times (around saturation time), where particles with high transversal momentum have lower longitudinal momentum. Overall this leads to a curved shape (like a boomerang).

5.6.1 Opening Angle

If one is interested in how the change of the momentum distribution affects the astrophysical gamma ray flux measured by a telescope the most interesting quantity is not the increase in width or the change of the mean but the change of the opening angle. In general both the increase of the transversal momentum spread and the decrease of the mean longitudinal momentum will lead to an increased opening angle with respect to the beam axis (the direction where mean transversal momentum is zero). We can estimate the opening angle θ only using these two quantities:

$$\langle\theta\rangle = \frac{\mathrm{Std}\left(p_\perp\right)}{\mathrm{Mean}\left(p_\parallel\right)} = \frac{\mathrm{Std}\left(p_\perp\right)}{\beta\gamma m_e - \Delta\mathrm{Mean}\left(p_\parallel\right)} \approx \frac{\mathrm{Std}\left(p_\perp\right)}{\gamma m_e} \tag{5.16}$$

The last approximation assumes that the beam is highly relativistic and dilute. It is only used for the following analytical estimate of the scaling:

$$\delta_{\langle\theta\rangle} \propto \delta_{\mathrm{Std}}\left(p_\perp\right) \propto \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \tag{5.17}$$

$$\Lambda_{\langle\theta\rangle} \propto \gamma^{-1}\Lambda_{\mathrm{Std}}\left(p_\perp\right) \propto \gamma^{-1}\gamma\delta_{\mathrm{Std}}\left(p_\perp\right) = \delta_{\mathrm{Std}}\left(p_\perp\right) \propto \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \tag{5.18}$$

The factor of γ does not feature into the growth rate since dividing by a constant in logarithmic scale will only give an offset. Treating the evolution of the mean value of longitudinal momentum will give a (relatively) stronger contribution at low Lorentz boost leading to a stronger reduction with higher Lorentz boost on the saturation expected here. Figures 5.30 and 5.31 show the evolution of the opening angle with time. It undergoes a similar evolution to the transverse momentum and can be described with the same analytical model (3.43). Figure 5.32 shows the extracted rates for evolution of the opening angle and they closely confirm to the scaling we predicted.

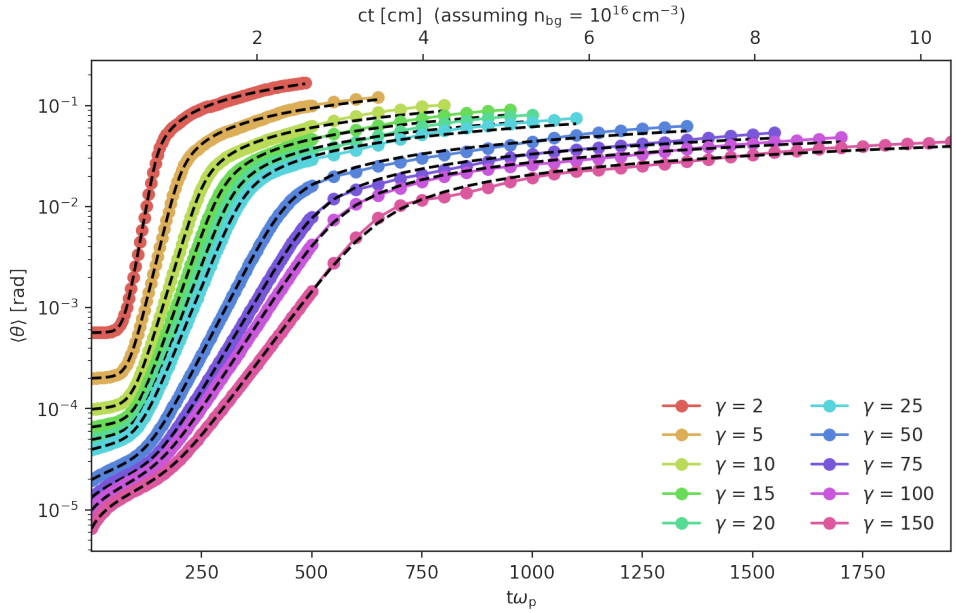

Figure 5.30: The evolution of the beam opening angle $\langle\theta\rangle$ is shown for runs with varying Lorentz boost γ and $\alpha = 10^{-3}$. The opening angle is calculated from the Mean($p_\parallel$) and Std($p_\perp$) at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $3t_{\mathrm{final}}$ (fitted value), data points are not shown either since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data, where the opening angle is substituted for the width of the momentum distribution, is shown.

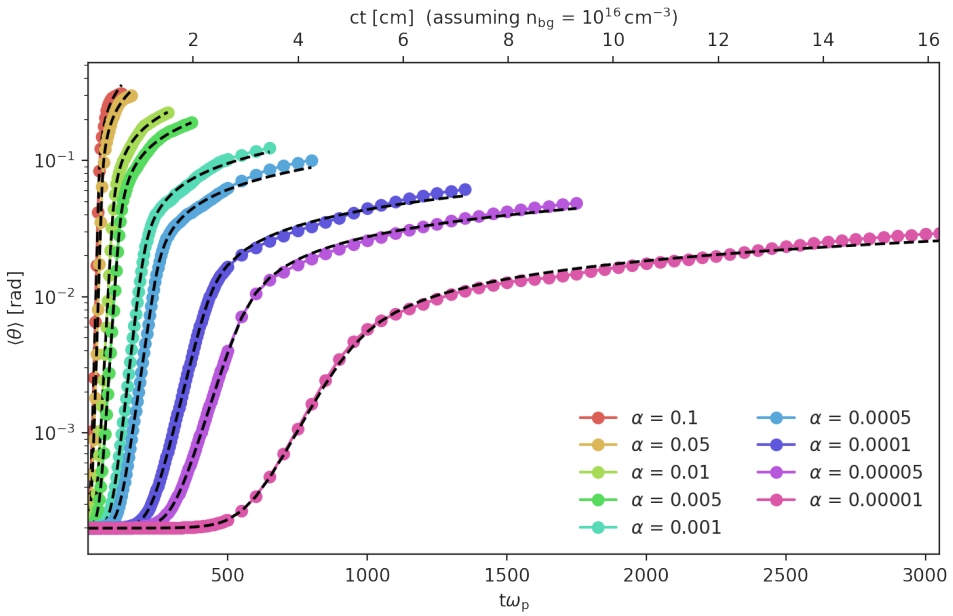

Figure 5.31: The evolution of the beam opening angle $\langle\theta\rangle$ is shown for runs with varying density ratio α and $\gamma = 5$. The opening angle is calculated from the Mean($p_\parallel$) and Std($p_\perp$) at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $3t_{\text{final}}$ (fitted value), data points are not shown either since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data, where the opening angle is substituted for the width of the momentum distribution, is shown.

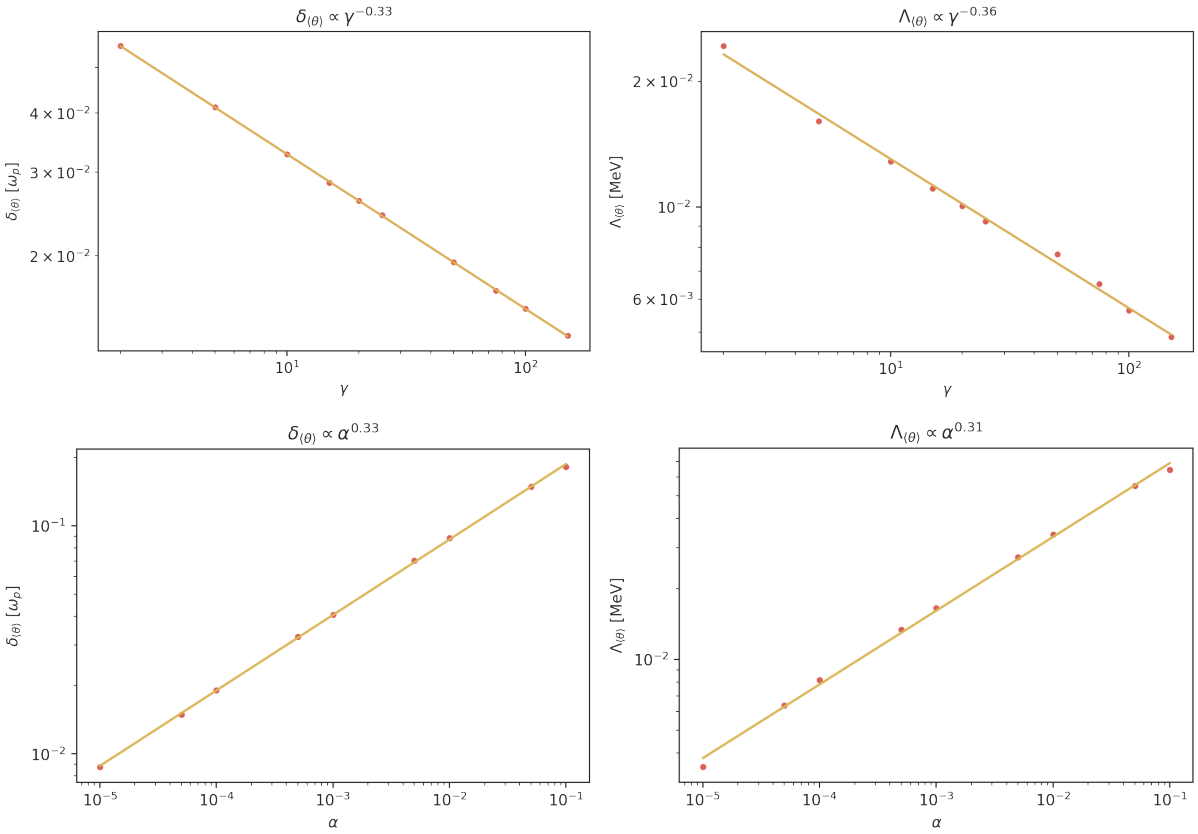

Figure 5.32: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the opening angle $\langle\theta\rangle$ with γ (**top**) and α (**bottom**). In all cases the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.43) as shown in figures 5.30 and 5.31.

5.7 Background Momentum Distribution

The instable behavior of the neutral pair beam leads to the drain of energy and momentum into an electric field which in turn can transfer energy into the background plasma. In figures 5.7 and 5.8 we can see that in some cases almost all of the energy ends up in the background plasma. In the case of very dilute beam we will roughly achieve equipartition between the fields and the background plasma.
The transfer of energy can either take the form of heating the background plasma by increasing the momentum spread isotropically without transferring bulk momentum (which means that the plasma is still at rest with respect to the laboratory frame) or by transferring bulk momentum accelerating the background plasma collectively in the direction of the beam. In the first case the entire loss of longitudinal momentum would have to be transfered to the electric field.
Figures 5.35 and 5.34 show the evolution of the mean and width of the longitudinal momentum distribution for the background plasma electrons overlaid for different combinations of α and γ. We find that even in cases where the beam is not very dilute the transfer of bulk momentum is small compared to the widening of the momentum distribution indicating that the background

plasma gains energy in the form of heat. If the beam is very dilute, meaning $\alpha\gamma \ll 1$, there is no transfer of bulk momentum beyond machine precision even though the background plasma experiences heating.
The heating process of the background plasma takes a similar form to the momentum widening of the beam during the instability. In the beginning the width of the momentum distribution remains at its initial value for some time until an exponential rise sets in, which is followed by a saturation regime. The figures 5.35, 5.36, 5.38 and 5.39 show the evolution of $\mathrm{Std}(p_\parallel)$ and $\mathrm{Std}(p_\perp)$ for the background plasma electrons. We can use the same model (3.43) to describes the heating process. We find that growth rate and the saturation level roughly follow the same scaling (see figures 5.37 and 5.40) as for the beam momentum distribution. However the "coupling constants" D_0 and D_1 are different. This is likely because of the difference in density between the beam and the background plasma. This also leads to final saturation level that is much smaller and thus significantly less heating than in the beam. We do not observe any effect of the initial fluctuations leading to a square root diffusion regime.

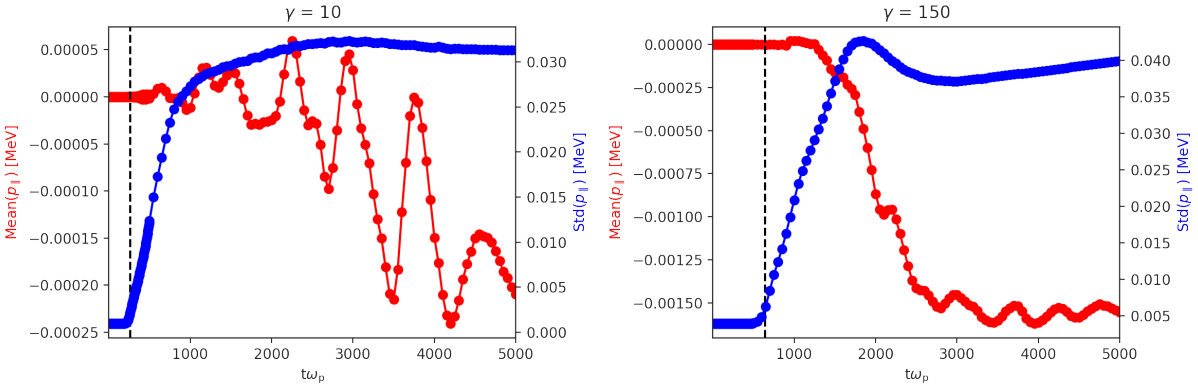

Figure 5.33: The evolution of the mean and width of the longitudinal momentum distribution for two runs with $\alpha = 10^{-3}$ and $\gamma = 10$ (**left**) and $\gamma = 150$ (**right**). The dashed vertical line shows the fitted saturation time t_{final}. The change in mean momentum is small compared to the increase in width. The strength of widening is only slightly correlated with γ, whereas the transfer of bulk momentum is roughly proportional to γ.

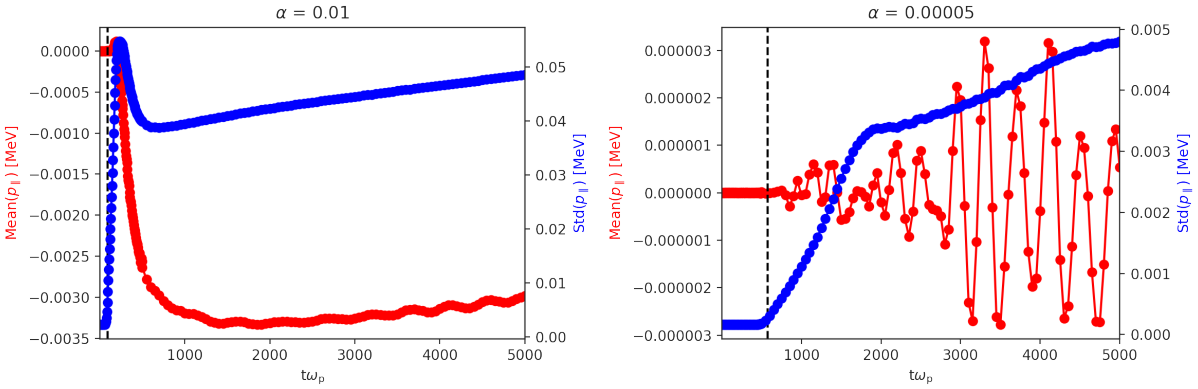

Figure 5.34: The evolution of the mean and width of the longitudinal momentum distribution for two runs with $\gamma = 5$ and $\alpha = 10^{-2}$ (**left**) and $\alpha = 5 \cdot 10^{-5}$ (**right**). The dashed vertical line shows the fitted saturation time t_{final}. The change in mean momentum is small compared to the increase in width. The strength of widening is only slightly correlated with γ, whereas the transfer of bulk momentum is roughly proportional to γ.

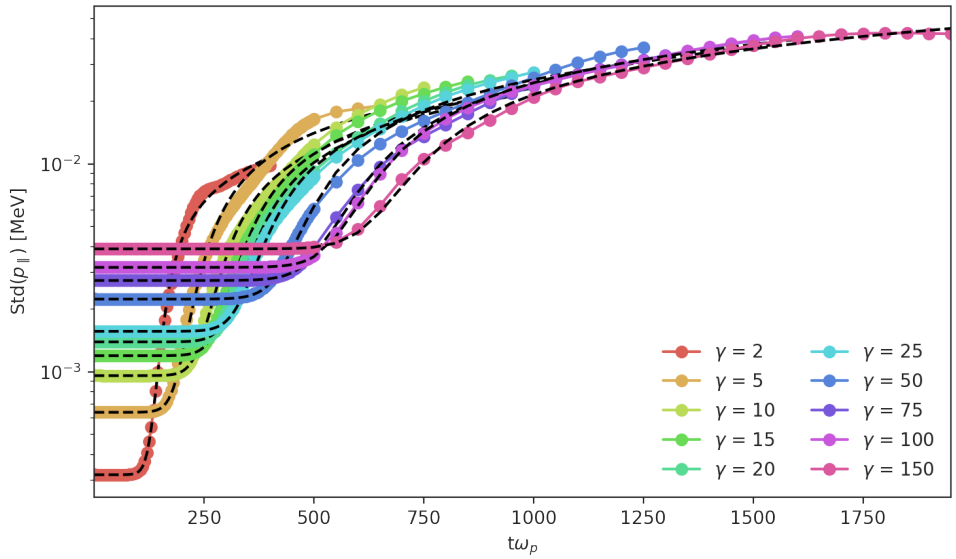

Figure 5.35: The evolution of the width of the longitudinal momentum $\text{Std}\left(p_{\parallel}\right)$ of the background plasma is shown for runs with varying Lorentz boost γ and $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

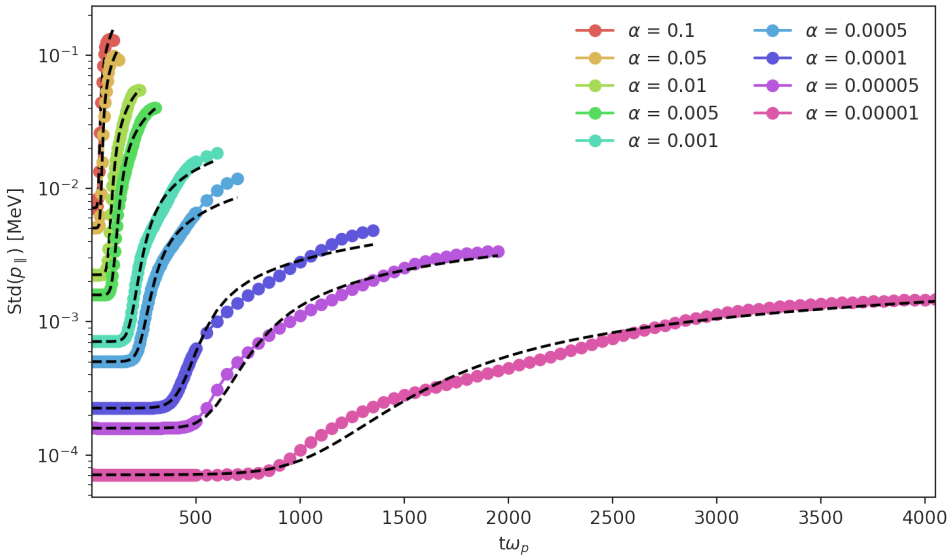

Figure 5.36: The evolution of the width of the longitudinal momentum Std $(p_\parallel)$ of the background plasma is shown for runs with varying density contrast α and $\gamma = 5$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

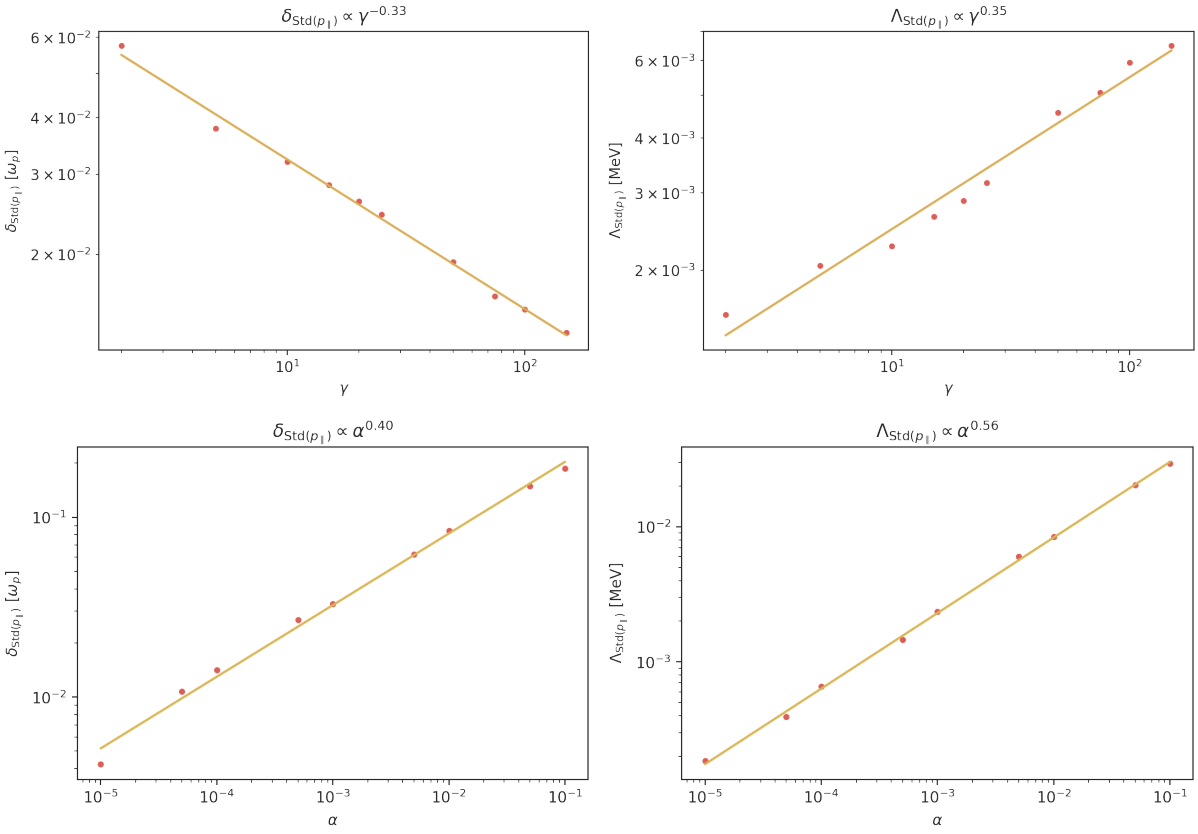

Figure 5.37: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the background plasma longitudinal momentum width $\mathrm{Std}(p_\parallel)$ with γ (**top**) and α (**bottom**). In all cases the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.43) as shown in figures 5.35 and 5.36.

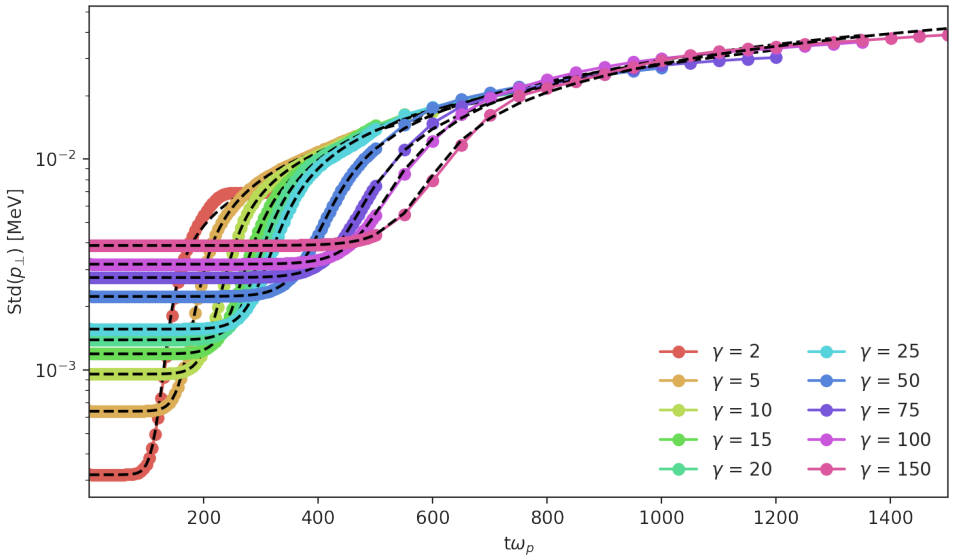

Figure 5.38: The evolution of the width of the transversal momentum Std $(p_\perp)$ of the background plasma is shown for runs with varying Lorentz boost γ and $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

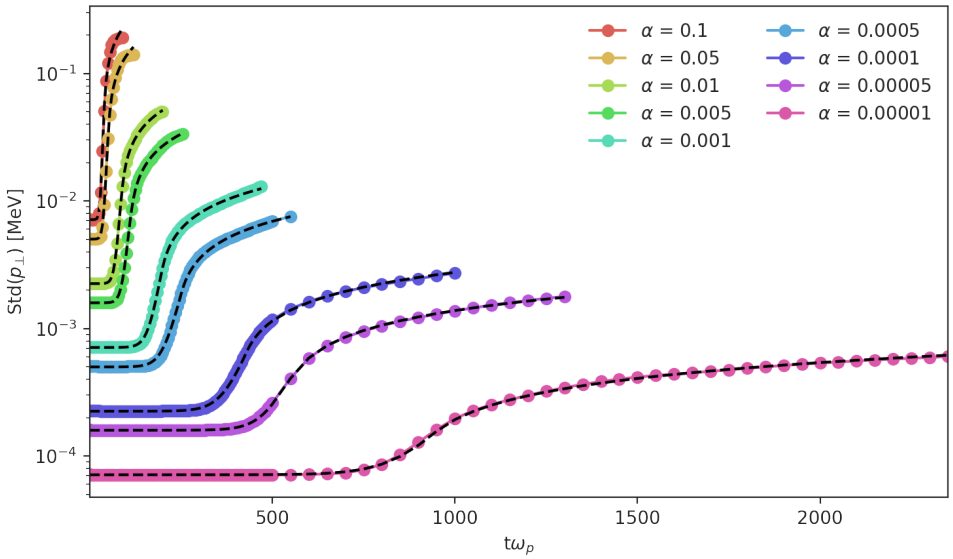

Figure 5.39: The evolution of the width of the transversal momentum Std $(p_\perp)$ of the background plasma is shown for runs with varying density contrast α and $\gamma = 5$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$ until a time of $500\omega_p^{-1}$ and afterwards at intervals of $50\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

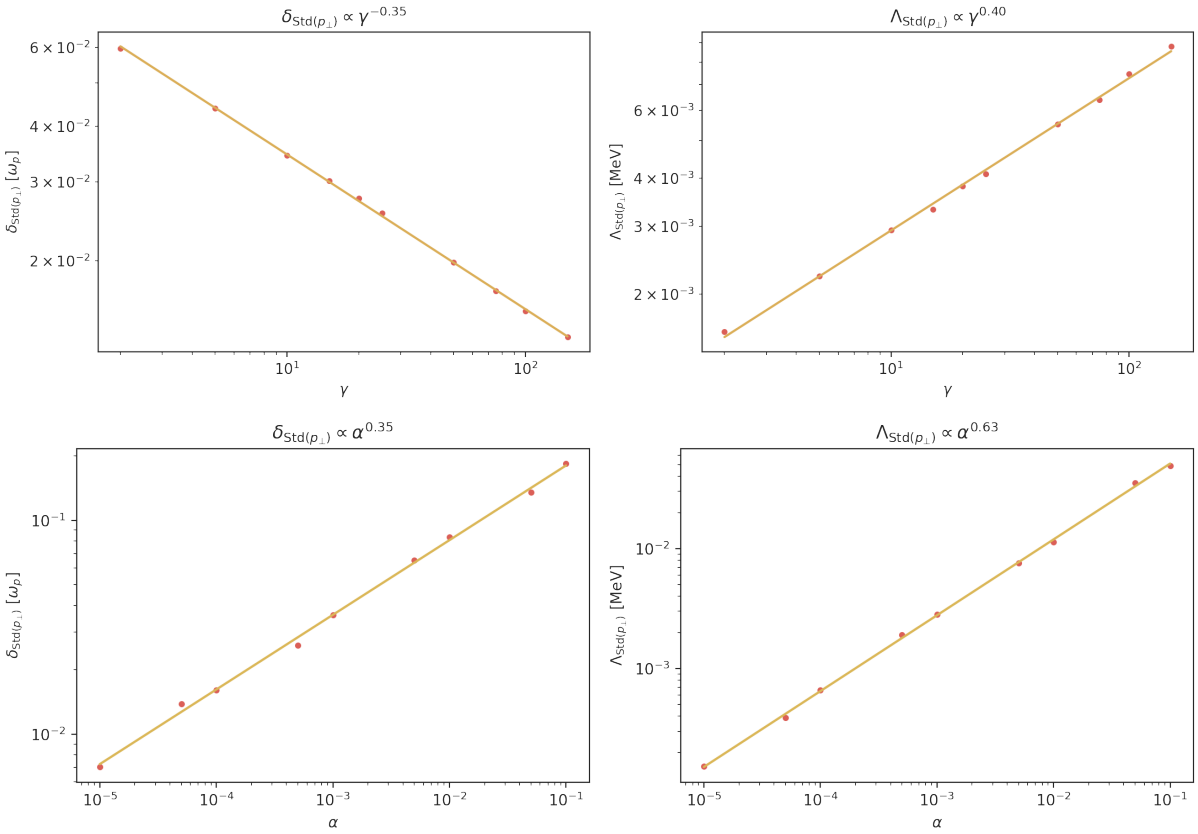

Figure 5.40: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the background plasma transversal momentum width $\mathrm{Std}(p_\perp)$ with γ (**top**) and α (**bottom**). In all cases the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.43) as shown in figures 5.38 and 5.39.

5.8 Initial Noise

At the start of a simulation run the particles of each species are randomly placed in the simulation box with a momentum randomly drawn from a predefined momentum distribution (thermal in the case of background species and shifted Gaussian in the case of beam species) leading to an almost uniform initial number density of each species of particles and current densities that cancel each other almost perfectly[2]. In the limit of infinite particles the initial charge and current density at each grid point of the simulation box would vanish exactly. The finite number of simulation particles leads to deviations from a homogeneous initial state. This noise produces a small charge and current density at each grid point leading to an initial electromagnetic field. The initial field serves the important purpose of seeding the instability. Without these initial fluctuations the pair beam system would be stable. The initial noise can further influence the evolution of the system in significant ways. One example we saw was the broadening of the beam momentum distribution due to diffusion caused by the initial fields.

[2]In the case of background species this is because of the isotropy of the sampled momentum distribution and in the case of beam species this is because positrons and electrons have the same momentum distribution

Since this noise is caused by the finite number of particles it follows Poisson statistics (shot noise). The fluctuations in charge density and current density thus scale inversely with the square root of the number of particles. From the Maxwell equations we can easily deduce that a square root dependence of the charge and current densities lead to a square root dependence of the field amplitude. Since the energy density of a field scales with the square of the field amplitude we can deduce that the energy density of the noise scales inversely with the number of simulation particles (see figure 5.41). It is important to note that each simulation particle stands for a vast number of physical particles and thus the initial noise in a simulation is most likely severely overestimated compared to a real physical system.
The initial thermal motion of the background particles contributes to the noise as the current produced a charged particle is proportional to its velocity. We can adjust the thermal velocity by changing the energy density ratio of beam to background plasma ϵ. Larger epsilon for constant α & γ corresponds to a colder background. The (mean) thermal velocity scales with the square root of the energy. Since the current density scales linearly with the velocity and the energy density of the noise scales inversely with the square of the current density, we can conclude that the initial noise due to the thermal motion scales inversely with ϵ. In figure 5.41 we can in fact see this scaling of the total noise with ϵ. From the scaling with ϵ being almost exactly linear we can conclude that almost the entire initial noise in our PIC simulations is caused by the thermal motion of the background plasma. We note however that this does not necessarily translate to the physical world. One important difference would be that we use an equal number of macro particles for each species, meaning that the background plasma and the beam have the same relative fluctuation. In the physical world the dilute beam consist of far fewer particles than the much denser background plasma leading to larger fluctuations in the beam.

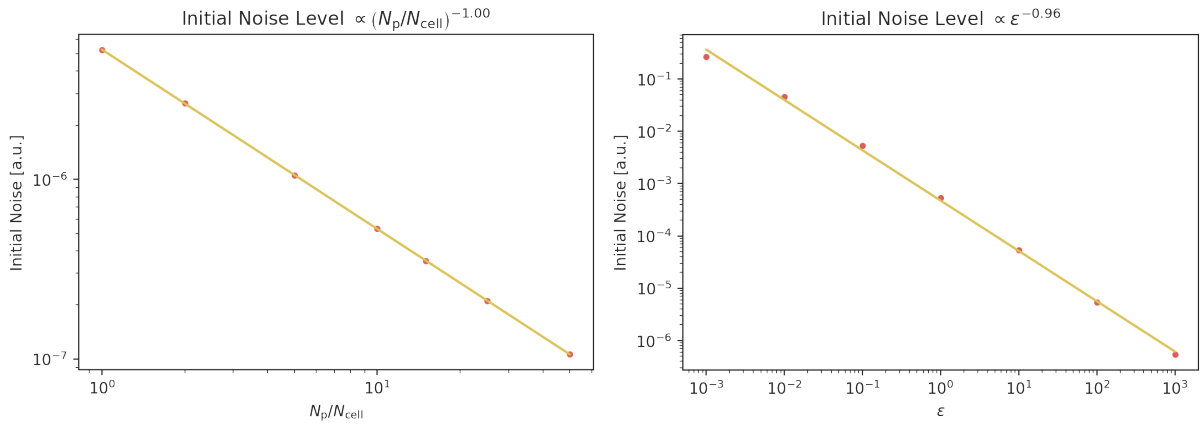

Figure 5.41: The energy density of the initial electromagnetic field noise as a function of the number of simulation particles per simulation cell (**left**) and as a function of the energy density ratio ϵ (**rigth**). ϵ is inversely proportional to the thermal energy of the plasma.

5.9 Cold Beam in 3D

The 2D simulations studied in this chapter should in principle be able to capture all relevant physical effects while requiring orders of magnitude less computational effort than 3D simulations at comparable resolution. Nevertheless we performed 3D simulations with reduced resolution to

confirm that no qualitative difference exist. An overview of the simulation result can be found in appendix B.
In principle the 2D simulations resolve momentum in three dimensions, however the third dimension, which is orthogonal to the simulation plane and also orthogonal to the beam direction, is completely unaffected by the instability and shows no broadening. This is most likely because the instable field modes have no component in this direction because the electromagnetic fields are only resolved in two spatial dimensions. We can confirm that in three spatial dimensions the momentum distribution evolves as expected in all three dimensions.
Even thought the 3D simulations where performed at reduced resolution, maybe most importantly the number of particles per cell was drastically reduced, one can still reduce the observed noise by averaging over the new dimension. For example when calculating the 2D growth rate maps, we calculate the growth rate for each mode resolved in a two dimensional plane. In a 3D simulation where the third dimension is resolved with N cells, we produce N independent realizations of this 2D plane, a noise reduction that is equivalent to increasing the number of particles per cell by a factor of N.
Additionally we can study the spectral growth rate in the plane orthogonal to the beam direction (figures 5.42 and 5.43). Here we find the most dominant modes in a ring centered around the origin with a radius of roughly $\pi\frac{\omega_p}{c}$. Neither the radius nor the width of the ring seem to be strongly affected by the choice of the Lorentz boost γ or the density ratio α. Since the most dominant modes in parallel to the beam direction can be found for $k_\parallel \approx \frac{\omega_p}{c}$ we can conclude the most dominant modes can be found at an oblique angle of roughly 72°.

$$\theta = \arctan\left(\frac{k_{\perp,\max}}{k_{\parallel,\max}}\right) \approx \arctan(\pi) \approx 72° \tag{5.19}$$

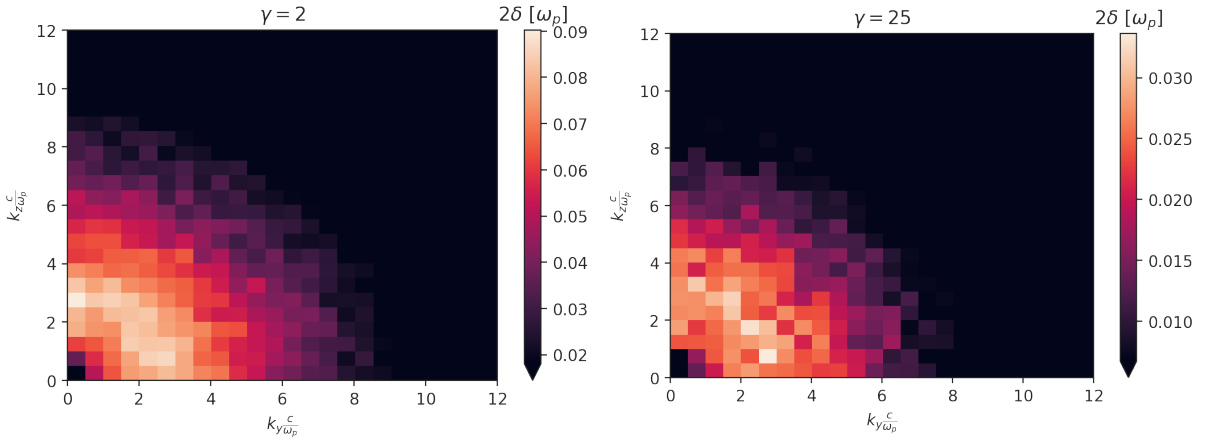

Figure 5.42: Spectral growth rate map of the transverse plane, where k_y and k_z are two components of the wave vector that are orthogonal to the direction of the beam. Shown are two simulation runs with $\alpha = 10^{-3}$ and $\gamma = 2$ (**left**) and $\gamma = 25$ (**right**) respectively.

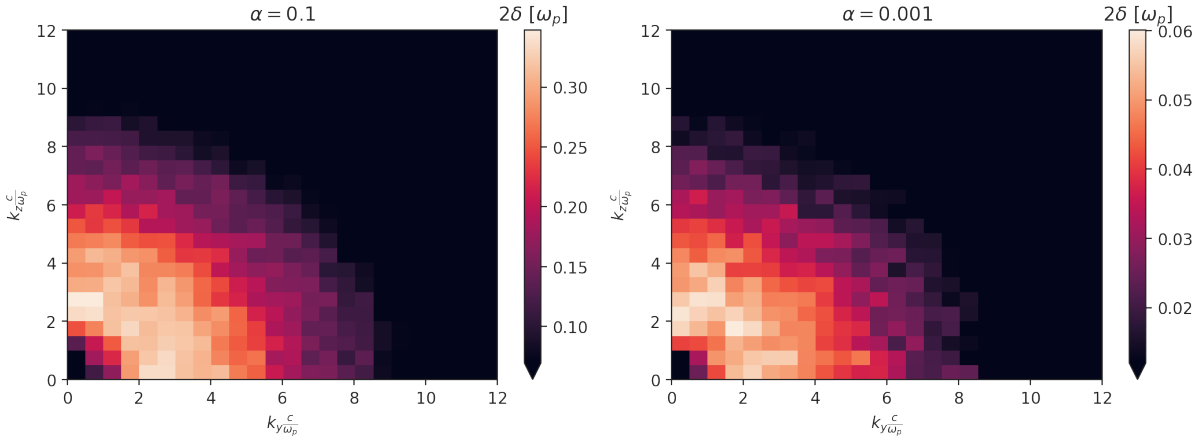

Figure 5.43: Spectral growth rate map of the transverse plane, where k_y and k_z are two components of the wave vector that are orthogonal to the direction of the beam. Shown are two simulation runs with $\gamma = 5$ and $\alpha = 0.1$ (**left**) and $\alpha = 10^{-4}$ (**right**) respectively.

6 | In the Cosmos

In this chapter we want to comment on the impact of plasma instabilities on the phenomenology of TeV blazars. As a first step we need to match the physical characteristics of TeV blazars to the dimensionless parameters α and γ. The parameter γ describes the Lorentz boost of the beam electrons and positrons, which can be trivially calculated from their energy for an individual particle. For a spectrum of energies the exact choice of γ is not as obvious[1] but as an order of magnitude estimate for TeV blazars $\gamma = 10^6 - 10^7$ is a safe choice, implying that the electrons and positrons are also at TeV energies.
The density ratio α is not as straight forward to model. One ingredient is of course the density of the IGM, which is known to be $n_{\mathrm{IGM}} \sim 10^{-7}\,\mathrm{cm}^{-3}$ [15]. Here we neglected the evolution with redshift or the effect of density perturbations. The density of the background medium will be orders of magnitude larger than the dilute particle beam and set the fundamental time and length scales for the plasma.

$$\omega_{p,\mathrm{IGM}} \approx 20\,\mathrm{s}^{-1} \tag{6.1}$$

$$\lambda_{p,\mathrm{IGM}} = \frac{2\pi c}{\omega_p} \approx 3.6 \cdot 10^{-9}\,\mathrm{pc} \tag{6.2}$$

We note that the condition for plasma-like behavior, namely that the size of the system exceeds λ_p, is easily fulfilled.
The electrons and positrons of the beam are produced via pair-production from TeV photons emitted by the blazar and photons from the EBL. It is easy to convince ourself that the density depends on the luminosity L of the blazar and the energy E of the photons as well as the precise spectrum of the EBL. A second influence on the beam density are the cooling mechanisms. The best understood cooling mechanism is ICS on the CMB. [36] estimates the beam density by assuming a balance between pair injection and cooling due to ICS. This derives an upper bound for the beam density, since a more efficient cooling mechanism (for example due to plasma instabilities) would reduce the beam density.

$$n_{\mathrm{b}} \approx 3.7 \cdot 10^{-22} \left(\frac{1+z}{2}\right)^{3\xi-4} \left(\frac{L}{10^{45}\,\frac{\mathrm{erg}}{\mathrm{s}}}\right) \left(\frac{E}{\mathrm{TeV}}\right)\,\mathrm{cm}^{-3} \tag{6.3}$$

$$\Rightarrow \alpha = \frac{n_{\mathrm{b}}}{n_{\mathrm{IGM}}} \approx 3.7 \cdot 10^{-15} \left(\frac{1+z}{2}\right)^{3\xi-4} \left(\frac{L}{10^{45}\,\frac{\mathrm{erg}}{\mathrm{s}}}\right) \left(\frac{E}{\mathrm{TeV}}\right) \tag{6.4}$$

[1] But since the momentum inversion is a key requirement for driving the instability in the first place a good suggestion would be to use the γ value that corresponds to the peak of the spectrum.

The term ξ is associated with the mean free path of VHE gamma rays which depends on the EBL and can be inferred to be $\xi = 4.5$ for $z < 1$ and $\xi = 0$ for $z > 1$. For a typical TeV gamma ray emitter[2] the redshift is $z < 0.5$, the luminosity is $L = 10^{41} - 10^{47}\,\frac{\text{erg}}{\text{s}}$ and the energy is $E = 1 - 10\,\text{TeV}$. This implies an upper bound for the density ratio $\alpha = 10^{-22} - 10^{-13}$. We note that the beam density is inversely proportional to the cooling rate. If an alternative cooling mechanism should be dominant that is for example one order of magnitude more efficient, then the beam density would likewise be reduced by an order of magnitude.
In order to drive the instability the beam particles need to experience many instability growth cycles before they are lost from the beam due to ICS cooling. This sets the condition that the growth rate of the instability has to be much larger than the ICS cooling time [36, 94], which can be calculated with the Thompson cross section σ_T and the energy density of the CMB u_CMB, which can be inferred from the CMB temperature.

$$\delta_\text{ICS} = \frac{4\sigma_\text{T} u_\text{CMB}}{3 m_e c}\gamma \approx 1.4 \cdot 10^{-20}\,(1+z)^4\,\gamma\,\text{s}^{-1} \tag{6.5}$$

We compare this with the maximum growth rate of the oblique mode for a cold beam (3.21).

$$\delta_\text{oblique} \approx 7.8 \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \text{s}^{-1} \tag{6.6}$$

$$\Rightarrow \frac{\delta_\text{oblique}}{\delta_\text{ICS}} = 5.6 \cdot 10^{20}\,(1+z)^{-4}\,\frac{\alpha^{\frac{1}{3}}}{\gamma^{\frac{4}{3}}} \tag{6.7}$$

In figure 6.1 we plot (6.7) assuming a nearby source with $z \ll 1$. Even for very unfavorable values of α and γ the instability growth rate is still many orders of magnitude larger than the ICS cooling rate. We can safely conclude that the instability growth will not be suppressed by ICS.
The instability is described by an exponential growth of the energy density of the electric field modes until a saturation regime is reached. If the saturation regime is reached over distances shorter than the distance between source and observer, we can use relationships we derived for the saturation levels to make inferences. From the 2D PIC simulations we can derive the saturation time t_final as a function of α and γ, as well as the distance the particles propagated until saturation d_final. We use the the values of t_final listed in table 5.4 and assume the inverse scaling of oblique growth rate to fit the numerical value of the pre-factor.

$$t_\text{final} = \frac{522}{20}\left(\frac{\gamma}{100}\right)^{\frac{1}{3}}\left(\frac{\alpha}{10^{-3}}\right)^{-\frac{1}{3}}\text{s} = 56\left(\frac{\gamma}{\alpha}\right)^{\frac{1}{3}}\text{s} \tag{6.8}$$

$$d_\text{final} = c t_\text{final} = 5.5 \cdot 10^{-7}\left(\frac{\gamma}{\alpha}\right)^{\frac{1}{3}}\text{pc} \tag{6.9}$$

The saturation length d_final is plotted in figure 6.2 as a function of the parameters α and γ. Under any realistic circumstances d_final is small compared to the distance from the source to the observer.
The instability grows from an initial noise level until the saturation level is reached. The saturation level is not sensitive to the initial noise and neither is the growth rate. As a consequence the

[2] A non-comprehensive list of known TeV gamma ray sources can be found in [36].

time until saturation increases when the initial noise decreases. This effect has to be considered when scaling the saturation time to astrophysical parameters. Since the beam is extremely dilute ($\alpha\gamma \ll 1$) we assume that the entire initial noise can be attributed to the background plasma. In section 5.8 we have shown how the initial noise then scales with ϵ and the number of (macro) particles. The temperature of the background plasma in galactic voids is usually taken to be $T_{\mathrm{IGM}} \sim 10^4\,\mathrm{K}$ [16].
We express the saturation time t_{final} for arbitrary ϵ and N_p with respect to a reference time $t_{\mathrm{final},0}$ with specific ϵ_0 and $N_{p,0}$ taken from the 2D PIC simulations, like (6.8).

$$u(t) = u(t=0; N_p, \epsilon)\exp(2\delta t) \tag{6.10}$$

$$\Rightarrow t = \frac{1}{2\delta}\log\left(\frac{u(t)}{u(t=0; N_p, \epsilon)}\right) \tag{6.11}$$

$$\Rightarrow t_{\mathrm{final}} = t_{\mathrm{final},0} + \frac{1}{2\delta}\log\left(\frac{N_p\epsilon}{N_{p,0}\epsilon_0}\right) \tag{6.12}$$

For the PIC simulation we find that the instability saturates after approximately ~ 10 e-folding times. Decreasing the initial noise by a factor of e^2 increases the saturation time by one growth period. Even for substantive differences in the initial noise the impact on the saturation time will not be large enough to change to overall picture of the instability saturating over very short distances.

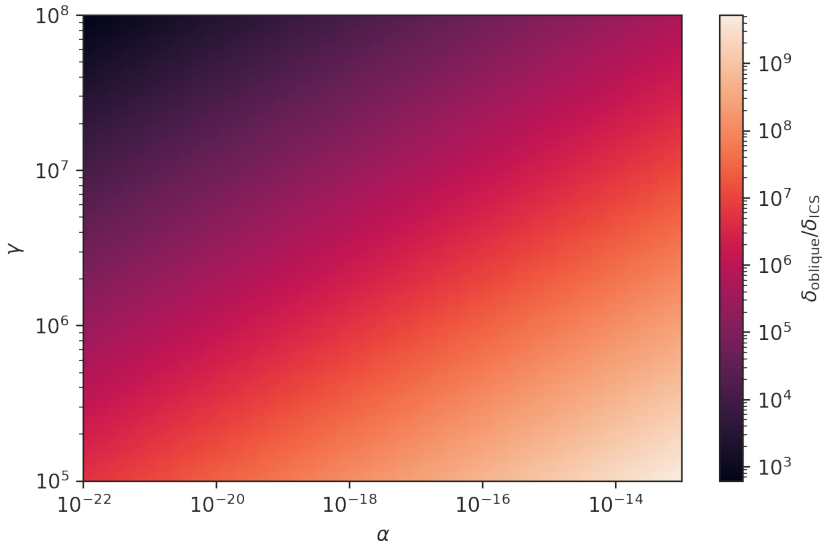

Figure 6.1: The ratio of the oblique instability growth rate $\delta_{\mathrm{oblique}}$ and the ICS cooling rate δ_{ICS} given in (6.7) assuming a source redshift $z \ll 1$. Only if the ratio is significantly larger than 1 can instability growth occur.

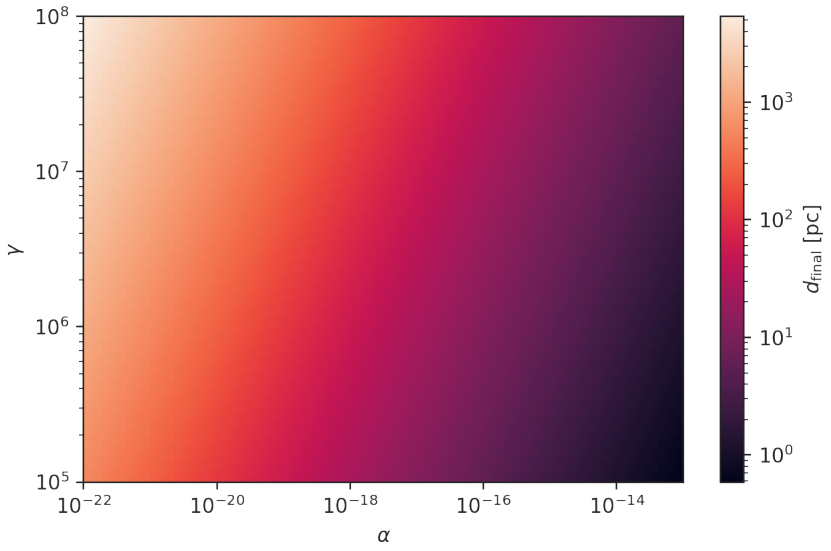

Figure 6.2: The distance a pair beam propagates until the instability reaches saturation. Relative to the distance between an earth and even the closest blazars this distance is extremely small. The effect of the different initial noise has not been taken into account here.

6.1 Instability Cooling

The growth of the field energy density with a rate δ will necessarily lead to the decay of the beam energy with the same rate. Consequently the instability cooling will be the dominating energy loss mechanism if $\delta_{\mathrm{oblique}} \gg \delta_{\mathrm{ICS}}$. Since the cooling via the instability does not produce gamma rays the secondary flux would be suppressed without the need for an external magnetic field. In this case also no extended gamma ray halo or time delayed gamma ray signal would be produced.
However in section 5.2.4 we showed that for a highly relativistic, very dilute particle beam the oblique instability saturates after only draining a minuscule fraction of the beam's energy. If we assume that the astrophysical pair beam also saturates at this level the ICS cooling would still be the dominant energy loss mechanism. Whether the astrophysical pair beam shows the same saturation behavior depends on the exact mechanism that leads to the stop of the instability growth.
Should the instability be indeed the dominating energy loss mechanism the beam density would need to be adjusted accordingly to reflect the faster loss of beam particles shifting α to smaller values.

6.2 Gamma Ray Halos

The deflection of electrons and positrons in the cascade due to IGMF has been suggested to produce a gamma ray halo that could be used to constrain the IGMF [15, 43, 17, 58, 6, 83]. Similarly the transversal heating of the beam due to plasma instabilities also deflects beam particles and could lead to a gamma ray halo. We estimate the deflection angle of the instability by scaling the saturation level $\Lambda_{\langle\theta\rangle}$ from the 2D PIC simulations (section 5.6.1) to astrophysical

parameters. We note that this angle is not the size of the gamma ray halo, which would need to be calculated from this angle in a source dependent way (see for example [15]).

$$\Lambda_{\langle\theta\rangle} = 0.5 \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \text{ rad} \tag{6.13}$$

The intrinsic opening angle of the pair beam is given by γ^{-1}. If the deflection of the beam does not clear this threshold and then it will not be observable even with perfect instrument. This necessarily sets a lower bound on the beam density ratio to produce a pair halo effect.

$$\alpha \gtrsim \frac{10}{\gamma^2} \tag{6.14}$$

Even for a high Lorentz boost value of $\gamma = 10^7$ the density ratio would need to be larger than $\alpha > 10^{-13}$, which we derived as an upper bound for the possible beam density ratios. At least by the end of the linear growth regime the opening angle of the beam would not increase enough to produce a halo.
If one were to observe gamma ray pair halos and would want to establish if they are produced by magnetic fields or plasma instabilities in a way that is independent on the modeling of the instability growth process, one could investigate the size of halo as a function of the beam density ratio α. Magnetic fields act on each individual particles and thus a denser beam would not produce a different halo, but plasma instabilities are the product of collective behavior and thus very sensitive to the beam density ratio with regard to their growth rates and saturation levels. In the case of a very coherent magnetic field the angular shape of the halo could also be used to distinguish these two possible explanations [38, 119].

6.3 Variable Sources

TeV gamma ray emitters are know to show variability on the time scale of seconds for GRBs and much longer time scales for blazars. For variable sources it has been suggested that the same deflections that lead to a secondary gamma ray halo would also lead to a time delay of the secondary gamma ray contribution, as they travel a longer path than the primary gamma rays. This would lead to a gamma ray echo which could be used to constrain the IGMF [15, 78, 52, 83, 56]. Deflections due to the plasma instabilities would lead to the same time delay effect and produce a gamma ray echo. However as we calculated for the gamma ray halo the deflections due to oblique plasma instability saturate at too small values to outright produce a significant effect.
One important difference between the deflection in the magnetic field and deflection in the plasma instability induced electric field is that the electric fields need some time to grow in the case of a variable beam. In section 7.1.1 we will perform PIC simulation of pair beams with finite length and find that the particles at the front of the beam experience little to no deflection and the full deflection is only experienced by particles after at least one growth cycle given by the inverse growth rate δ^{-1}. Since only part of the secondary gamma ray signal would be delayed, the observation might be more precisely described as a gamma ray afterglow instead of a gamma ray echo. Furthermore this sets a bound on the length of the emission τ_{emission} of the primary

gamma ray signal required to induce instability related effects.

$$\tau_{\text{emission}} \gg \delta^{-1}_{\text{oblique}} = 0.13 \left(\frac{\gamma}{\alpha}\right)^{\frac{1}{3}} \text{ s} \tag{6.15}$$

In figure 6.3 we show the inverse growth rate as a function of α and γ. For typical blazar parameters the emission time would need to be at least of the order of weeks to produce a significant effect. This does not consider the contribution to instability growth due to a previous low-flux state.
Likewise for a GRB that shows variability on the time scale of usually less than 100 seconds this sets a lower bound on the beam density (and thus the luminosity of the GRB). We assume $\tau_{\text{emission}} \approx 100\,\text{s}$.

$$\tau_{\text{emission}} \approx 100\,\text{s} \gg 0.13 \left(\frac{\gamma}{\alpha}\right)^{\frac{1}{3}} \text{ s} \tag{6.16}$$

$$\Rightarrow \alpha \gtrsim \frac{\gamma}{10^9} \tag{6.17}$$

For a typical $\gamma = 10^6 - 10^7$ this would yield a beam density ratio of $\alpha = 10^{-3} - 10^{-2}$ which would require GRB luminosities on the order of $L = 10^{56} - 10^{58}\,\frac{\text{erg}}{\text{s}}$. This means that GRBs are too short to be affected by plasma instabilities.

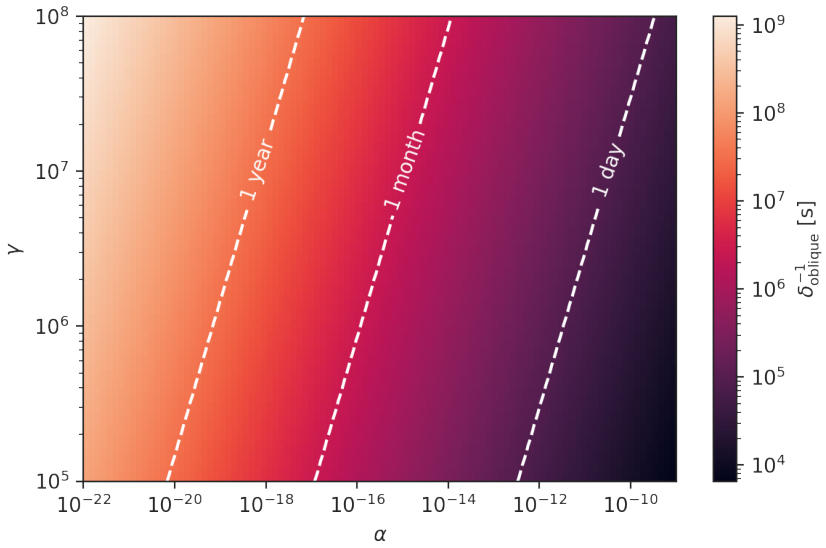

Figure 6.3: The inverse growth rate of the oblique instability as a function of γ and α. Only when the emission time τ_{emission} exceeds this value can a significant instability effect on the gamma ray signal be observed.

6.4 Saturation of a Replenishing Beam

So far we considered the effect of the oblique plasma instability by extrapolating the saturation levels we found in 2D PIC simulations to astrophysical scales. However as described in (3.42) and (3.43) the evolution of the beam momentum distribution does not stop when the fields

saturate. In fact after saturation the effect of the energy drain and the momentum diffusion is the strongest and can continue indefinitely. In the 2D PIC simulation the feedback of the fields on the beam leads to a growth of the transversal momentum spread until the beam becomes too hot and the instability growth stops. Then the electric field still continues to evolve, mainly in the form of transferring energy to modes with longer wavelengths. This leads to the reduction of the diffusion and drift terms and eventually the evolution of the momentum distribution stops. In the astrophysical scenario however the beam is constantly fed new, cold particles via pair production of the primary gamma rays on the EBL. These new, cold particles can continue to drive the instability and eventually an equilibrium will be reached where the diffusion and drift term take on a constant value. We will try to estimate this constant value by assuming this equilibrium will be reached at the saturation level of our PIC simulations where D_1 and V_1 take on their maximum values henceforth called $D_{\max}$ and $V_{\max}$.

$$D_{\max} = 6.6 \cdot 10^{-3} \gamma^{1.1} \alpha^{0.93} \,\mathrm{MeV}^2 \omega_p \tag{6.18}$$

$$2\delta V_{\max} = 0.74 \gamma^{0.28} \alpha^{0.62} \,\mathrm{MeV} \omega_p \tag{6.19}$$

It is not at all clear that the equilibrium between dissipation of the fields and new growth would be reached at this level. Since the injection of new pairs is determined by the mean free path of TeV gamma rays and the relevant scale for the evolution of the electric fields in a plasma is λ_p, which is much shorter, it is likely that our values for $D_{\max}$, $V_{\max}$ only give a vague upper limit. Using the value for the drift term we can derive an energy loss time τ_{loss} (or rather an upper limit).

$$\tau_{\mathrm{loss}} = \frac{E}{-\frac{\mathrm{d}E}{\mathrm{d}t}} = \frac{\gamma m_e}{2\delta V_{\max}} = 0.7 \frac{\gamma^{0.72}}{\alpha^{0.62}} \omega_p^{-1} \tag{6.20}$$

In figure 6.4 we compare this rate with the loss rate due to ICS. For reasonable parameters it is possible for energy loss due to the instability in the saturation regime to be the dominant energy loss mechanism.
Since the saturation length of the oblique instability is so short, one can try to model the entire effect of the instability always assuming the saturation level and taking $D(t) = D_{\max}$ and $V(t) = V_{\max}$ to be constants. These two constants would be enough to model the feedback of the collective effects on the beam particles. Using state of the art cosmic ray propagation codes like CRPropa [13, 14] one can simulate the full impact of the instabilities on the blazar spectrum, the development of gamma ray halos and gamma ray echos/afterglows. Nevertheless we stress that it would take further studies of the plasma instabilities in the context of a beam that replenishes itself constantly with new, cold beam particles, to give robust estimates of the energy loss and diffusion rates.

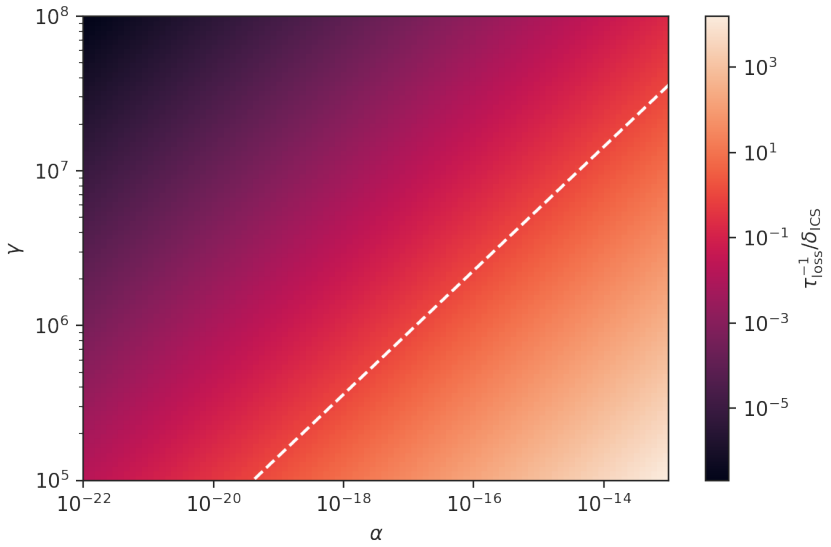

Figure 6.4: Comparison of the energy loss rate $\tau_{\rm loss}^{-1}$ in the saturation regime with the Inverse Compton rate $\delta_{\rm ICS}$. The dashed line indicates where the rates are equal.

7 | Beyond the Cold Beam

In chapter 5 we took an in-depth look at an idealized pair beam undergoing unstable behavior. We found that beyond draining energy from the pair beam into the electric field and heating the background plasma, momentum diffusion also leads to heating of the beam. In this chapter we will look beyond this idealized scenario by considering key effects that violate the idealized scenario. Namely we will consider pair beams that are not infinitely extended but have a finite length and/or width. We will consider scenarios where the beam momentum distribution can not be considered cold anymore and scenarios where the neutrality of the pair beam is violated and thus the beam carries a current. Furthermore we will consider changing the energy density hierarchy of beam plasma and background plasma to understand how the behavior of a laboratory system is related to the astrophysical system.
we will also consider introducing an inhomogeneous background medium and external magnetic fields. These two effects are in general quite complex and we will only address scenarios where the coherence length of density perturbations and magnetic fields is large.
In all of these scenarios one key question we would like to address is quantifying the limit where the behavior of the idealized pair beam is recovered (if it exists).

7.1 Finite Extension

Previously we used periodic boundary conditions to simulate a beam that is infinitely extended in length and width. In a laboratory experiment the particle beam will necessarily have a finite length and width. This means that the density of beam particles tapers off from the center of the beam outwards. This changes the scenario radically. Not only is the density ratio α now a function of the position within the beam, but also the energy density of each mode and thus the effect on the momentum distribution is a function of position. We postulate that in the limit of very high length or width we should approximately recover the behavior of the infinitely extended system.
To investigate the effect of finite size we will consider the effect of length and width separately. We will use open boundary conditions with a Gaussian density profile while keeping the periodic boundary conditions with a constant density profile in the other direction. The length $L_\parallel$ or width $L_\perp$ will be defined as the standard deviation of the Gaussian density profile, that is the central interval that contains 68% of the beam particles. The density ratio α will be defined as the ratio between the peak beam density and the background density.

7.1.1 Finite Length

We consider beam of variable length $L_\parallel$ with a Gaussian density profile in a simulation box with length 500 $\frac{c}{\omega_p}$ and open boundary conditions in longitudinal direction. Figure 7.1 shows the evolution of the total field energy density for simulation runs with varying beam length. Since the energy density is normalized to the total initial beam energy, the initial noise regime increases for a shorter beam that contains overall less particles. Furthermore we can see that for shorter beam the slope of exponential growth phase is reduced. By fitting (5.5) we extracted the growth rate for each simulation run. Shown in figure 7.2 is the growth rate as a function of beam length for low value of the Lorentz boost γ and a high value respectively. In both instances we find that for very short beams the extracted growth rate scales almost proportionally to the beam length and for very long beams it asymptotically approaches the growth rate of an infinitely long beam we will call δ_0. For the high Lorentz value the transition into the asymptotic regime happens later. Heuristically we find a formula that describes the growth rate as a function of $L_\parallel$ and δ_0.

$$\delta(L_\parallel) = \delta_0 \tanh\left(2\delta_0 \frac{L_\parallel}{c}\right) \tag{7.1}$$

Leaving δ_0 as a free parameter we performed a fit of (7.1) to the data points in figure 7.2 which mostly agrees with the data. We can now derive a condition for which we recover the growth rate of the infinitely long beam.

$$\frac{L_\parallel}{c} \gg \delta_0^{-1} \Rightarrow \delta(L_\parallel) \approx \delta_0 \tag{7.2}$$

The modes that are excited by the instability process (see differential growth rate maps in figure 7.3) are the same for long and short beams. Beyond the overall reduction of the growth rate for a shorter bunch we also observe a sort of interference pattern reducing the extracted differential growth rate in regular intervals of $k_\perp$. However this could be a numerical artifact of the discrete Fourier transform with open boundaries.
A major difference in the case of a finitely long beam is that the front of the beam continuously interacts with fresh, unperturbed background plasma. Figure 7.4 shows the density of the beam as a function of the position at multiple time steps. At early times the entire beam comforts to the Gaussian shape the beam was initialized with. At late times this shape becomes strongly perturbed by the instability process, but only at the back of the particle beam. The front of the bunch still closely resembles the initial Gaussian profile. We can conclude that only a subset of the particles at the back of the bunch significantly experience the effects of the instability mechanism. Because the excitation of unstable modes is now not only a function of time but also space this instability mechanism has also be called *spatiotemporal* [97].
The particle beam momentum distribution changes as a consequence of instability growth. If different parts of the beam experience a different amount of instability growth then the momentum distribution will be affected differently as a function of position. More precisely the shift of the mean longitudinal momentum and the widening of the momentum distribution in all directions will be less severe in the front of the bunch. The total momentum distribution integrated over position will thus be a combination of a narrow Gaussian that reflects the initial momentum distribution of the particles at the front of the beam and much wider distribution of the particles at the back of the beam that is shifted to lower energies. In figure 7.5 we show the longitudinal

beam momentum distribution at a number of timesteps near the end of the instability growth for a short and a long beam. In both instances we find that part of the distribution remains in the narrow initial state and part of the distribution is broadened and shifted to lower energies. For the transversal momentum distribution (figure 7.6) for the same simulation runs we similarly find that a subsection of the beam particles remain with narrow distribution while another part experiences significant broadening. We also find that for both the longitudinal and transversal momentum distribution the fraction of the beam particles that makes up the narrow part of the momentum distribution decreases for a longer beam.
Due to the more complex evolution of the beam momentum distribution the evolution of $\Delta\text{Mean}(p_\parallel)$ (figure 7.7), $\text{Std}(p_\parallel)$ (figure 7.8) and $\text{Std}(p_\perp)$ (figure 7.9) is also changed. More precisely for a shorter beam the rate of the shift and widening is reduced and the final saturation value is lower. For each simulation run we extracted the growth rates by fitting (3.42) and (3.43) respectively. The growth rates (figure 7.10) all conform to (7.1) like the integrated growth rate of the field energy density.

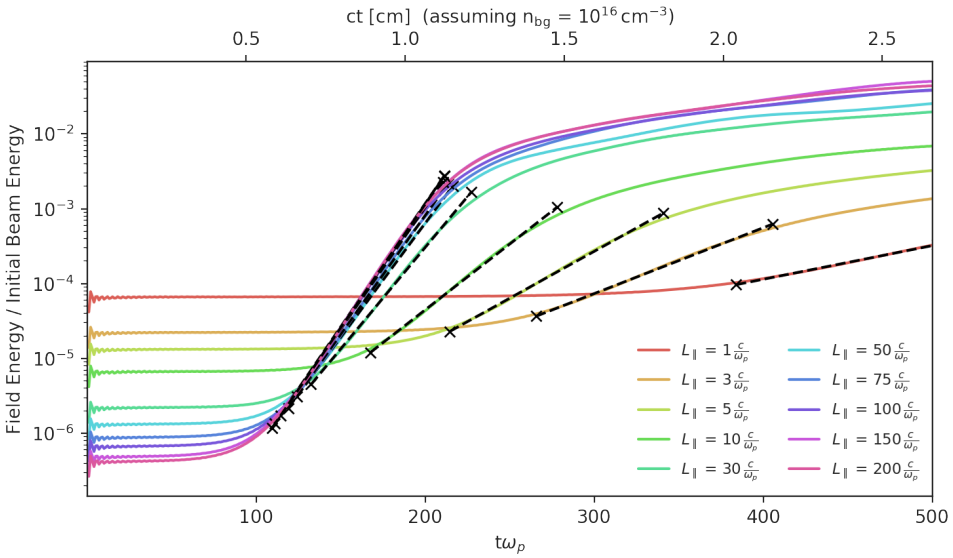

Figure 7.1: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time. Each simulation run has a density contrast of $\alpha = 10^{-3}$, Lorentz boost $\gamma = 5$ and a varying beam length $L_\parallel$. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line. The initial noise regime is mostly driven by the background plasma and increases for a shorter beam (less particles) compared to the beam energy.

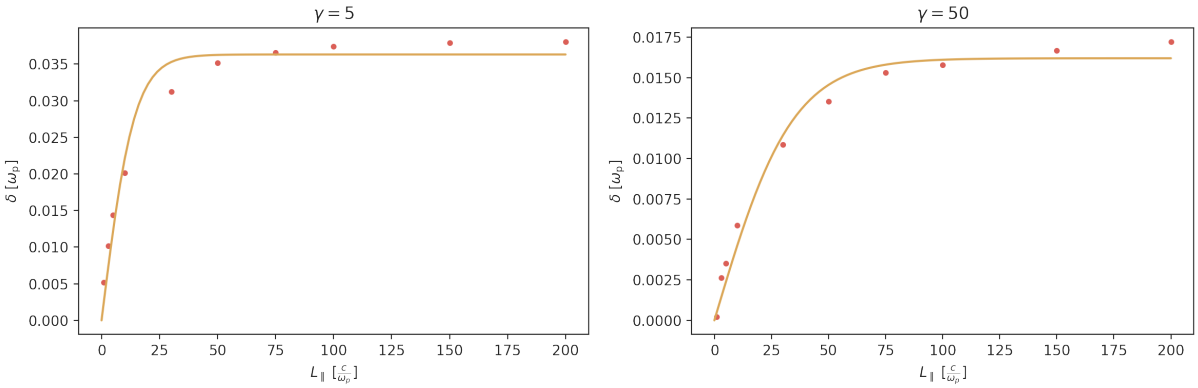

Figure 7.2: The integrated growth rate as a function of beam length for a density ratio $\alpha = 10^{-3}$ and Lorentz boost $\gamma = 5$ (**left**) and $\gamma = 50$ (**right**). In both cases the data conforms to (7.1) which has been fitted with only δ_0 as a free parameter.

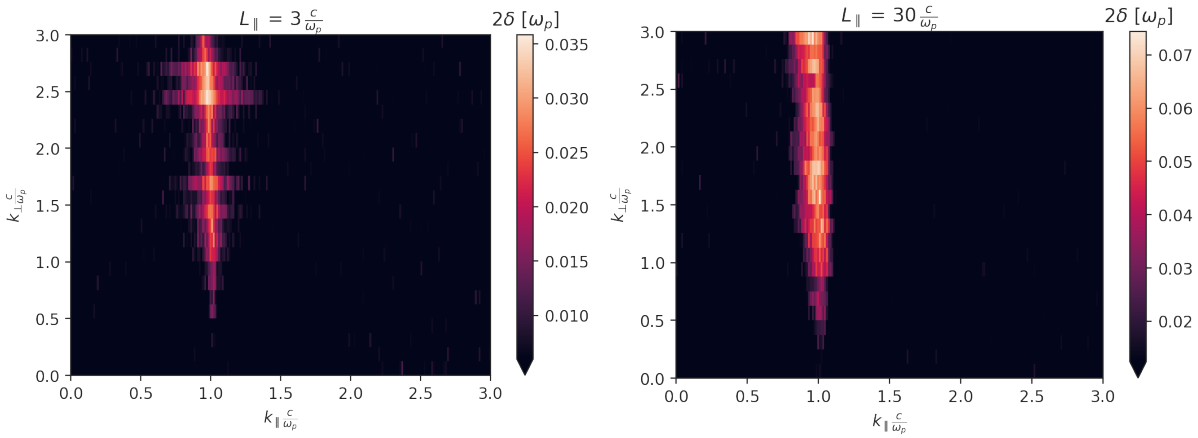

Figure 7.3: Map of the growth rate as a function of wavevector for each mode for two simulation runs. On the **left** the beam length $L_\parallel = 3\,\frac{c}{\omega_p}$ is lower than on the **right** where $L_\parallel = 30\,\frac{c}{\omega_p}$. For both $\gamma = 5$ and $\alpha = 10^{-3}$ are the same. For a shorter beam length an interference pattern in $k_\perp$ is present, which becomes less pronounced for longer beams.

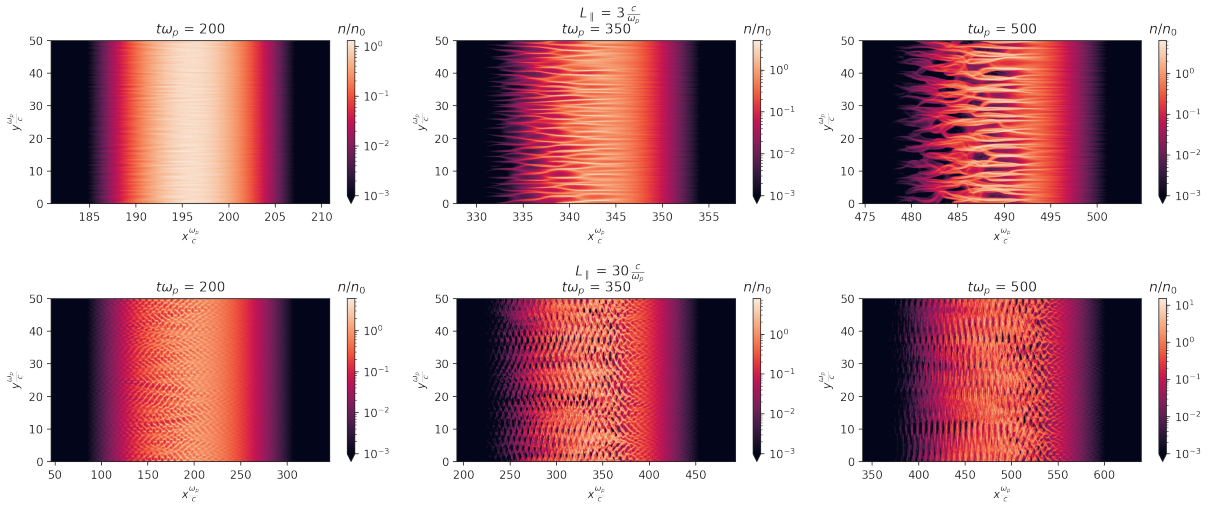

Figure 7.4: The density (normalized to the initial peak density n_0) as a function of position at three different time steps for a beam length $L_\parallel = 3\,\frac{c}{\omega_p}$ (**top**) and $L_\parallel = 30\,\frac{c}{\omega_p}$ (**bottom**). The beam moves towards the right.

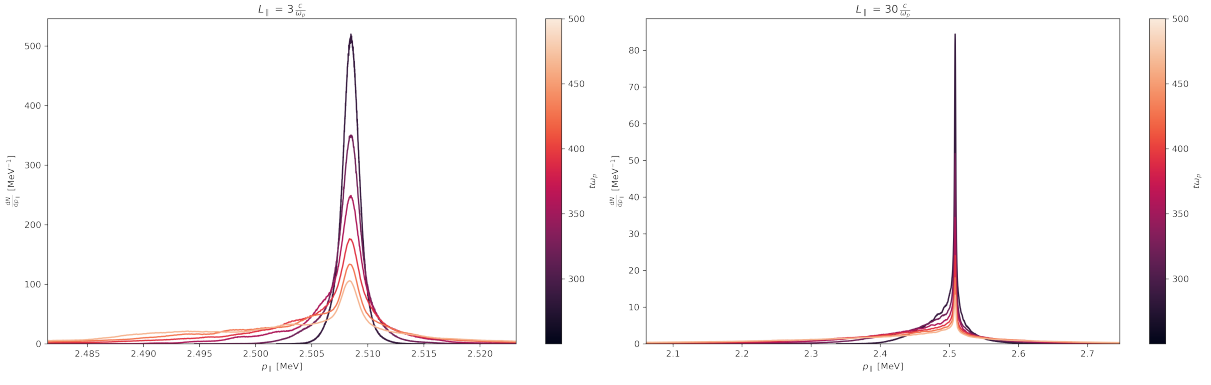

Figure 7.5: The longitudinal beam momentum distribution, while integrating over all other momentum dimension, for a beam with $L_\parallel = 3\,\frac{c}{\omega_p}$ (**left**) and $L_\parallel = 30\,\frac{c}{\omega_p}$ (**right**) at regular intervals near the end of the linear growth phase.

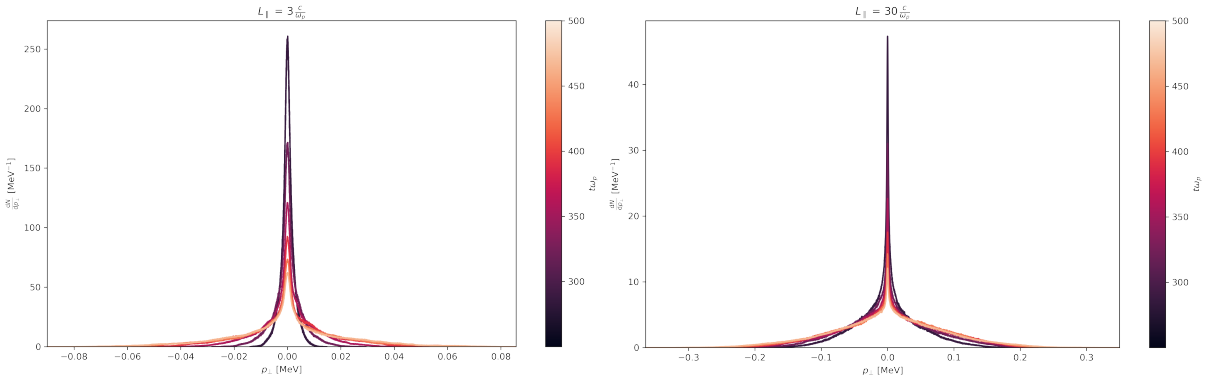

Figure 7.6: The transversal beam momentum distribution, while integrating over all other momentum dimension, for a beam with $L_\parallel = 3\,\frac{c}{\omega_p}$ (**left**) and $L_\parallel = 30\,\frac{c}{\omega_p}$ (**right**) at regular intervals near the end of the linear growth phase.

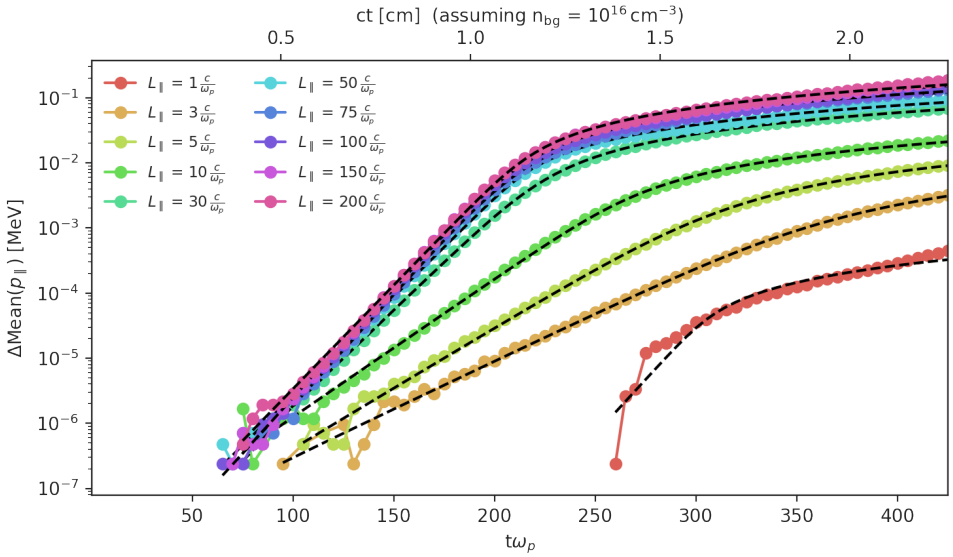

Figure 7.7: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ is shown for runs with varying beam length $L_\parallel$, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The mean value is calculated at regular intervals of $5\omega_p^{-1}$. However for early times the shift of the mean momentum is dominated by numeric effects (finite binning) and not shown here. Furthermore for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown either since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.42) to the data is shown.

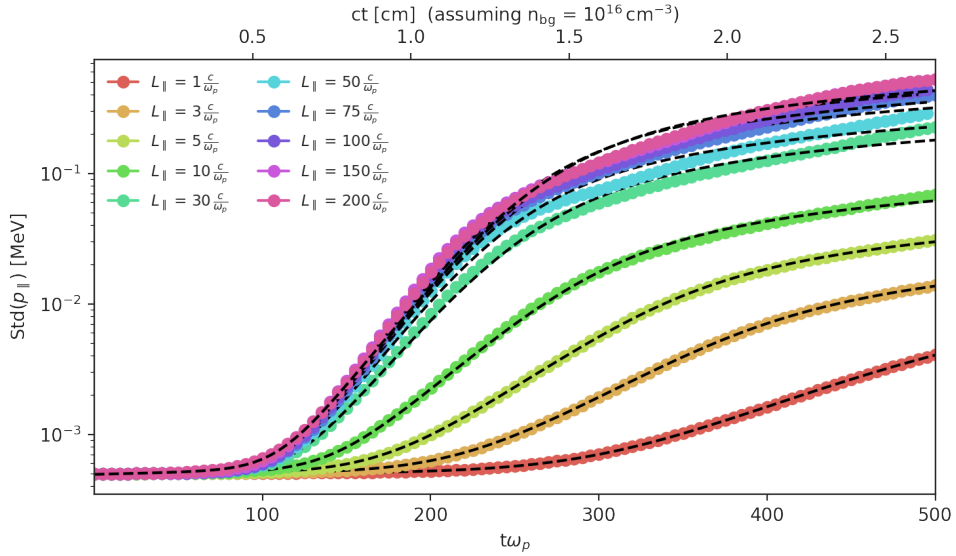

Figure 7.8: The evolution of the width of the longitudinal momentum Std $(p_\parallel)$ is shown for runs with varying beam length $L_\parallel$, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

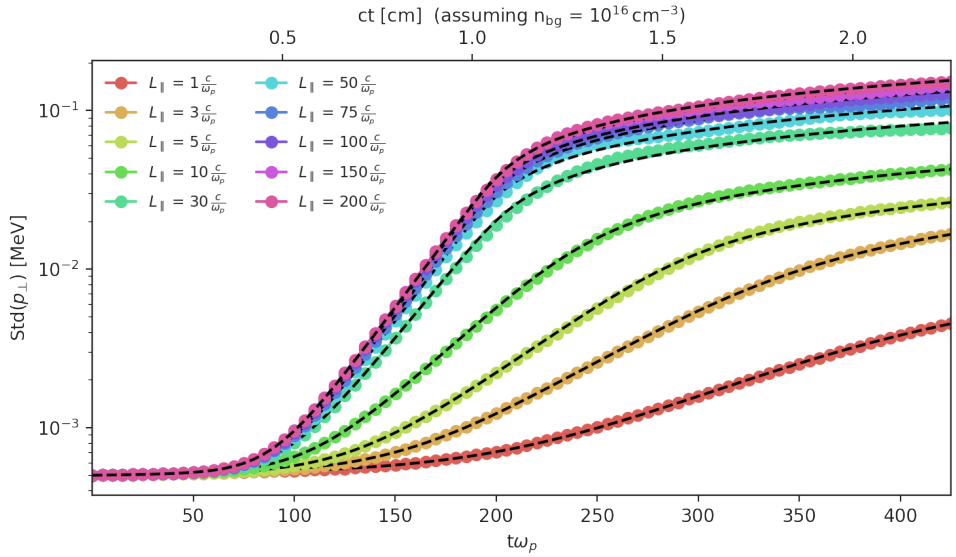

Figure 7.9: The evolution of the width of the transversal momentum Std $(p_\perp)$ is shown for runs with varying beam length $L_\parallel$, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

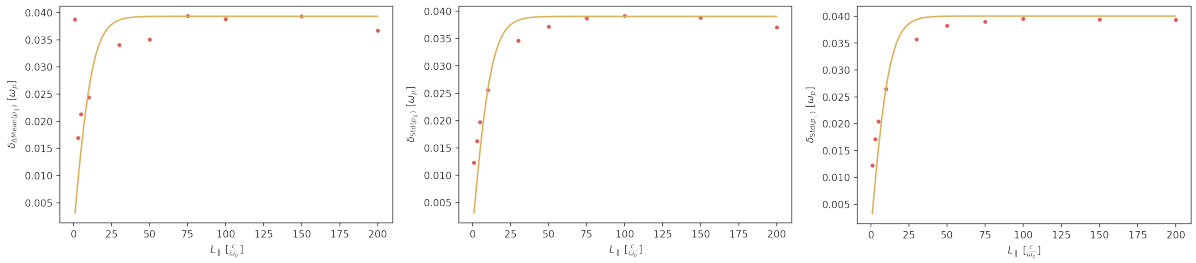

Figure 7.10: The growth rate of $\Delta\text{Mean}(p_\parallel)$ (**left**), $\text{Std}(p_\parallel)$ (**middle**) and $\text{Std}(p_\perp)$ (**right**) extracted from figures 7.7, 7.8 and 7.9 as a function of beam length $L_\parallel$. In all three cases the data conforms to (7.1) which has been fitted with δ_0 as a free parameter

7.1.2 Finite Width

We perform a study that is similar to the finite length also for the finite width. Here we consider a beam with a finite width $L_\perp$ with a Gaussian density profile in a simulation box that has a width of 50 $\frac{c}{\omega_p}$ and open boundary conditions in transverse direction. In figure 7.11 we show the evolution of the total field energy density for simulation runs with varying beam width. Again the field energy density is normalized to the initial beam energy and thus the energy density due to shot noise is larger for a narrower beam with less beam particles. The narrower beam also leads to a less steep slope of the exponential growth phase. The growth rate has been extracted by fitting (5.5) and plotted as a function of beam width in figure 7.12 for simulation runs with low and high Lorentz boost γ. We find that for very narrow beams the growth rate is almost proportional to the beam width $L_\perp$ and for very wide beam the growth rate asymptotically approaches the growth rate of an infinitely wide beam δ_0. Surprisingly the width at which the transition between these two regimes takes place is not dependent on γ and thus the growth rate. We find a heuristic formula that describes the growth rate as a function of $L_\perp$ and δ_0.

$$\delta(L_\perp) = \delta_0 \tanh\left(2L_\perp \frac{\omega_p}{c}\right) \tag{7.3}$$

While leaving δ_0 as a free parameter we performed a fit of (7.3) to the data points shown in figure 7.12. The fit describes the data quite well. This leads us to a condition on the beam width for which we recover the growth rate of the infinitely wide beam. This conditions essentially only requires that the beam is wide enough to show plasma-like collective behavior.

$$L_\perp \gg \frac{c}{\omega_p} \Rightarrow \delta(L_\perp) \approx \delta_0 \tag{7.4}$$

For the most dominant unstable modes (figure 7.13) we find that the beam width has no influence except for an interference pattern that suppresses the growth rate in regular intervals of $k_\parallel$. Again this could be an artifact of the discrete Fourier transform and the open boundary conditions.
For a beam with a density that tapers off transversally the density ratio of beam to background plasma α is a function of the position. Consequently the instable modes will evolve differently as a function of position as well. Figure 7.14 shows the beam plasma density as a function of position for a number of time steps. We can see that the initial Gaussian distribution evolves a structure. Furthermore at late times we observe that the density at the center of the beam is

depleted and we find two new maximums to the left and the right from the center in transversal direction. The beam particles gain transversal momentum due to the instability growth and start to move away from the center of the beam. This increases the beam plasma density at the outer edges of the beam and the instability can proceed there with increased growth rate. As a further consequence of the gain of transversal momentum the beam widens over time possibly posing a challenge for the further transport and measurement in a laboratory experiment. The widening of the beam will be separately discussed.
The splitting of beam particles moving in opposite transversal directions and the delayed instability growth in the outer edges of the beam can also be observed in the beam momentum distribution. For the longitudinal momentum distribution (figure 7.15) we see that a narrow peak on top of an already broadened momentum spectrum that disappears at a later point. For the transversal momentum distribution (figure 7.16) we can see this lingering peak as well. This peak consists of the particles at the outer edges that remain in their initial momentum distribution for a longer time. Furthermore in the transversal momentum distribution we can also observe the development of two peaks of left and right moving particles at later times.
Despite the more complex evolution of the momentum distribution the evolution of the moments $\Delta\mathrm{Mean}(p_\parallel)$ (figure 7.17), $\mathrm{Std}(p_\parallel)$ (figure 7.18), $\mathrm{Std}(p_\perp)$ (figure 7.19) can still be described with the equations derived from the Fokker-Planck equation albeit with a reduced overall growth rate. Figure 7.20 shows the growth rates as a function of the beam width, which all conform to the relationship (7.3) as confirmed by a fit.

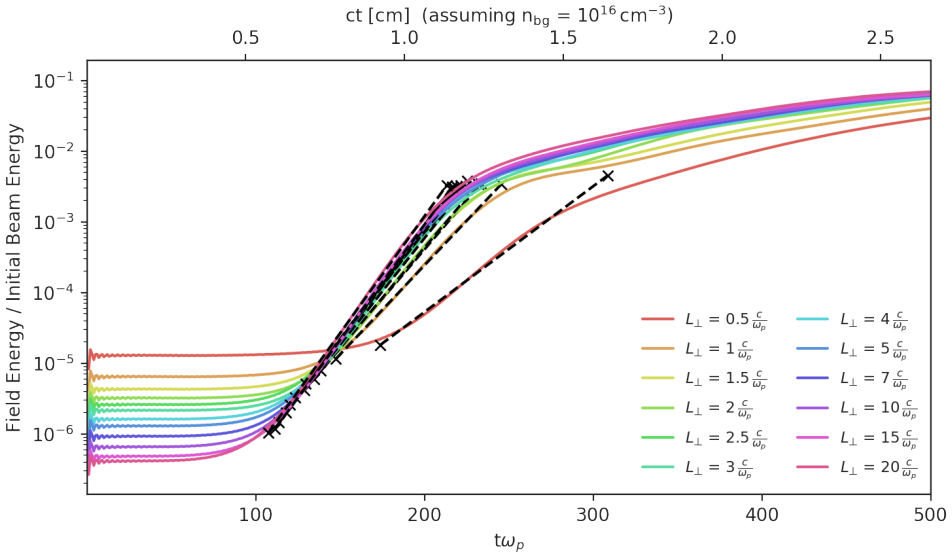

Figure 7.11: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time. Each simulation run has a density contrast of $\alpha = 10^{-3}$, Lorentz boost $\gamma = 5$ and a varying beam length $L_\perp$. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line. The initial noise regime is mostly driven by the background plasma and increases for a narrow beam (less particles) compared to the beam energy.

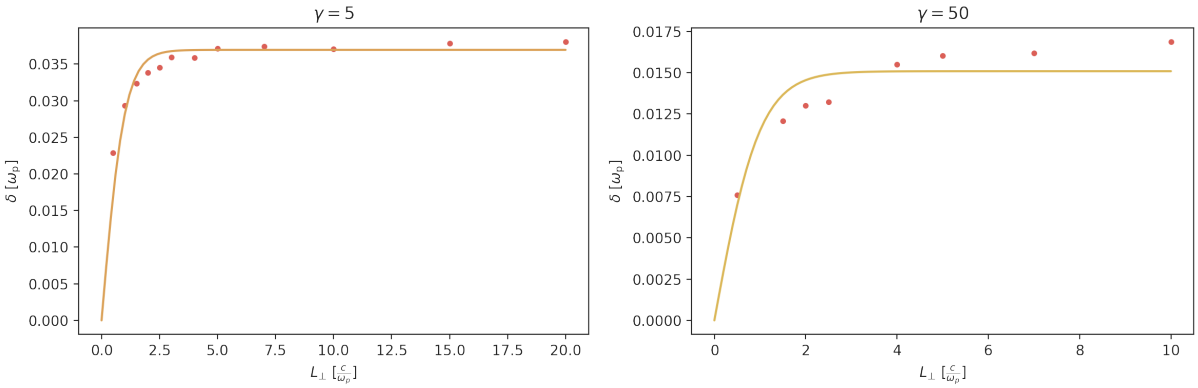

Figure 7.12: The integrated growth rate as a function of beam width for a density ratio $\alpha = 10^{-3}$ and Lorentz boost $\gamma = 5$ (**left**) and $\gamma = 50$ (**right**). In both cases the data conforms to (7.3) which has been fitted with only δ_0 as a free parameter.

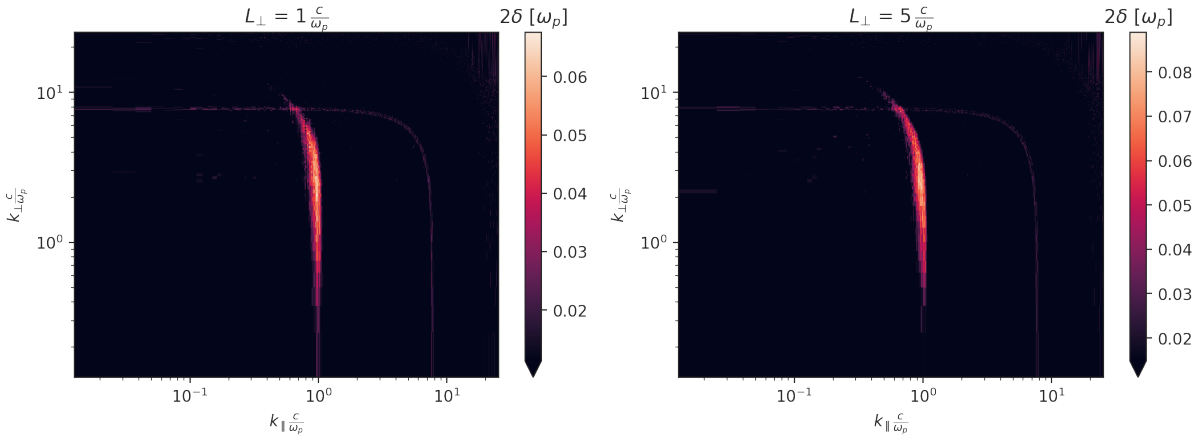

Figure 7.13: Map of the growth rate as a function of wavevector for each mode for two simulation runs. On the **left** the beam width $L_\perp = 1\,\frac{c}{\omega_p}$ is lower than on the **right** where $L_\perp = 5\,\frac{c}{\omega_p}$. For both $\gamma = 5$ and $\alpha = 10^{-3}$ are the same. For a narrower beam width an interference pattern in $k_\parallel$ is present, which becomes less pronounced for wider beams.

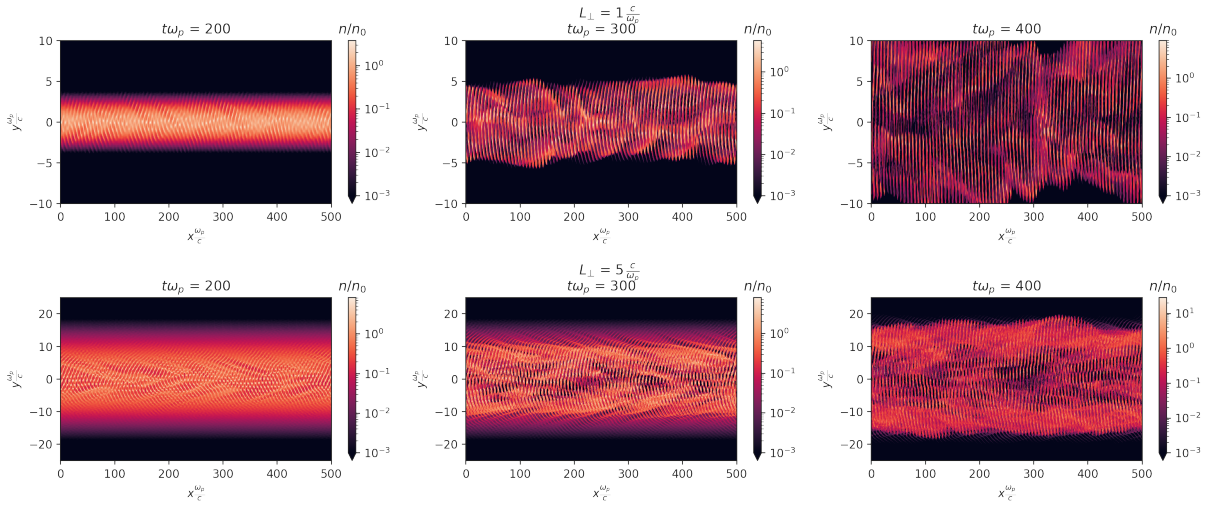

Figure 7.14: The density (normalized to the initial peak density n_0) as a function of position at three different time steps for a beam width $L_\perp = 1\,\frac{c}{\omega_p}$ (**top**) and $L_\perp = 5\,\frac{c}{\omega_p}$ (**bottom**). The beam moves towards the right.

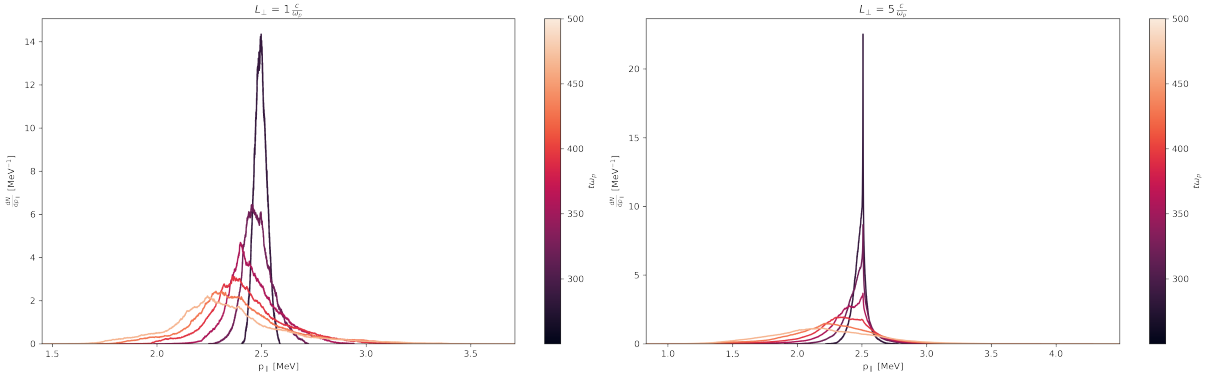

Figure 7.15: The longitudinal beam momentum distribution, while integrating over all other momentum dimensions, for a beam with $L_\perp = 1\,\frac{c}{\omega_p}$ (**left**) and $L_\perp = 5\,\frac{c}{\omega_p}$ (**right**) at regular intervals near the end of the linear growth phase.

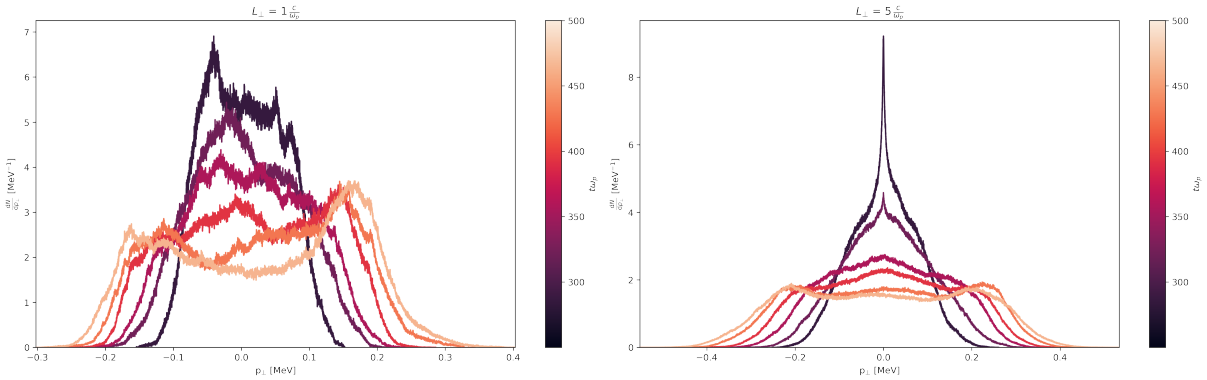

Figure 7.16: The transversal beam momentum distribution, while integrating over all other momentum dimensions, for a beam with $L_\perp = 1\,\frac{c}{\omega_p}$ (**left**) and $L_\perp = 5\,\frac{c}{\omega_p}$ (**right**) at regular intervals near the end of the linear growth phase.

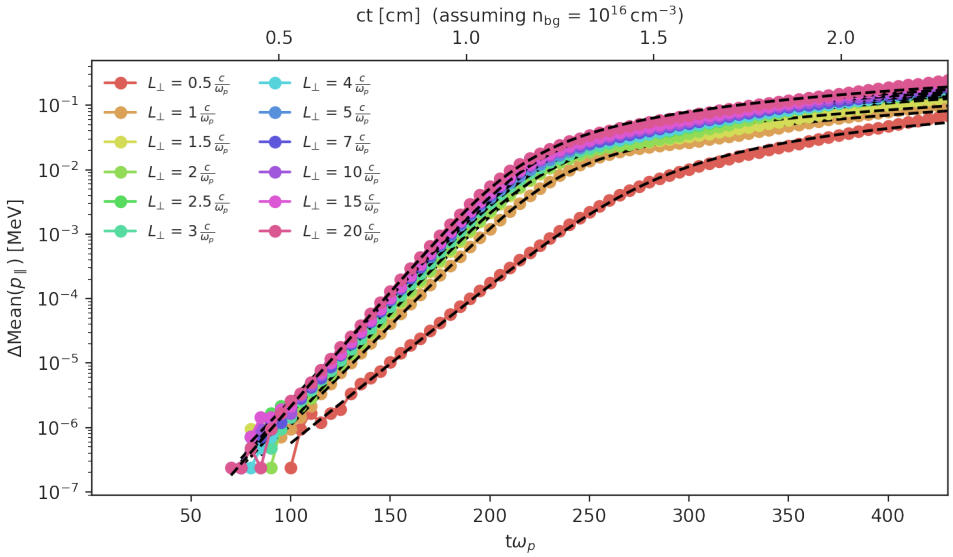

Figure 7.17: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ is shown for runs with varying beam width $L_\perp$, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The mean value is calculated at regular intervals of $5\omega_p^{-1}$. However for early times the shift of the mean momentum is dominate by numeric effects (finite binning) and not shown here. Furthermore late times, after $2t_{\text{final}}$ (fitted value), data points are not shown either since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.42) to the data is shown.

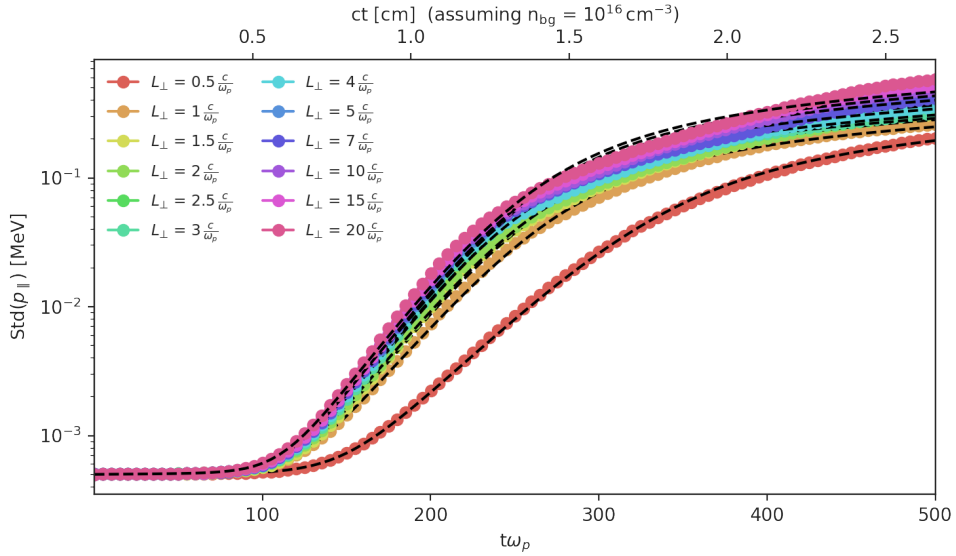

Figure 7.18: The evolution of the width of the longitudinal momentum Std $(p_\parallel)$ is shown for runs with varying beam width $L_\perp$, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

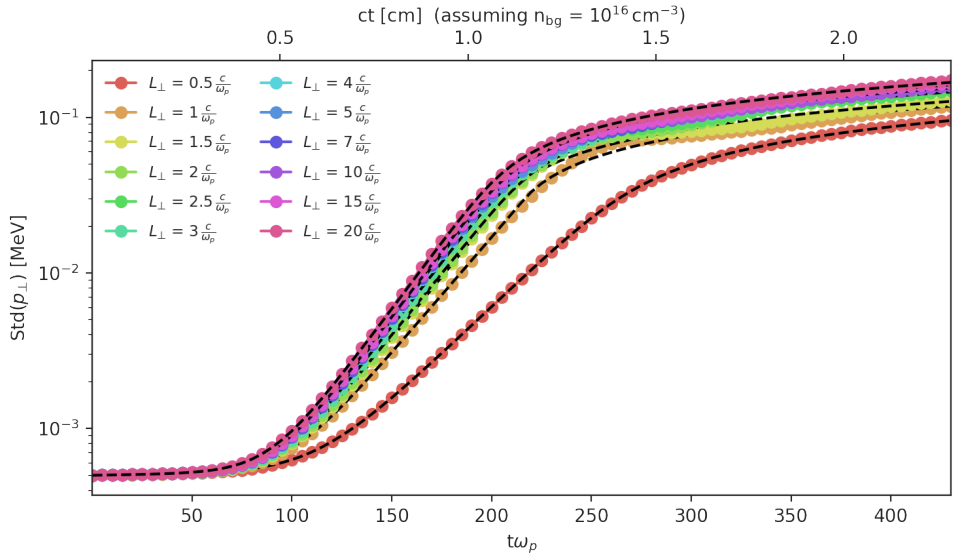

Figure 7.19: The evolution of the width of the transversal momentum Std $(p_\perp)$ is shown for runs with varying beam width $L_\perp$, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

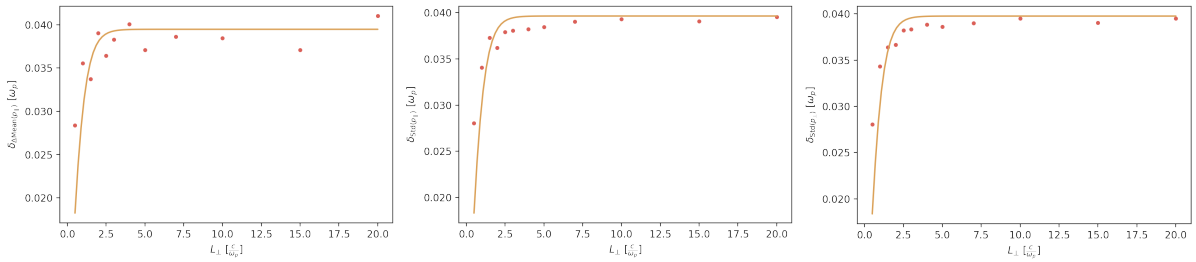

Figure 7.20: The growth rate of $\Delta\text{Mean}(p_\parallel)$ (**left**), $\text{Std}(p_\parallel)$ (**middle**) and $\text{Std}(p_\perp)$ (**right**) extracted from figures 7.17, 7.18 and 7.19 as a function of beam width $L_\perp$. In all three cases the data conforms to (7.3) which has been fitted with δ_0 as a free parameter

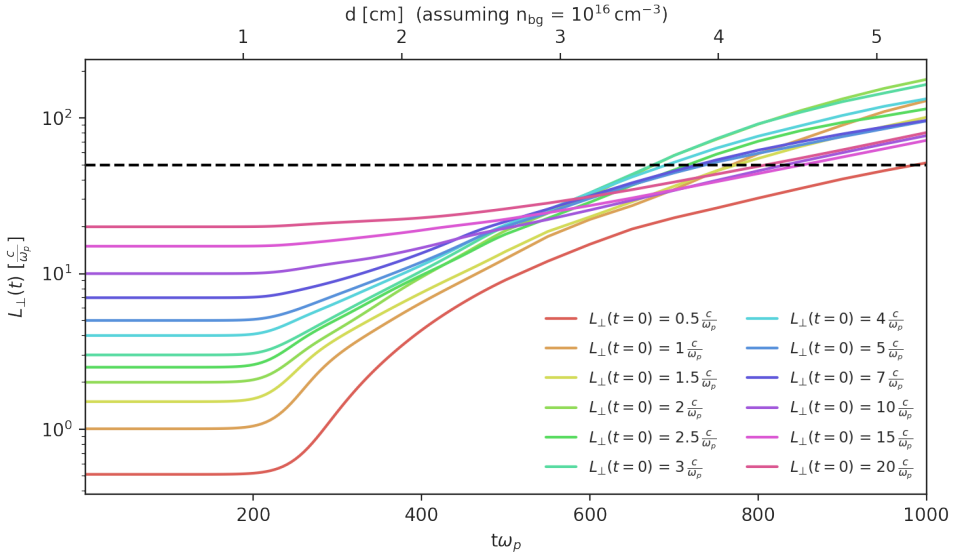

Figure 7.21: The evolution of the beam width $L_\perp$ as a function of time for runs with varying initial beam width $L_\perp\,(t=0)$. The dashed line represents the width of the simulation box ($50\,\frac{c}{\omega_p}$). Particles which stray too far from the center of the beam leave the simulation box through the open boundaries and are discarded.

Beam Broadening

Already in section 5.6.1 we showed that the transversal heating of the beam due to the instability growth can increase the beam opening angle of an initially cold beam by orders of magnitude. The final opening angle after the instability saturates depends on the values of α and γ, but for values that could be realistic in a laboratory experiment the final opening angle could exceed 1°. In the case of a beam with finite width this means a beam that becomes broader with time or distance. In a laboratory experiment using the momentum distribution as an observable the beam would need to propagate some distance before reaching the instrument used for the measurement and the increased width could spoil the measurement.
Figure 7.21 shows the evolution of the beam width $L_\perp$, as defined by the standard deviation of the transversal position distribution, with time. It is important to note that the simulation box

is note infinitely wide and that particles which stray too far from the center of the beam leave the simulation box and are discarded. Once the width of the beam becomes comparable to the simulation box width the evolution is not necessarily trustworthy anymore. After an early phase of exponential widening that is more pronounced for initially narrower beams follows a regime where the transverse momentum has saturated and the beam continues to widen with a constant rate, as expected of a beam with a constant opening angle.

7.2 Warmer Beam

Even thought the shifted Gaussian momentum distribution we use for the beam is not strictly a thermal distribution we can use the term "cold" to describe that the momentum spread is small compared to the total momentum. We will use the standard deviation σ of the momentum distribution to quantify the momentum spread. Since the beam plasma is not necessarily a thermalized system it is also possible for the momentum spread to be different in longitudinal and transversal direction. In that sense a beam could be longitudinally hot and transversally cold or vice versa. Furthermore we have not yet given a quantitative assessment under which conditions a beam is considered cold.
For the oblique instability one can estimate that the beam transitions from the reactive to the kinetic regime when the inequality (3.25) is no longer fulfilled. This indicates that only the transversal temperature is relevant for this transition. Using the oblique growth rate we can estimate the transversal momentum spread $\sigma_\perp$ above which the instability is kinetic.

$$\sigma_\perp \sim m_e \alpha^{\frac{1}{3}} \gamma^{\frac{2}{3}} \tag{7.5}$$

The scaling of this transition with α and γ also traces the scaling of the saturation level further indicating that the saturation occurs when the instability turns kinetic.
The scenario of a longitudinally hot and transversally cold, or at least mildly warm, beam is the one that probably most closely resembles the astrophysical pair beam. The astrophysical pair beam is created from gamma rays that follow a power-law distribution over many orders of magnitude in energy. The resulting electron positron pair will come with an opening angle of roughly γ^{-1}. The highly relativistic particle will thus have a broad longitudinal momentum spectrum and a transversal momentum spread on the order of the electron mass. This indicates that for the parameters α, γ that describe an astrophysical pair beam the system falls close to the transition regime between cold and hot beam.

$$\sigma_\perp \approx \sin\left(\gamma^{-1}\right) p_{\text{tot}} = \sin\left(\gamma^{-1}\right) \gamma m_e \approx m_e \tag{7.6}$$

In the previous PIC simulations we assumed that a momentum spread of 0.5 keV (all directions) is sufficiently small for a beam to be considered cold under any circumstances. We will find that this is correct.

7.2.1 Longitudinal Temperature

We run multiple simulations of a neutral pair beam with $\gamma = 10$, $\alpha = 10^{-3}$ and $\sigma_\perp = 0.5\,\text{keV}$ where we vary the initial longitudinal momentum spread from very cold (0.5 keV) to very hot (5 MeV). For the largest choice of $\sigma_\parallel$ the longitudinal momentum spread is in fact so large

that the momentum distribution has a significant fraction of particles with negative longitudinal momentum, that means they are in fact counter-streaming.
We find that the initial longitudinal momentum spread has little to no influence on the evolution of the instability. In the evolution of the field energy density (figure 7.22), the mean longitudinal momentum shift (figure 7.23) and the transversal width of the momentum distribution (figure 7.25) we can only see a small reduction in the growth rate once the initial longitudinal width of the momentum distribution becomes large enough that a significant fraction of the particles have negative momentum.
In figure 7.24 we can see that if the initial width of the momentum distribution is already larger than the broadening we expect from the instability the effect remains invisible, even though the instability can proceed as usual in field energy density or transversal momentum.
The increase of the initial longitudinal momentum spread only reduces the overall growth rate if some particle start with negative parallel momentum and do not participate in the instability process. The particles with $p_{\parallel} < 0$ do not participate in the instability because they do not experience a momentum inversion. For this particle population the slope of the momentum distribution in the direction of movement is always negative. This can be observed for simulation runs with $\sigma_{\parallel} = 5\,\mathrm{MeV}$ in figure 7.27 where the momentum distribution of particles with negative parallel momentum does not evolve and in figure 7.28 where a subset of particles remain in their initially narrow momentum distribution.

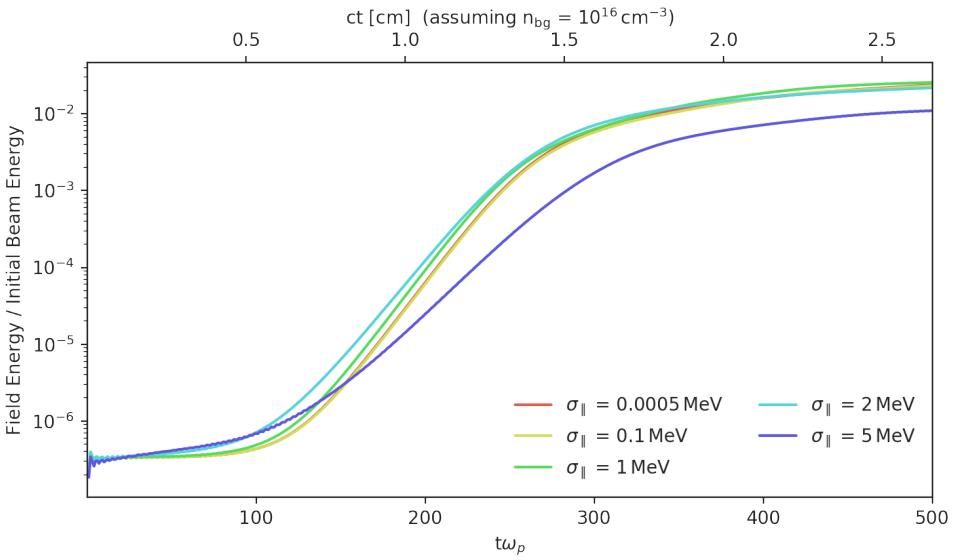

Figure 7.22: The energy density of the electromagnetic fields of various simulation runs with varying initial longitudinal momentum spread $\sigma_{\parallel}$ as a function of time.

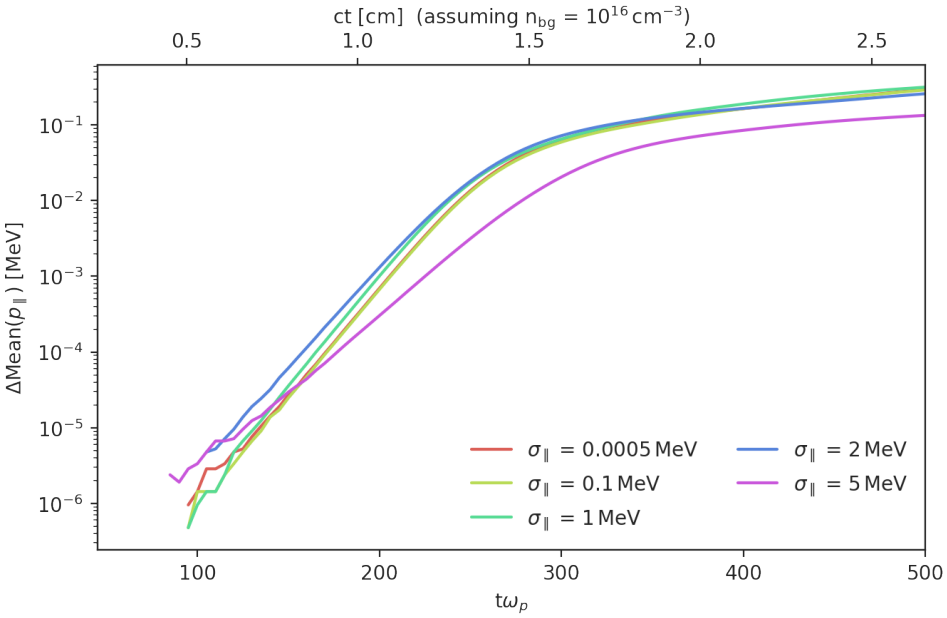

Figure 7.23: The shift of the mean of the longitudinal momentum distribution $\Delta\text{Mean}(p_\parallel)$ as a function of time for simulation runs with a varying initial longitudinal momentum spread $\sigma_\parallel$. The mean value is calculated at regular intervals of $5\omega_p^{-1}$. However for early times the shift of the mean momentum is dominated by numeric effects (finite binning) and not shown here.

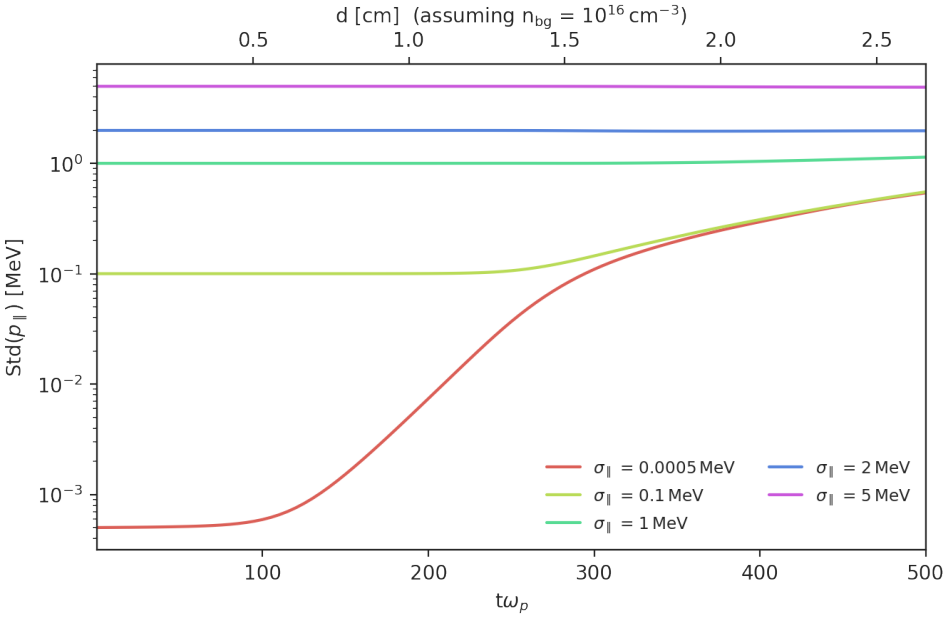

Figure 7.24: The longitudinal momentum spread $\text{Std}(p_\parallel)$ as a function of simulation time for various simulation runs with varying initial longitudinal momentum spread $\sigma_\parallel$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

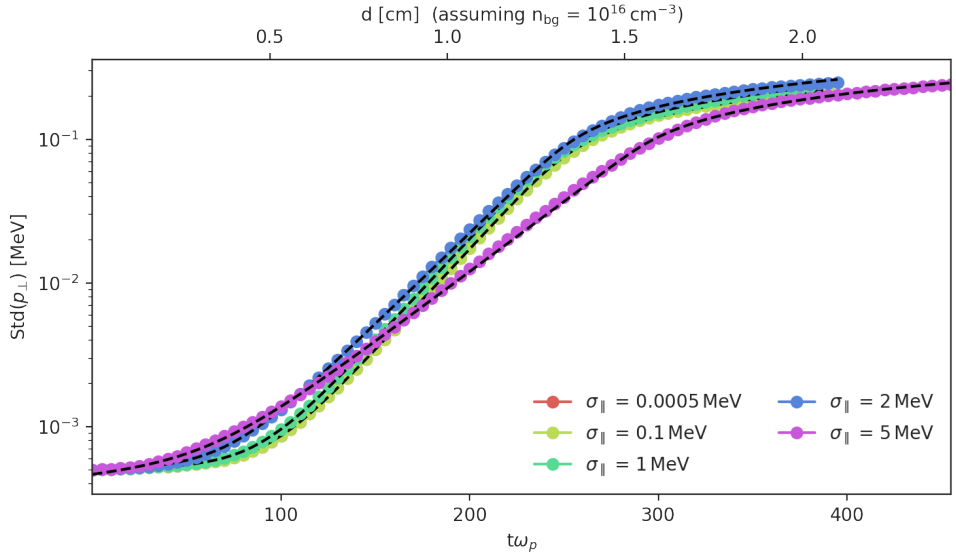

Figure 7.25: The transversal momentum spread $\mathrm{Std}(p_\perp)$ as a function of simulation time for various simulation runs with varying initial longitudinal momentum spread $\sigma_\parallel$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However for late times, after $2t_{\mathrm{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. The dashed lines represent a fit of (3.43 to the data). The extracted growth rate is shown in figure 7.26.

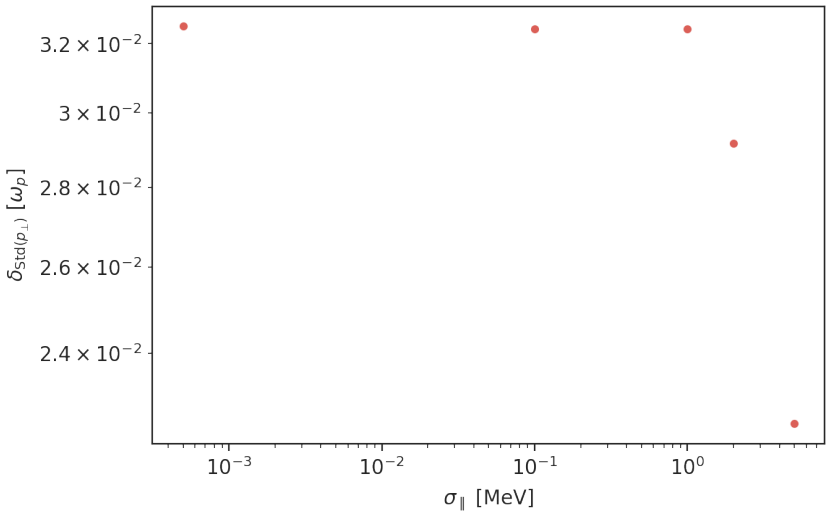

Figure 7.26: The growth rate of the transversal momentum spread $\mathrm{Std}(p_\perp)$ extracted from figure 7.25 as a function of initial longitudinal momentum spread $\sigma_\parallel$.

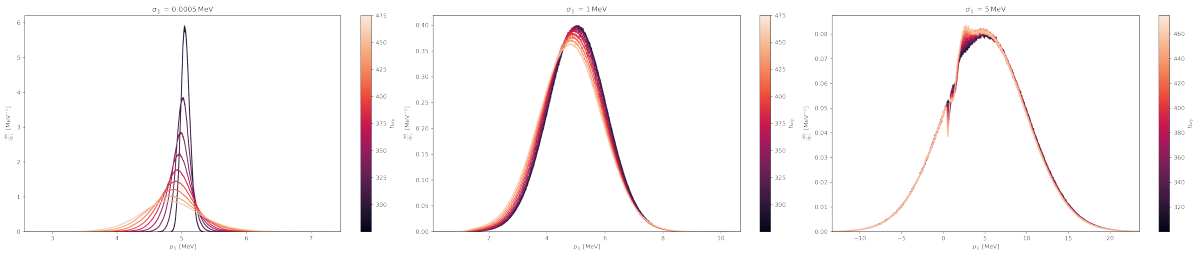

Figure 7.27: The longitudinal beam momentum distribution at a number of time steps near the end of instability growth for simulation runs with varying initial longitudinal momentum spread $\sigma_\parallel$. In the **right** plot the momentum distribution only evolves for particles with $p_\parallel > 0$.

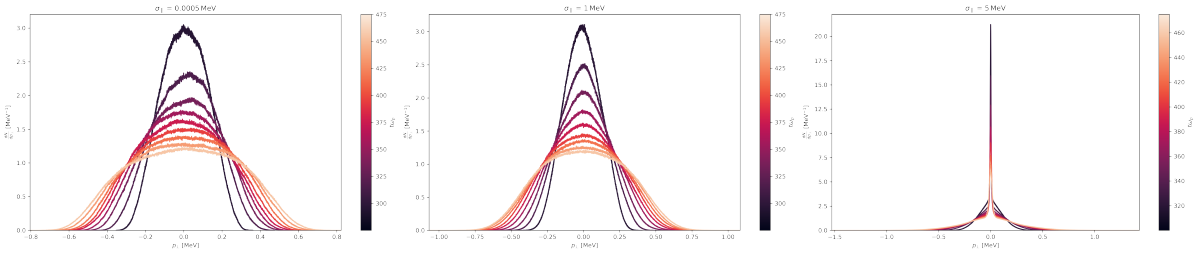

Figure 7.28: The transversal beam momentum distribution at a number of time steps near the end of instability growth for simulation runs with varying initial longitudinal momentum spread $\sigma_\parallel$. In the **right** plot a fraction of the particles does not undergo the momentum diffusion process, likely the particles with negative parallel momentum.

7.2.2 Transversal Temperature

For the transversal temperature we perform the same scan of the initial transversal momentum spread $\sigma_\perp$ from very cold (0.5 keV) to hot (larger than m_e) while keeping the Lorentz boost γ, the density ratio α and the initial longitudinal momentum spread $\sigma_\parallel = 0.5\,\text{keV}$ constant. We define a beam as cold when its momentum spread is small compared to its total momentum and thus the definition of coldness is intimately linked to the total momentum of the beam. Whether a beam is cold or not has to be answered in context of its Lorentz boost γ. More precisely a higher value means that the threshold for the transversal energy spread above which a beam is no longer cold increases and if a certain momentum spread leads to a cold beam at low γ then it will be cold for all higher values as well. For the parameters $\alpha = 10^{-3}$ and $\gamma = 10$ we expect the transition from the reactive (cold) beam to the kinetic (hot) beam for $\sigma_\perp \sim 0.2\,\text{MeV}$ according to (7.5). This criterion is more strict than simply requiring that the transversal momentum spread is small compare to its total momentum as the growth rate can become arbitrarily small.
The time evolution of the field energy density (figure 7.29) and the beam energy loss (figure 7.30) show that increasing the initial transverse momentum spread quickly diminishes the growth of the instability. Surprisingly a larger initial transverse momentum spread leads to a very fast buildup of weak electromagnetic fields that goes hand in hand with a very rapid but overall minor loss of beam energy. A larger transverse momentum spread leads to stronger fields and larger energy loss of the beam. This early electromagnetic field also leads to an early diffusion of longitudinal beam momentum (figure 7.31) where the longitudinal momentum spread increases

faster with larger $\sigma_\perp$ before the stronger diffusion in the instability created fields takes over and the momentum spread grows faster with lower $\sigma_\perp$.
Even though a higher initial transversal momentum spread leads to a slower instability growth the saturation levels of the field energy density, the beam energy loss and longitudinal momentum spread are unaffected. They only reach the same level at a later saturation time. For the evolution of the transversal momentum spread $\mathrm{Std}(p_\perp)$ (figure 7.32) we similarly find that as long as the initial momentum spread is not larger than the expected saturation level the instability mechanism will eventually broaden the momentum distribution to the saturation level spread although here too a higher initial spread $\sigma_\perp$ leads to a slower growth rate. This can also be explicitly seen in figure 7.33 where the growth rate of the beam energy loss and the growth rate of the transversal momentum spread are shown as a function of $\sigma_\perp$. Increasing the initial momentum spread at first only leads to a reduction of the growth rate and only once $\sigma_\perp$ exceeds the expected saturation level of $\mathrm{Std}(p_\perp)$ do the rates vanish. The transversal momentum distribution (figure 7.35) remains unchanged for $\sigma_\perp$ larger than the saturation value, whereas the longitudinal momentum distribution (figure 7.34) still undergoes some minor evolution. Increasing the initial transversal momentum spread does not only affect the growth rate of unstable modes, but also alters which modes are excited in the first place. Figure 7.36 shows the differential power spectrum at different time steps. We know that for a cold beam the oblique instability excites modes with $k_\parallel \sim \frac{\omega_p}{c}$ and a wide spectrum in $k_\perp$ with the most dominant modes found at an oblique angle. A larger initial transversal momentum spread reduces the range of $k_\perp$ for which the modes experience significant growth and the most dominant modes are found at lower $k_\perp$ as well. This can be attribute to the fact that according to (3.24) modes with higher transverse wavenumber turn kinetic with lower transverse momentum spread. In figure 7.37 we show the power of the most excited mode $500\,\omega_p^{-1}$ after simulation start, which is long after the saturation time of the cold beam. Even for $10^{-2}\,\mathrm{MeV} < \sigma_\perp < 10^{-1}\,\mathrm{MeV}$ these modes saturate at the same level, even though we find that for $\sigma_\perp > 10^{-2}\,\mathrm{MeV}$ the growth rate is already reduced. For $\sigma_\perp > 10^{-1}\,\mathrm{MeV}$ the power of the most dominant mode reduces with larger $\sigma_\perp$. This suggests that even for very large $\sigma_\perp$ the instability can still grow some modes that are parallel to the beam direction, albeit with drastically reduced growth rate.

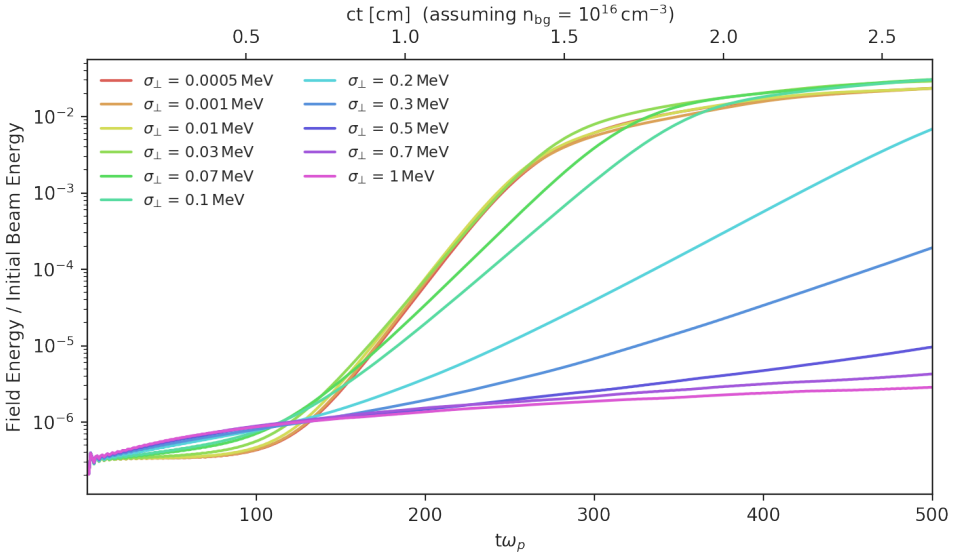

Figure 7.29: The energy density of the electromagnetic fields of various simulation runs with $\gamma = 10$, $\alpha = 10^{-3}$ and varying initial transversal momentum spread $\sigma_\perp$ as a function of time.

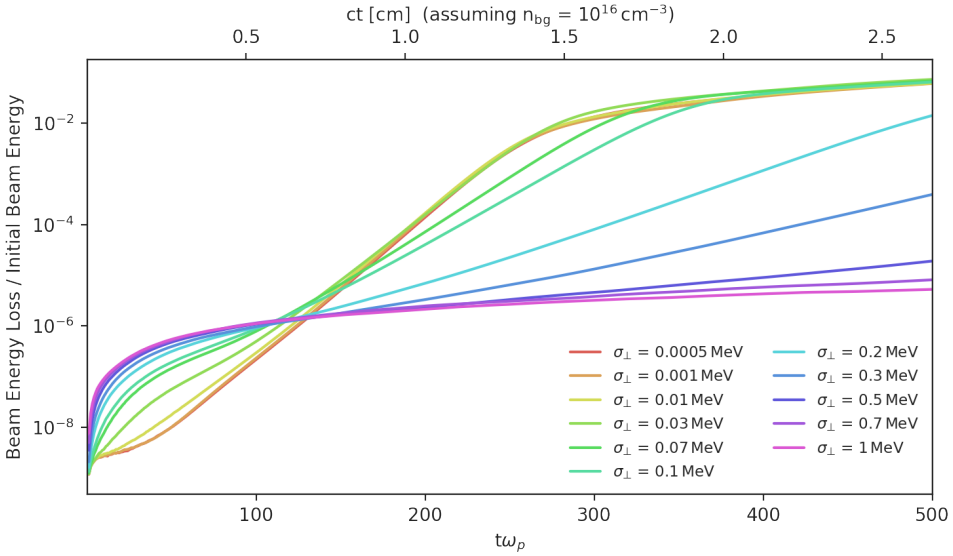

Figure 7.30: The energy loss of the beam particles normalized to the initial beam energy of various simulation runs with $\gamma = 10$, $\alpha = 10^{-3}$ and varying initial transversal momentum spread $\sigma_\perp$ as a function of time.

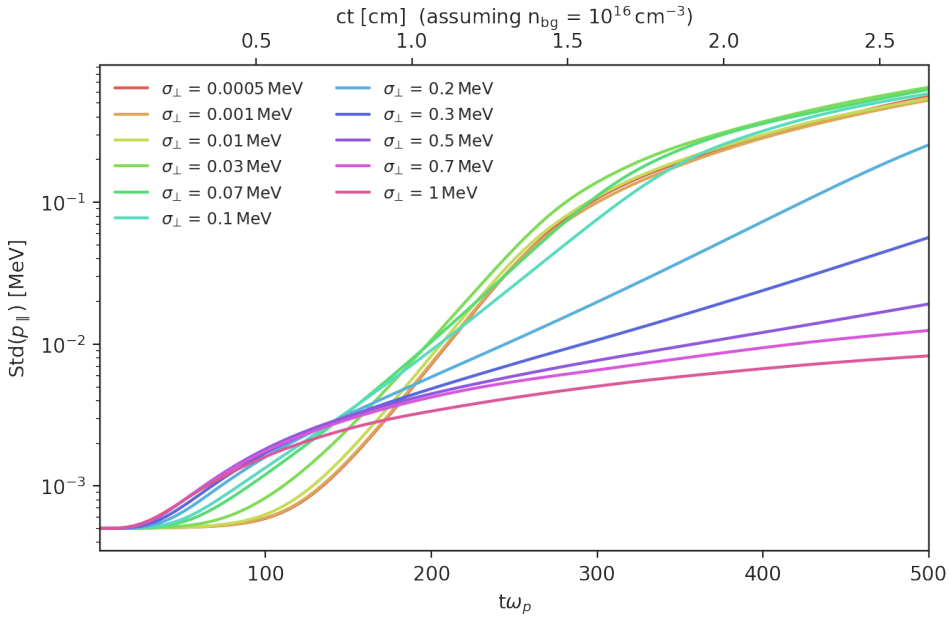

Figure 7.31: The longitudinal momentum spread $\mathrm{Std}(p_\parallel)$ as a function of simulation time for various simulation runs with varying initial transversal momentum spread $\sigma_\perp$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

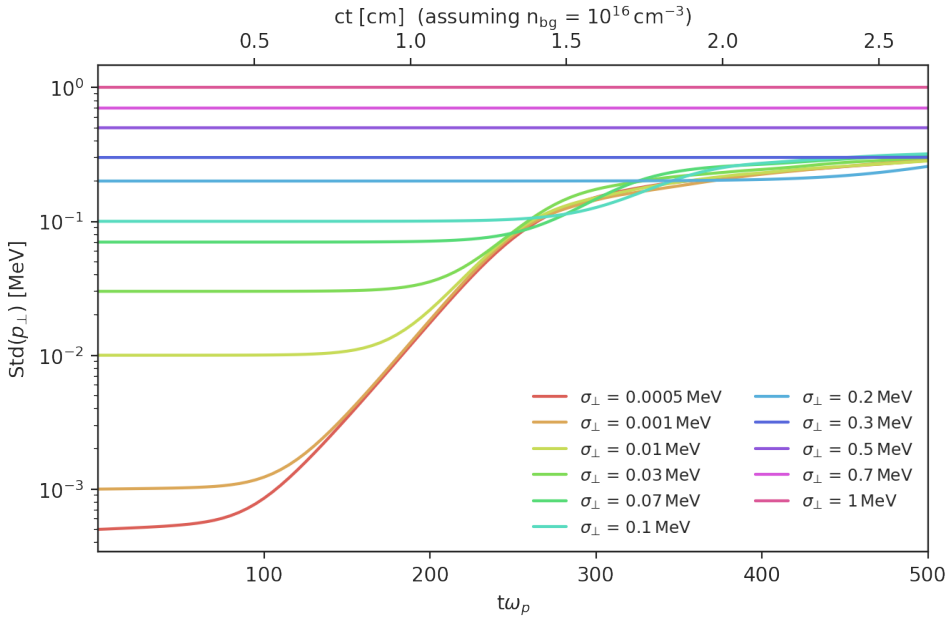

Figure 7.32: The transversal momentum spread $\mathrm{Std}(p_\perp)$ as a function of simulation time for various simulation runs with varying initial transversal momentum spread $\sigma_\perp$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

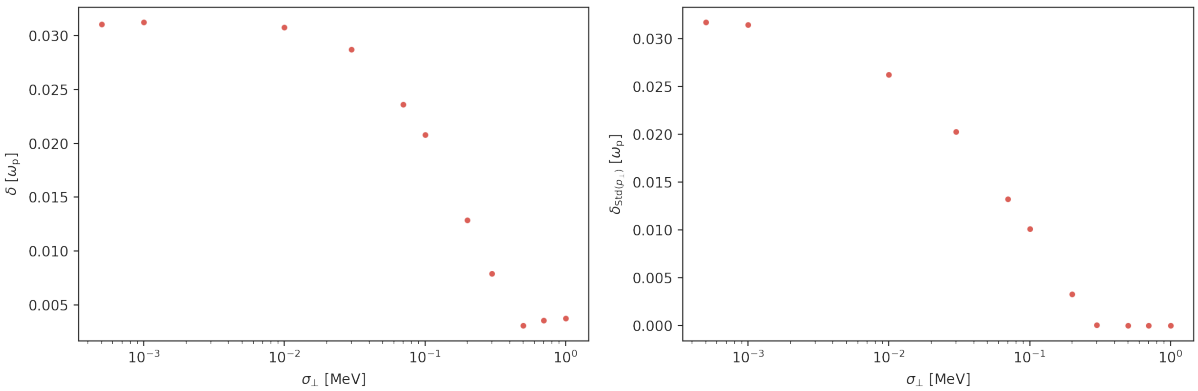

Figure 7.33: The beam energy loss rate (**left**) extracted from figure 7.30 and the growth rate of the transverse momentum spread $\mathrm{Std}(p_\perp)$ extracted from figure 7.32 as a function of the initial transverse momentum spread $\sigma_\perp$. For a beam with $\gamma = 10$ the total momentum is $p \approx 5\,\mathrm{MeV}$.

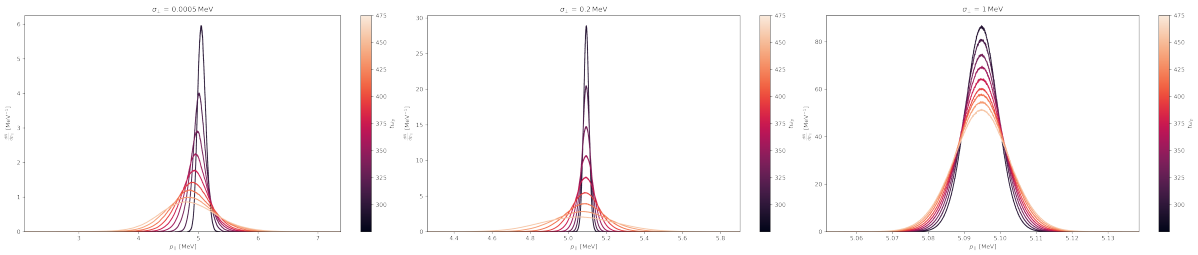

Figure 7.34: The longitudinal beam momentum distribution at a number of time steps near the end of instability growth for simulation runs with varying initial transversal momentum spread $\sigma_\perp$.

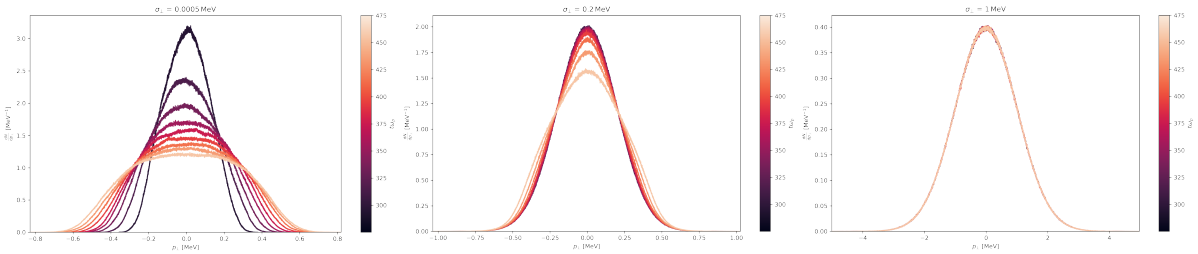

Figure 7.35: The transversal beam momentum distribution at a number of time steps near the end of instability growth for simulation runs with varying initial transversal momentum spread $\sigma_\perp$.

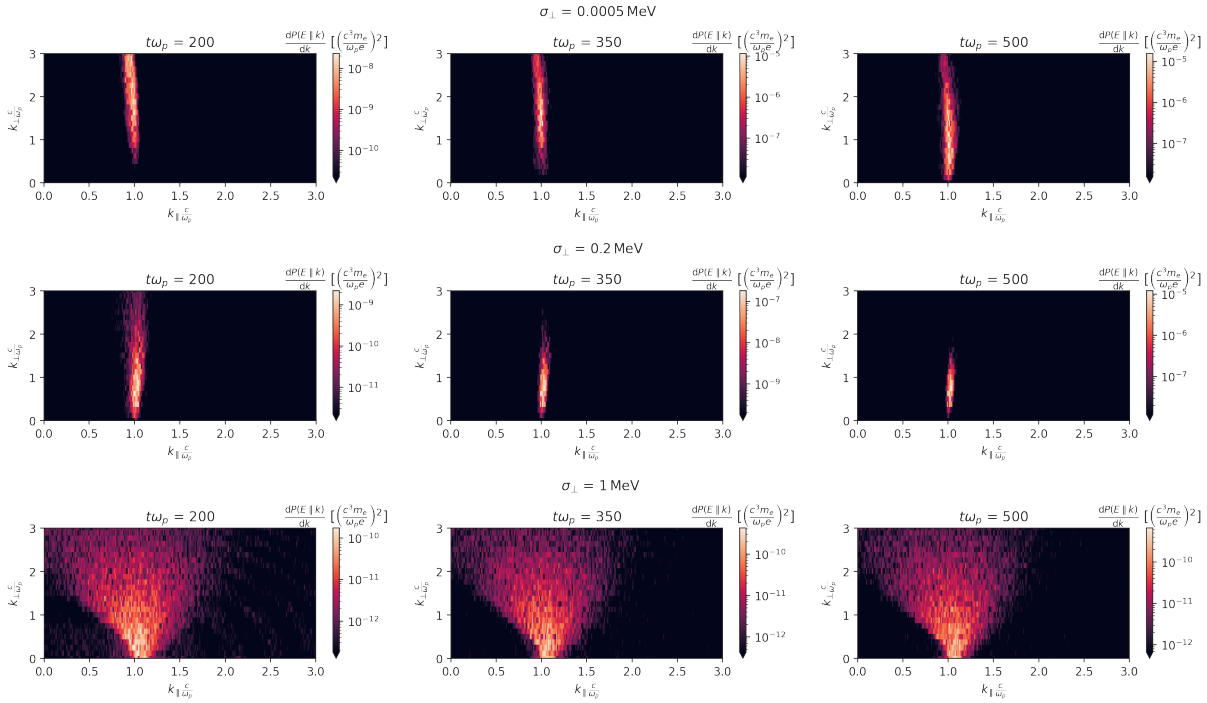

Figure 7.36: The differential power spectrum for a simulation runs with transversally cold (**top**), mildly warm (**middle**) and hot (**bottom**) momentum distribution and $\gamma = 5$, $\alpha = 10^{-3}$ at three different time steps. The most dominant modes are always found at $k_\parallel \sim \frac{\omega_p}{c}$, but for the hot beam other modes also contribute significantly. For higher $\sigma_\perp$ dominating modes becomes confined to smaller region in $k_\perp$.

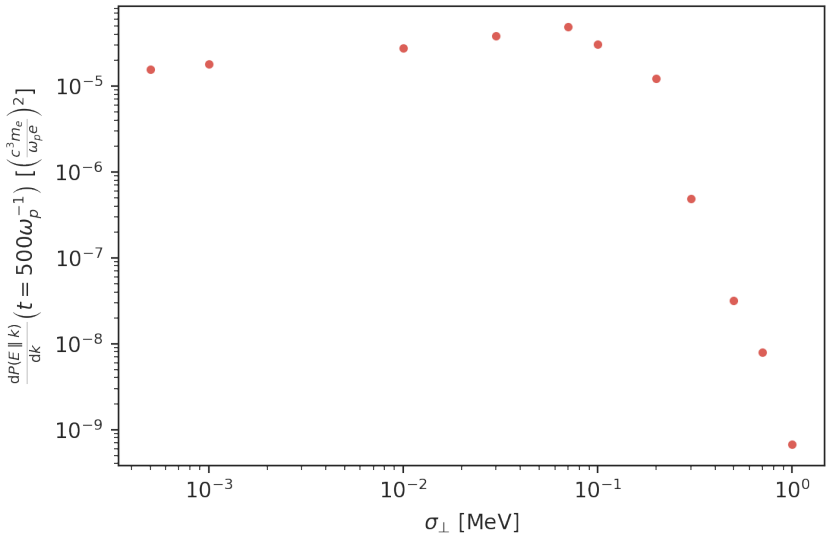

Figure 7.37: The power of the most dominant Fourier mode at time $t = 500\omega_p^{-1} > t_{\text{final}}$ as a function of the initial transverse momentum spread $\sigma_\perp$ for simulation runs with $\gamma = 10$, $\alpha = 10^{-3}$.

7.3 Charge Asymmetry

A neutral electron-positron beam consists in equal parts of positrons and electrons and thus carries neither a net current nor a net charge that could produce a separate magnetic or electric field. In a laboratory experiment it would only be possible to produce an approximately neutral beam that still caries some excess of electrons. We define the parameter χ as the ratio of positron density to electron density in the beam.

$$\chi = \frac{n_{e^+}}{n_{e^-}} \tag{7.7}$$

For a neutral beam χ is equal to 1 and for a pure electron beam χ is equal to 0. The case of χ larger than 1, where the number of positrons exceeds the number of electrons, is practically irrelevant and would (up to a sign) carry the same charge and current as a beam with χ^{-1}.
A beam with χ smaller than 1 carries a current density j.

$$j = \frac{1-\chi}{1+\chi} j_0 \tag{7.8}$$

Here j_0 is the current density of a pure electron beam with the same density ratio α and Lorentz boost γ as the pair beam.

$$j_0 = e\left(1 - \frac{1}{\gamma^2}\right) c\alpha n_{\text{bg}} \approx ec\alpha n_{\text{bg}} \tag{7.9}$$

For a relativistic beam the velocity can be assumed to be equal to the speed of light.
Figure 7.38 shows the evolution of the field energy density for simulation runs with varying beam charge ratio χ. Right from the start of the simulation runs the presence of a net current leads to an oscillating energy density with an amplitude that increases with the beam's current. Nevertheless the instability proceeds entirely unaffected by this oscillation and eventually the energy density of the instability fields exceed the field oscillation induced by the current. Even in the case of a pure electron beam ($\chi = 0$) the system shows the same late time behavior. This is because the instability only couples to the square of the particle charge and thus in the absence of external magnetic fields a pure electron beam shows the same instable behavior as a neutral pair beam. We find the same qualitative behavior for the energy loss of the beam particles (see figure 7.39) with one major difference: the period of the oscillation is twice as long as for the energy density of the electromagnetic fields. This is more easily visible when zooming in on early times, as in figure 7.40. There we can easily tell that the electromagnetic field[1] oscillates with the plasma frequency[2].
From the same late time behavior we have already concluded that the unstable behavior is independent of the beam charge ratio χ. An even more conclusive proof of this can be found in the spectral growth rate (figure 7.41). We find that the growth rates of the instable modes are completely the same for a neutral pair beam or a pure electron beam.
A key for experimentally probing the instable behavior is the effect of the instability growth on the beam momentum distribution. Unsurprisingly the oscillation of the beam energy due to the presence of a current also leads to the same oscillation behavior of the mean longitudinal

[1] The energy density is given by the square of the field amplitude and thus oscillates with twice the frequency
[2] When expressing time in units of the inverse plasma frequency ω_p^{-1} this is equal to a period of 2π

momentum. In figure 7.23 we can see that $\Delta\text{Mean}\left(p_{\parallel}\right)$ oscillates with the same frequency as the beam energy loss at early times and late times the instability takes over.
The widening of the momentum distribution is entirely driven by growth of field modes due to the instability. Since the growth of these modes is completely unaffected by the beam charge ratio χ it comes at no surprise that the evolution of $\text{Std}\left(p_{\parallel}\right)$ and $\text{Std}(p_{\perp})$ are in turn not influenced by the non-neutrality of the beam (see figure 7.43 and 7.44).
An important insight from this is that deviating from a perfectly neutral beam will not spoil an experiment that tries to observe the instable behavior via the change of the momentum distribution. In fact even a pure electron beam could be a viable experimental probe for the instable behavior.

Oscillation Amplitude

The oscillation of the field energy is caused by the build up of a magnetic field due to the current carried by the charged beam. The field then induces a return current that counteracts itself by reducing the longitudinal momentum of the beam. Since the magnetic field induced by a current is proportional to the current, the amplitude of the magnetic field oscillation A_{field} induced by a charged bunch is proportional to its current density.

$$A_{\text{field}} \propto j \propto \alpha \tag{7.10}$$

The energy density of a field is given by the square of the amplitude and thus the amplitude of the energy density oscillation A_{energy} scales with the square of the current density.

$$A_{\text{energy}} \propto j^2 \propto \alpha^2 \tag{7.11}$$

In figure 7.45 we confirmed this scaling. Usually we express the energy density of the fields or the beam with respect to the initial energy of the beam E_{beam}. In this case the normalized amplitude A_{norm} scales differently with α and γ.

$$A_{\text{norm}} \propto \frac{j^2}{E_{\text{beam}}} \propto \frac{\alpha}{\gamma} \tag{7.12}$$

Figure 7.46 confirms this scaling of the amplitude of the field energy density oscillation for pure electron beams. The relative amplitude of the oscillation can be reduced by increasing the Lorentz boost or reducing the density ratio.
It is important to note that here we only considered infinitely extended beams. In the case of a beam with a finite extension repulsive effects due to a net space charge of the beam or focusing effects due to a toroidal magnetic field might play a role.

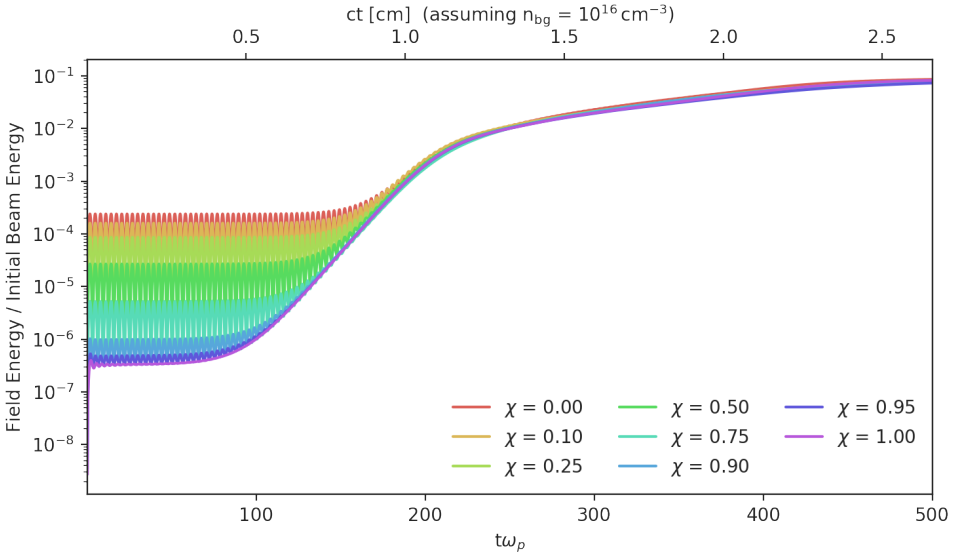

Figure 7.38: The energy density of the electromagnetic fields as a function of time for simulation runs with varying beam charge ratio χ. The presence of a net current leads to an oscillating energy density at early times with an amplitude that increases with the beam's current. At late times the energy density of the instability exceeds the oscillation in all cases and the energy density saturates at the same level.

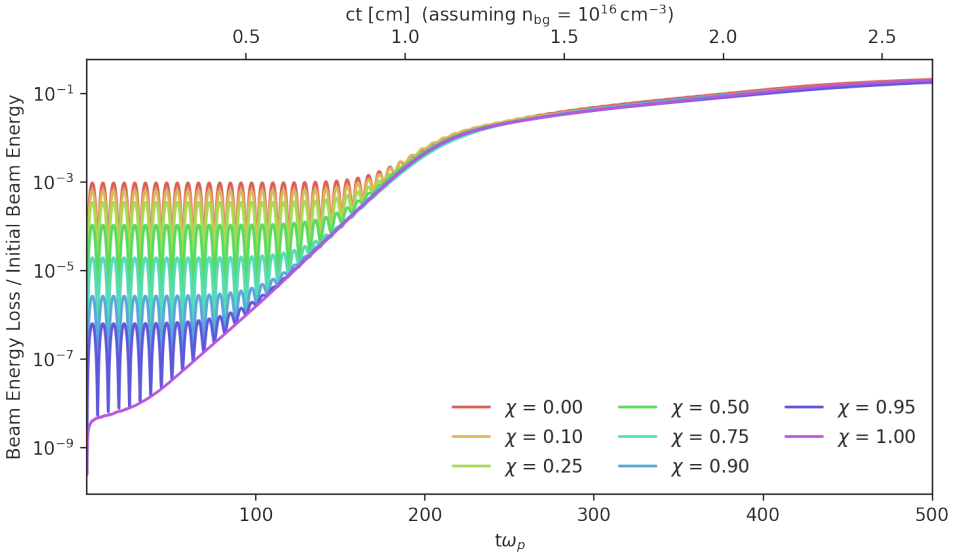

Figure 7.39: The energy loss of the particle beam as a function of time for simulation runs with varying beam charge ratio χ. Similarly to figure 7.38 at early times we find an oscillation induced by the presence of a current but the instability proceeds unaffected and at late times the system reaches the same final state independent of beam charge ratio.

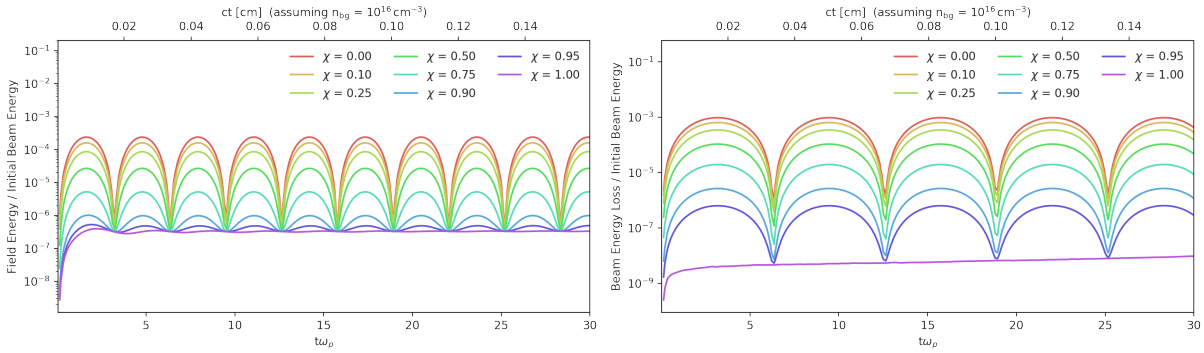

Figure 7.40: The field energy density (**left**) and the beam energy loss (**right**) as a function of simulation time for very early times only. Here it is very apparent that the amplitude of the oscillation increases with a stronger current (smaller χ) and that field energy oscillates twice as fast as the beam energy. The oscillating electric field also perturbs the background plasma. The energy of the background plasma oscillates at the same frequency as the beam plasma but shifted by a quarter period.

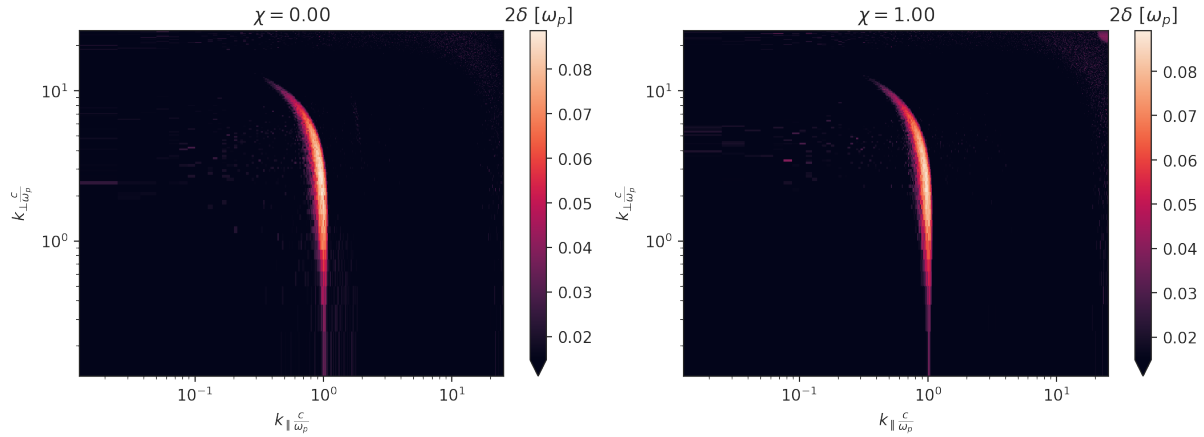

Figure 7.41: The spectral growth rate for two simulations runs. On the **left** for a beam that purely consists of electrons. On the **right** a neutral pair beam. The growth rate is completely identical in both cases.

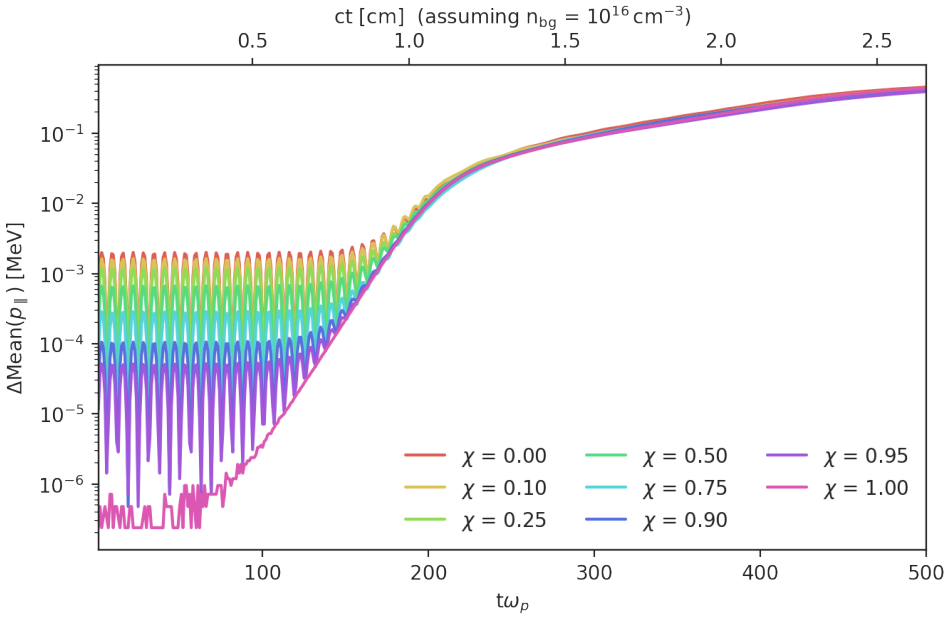

Figure 7.42: The shift of the mean longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ of the beam particles as a function of time for simulation runs with varying beam charge ratio χ. Similarly to beam energy loss we find that a net current leads to an initial oscillation that is eventually outgrown by the exponential growth of the instability leading to a late time behavior that is independent of χ.

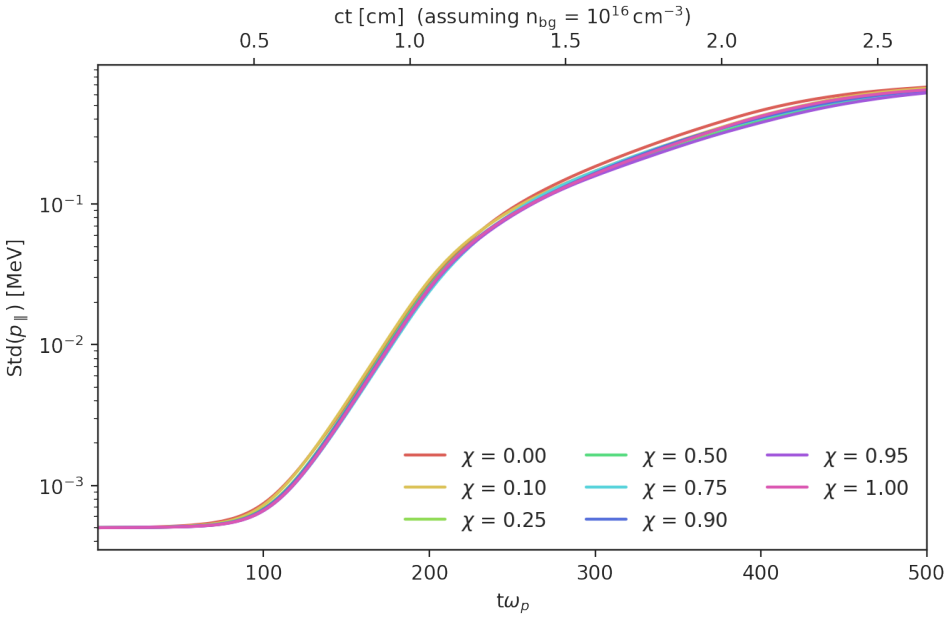

Figure 7.43: The standard deviation of the longitudinal momentum distribution $\text{Std}\left(p_\parallel\right)$ as a function of time for simulation runs with varying beam charge ratio χ. The widening of the momentum distribution is completely unaffected by χ

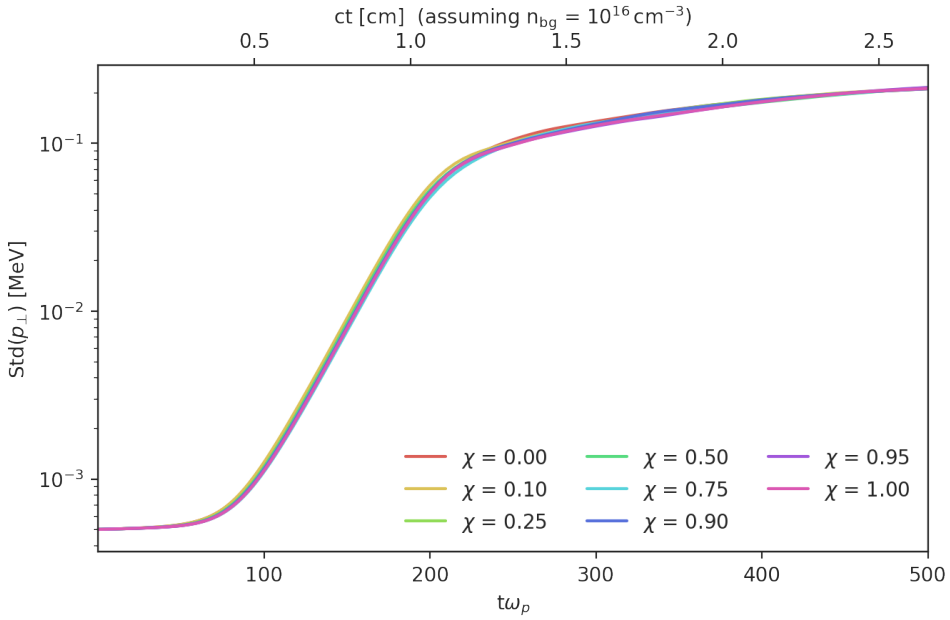

Figure 7.44: The standard deviation of the transversal momentum distribution $\mathrm{Std}(p_\perp)$ as a function of time for simulation runs with varying beam charge ratio χ. The widening of the momentum distribution is completely unaffected by χ

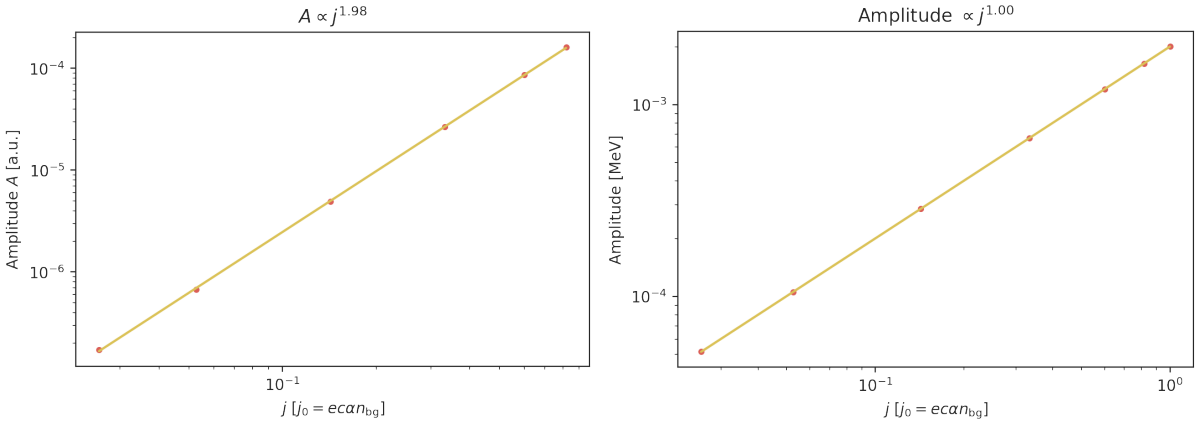

Figure 7.45: The amplitude of the field energy density oscillation (**left**) and the oscillation of $\Delta\mathrm{Mean}\left(p_\parallel\right)$ (**right**) as a function of the beam's current.

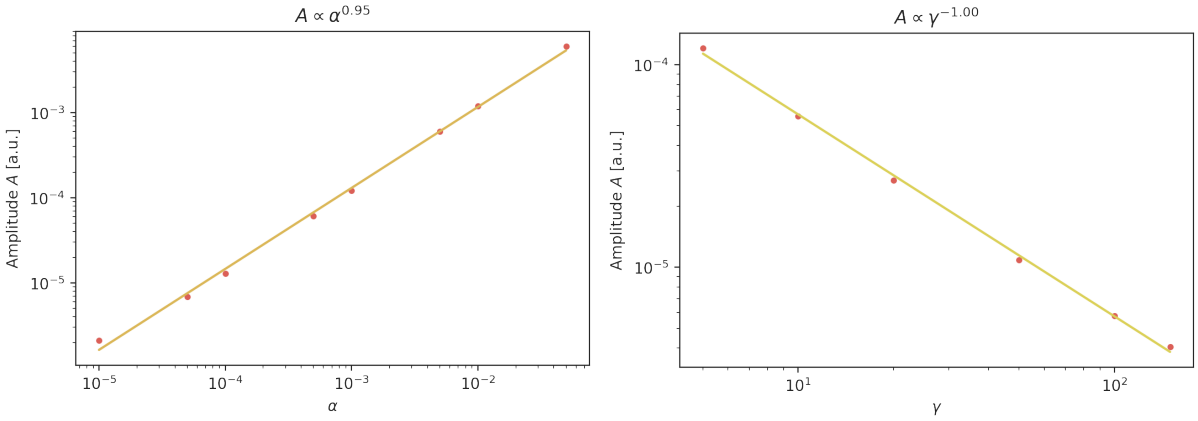

Figure 7.46: The amplitude of the field energy density oscillation of a pure electron beam ($\chi = 0$) scaling with α (**left**) and γ (**right**). The amplitude was normalized by the initial energy of the beam. The scaling follows the expectation outlined in (7.12).

7.4 Kinetic Energy Density Ratio

The physical system we want to study is a pair beam propagating through a much denser background plasma. The beam is only a tiny perturbation to the whole system. One can impose a number of more or less stringent conditions to determine if contribution of the beam to the total system is sufficiently small. The most obvious and least stringent is for the density ratio α to be much smaller than 1. A stronger condition would be that the energy density of the highly relativistic beam is smaller than the energy density of the background plasma[3]. This leads to the condition $\alpha\gamma \ll 1$. Stronger still would be the requirement that the kinetic energy of the beam is smaller than the kinetic energy of the background plasma. Their ratio is given by the parameter ϵ.

$$\epsilon = \frac{n_\mathrm{b}(\gamma - 1)m_e}{\frac{3}{2}n_\mathrm{bg}k_bT} \tag{7.13}$$

In a laboratory experiment the background plasma will be produced in a gas discharge with typical temperatures on the eV scale [26, 128]. Considering realistic values for the density ratio and the Lorentz boost this will lead to a background plasma with much less kinetic energy than the beam ($\epsilon \gg 1$). In an astrophysical scenario on the other hand the density ratio α is typically so low that the kinetic energy of the background plasma dominates ($\epsilon \ll 1$).
Furthermore we have shown in section 5.8 that the kinetic energy of the background plasma is proportional to the initial noise of the system.

$$\text{Initial noise} \propto \frac{1}{N_p\epsilon} \tag{7.14}$$

The large noise due to high plasma temperature would have to be compensated by increasing the number of simulation particles accordingly making simulations with small ϵ computationally very expensive.

[3] The energy density of the background plasma is almost entirely given by its rest mass energy.

In section 5.2.4 we saw that once we start to break the condition $\alpha\gamma \ll 1$ the behavior of the system starts to change. Namely the equipartition between beam energy being converted into field energy and heating the background plasma is lost. In this section we want to study the effect of the ϵ parameter on the instable system to make sure that we can scale what we learn from studying a laboratory system to astrophysical scales.
Figure 7.47 shows that raising the background plasma temperature, to reach a lower kinetic energy density ratio ϵ, leads to larger initial noise. As long as the energy density of the initial noise is lower than the final saturation level of the instability the system saturates at roughly the same level, indicating that the instability proceeds uninhibited by the value of ϵ. Only if the initial noise is already larger than the expected final saturation level of the instability do we not see a change in the energy density with time. However if we consider the energy loss of the beam (figure 7.48) we can see that even for the lowest ϵ the energy drain from the beam reaches the same level. Here the main difference is at early times: for small ϵ there is an initial phase where the beam does not lose energy, in fact a small amount of energy is even gained from the background plasma. Only when the energy loss of the beam due to the instability is large enough does the total energy budget turn over and we see significant energy loss.
Due to the higher background temperature the initial noise in the system is larger. This noise seeds the instability and thus for more noise the instability starts from a higher level already and thus also saturates at earlier time. In both figure 7.47 and figure 7.48 we can clearly see that at the same time simulation runs with lower ϵ take on a higher value, since lower ϵ means starting from a higher initial field amplitude.
In figure 7.49 we extracted the slope of the linear growth phase from figure 7.48. Even though the energy drain always saturates at roughly the same level and the linear growth phase looks similar we find that the extracted growth rate reduces with smaller ϵ, possibly indicating a suppression of the instability.
In the spectral growth rate maps (figure 7.50) we find that beyond a reduction of the maximum spectral growth rate for simulation runs with smaller ϵ, there are only minor effects. For beams with large ϵ the region of maximum growth curves towards smaller $k_\parallel$ for high $k_\perp$. This is likely a computational effect that becomes stronger when the plasma frequency is resolved with less grid points (the influence of the resolution of the plasmawavelength on the spectral growth map is discussed in appendix A.3). Nevertheless when reducing ϵ by heating the background plasma this effect subsides as well.
In a similar vein we can investigate the effect of adjusting ϵ on the change of the momentum distribution. For $\Delta\text{Mean}\left(p_\parallel\right)$ (figures 7.51 and 7.52), $\text{Std}\left(p_\parallel\right)$ (figure 7.53 and 7.54) as well as $\text{Std}(p_\perp)$ (figure 7.55 and 7.56) we again find the picture that the late time behavior is largely unaffected by changing ϵ whereas the early time behavior can change strongly due to the larger noise for lower ϵ leading to an earlier saturation of the instability and for particularly small ϵ even an evolution which is almost entirely influenced by the noise. The proportionality constants V_1, D_0 and D_1 that govern how the energy density of the instability couples to the moments of the momentum distribution are strongly dependent on the value of ϵ.
The growth rate δ does not show a clear dependence on ϵ. For the widening of the momentum distribution larger ϵ seem to lead to a slightly reduced growth rate, whereas for the shift of the mean momentum and the energy density larger ϵ leads to a slight increase of the growth rate. Our physical understanding is that the growth rate should be the same for all of these processes. The most likely explanation is that the change in extracted growth rate is only artificial and the

underlying instability grows with a rate that is independent of ϵ. However when ϵ becomes too small the large (unphysical) noise leads to behavior that seemingly changes the instable growth behavior.
We draw the conclusion that if the numerical noise is kept to a sensibly low level then there is no reason why the system with very low ϵ like an astrophysical pair beam should behave differently from a similar system with much larger ϵ. The very different kinetic energy density ratios do not hinder one from using a laboratory system to make inferences of the astrophysical system.

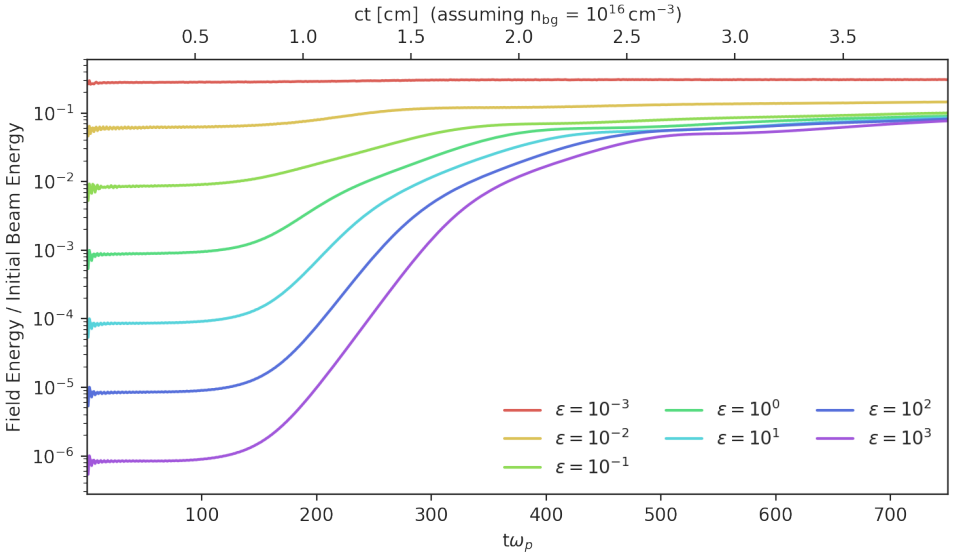

Figure 7.47: The field energy density as a function of time for simulation runs with $\gamma = 2$, $\alpha = 10^{-4}$ and varying kinetic energy density ratio ϵ. ϵ is varied by adjusting the temperature of the background plasma. Low ϵ leads to more initial noise. Once the initial noise exceeds the expected saturation level of the instability no instable behavior is visible anymore.

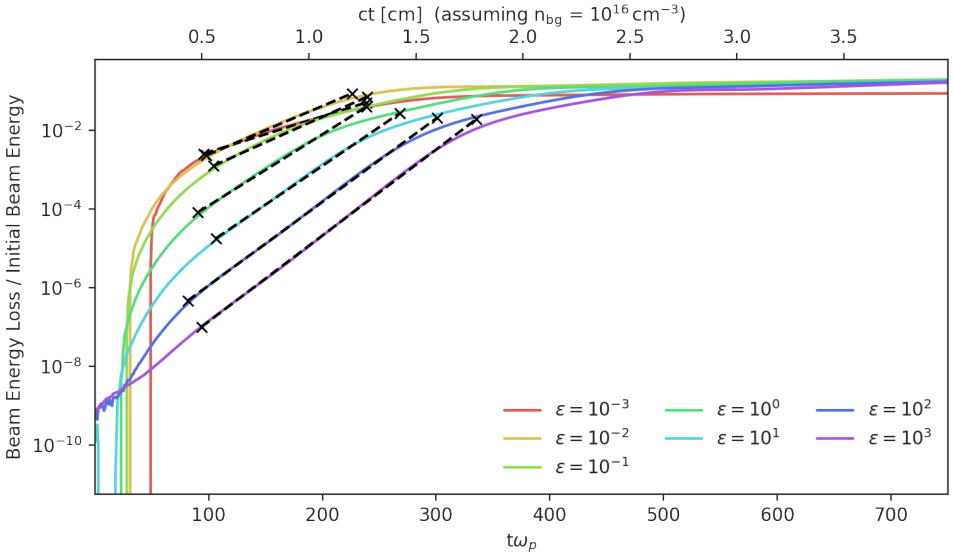

Figure 7.48: The energy loss of the pair beam as a function of time for simulation runs with $\gamma = 2$, $\alpha = 10^{-4}$ and varying kinetic energy density ratio ϵ. The dashed line represent a fitted exponential growth regime, the slopes of which are plotted in figure 7.49.

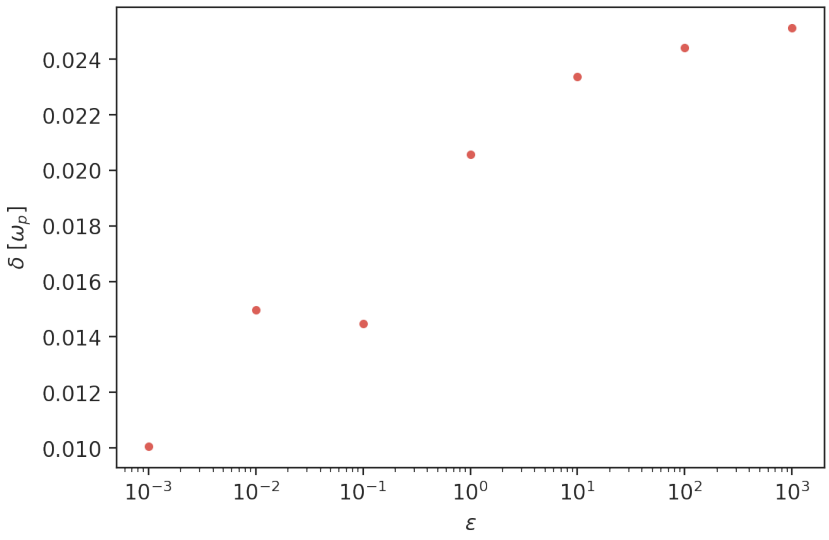

Figure 7.49: The slopes of the fitted exponential regimes in figure 7.48. For smaller ϵ the fitted slope becomes less step.

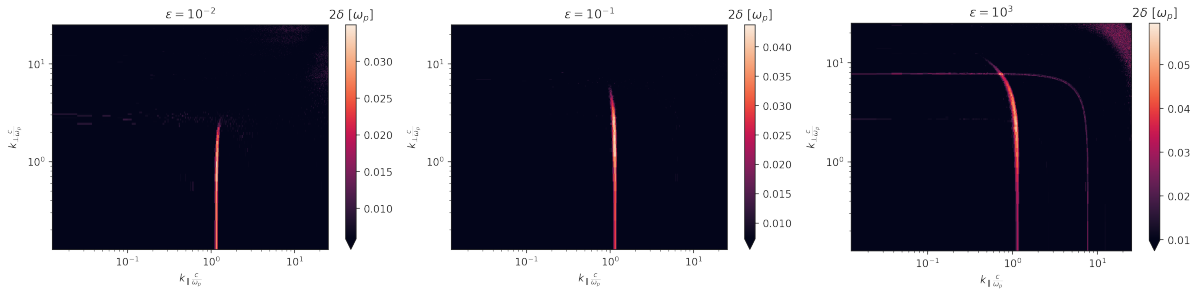

Figure 7.50: The spectral growth rate for simulation runs with varying kinetic energy density ratio ϵ. On the **left** ($\epsilon = 10^{-2}$) and in the **middle** ($\epsilon = 10^{-1}$) the kinetic energy density of the background plasma is larger than of the kinetic energy density of the pair beam. On the **right** ($\epsilon = 10^{3}$) the hierarchy is reversed. In all cases the most dominant modes can be found around $k_{\parallel} \sim \frac{\omega_p}{c}$. However for very cold background plasma the modes of the resonance with the largest transverse component can be found at slightly lower $k_{\parallel}$ whereas for smaller ϵ this trend is reduced and at least in the case of $\epsilon = 10^{-2}$ even slightly reversed.

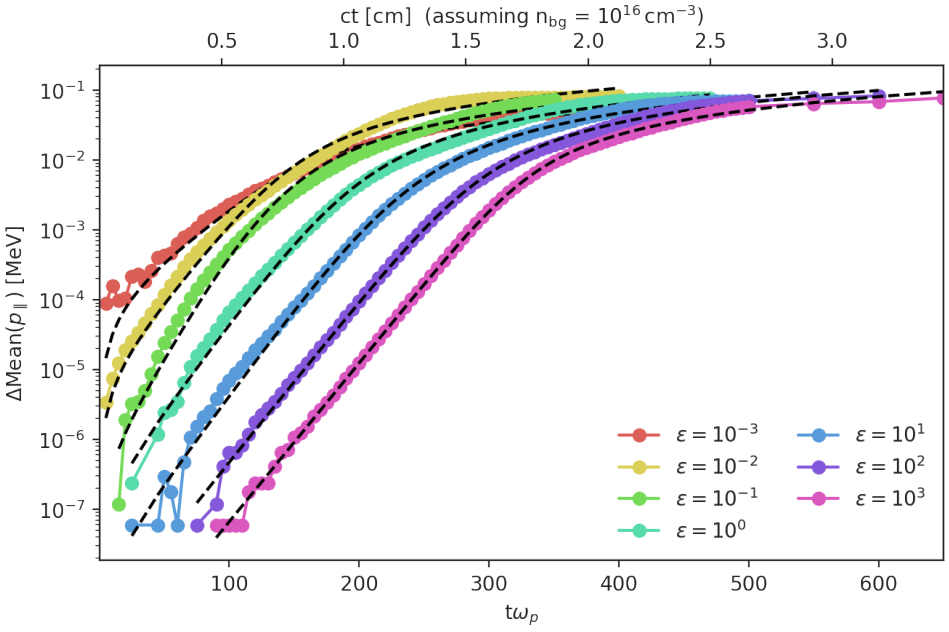

Figure 7.51: The shift of the mean of the parallel momentum distribution $\Delta\text{Mean}\left(p_{\parallel}\right)$ as a function of time for simulation runs with $\gamma = 2$, $\alpha = 10^{-4}$ and a varying kinetic energy density ratio ϵ. In dashed line a fit of the analytical model (3.42) is shown. The extracted fit parameter δ and V_1 are plotted in figure 7.52.

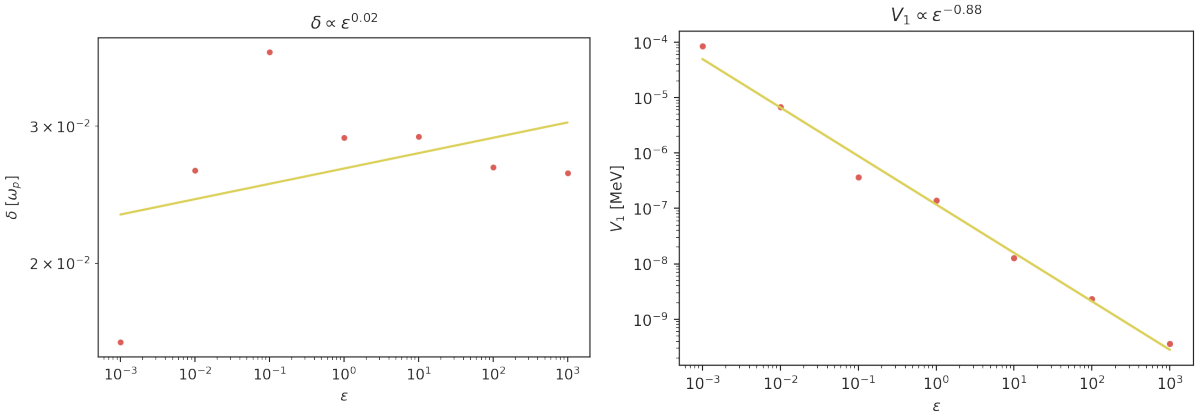

Figure 7.52: The growth rate δ (**left**) and the parameter V_1 (**right**) of the fit to (3.42) shown in figure 7.51 each as a function of ϵ. While V_1 shows a pretty clear dependence on ϵ the influence on the growth rate is not so clear. For large ϵ the growth rate seems to take on a constant value, while for low value deviations with no clear systematic are observed.

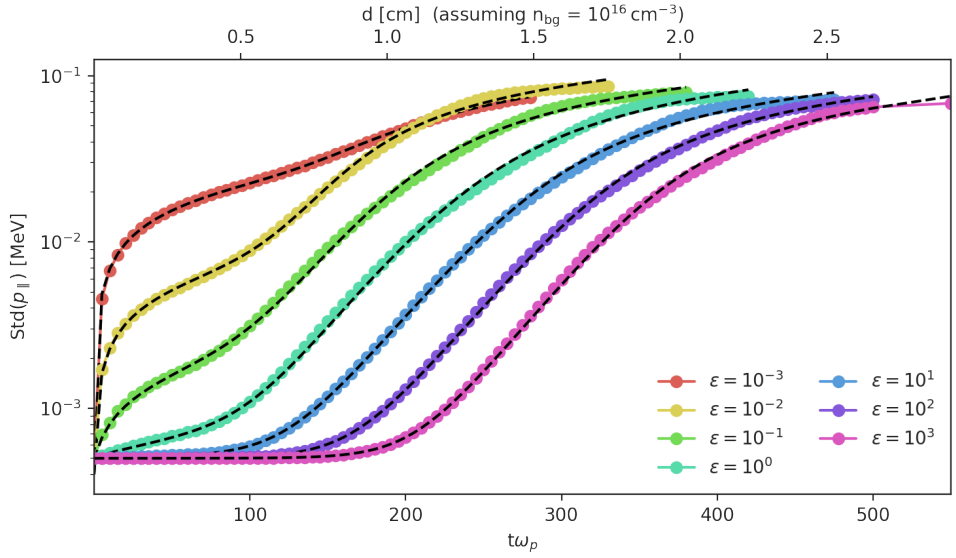

Figure 7.53: The evolution of the longitudinal beam momentum spread $\mathrm{Std}(p_\parallel)$ with time for simulation runs with $\gamma = 2$, $\alpha = 10^{-4}$ and a kinetic energy density ratio ϵ that is varied by adjusting the temperature of the background plasma. In dashed lines a fit to (3.43) for each run is shown. The extracted fit parameters δ, D_0 and D_1 are plotted in figure 7.54.

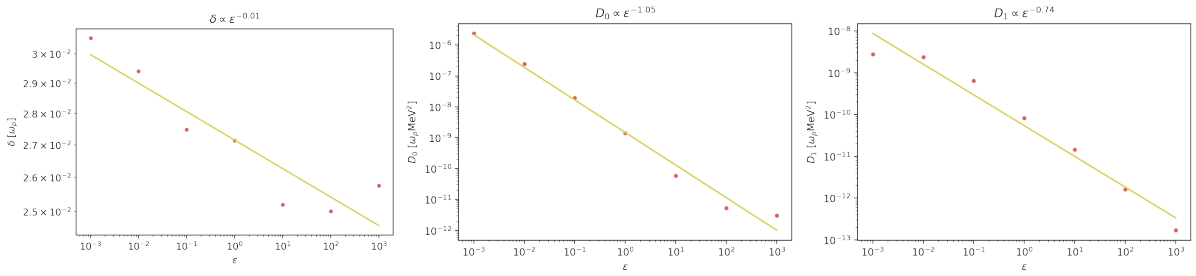

Figure 7.54: The growth rate δ (**left**), and the parameters D_0 (**center**) and D_1 (**right**) of the fit to (3.43) shown in figure 7.53 each as a function of ϵ. Both D_0 and D_1 show a clear dependence on ϵ, whereas the growth rate δ only shows a very slight increase with small ϵ.

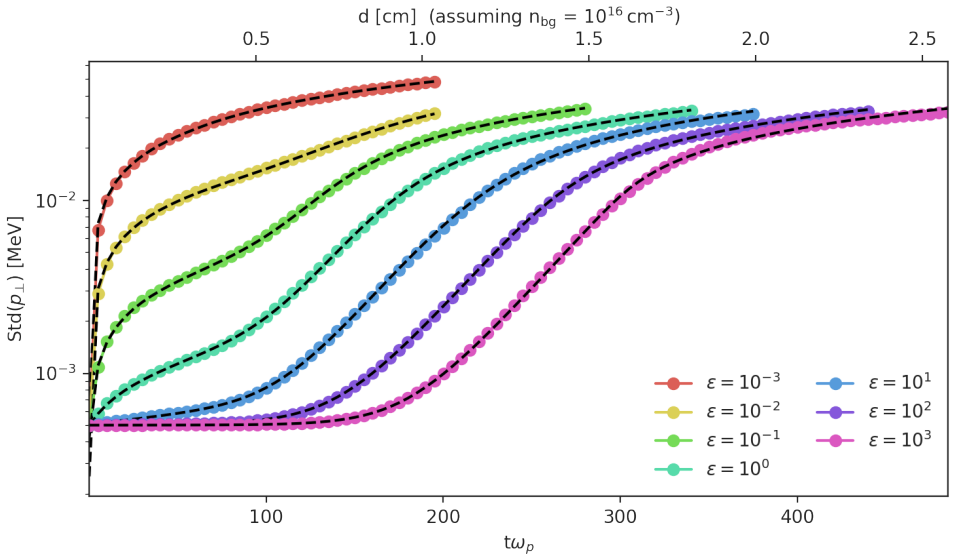

Figure 7.55: The evolution of the transversal beam momentum spread Std($p_\perp$) with time for simulation runs with $\gamma = 2$, $\alpha = 10^{-4}$ and a kinetic energy density ratio ϵ that is varied by adjusting the temperature of the background plasma. In dashed lines a fit to (3.43) for each run is shown. The extracted fit parameters δ, D_0 and D_1 are plotted in figure 7.56.

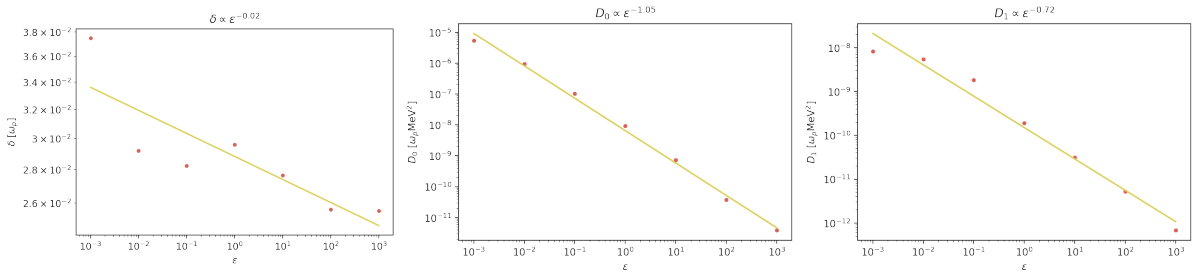

Figure 7.56: The growth rate δ (**left**), and the parameters D_0 (**center**) and D_1 (**right**) of the fit to (3.43) shown in figure 7.55 each as a function of ϵ. Both D_0 and D_1 show a clear dependence on ϵ, whereas the growth rate δ only shows a very slight increase with small ϵ.

7.5 Magnetic Fields

In a magnetic field moving charged particles are deflected without changing energy due to the Lorentz force.

$$\vec{F} = q\vec{v} \times \vec{B} \tag{7.15}$$

In a constant magnetic field this will lead to charged particles moving on a spiral trajectory. On the other hand in a stochastic magnetic field the particle will undergo many different deflections leading to a broadening of the beam opening angle similar to what we see from instabilities. It has been shown by Alawashra & Pohl [9] that this broadening will suppress the instability growth for a strong enough magnetic field and that the required field strength is smaller for more coherent fields. However due to the nature of the Lorentz force, a simulation that is capable of studying the effect of the magnetic deflections needs to resolve three spatial dimensions. The single exception that can be studied in a two dimensional simulation is an external, static magnetic field that is aligned orthogonal to the plane of the simulation. We will have a look at this scenario to derive some understanding of the influence of external magnetic fields that could serve as a starting point for further investigations.
A constant magnetic field forces charged particles on a spiral trajectory with the radius called the gyro radius r_g with the associated frequency called the gyro frequency ω_g. We will use the gyrofrequency to parameterize the magnetic field by using the magnetic field at which the gyro frequency equals the plasma frequency as a reference.

$$|B| = \frac{\omega_g m_e}{|e|} \Rightarrow B_{\text{ref}} = \frac{\omega_p m_e}{e} \tag{7.16}$$

Then the gyro radius of an ultra-relativistic particle will be given in units of the plasma wavelength.

$$r_g = \frac{m v_\perp}{eB} = \frac{v_\perp}{\omega_g} \approx \frac{c}{\omega_g} \tag{7.17}$$

Figure 7.57 shows the evolution of the beam energy loss as a function of the external magnetic field (orthogonal to the plane of simulation). For low magnetic fields the typical evolution of an exponential growth phase can be observed. As the magnetic field increases the growth rate of this exponential phase is reduced (compare figure 7.58). Once the magnetic field becomes strong

enough the instability growth is suppressed entirely and a new oscillation behavior becomes apparent.
Applying an external magnetic field orthogonal to the beam direction deflects the particles sideways but in opposing directions for electrons and positrons. Figures 7.59 and 7.60 show how an external magnetic field shifts the momentum distribution of the beam electrons continuously over time. This leads to a misalignment of electron and positron component of the beam and suppresses the instability growth if the magnetic field is strong enough. The suppression also changes the power spectrum (figure 7.61). The instability excites modes with $k_\parallel \frac{c}{\omega_p} \sim 1$ and a broad range in $k_\perp$. When applying a strong external magnetic field the power spectrum instead transitions to a circle of excited modes with $k_\parallel^2 + k_\perp^2 = \left(\frac{\omega_p}{c}\right)^2$. Furthermore the power of the maximum mode (figure 7.62) is reduced by many orders of magnitude.
The transition from an instability dominated regime to a regime that is dominated by the external magnetic field takes place when the period of the magnetic gyration becomes comparable to the time of instability growth. We use the inverse growth rate to estimate the time of instability growth and thus roughly estimate for the magnetic field required to suppress the instability:

$$B_z \frac{e}{\omega_p m_e} \gtrsim \delta \tag{7.18}$$

We note that Alawashra & Pohl [9] find that for incoherent fields the required magnetic field to suppress the instability is even higher.
It is important to note that this 2D treatment can only give us a rather qualitative intuition of the influence of an external magnetic field. Especially we have to note that the magnetic field of course affects electron and positrons very differently (namely deflecting them in opposite directions) and thus any similarities between a neutral pair beam and a pure electron beam could be spoiled by applying an external magnetic field. As a minor note we mention that in the case of very large magnetic fields, and thus small gyroradii, it might be necessary to include radiative effects in the simulation code[4] to correctly describe the energy of the beam particles.

[4] Radiative effects can be treated with EPOCH, however they are turned of by default

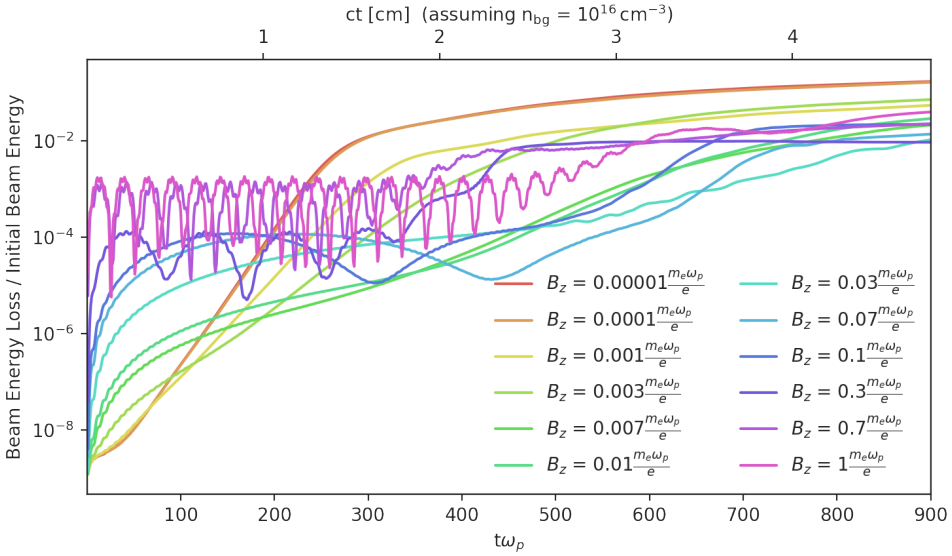

Figure 7.57: The energy loss of the pair beam as a function of time for simulation runs with varying external magnetic field strength B_z. Increasing the magnetic field leads to a reduction of the beam energy loss rate (figure 7.58). For very strong external magnetic fields the entire evolution changes with very fast initial energy loss followed by an oscillation regime.

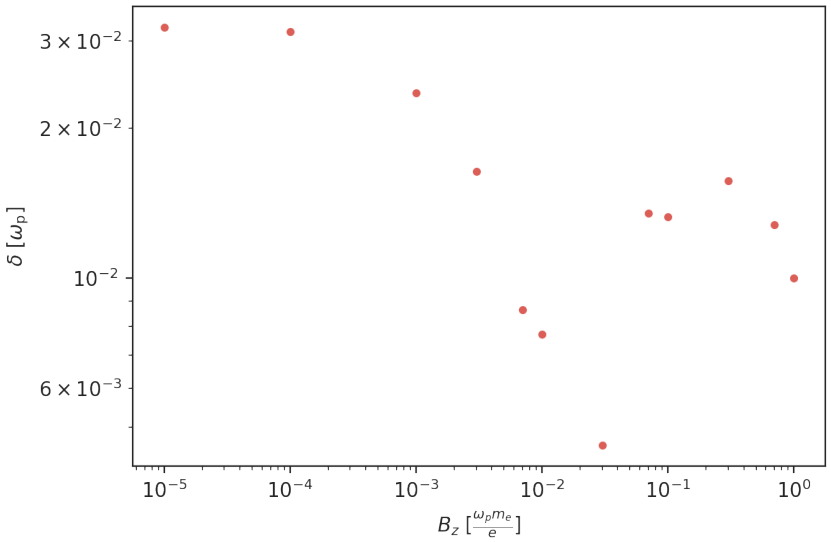

Figure 7.58: The extracted rate of the energy loss of the pair beam (figure 7.57). Up until $B_z = 0.03\,\frac{\omega_p m_e}{e}$ the fitted growth rate reduces with larger magnetic field. For higher magnetic fields the entire evolution changes and exponential growth does no longer describe the system very well.

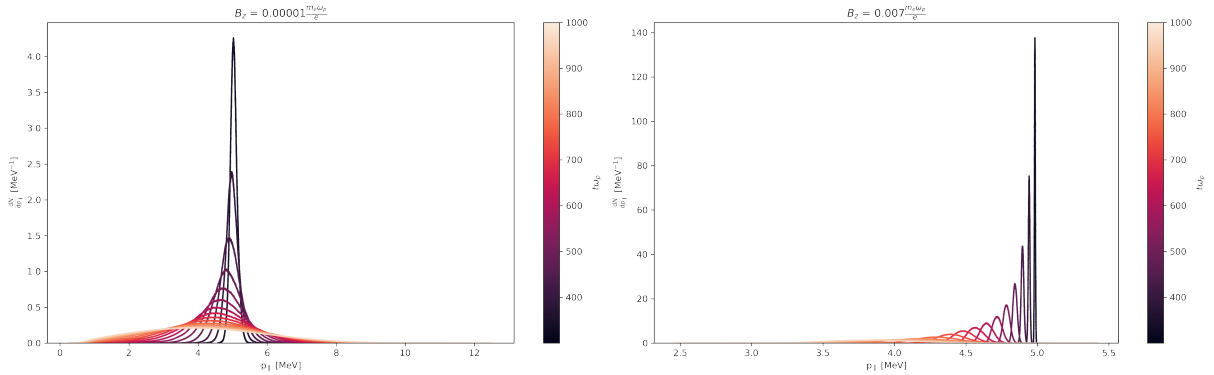

Figure 7.59: The longitudinal beam momentum distribution for two simulation runs with different external magnetic field strengths at a number of time steps near the end of instability growth. On the **left** the external magnetic field is very weak and does not alter the momentum distribution. The normal broadening and shift behavior can be observed. On the **right** the external magnetic field is stronger and bends the beam trajectory sideways, which leads to a continuous reduction of longitudinal momentum.

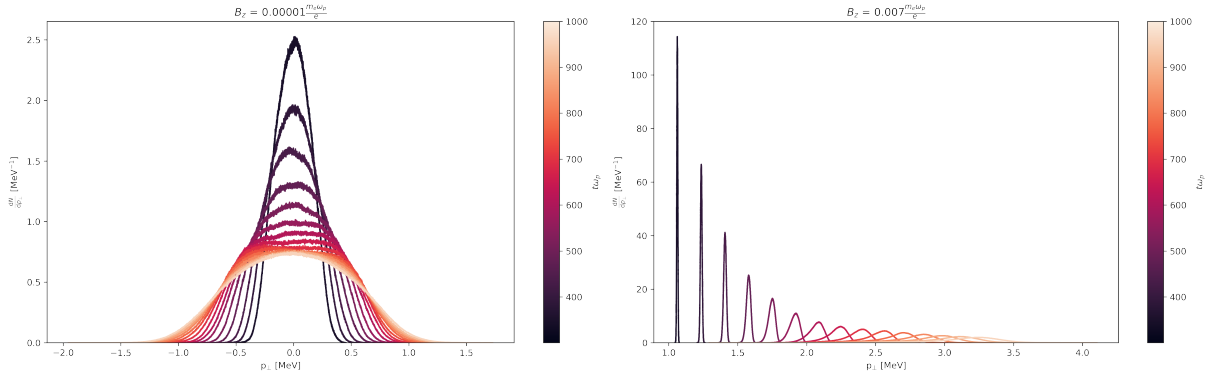

Figure 7.60: The transversal beam momentum distribution (for the electron component) for two simulation runs with different external magnetic field strengths at a number of time steps near the end of instability growth. On the **left** the external magnetic field is very weak and does not alter the momentum distribution. The normal broadening behavior can be observed. On the **right** the external magnetic field is stronger and bends the beam trajectory sideways, which leads to a continuous gain of transversal momentum.

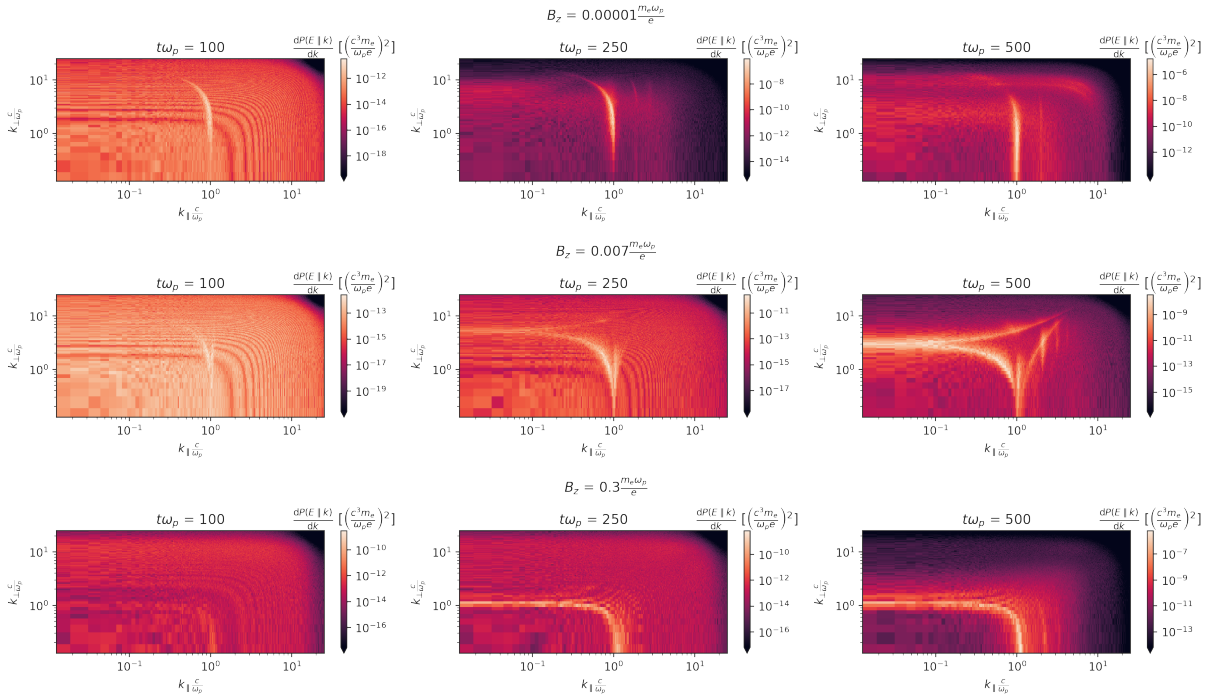

Figure 7.61: The differential power spectrum for three simulation runs taken at different times during the simulation. The **top** shows the power spectrum for a very weak external magnetic field. The **middle** shows a moderate magnetic field where the instability growth is reduced due to the external field. The **bottom** shows a very strong external magnetic field where no instability growth takes place.

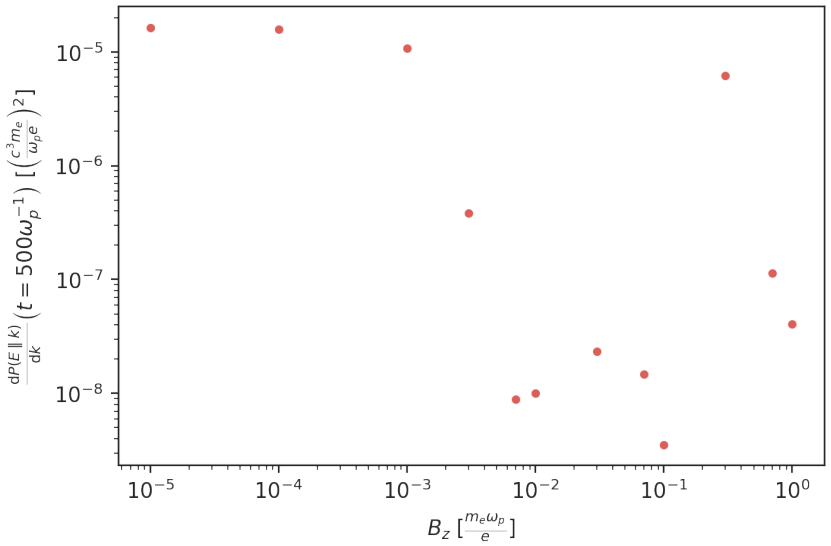

Figure 7.62: The power of the strongest mode at time $500\omega_p^{-1}$ as a function of the external magnetic field strength. Applying a strong external field suppresses the instability and leads to much smaller excitations of the parallel electric field.

7.6 Inhomogeneous Background Medium

The plasma frequency ω_p determines the most important time and length scales of the instable system. The most dominant instable modes are found at $k_\parallel \sim \frac{\omega_p}{c}$ and the growth rate of the instability is proportional to the plasmafrequency. The plasmafrequency in turn only depends on physical constants and the density of the plasma. For a very dilute beam the total plasma density is approximately the density of the background plasma. If the background plasma is homogeneous then the ω_p is a constant. However if we relax the assumption that the background plasma is homogeneous, the plasma frequency becomes a function of position in space. Consequently a beam propagating through a background plasma with changing density will excite different modes at different times. The growth of the instability is a resonant effect and thus a changing plasma frequency could suppress the instable behavior.
The inhomogeneity of the background plasma can take the form of a rapid oscillation (short coherence length) or a gradual change over larger distances (long coherence length). Here we will only focus on a density change with long coherence length that could for example occur in the plasma cell of a laboratory experiment. However inhomogeneities with very short coherence length could also be a mechanism for the suppression of the instability in astrophysical pair beams [106, 107]. We will assume that the inhomogeneity is described by a single sine wave and expand for distances that are much smaller than the wavelength. Then the inhomogeneity of background medium takes on the form a density ramp with constant slope [105].

$$n(x) = n_{\text{bg}}\left(1 + A\sin\left(\frac{2\pi}{\lambda}x\right)\right) \approx n_{\text{bg}}\left(1 + A\frac{2\pi}{\lambda}x\right) \tag{7.19}$$

$$\Rightarrow \dot{n} = \frac{\partial n}{\partial x} = n_{\text{bg}} A\frac{2\pi}{\lambda} \tag{7.20}$$

Here n_{bg} is the average or reference density of the background medium, A is the amplitude of the density oscillation and λ is the length scale of the density oscillation. The evolution of the density takes on the form of a constant change with a rate that is determined by the amplitude divided by the wavelength of the density inhomogeneity. We normalize $\dot{n}$ by expressing it as the density change in units of the reference density n_{bg} over a plasma wavelength $\frac{c}{\omega_p}$.
The scenario essential becomes a constant slope of density increase (or decrease) along the path of the particle beam. The plasma frequency, which depends on the square root of the plasma density, grows with the square root of the distance covered. This implies that the modes that are excited by the instability are a function of position and thus time. At late times of the instability development modes with higher $k_\parallel$ (for growing density) are excited.
Figure 7.63 shows increasing the slope of the density reduces the rate of beam energy loss and the final beam energy loss. More precisely once the relative rate of density change per plasma wavelength reaches the permille level, roughly the growth rate in units of the plasma frequency, the growth rate and the saturation level starts to reduce following a power-law relationship (see figure 7.64). The Fourier plot (figure 7.65) shows a clear smearing out of the excited modes in $k_\parallel$ for a large background density slope.
In the evolution of the longitudinal and transversal momentum distribution (figures 7.66 and 7.68 respectively) we also observe a suppression of the instability for a high background density slope. Both high and low density slope can be well described by the analytical model (3.43) derived from the Fokker-Planck equation. Some of the fit parameters are shown in figures 7.67

and 7.69. The transition from a constant regime for low density slopes to a power-law decay for high slopes takes place at a similar threshold as for the beam energy loss. While the power-law index is the same for the growth rate for the beam energy loss and the momentum broadening, this is not the case for the saturation level. Increasing the slope of the background density also leads to an exponential increase of the fit parameter D_1 which is closely associated with the initial noise of a simulation run. This leads to stronger initial momentum broadening for high density slopes.
We conclude that the suppression of the instability growth due to the density slope already starts when the relative rate of change is still small compared to the instability growth rate δ and it is not enough that the density change is small over one growth cycle. Instead the density change has to be small over the whole time it takes to grow the instable modes to saturation. This conclusion is supported by the observation from figure 7.65 that only a small rate of density change can ensure that the same modes are excited at different times during the instability growth. Furthermore this implies that a configuration with a narrower band in $k_{\parallel}$ with significant instability growth (due to a high Lorentz boost γ or a low density ratio α) would be more susceptible to disruption due to a density gradient.

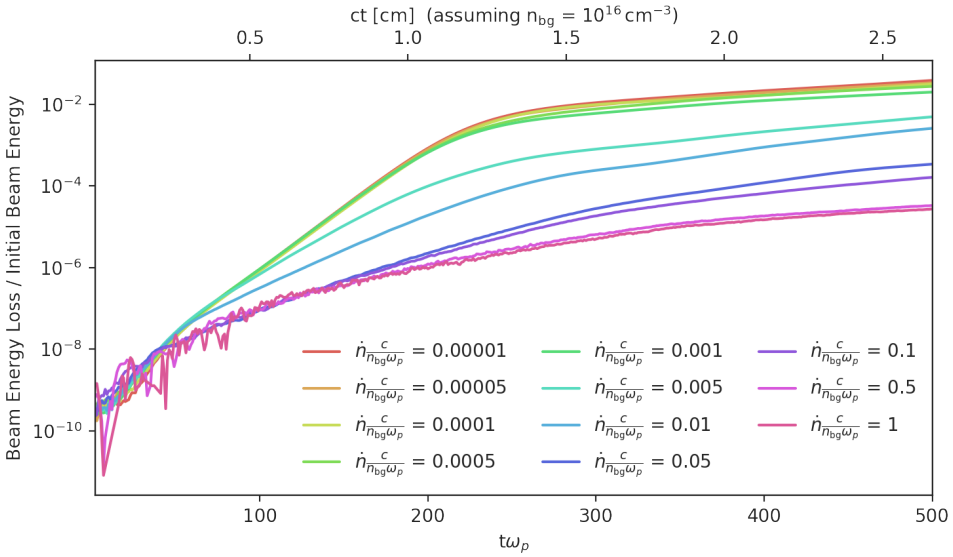

Figure 7.63: The energy loss of the pair beam as a function of time for simulation runs with varying background density slope $\dot{n}$. A stronger slope reduces the growth rate and the saturation level of the energy drain due to the instability. The extracted growth rate and saturation level are shown in figure 7.64. The time was normalized using the plasmafrequency derived from the reference background density n_{bg}.

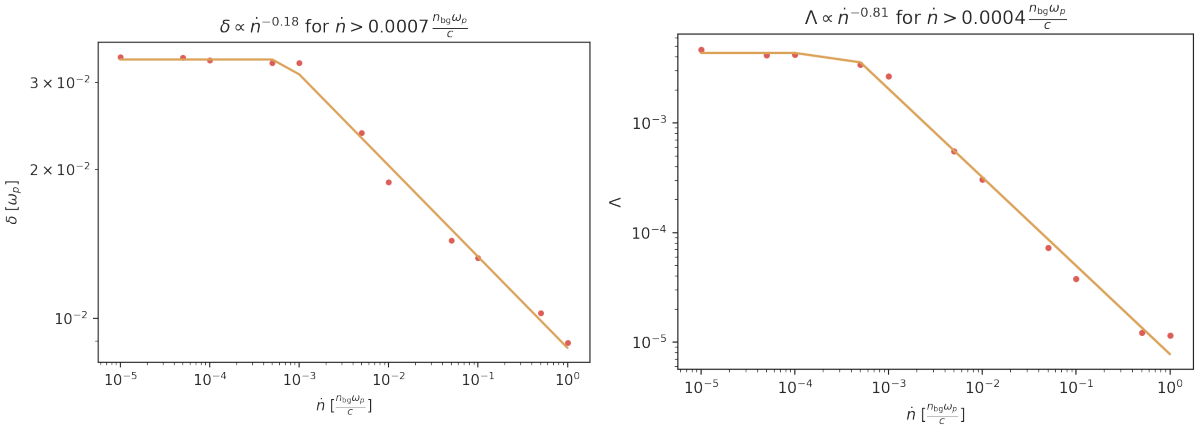

Figure 7.64: The growth rate δ (**left**) and the saturation level Λ (**right**) as a function of the background density slope $\dot{n}$ as extracted from the beam energy loss shown in figure 7.63. The growth rate and saturation level take on a constant value for a slowly changing background plasma density and above a critical value they start to decrease following a power law behavior.

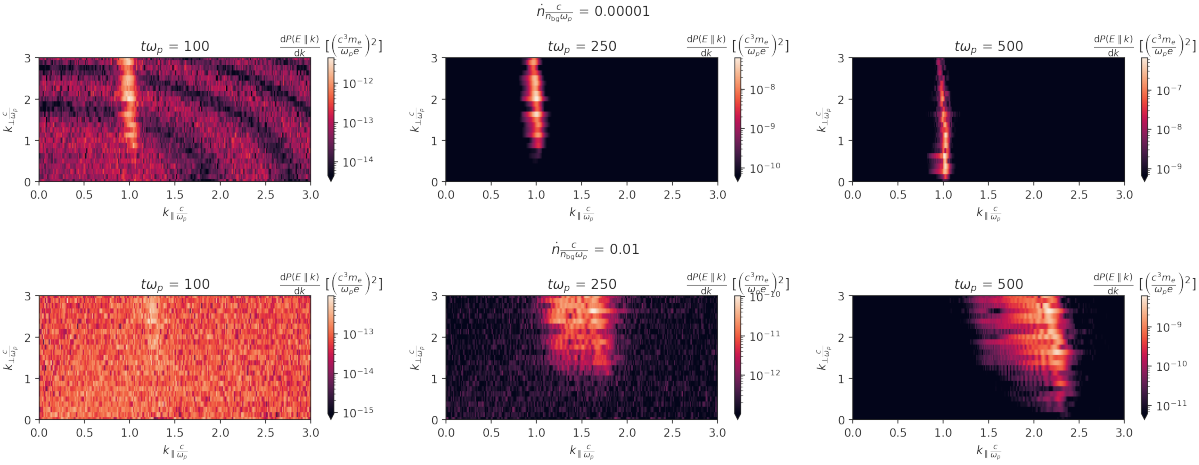

Figure 7.65: The differential power spectrum for two simulation runs with low rate of density change (**top**) and high rate of density change (**bottom**) at different time steps during a simulation run. The gradual increase of the local background plasma density shifts the plasmafrequency ω_p to higher value and leads to a smearing of the power spectrum. For a stronger density change the smearing of the power spectrum is stronger.

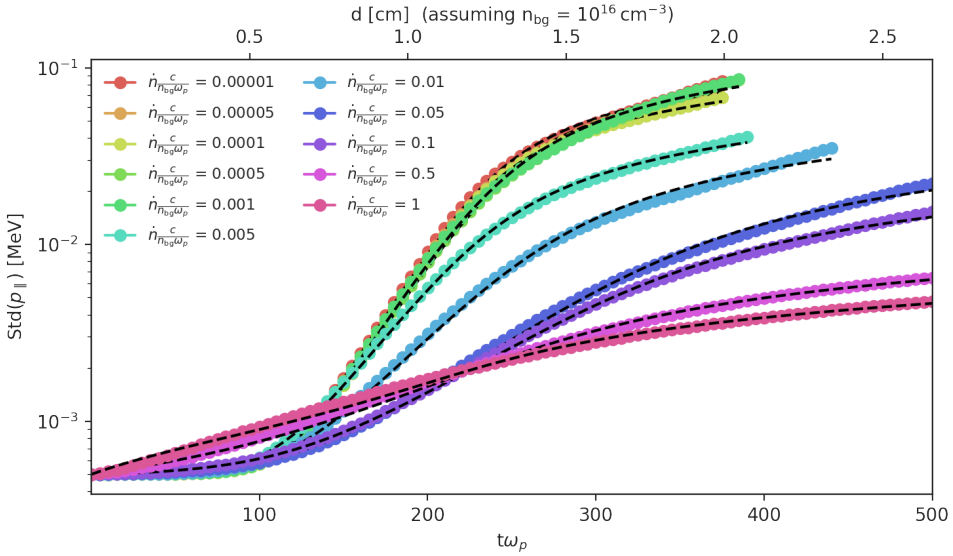

Figure 7.66: The evolution of the width of the longitudinal momentum $\mathrm{Std}(p_\parallel)$ is shown for runs with varying background density slope $\dot{n}$. The standard is calculated at regular intervals of $5\,\omega_p^{-1}$ until a time of $500\,\omega_p^{-1}$ and afterwards at intervals of $50\,\omega_p^{-1}$. Data points are shown up to times of twice the fitted value for the time of instability saturation. In dashed lines a fit of (3.43) to the data is shown. The extracted values for the growth rate δ, the noise parameter D_1 and the saturation level Λ can be found in figure 7.67.

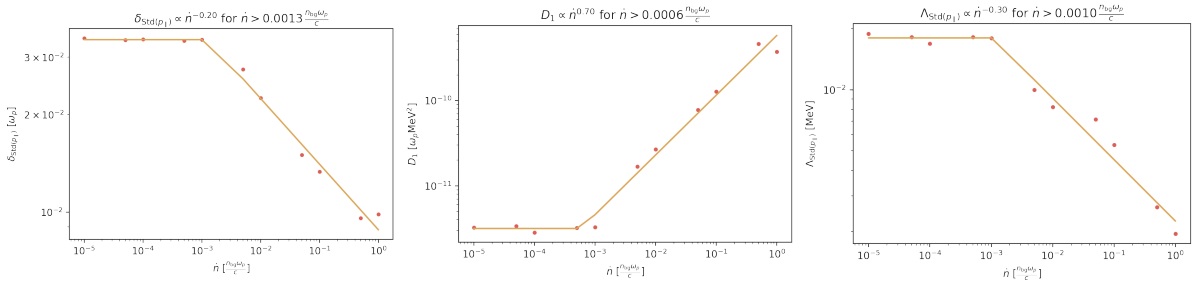

Figure 7.67: The extracted parameters for the growth rate δ (**left**), the noise parameter D_1 (**middle**) and the saturation level Λ (**right**) from the fit in figure 7.66 for the evolution of the width of the longitudinal momentum. For high density slope the growth rate and the saturation level are reduced, whereas the noise level D_1 increases.

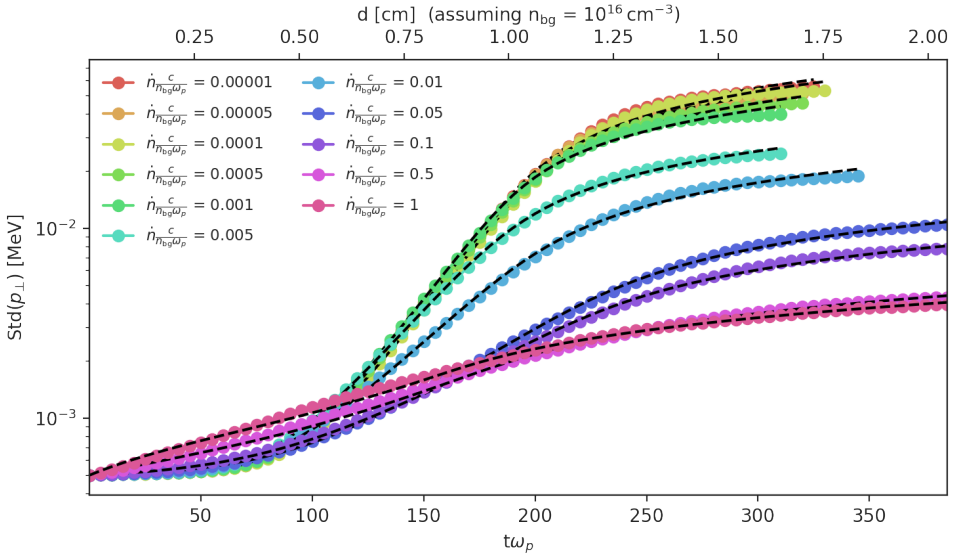

Figure 7.68: The evolution of the width of the transversal momentum $\mathrm{Std}(p_\perp)$ is shown for runs with varying background density slope $\dot{n}$. The standard is calculated at regular intervals of $5\,\omega_p^{-1}$ until a time of $500\,\omega_p^{-1}$ and afterwards at intervals of $50\,\omega_p^{-1}$. Data points are shown up to times of twice the fitted value for the time of instability saturation. In dashed lines a fit of (3.43) to the data is shown. The extracted values for the growth rate δ, the noise parameter D_1 and the saturation level Λ can be found in figure 7.69.

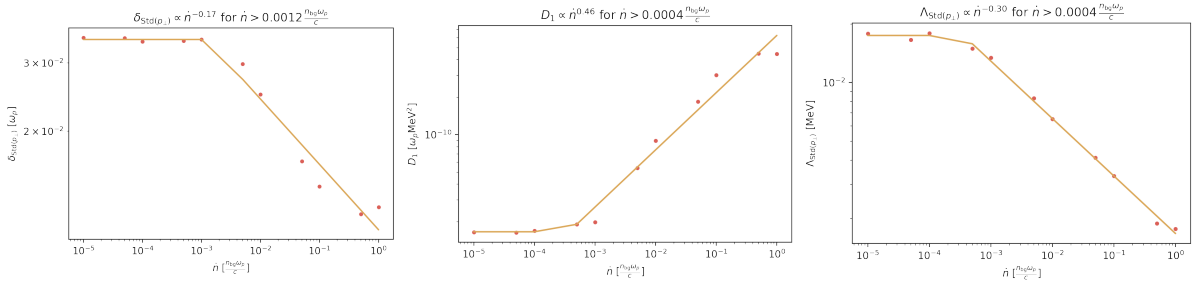

Figure 7.69: The extracted parameters for the growth rate δ (**left**), the noise parameter D_1 (**middle**) and the saturation level Λ (**right**) from the fit in figure 7.68 for the evolution of the width of the transversal momentum. For high density slope the growth rate and the saturation level are reduced, whereas the noise level D_1 increases.

8 | In the Laboratory

A laboratory experiment would be a great opportunity to compare the results of our simulations with reality in a situation that allows for in situ measurements of the fields or particle beam. In particular the behavior of a pair beam in the non-linear post saturation regime and the unstable behavior of pair beams in the transition region of a mildly hot beam would be interesting regimes. In figure 8.1 we sketch a possible setup for such a laboratory experiment. In this case an electron beam is created by a LWFA (laser wakefield accelerator), however depending on the requirements this could also be a conventional accelerator. This electron beam needs to be converted into a beam that ideally contains an equal number of electrons and positrons with equal momentum distribution. This has been demonstrated to be possible [127, 100, 99, 101] by hitting a high density target where the electrons undergo a cascade of Bremsstrahlung and pair production. The resulting pair beam will have a broad energy spectrum, high divergence and an overabundance of electrons which makes it necessary to apply clever beam optics that can focus both charges simultaneously.
The produced, hopefully high quality, pair beam will then be led through a plasma cell where the instability occurs leading to the build up of electromagnetic fields in the plasma cell as well as heating the background plasma and the pair beam. The modified pair beam will then leave the plasma cell where its momentum spectrum can be measured with a magnetic spectrometer.

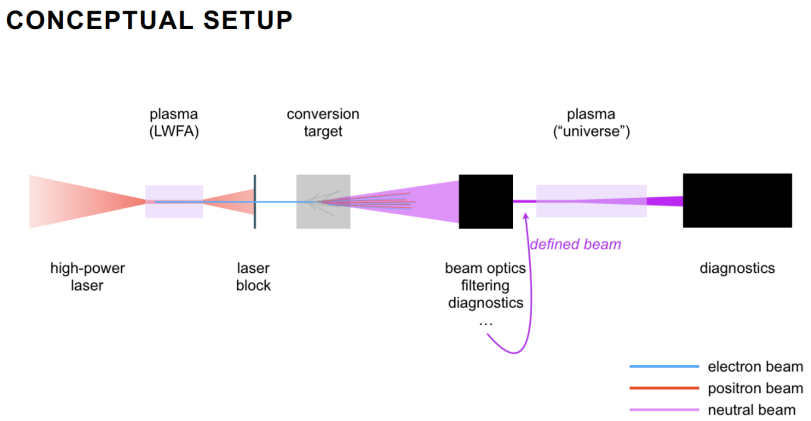

Figure 8.1: Sketch of a possible laboratory setup where an electron beam is produced by a laser wakefield accelerator (LWFA) and converted into a neutral pair beam in a high density target. Courtesy: Benno Zeitler

8.1 Initial Electron Beam

The quality of modern accelerators are often judged on many different aspects of their performance: Their maximum energy, their energy spread, the number of particles per bunch (which is equivalent to the beam charge), their emittance [62], the length of their bunches, etc. However when we want to produce a pair beam by hitting a high density target many of these qualities will be spoiled as a result of the radiative interactions and scattering processes. The resulting beam will be highly divergent and have a wide energy spectrum. We will only consider three quantities of the incoming electron beam to be relevant for the output pair beam: The initial mean energy of the beam particles, the initial beam charge and the initial bunch length.
The number of output particles, and thus the output beam density, obviously scales proportionally to the number of input particles. The input beam energy will define the maximum energy for output particles. Even though most particles will be close to the critical energy of the converter material, a higher energy input beam will increase the mean energy of the output beam. Since in a cascade the number of particles increases exponentially with each generation increasing the energy of the input beam also leads to more output particles and thus a higher output density. Lastly the length of the beam will increase only by a few femtoseconds due to interactions and deflections in the converter target and thus the beam length of the output beam will be approximately equal to the length of the input beam. A longer beam with the same number of particles will have a lower density, but can drive the instability longer. LWFA can produce electron beams with up to nanocoulomb charge, hundreds of MeV in energy and a bunch length of the order of tens of femtoseconds [47, 79].
We have shown in section 7.1.1 that the length of the beam needs to be longer than the inverse growth rate to drive the oblique instability. In figure 8.2 we show the inverse growth rate as a function of the energy and density of the output beam for different background plasma densities. For energies and densities that could be reasonably achieved with a LWFA the length of the beam would have to exceed tens of picoseconds, about three orders of magnitudes larger than what these kind of machines usually produce. The required beam length can be reduced by increasing the density of the background plasma, however an increase of the growth rate by three orders of magnitude would require an increase of the background density by 18 orders of magnitude.

$$\delta \propto \left(\frac{\alpha}{\gamma}\right)^{\frac{1}{3}} \omega_p \propto= n_{\mathrm{bg}}^{\frac{1}{6}} \tag{8.1}$$

If one were to stretch the input beam, for example with a magnetic chicane, the length would increase at the cost of reducing the density of the beam. Since the beam density scales linearly with the length of the beam and the growth rate scales with the third root of beam density stretching the beam enough would eventually lead to a beam that can drive the instability. However for a beam with a length that is three orders of magnitude too short, one would need to stretch the beam by four orders of magnitude.

$$\frac{L_\parallel}{c} \gg \delta^{-1} \propto \omega_p^{-1} \left(\frac{\gamma}{\alpha}\right)^{\frac{1}{3}} \propto \left(L_\parallel L_\perp^2\right)^{\frac{1}{3}} \propto L_\parallel^{\frac{1}{3}} \tag{8.2}$$

An alternative proposal for the creation of pair beams is to use high energy proton beams to induce the pair cascades [19]. Here the proton could produce neutral pions via strong interaction which can decay into two photons that then each start an electromagnetic cascade. Protons can

be accelerated with much less effort due to their higher masses and those high energy protons would produce many more secondary particles to increase the charge of the output beam. Facilities like AWAKE (bunches of $3 \cdot 10^{11}$ protons with 400 GeV energy and 400 ps bunch length) [66] or HiRadMat (bunches of $1.2 \cdot 10^{11}$ protons with 440 GeV energy and 375 ps bunch length) [57] could produce output beam consisting of electrons and positrons each with a density in excess of $10^{13}\,\mathrm{cm}^{-3}$ [19] with lengths of hundreds of picoseconds. These beams could be suitable for studying the oblique instability.

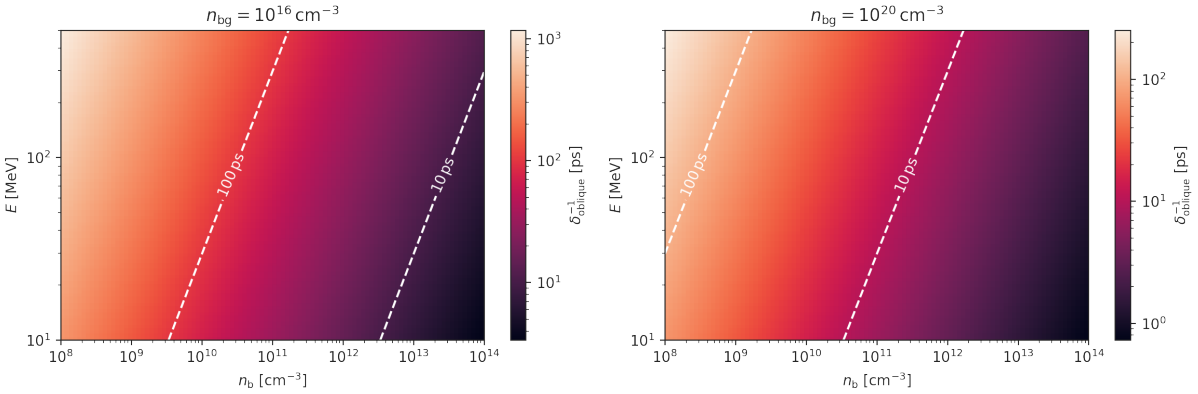

Figure 8.2: Inverse growth time as a function of output beam mean energy E and output beam density n_b.

8.2 Pair Beam Production

To produce a beam of electron and positron pairs we want to induce electromagnetic cascades in a high density target. In an electromagnetic cascade the incident electron produces a high energy photon via Bremsstrahlung in a material. The photon can then go on to produce a pair of electron and positron via the pair production process while the incident electron can again emit a Bremsstrahlung photon. The mean distance a particle travels before it interacts is a material constant called the radiation length X_0. This mechanism can continue until all particles carry energy below the critical energy E_c, another material constant.
In the most simple analytical model called the Heitler model one can assume that after each interaction the two outgoing particles carry half the energy of the incoming particle [108]. Then the cascade produced by an incident particle with energy E_{in} can continue for N generations, where N_p particles are produced at the end of the cascade that penetrates a depth of X_{max} of the material.

$$N = \log_2\left(\frac{E_{\mathrm{in}}}{E_{\mathrm{c}}}\right) \tag{8.3}$$

$$N_p = 2^N = \frac{E_{\mathrm{in}}}{E_{\mathrm{c}}} \tag{8.4}$$

$$X_{\mathrm{max}} = NX_0 = X_0 \log_2\left(\frac{E_{\mathrm{in}}}{E_{\mathrm{c}}}\right) \tag{8.5}$$

Below the critical energy electrons and positrons mainly lose energy due to ionization producing additional electrons (see figure 8.3). As a consequence below the critical energy a vast over-

abundance of electrons is produced. Ideally the material for the converter target should have a low critical energy to maximize the particle yield and a short radiation length to minimize the influence of scattering in the material. In table 8.1 we summarized a number of materials that could be suitable. In particular lead and tungsten are widely available and offer a low critical energy and short radiation length.

Table 8.1: Radiation length and critical energy of a number of dense materials [130]. The critical energy is given for electrons, which is slightly larger than the critical energy for positrons

Element	Radiation Length X_0 [cm]	Critical Energy E_c [MeV]
Cu	1.436	19.42
Sn	1.206	11.86
Ta	0.4094	8.09
W	0.3504	7.97
Pb	0.5612	7.43

To study the conversion of the input electron beam in the converter target, we use the Geant4 [7, 11, 12] simulation framework to simulate a monoenergetic beam of 10^6 electrons (equivalent to a charge of 0.16 pC) with zero emittance and zero timespread hitting targets of various materials and thickness. We adopt a reference energy of 500 MeV for the input electrons. In figure 8.4 we show the ratio of the produced positrons to electrons as a function of the converter material and length. For a very thin converter increasing the thickness produces more positrons, but after a couple radiation lengths the ratio begins to shrink again if we consider particles of all energies. This can be attributed to the ionization electrons produced at low energies. If we only consider particles well above the critical energy of the material the ratio converges towards charge parity after a few radiation lengths. We conclude that it is only possible to produce a beam that is close to charge neutrality we apply some filtering mechanism to remove part of the produced particles.
In figure 8.5 we show the spectrum of all produced particles for a 20 mm tungsten converter target. The by far most commonly produced particles are photons, followed by electrons and positrons produced in almost equal amount and then neutrons. The neutrons could become a problem with regard to radiation safety in an experiment. Other particles are produced in negligible amounts. All produced particles very roughly follow a power-law spectrum that reaches up to the energy of the incident particles. A requirement to drive the instability growth is the inversion of the momentum distribution. This again requires some form of intervention to remove part of the lower energy particles. Since it is below the critical energy where electrons are produced in much larger numbers than positrons this would also address our problem with the charge neutrality.
We can use to our advantage that low energy particles in the cascade are produced with a larger intrinsic opening angles and are more susceptible to scattering. If we consider the deflection angle with respect to the incident beam axis θ, then particles with lower energies will on average have larger θ. If we introduce an aperture of some kind that only allows particles with small θ to pass, we will produce a beam that is mostly charge neutral and has a momentum inversion. The very forward going particles is also what would arrive in a plasma cell located at some distance from the target. An artificial cut of $\theta < 5°$ approximately replicates the particles that arrive in a plasma cell with a radius of 1.75 mm at a distance of 2 cm from the converter target. In figure

8.6 we show the influence of introducing a 5° degree cut applied to θ on the longitudinal momentum spectrum. The resulting spectrum has a momentum inversion and particles below the critical energy are almost entirely removed. Additionally we point out that the number of high energy particles decreases for a thicker absorber even though more particles across all energies are produced. One may consider using a thinner converter target to produce a non-neutral, but high energy and high density pair beam. However for a thin target the momentum distribution of the electrons and positrons is quite different (see figure 8.7). The excess of electrons for a thin target manifests mainly at high energies where no positrons are present. Only for a target thickness that also leads to a neutral beam above the critical energy does the momentum distribution of both particle species becomes similar. Even though the input particle beam consists of monoenergetic electrons with an energy of 500 MeV the energy spectrum of the output pair beam peaks at much lower energies. The exact value depends on the thickness of the converter target.
The transversal momentum distribution (figure 8.8) shows most particles with transverse momentum of the order of some MeV, leading to a beam that can certainly not be described as cold. A thicker converter target leads to a beam with less transversal momentum overall, although this is likely offset by the reduction in longitudinal momentum as well. Introducing a cut on the angle θ does not substantially change the momentum distribution. This introduces the further challenge of cooling the beam in order to drive the oblique instability efficiently.
One of the most important quantities for the instability growth is the density of the pair beam. A denser pair beam leads to faster instability growth making a laboratory experiment more feasible. Due to the inhomogeneous distribution of particles with respect to time and position in space the density of the pair beam is not an immediately well defined quantity. We define the density as the number of particles divided by the 68th percentile in distance from the beam axis in both transverse directions (r_{68}) and the 68th percentile in arrival time times the speed of light (ct_{68}).

$$n_\mathrm{b} = \frac{N_{e^-} + N_{e^+}}{ct_{68} r_{68}^2} \tag{8.6}$$

At the backside of the converter target (figure 8.9) the output particles are focused in a small spot around the beam axis. For a thicker target the particles experience more deflection and the spot size (defined as the area that contains 68% of the outgoing positrons) grows from 0.15 mm for a 5 mm lead converter to 1.2 mm for a 20 mm converter target. Most particles, especially those with high energies, leave the back of the converter target within a few femtoseconds (see figure 8.10). Nevertheless the time distribution of the produced particles follows an exponential decay with few particles experiencing larger delays. 68% of particles arrive within 122 fs for a 20 mm lead target (for a 5 mm target within 32 fs).
Since the pair beam is diverging the density is also a quantity that rapidly evolves with distance from the converter target. In figure 8.11 we show the density as it evolves with distance for different converter thicknesses for an input beam with 10^6 electrons of 500 MeV. Even though a thicker converter target produces more secondary particles, the density of the pair beam is reduced with increasing thickness. Within the first few millimeters the pair beam density already drops by orders of magnitude. When considering only the very forward going particles ($\theta < 5°$) the density does drop faster at first, but at distance larger five to ten millimeters the density it typically one order of magnitude larger than without this cut. For a 20 mm converter target

and a plasma cell at a distance of another 20 mm particle densities in excess of $10^7\,\mathrm{cm}^{-3}$ can be reached, for a beam with a charge of a nanocoulomb this would translate to a density of $10^{11}\,\mathrm{cm}^{-3}$.
In accordance with the Heitler model the number of produced secondary particles scales linearly with the energy of the incident particle. Including the effect that particles with higher energy experience less scattering in the material and are produced with smaller opening angle the density increases stronger than linear with energy.

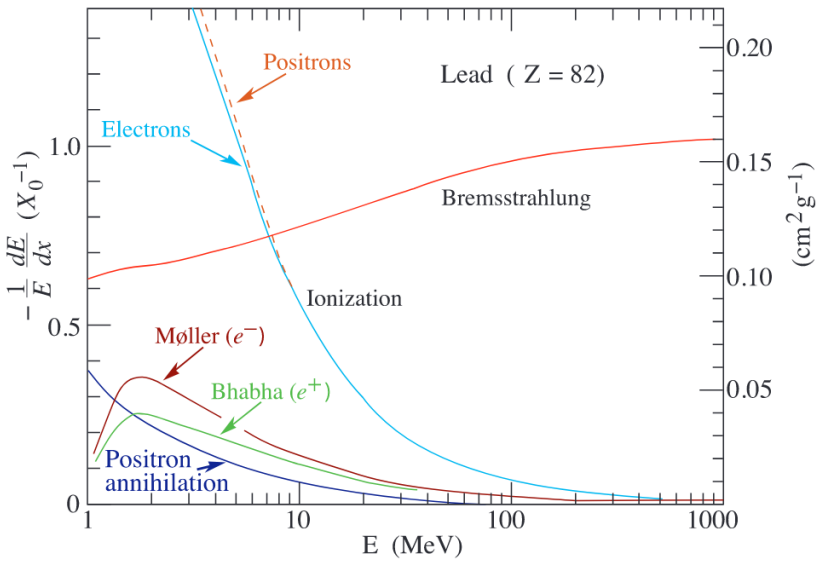

Figure 8.3: The fractional energy loss per radiation length X_0 in lead as a function of the energy of an electron or positron [130]. Above the critical energy E_c Bremsstrahlung is the dominant process leading to a cascade that produces electrons and positrons in equal number. Below the critical energy only electrons are produced via ionization.

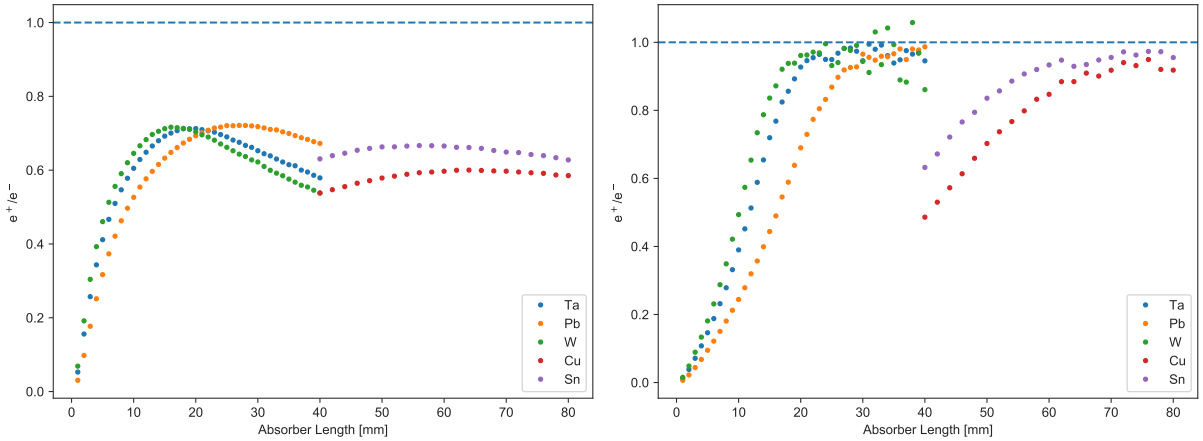

Figure 8.4: Ratio of the positrons to electrons produced from 500 MeV incident electrons as a function of converter material and length. The **left** shows all particles produced, the **right** shows only particles with energies above 100 MeV.

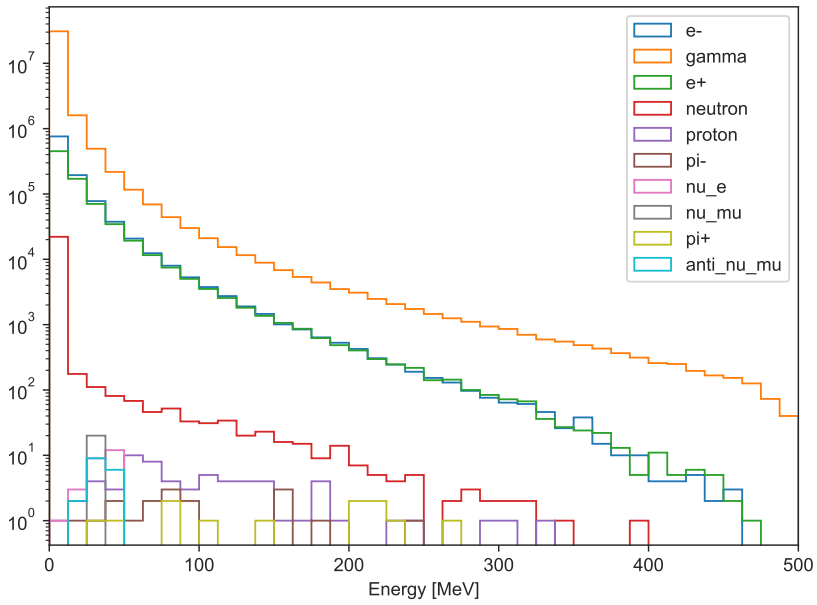

Figure 8.5: All Particle Spectrum for a 20 mm tungsten target and an incident beam of 10^6 electrons with an energy of 500 MeV. The overproduction of photons can be explained by the pair production length actually being $\frac{9}{7}X_0$ and thus slightly larger than the energy loss length Bremsstrahlung.

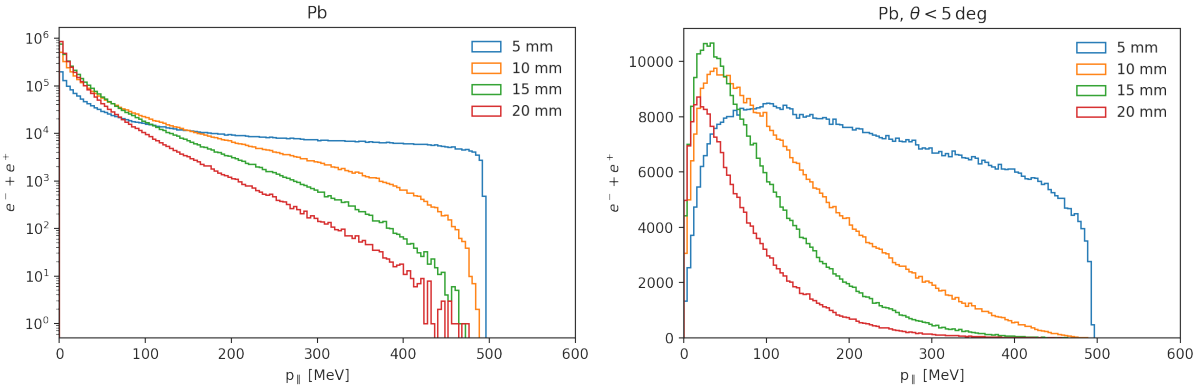

Figure 8.6: The electron and positron combined parallel momentum spectrum for lead converter targets of different thickness. The **left** shows all particles and the **right** only particles with an angle smaller than 5° with respect to the incident beam axis. $p_\parallel$ is the momentum component that is parallel to the incident beam.

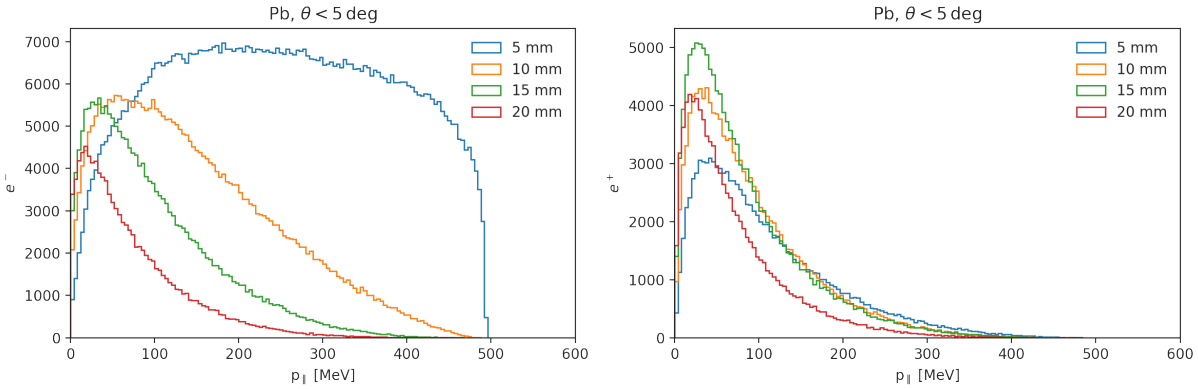

Figure 8.7: Electron (**left**) and positron (**right**) momentum spectrum for lead converter targets of different thickness considering only particles with deflection angle $\theta < 5°$. Only for a target that has a thickness of at least a few radiation lengths does the momentum distribution of both species become the same.

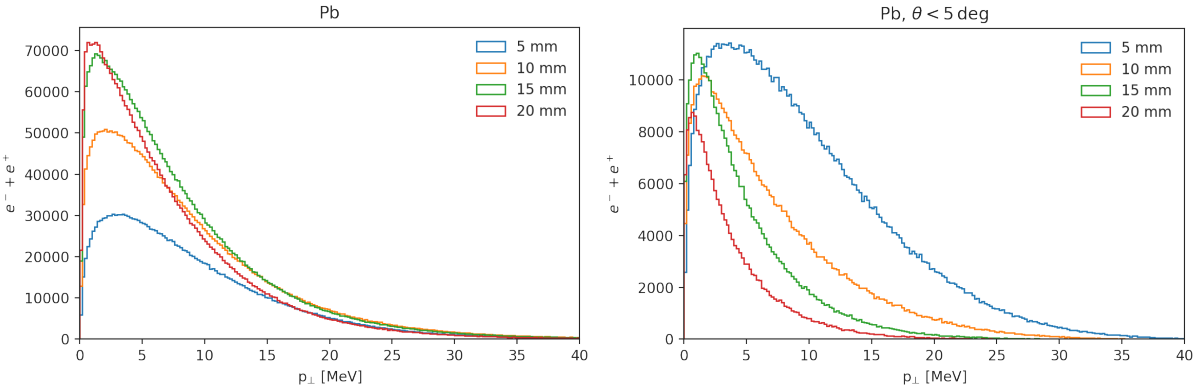

Figure 8.8: The electron and positron combined transversal momentum spectrum for lead converter targets of different thickness. The **left** shows all particles and the **right** only particles with an angle smaller than 5° with respect to the incident beam axis. $p_\perp$ is the momentum component that is orthogonal to the incident beam.

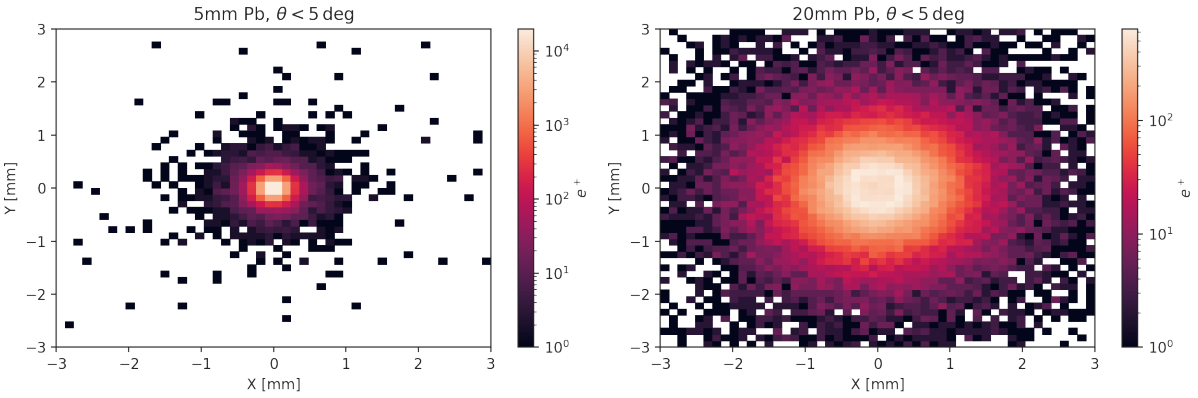

Figure 8.9: Position of the very forward going ($\theta < 5°$) output positrons as they leave a 5 mm (**left**) and a 20 mm (**right**) lead converter target. Most of the positrons are located in the central region with a diameter of 0.15 mm for 5 mm and 1.5 mm for 20 mm of lead.

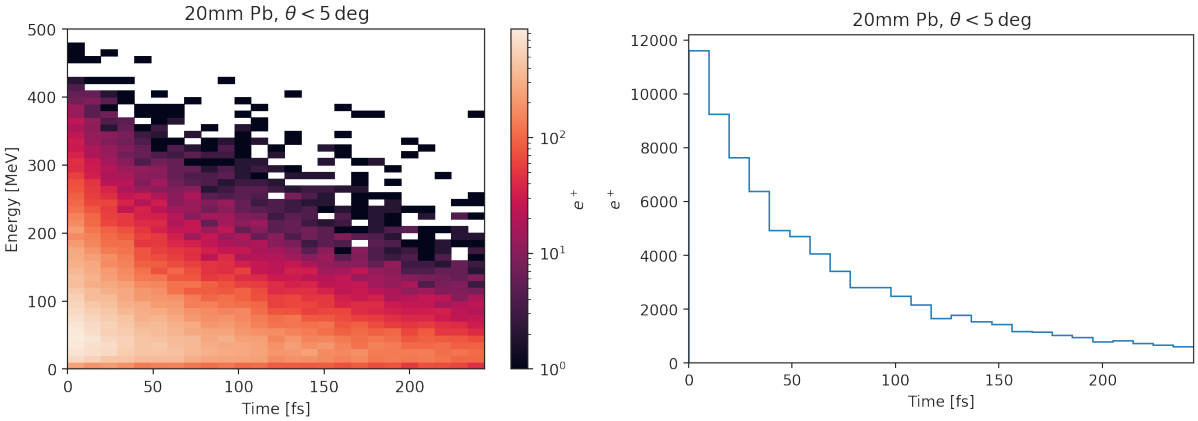

Figure 8.10: Time of leaving the conversion target for very forward going ($\theta < 5°$) output positrons. The **left** shows the correlation between time and energy while the **right** integrates over energy.

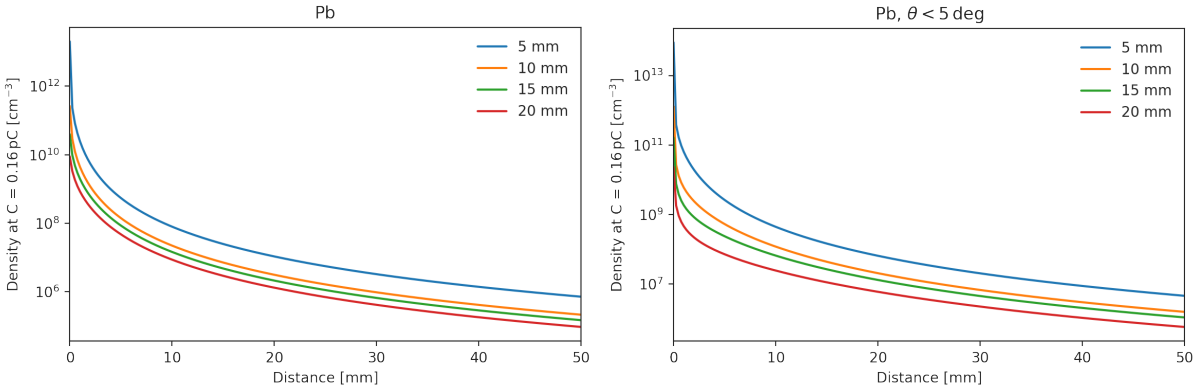

Figure 8.11: The density of the pair beam for lead converters of different thickness as a function of distance from the converter target for all particles (**left**) and very forward going ($\theta < 5°$, **right**). Within the first millimeters the density already drops by multiple orders of magnitude.

8.3 Beam Optics

In the previous section we described the conversion of the electron beam into a pair beam of electrons and positrons in a converter target. We found that the result is a highly divergent beam with a falling energy spectrum with a density that rapidly reduces over short distances. On top of that we only attain desired features like a momentum inversion and charge neutrality if we introduce an additional cut on the divergence angle, which would be equal to a small plasma cell some distance away. Finally the beam is also top transversally hot. All of these problem could potentially be addressed with clever beam optics that filter or focus the beam to transport it over large distances or even increase its density and decrease its transversal momentum at a focal point where the plasma cell could be placed.
A charged particle beam could be focused using magnetic beam elements like a solenoid magnet [69], quadrupole magnets [50] or active plasma lenses (APLs) [111, 132, 122]. However a beam

consisting of positively and negatively charged particles faces the additional challenge of finding beam optics that apply the same or similar focusing to both particle species. This is the case for a single solenoid magnet, where the Lorentz force acts twice on a particle to achieve a focusing effect. Inside the solenoid the magnetic field lines are parallel to the beam axis. A particle with a transverse momentum component will then experience a Lorentz force that leads to an azimuthal acceleration. The azimuthal momentum component will then likewise lead to transverse acceleration that counteracts the initial transverse momentum component independent of the charge of the particle. The twofold application of the Lorentz force leads to an energy dependent focusing that scales with γ^{-2}. The focusing strength of a single solenoid with magnetic field B and length L can be given in the thin lens approximation [69] as follows:

$$\frac{1}{f} \approx \left(\frac{e}{2m_e c\beta\gamma}\right)^2 B^2 L = 8.6 \left(\frac{\gamma}{100}\right)^{-2} \left(\frac{B}{\mathrm{T}}\right)^2 \left(\frac{L}{\mathrm{m}}\right) \mathrm{m}^{-1} \tag{8.7}$$

Using a solenoid to focus a pair beam would be conceptually very simple, because it focuses the beam symmetrical in both planes simultaneously, but the challenge would be to achieve magnetic fields that are sufficiently strong. A solenoid of length 10 cm would need a field strength exceeding 1 T to focus 50 MeV electrons with a focal length of 1 m. On the other hand quadrupole magnets focus a particle beam (depending on the charge) in one transverse plane and defocus in the other, making it necessary to employ more complex layouts of three quadrupoles [50] to achieve net focusing in both planes. APLs can symmetrically focus one particle species, while defocusing the other, making it again possible that a layout of three APLs could achieve a focusing effect for both species.
The energy dependence of the focal length could be used as an energy filter. Only particles of a narrow energy range will be focused at the position of the plasma cell. The unfocused particles with higher and lower energies will overall constitute a low density plasma beam that will not perturb the plasma cell much, whereas the focused particles constitute a much higher density plasma with a narrow energy spread that could be similar to a cold beam and drive significant instability growth.

8.4 Plasma Cell

The plasma cell will be at the heart of a beam line where the plasma instability growth takes place. The most important quantity of the background plasma will be its density that sets the time and length scales by way of the plasma frequency and directly influences the growth rate. The plasma cell needs to be long enough to contain the full instability growth. Background plasma needs to be stable over the entire instability growth time and have a density profile that is homogeneous enough (compare section 7.6). The density of the plasma also sets the transverse size of the beam for plasma like collective behavior. A plasma can be created by ionizing a gas with a current discharge or a laser [114, 128, 91]. Fully ionizing a Hydrogen gas of number density n_H will produce an electron-proton plasma of the same number density. The density of a Hydrogen plasma n_e is thus given by the ionization fraction x_e as well as the temperature and pressure of the original gas.

$$pV = N_\mathrm{H} k_B T \Rightarrow n_e = x_e \frac{p}{k_B T} = 2.4 \cdot 10^{20} x_e \left(\frac{p}{\mathrm{bar}}\right) \left(\frac{T}{300\,\mathrm{K}}\right)^{-1} \mathrm{cm}^{-3} \tag{8.8}$$

For ordinary conditions in a laboratory environment fully ionizing a Hydrogen gas will produce a plasma with a density of the order of $10^{20}\,\mathrm{cm}^{-3}$. For a density of this magnitude a beam with a transverse extent of more than $0.1\,\mu\mathrm{m}$ would already be enough to achieve full instability growth. Since the plasma frequency scales with the square root of the plasma density even for a much smaller density of $10^{14}\,\mathrm{cm}^{-3}$ a transverse extent in excess of $10\,\mu\mathrm{m}$ would still satisfy the requirement (compare figure 8.12).

$$L_\perp \gg \frac{2\pi c}{\omega_p} \tag{8.9}$$

The density of the background plasma needs to be stable over the entire distance of instability growth. For a configuration that needs of the order of 10 ps to reach saturation this would require a plateau of the density of the order of $300\,\mu\mathrm{m}$.

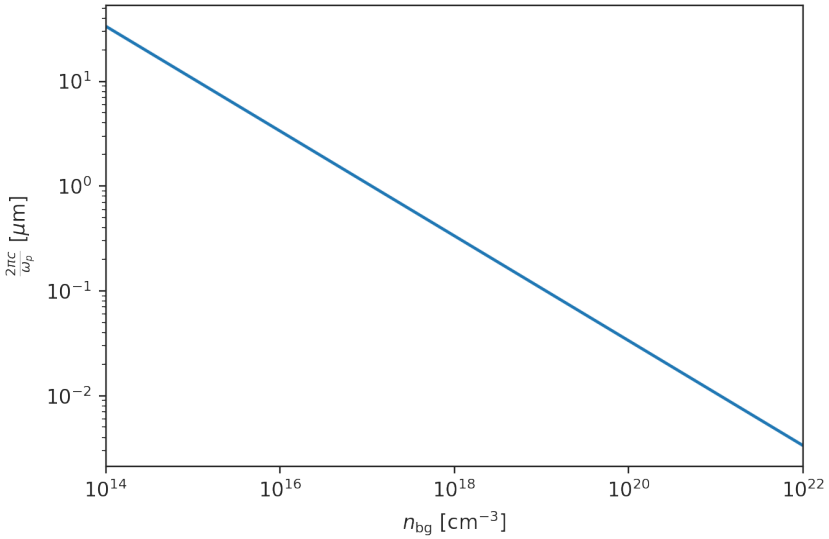

Figure 8.12: Minimum size of the plasma beam in each direction to achieve plasma-like behavior as a function of the background plasma density n_{bg}. For reasonable densities it is enough to exceed the threshold of $1\,\mu\mathrm{m}$.

8.5 Diagnostics

The instability mechanism leads to the growth of electromagnetic fields in the plasma cell. One possibility for measuring the instability growth would be directly measure these fields [128] for example via proton radiography [76, 75]. However we have seen in section 5.5 that the growth of the instability also leads to significant feedback on the momentum distribution of the pair beam. An initially cold beam with $\gamma = 150$ and $\alpha = 10^{-3}$ (figure 5.16) can develop a momentum spread of multiple MeV in both directions. Using a magnetic spectrometer, where particles are deflected in a magnetic field and the amount of their deflection is recorded with a detector and used to infer their momentum, the momentum distribution of the pair beam could be measured with sufficient resolution to detect this effect. The two parameters α and γ can be varied for example by adjusting the gas pressure in the plasma cell, the magnetic field of the solenoid or the

thickness of the converter target. The scaling of the change of the momentum distribution with α and γ could be probed for unequivocal proof that the effect is caused by the oblique instability. If the instability is grown up to the saturation level then the scaling can be simply compared with (5.15). Since the saturation level scaling follows from the growth rate scaling the change of the momentum distribution allows for direct probing of the instability growth rate itself.

8.6 Pure Electron Experiment

Performing a laboratory experiment with pair beams of electrons and positrons has shown itself to be very tricky with the key hurdle being the creation of handling of the pair beam itself. On top of that we have shown that the pair beams we could create would likely have a large overabundance of electrons, leading to a pair beam that is very different from the perfectly neutral astrophysical beams we aim to learn about. Luckily we have seen in section 7.3 that non-neutral beams undergo the same instability mechanism as neutral pair beams. In fact, even a beam that consists only of electrons undergoes the same unstable behavior. This opens up the exciting opportunity of performing the suggested experiment with a pure electron beam, sidestepping all the problems we have described about creating a pair beam in a conversion target and focusing the beam with complicated beam optics. In the next chapter we will investigate the unstable behavior of pure electrons in-depth.

9 | Pure Electron Beam

In section 7.3 we studied the influence of non-neutrality of the beam due to an over-abundance of electrons as a correction to the idealized scenario of a neutral pair beam, which we want to study. We realized that the instability can even proceed for beams with a very strong asymmetry between electrons and positrons. It is even possible to produce instable growth for a pure electron beam. In fact the momentum broadening of the beam is the same for a neutral pair beam and a pure electron beam. Since the production of a neutral pair beam of high density and high charge is a great experimental challenge, whereas producing pure electron beam of the same qualities is comparatively much easier with current technology, we want to explore the possibility of performing the experiment with a pure electron beam.
To this end we need to check that a pure electron beam comforts to the same scaling behavior as a neutral pair beam, that the key experimental observables are not spoiled by the effects of charge-carrying beam and that there are no cross effects between the non-neutrality and other corrections to the idealized neutral pair beam discussed in chapter 7. We know that a neutral pair beam and a pure electron beam start to behave much differently if an external magnetic field is applied. However also the scenario of finite beam extension needs to be investigated.

9.1 Growth Rates

As a first step we want to confirm that for an infinitely extend, cold, pure electron beam in a homogeneous environment the instability evolves with the same rate and leads to a diffusion of the beam momentum distribution with the same rate and saturation level. Only then we can conclude that using a charged beam as a proxy for a neutral pair beam will lead to the same result under optimal conditions. Up to the obvious absence of a positron species in the simulation and the total runtime of the simulation, which was reduced to $1500\omega_p^{-1}$, we use the same configuration for the simulation runs presented in this chapter as outlined in table 5.1.
The most straight-forward way to determine the rate of instability growth is the integrated growth rate that can be extracted from the evolution of the total field energy density evolution. In figures 9.1 and 9.3 we show the evolution of the total field energy density as a function of γ and α respectively. In both cases the initial noise regime exhibits a oscillation of the field energy density, that depends on the current density of the beam, until the instability takes over and the field energy density grows exponentially. We expect that the instability grows with the rate of the oblique instability and the time until saturation t_{final} scales like the inverse growth rate. Again we extract the growth rate and the time of saturation by fitting the logarithmic energy density to (5.5).
Figures 9.2 and 9.4 show that the growth rate is fitted to follow a power law as expected albeit

with a slightly smaller spectral index of -0.27 for the γ scaling and -0.24 for the α scaling. The inverse saturation time t_{final}^{-1} also follows a power law with a spectral index that is closer to the expected value of $\frac{1}{3}$ for both α and γ. The scaling of the inverse saturation time is much less sensitive to the initial oscillation of the field energy density at early times and could be the more reliable quantity.
To further investigate the growth rate of the instability we also consider the constrained growth rate that only considers the energy density of the most dominant modes. When only considering a small number of modes the influence of the initial oscillation due to the beam's current is strongly reduced. Figures 9.5 and 9.6 show the mode constrained field energy density evolution. For the energy density of the modes that are relevant to the instability we observe no oscillation at early times. Consequently the exponential rise of the field energy density becomes apparent at earlier times. At late times the total field energy density is dominated by the selected instable modes and thus the saturation behavior remains the same. The extracted growth rates have been plotted in figure 9.7 as functions of both γ and α and clearly conform to a power law with the exact scaling of the oblique instability. We can conclude that indeed the growth of the instable modes behaves the same for a pure electron beam and a neutral pair beam in the idealized scenario of an infinitely extended, cold beam.
For a laboratory experiment the change of the beam momentum distribution due to the growth of the instability is a key observable. We need to check that this observable remains experimentally accessible for a pure electron beam. The energy loss of the beam leads to the shift of the mean longitudinal momentum $\text{Mean}(p_\parallel)$ to lower values and the build up of the electric field leads to the diffusion of beam momentum that widens the beam momentum distribution ultimately growing $\text{Std}(p_\parallel)$ and $\text{Std}(p_\perp)$. However for a charged beam the presence of a current also leads to an oscillation of the beam energy and thus the mean longitudinal momentum. To reflect this we extend (3.42) with an additional oscillation term, where A is the amplitude of the oscillation that proceeds with a frequency of $\frac{\omega_p}{2}$.

$$\Delta\mu(t) = -2\delta V_1 \int_0^t \mathrm{d}t' \left(\kappa^{-\delta t'} + \kappa^{-\delta t_{\text{final}}}\right)^{-\frac{2}{\log\kappa}} + A\sin^2\left(\frac{\omega_p t}{2}\right) \tag{9.1}$$

The evolution of the longitudinal and transversal momentum spread is not significantly altered and we continue to use (3.43).
Figures 9.8 and 9.9 show the evolution of the mean longitudinal momentum shift for simulation runs with varying γ and α respectively. The extended model (9.1) can describe the total evolution quite well far into the non-linear regime. The parameters of each individual fit are listed in table 9.1 and the scaling of the growth rate and the saturation level with α and γ have been plotted in figure 9.10. Notably we find that the smoothness parameter κ is fitted with a higher value, meaning a faster transition into the non-linear regime, for higher saturation levels. For a neutral beam we found this behavior only for the evolution of $\text{Std}(p_\perp)$, which was best described by our model.
The evolution of the longitudinal momentum spread is shown in figure 9.11 and 9.12 for simulation runs with varying γ and α. The fitted diffusion model (3.43) describes the evolution quite well, but tends to underestimate the momentum spread far into the non-linear regime, a behavior we also observed for the neutral pair beam. The parameters of each individual fit are listed in table 9.2 and the scaling of the growth rate and the saturation level with α and γ have been plotted in figure 9.13.

Lastly, the evolution of the transversal momentum spread is shown in figures 9.14 and 9.14. Here the fitted model describes the data very well deep into the non-linear regime. The scaling of the growth rate and saturation level are plotted in figure 9.16. Again we find that a higher saturation level is associated with a less smooth transition from the linear into the non-linear regime (see table 9.3).

$$\delta_{\Delta\mathrm{Mean}(p_\parallel)} \propto \frac{\alpha^{0.27}}{\gamma^{0.33}} \tag{9.2}$$

$$\Lambda_{\Delta\mathrm{Mean}(p_\parallel)} \propto \alpha^{0.48}\gamma^{0.61} \tag{9.3}$$

$$\delta_{\mathrm{Std}(p_\parallel)} \propto \frac{\alpha^{0.33}}{\gamma^{0.33}} \tag{9.4}$$

$$\Lambda_{\mathrm{Std}(p_\parallel)} \propto \alpha^{0.36}\gamma^{0.59} \tag{9.5}$$

$$\delta_{\mathrm{Std}(p_\perp)} \propto \frac{\alpha^{0.34}}{\gamma^{0.32}} \tag{9.6}$$

$$\Lambda_{\mathrm{Std}(p_\perp)} \propto \alpha^{0.32}\gamma^{0.63} \tag{9.7}$$

For the widening of the momentum distribution in both directions we find that the growth rate replicates the rate of the oblique instability and that the saturation level also agrees with what has been found for the neutral beam (see chapter 5.5.1). For the change of the mean longitudinal momentum this is not strictly the case, as we find slightly different scaling relationships with respect to α, but this could again be influenced by the early oscillation leading to a biased fit. Nevertheless, if we discount the early oscillation of the mean longitudinal momentum, the evolution of the momentum distribution of an idealized pure electron beam is very similar to a neutral pair beam and we can conclude that at least in principle a pure electron beam would allow us to study the same physics.

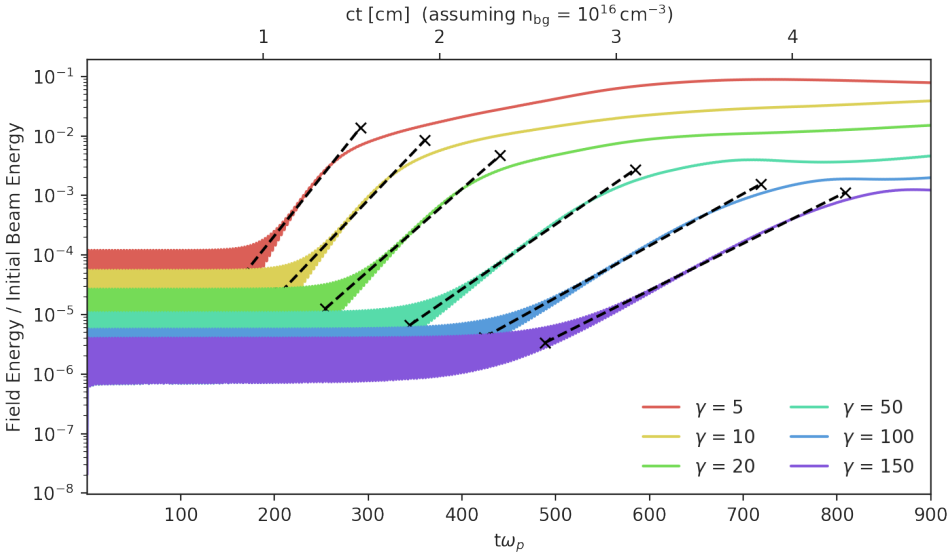

Figure 9.1: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the electron beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a density contrast of $\alpha = 10^{-3}$ and varying Lorentz boost γ. For each simulation run the start and end point of the linear growth phase is indicated by cross and the linear growth rate is indicated by a dashed line. The extracted growth rate and the time of the end of the linear growth phase are shown in figure 9.2.

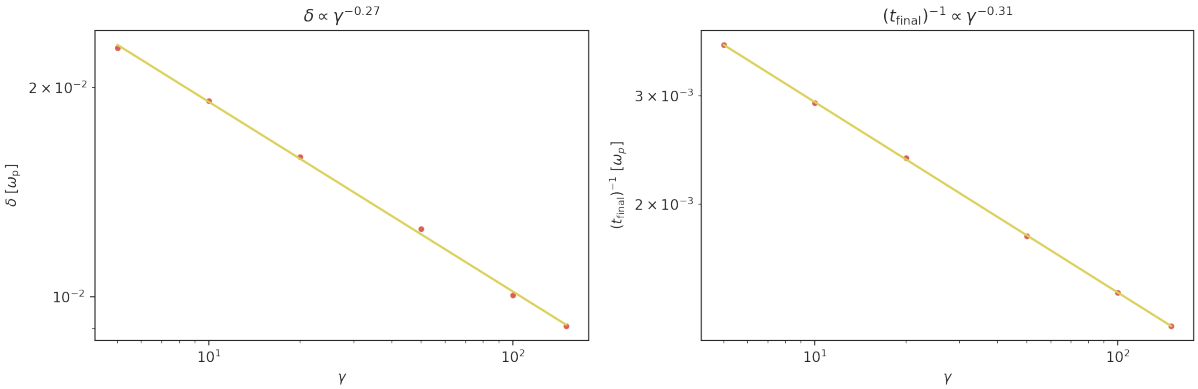

Figure 9.2: The integrated growth rate (**left**) and the inverse time at the end of the linear growth phase (**right**) as a function of γ extracted from total energy density shown in figure 9.1. Both follow a power law with a fitted index slightly below $\frac{1}{3}$.

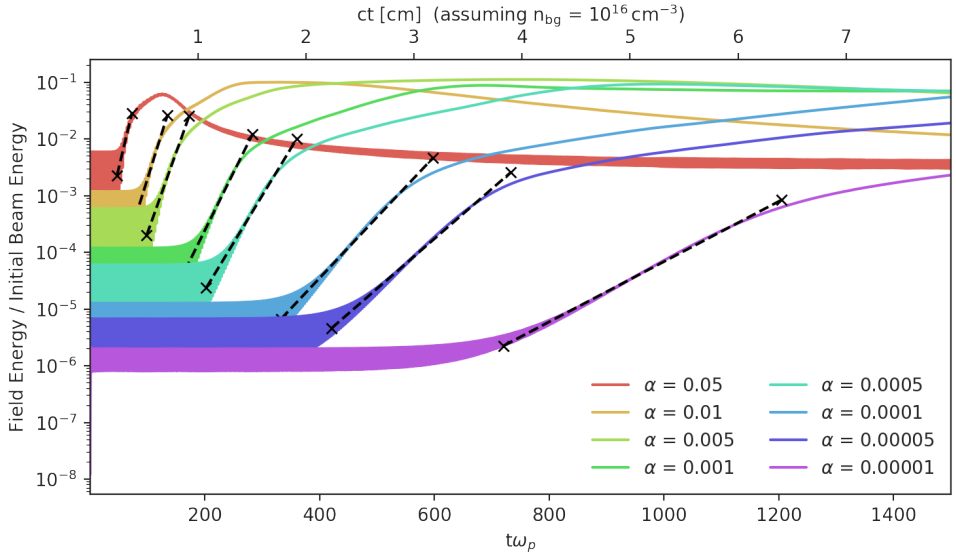

Figure 9.3: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the electron beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a Lorentz boost of $\gamma = 5$ and varying density contrast α. For each simulation run the start and end point of the linear growth phase is indicated by cross and the linear growth rate is indicated by a dashed line. The extracted growth rate and the time of the end of the linear growth phase are shown in figure 9.4.

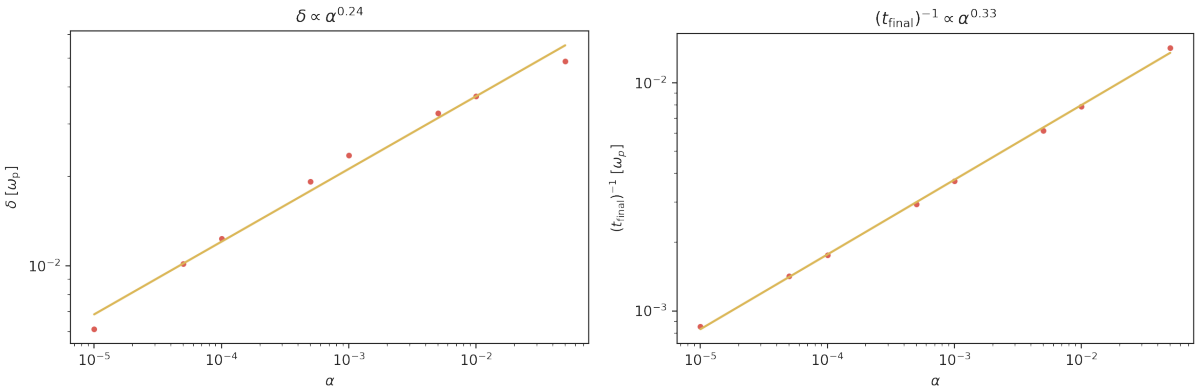

Figure 9.4: The integrated growth rate (**left**) and the inverse time at the end of the linear growth phase (**right**) as a function of α extracted from total energy density shown in figure 9.3. Both follow a power law with a fitted index of $\frac{1}{3}$ for the inverse saturation time t_{final}^{-1} and slightly lower for the growth rate.

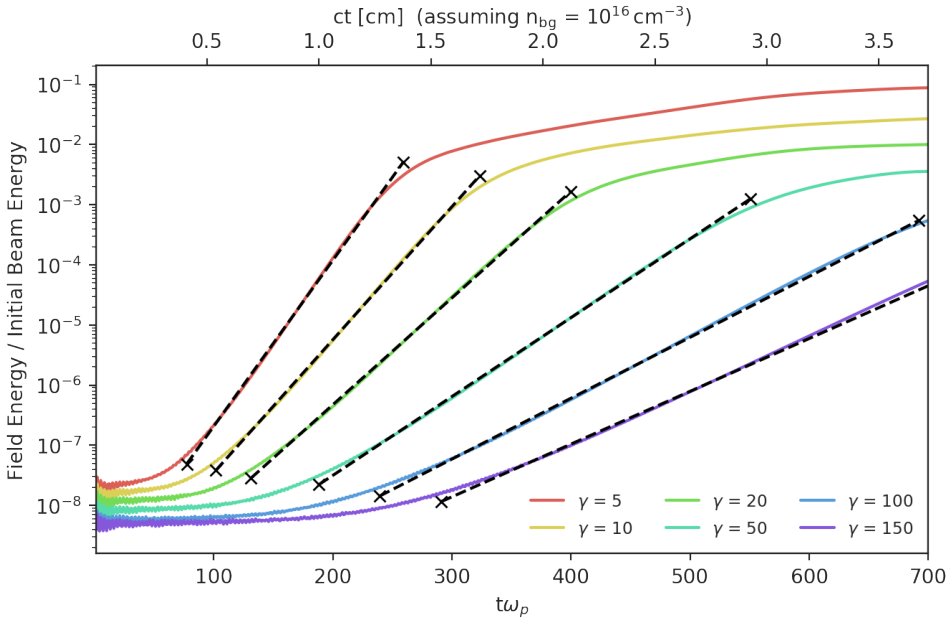

Figure 9.5: The energy density of the electric field that dominate the instability ($k_{\parallel} \sim \frac{\omega_p}{c}$) normalized to the initial energy density of the electron beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a density contrast of $\alpha = 10^{-3}$ and varying Lorentz boost γ. The dashed line indicates the fitted linear growth phase. The extracted growth rate is shown in the left of figure 9.7.

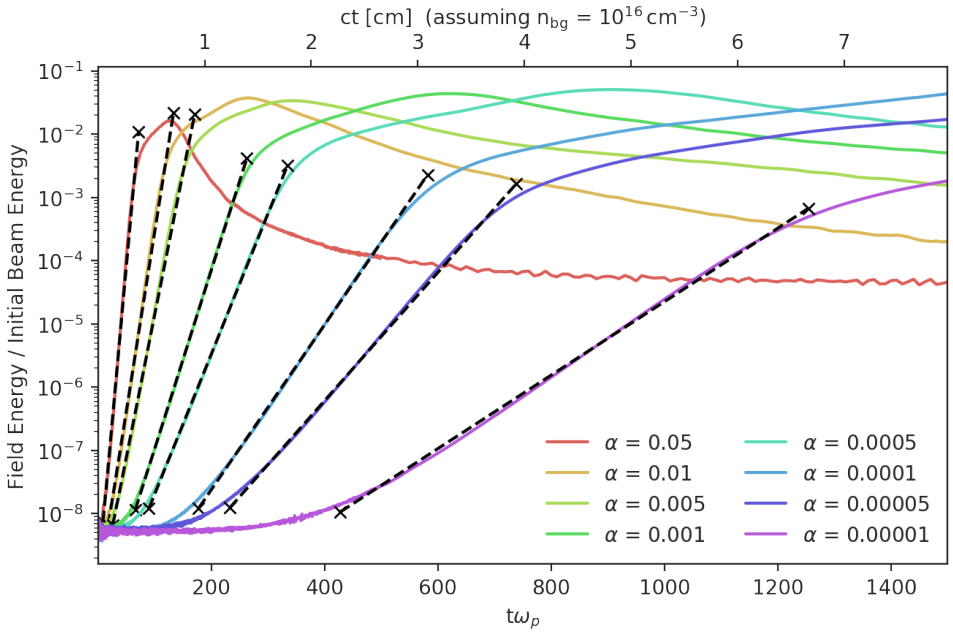

Figure 9.6: The energy density of the electric field that dominate the instability ($k_{\parallel} \sim \frac{\omega_p}{c}$) normalized to the initial energy density of the electron beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a Lorentz boost of $\gamma = 5$ and varying density contrast α. The dashed line indicates the fitted linear growth phase. The extracted growth rate is shown in the left of figure 9.7.

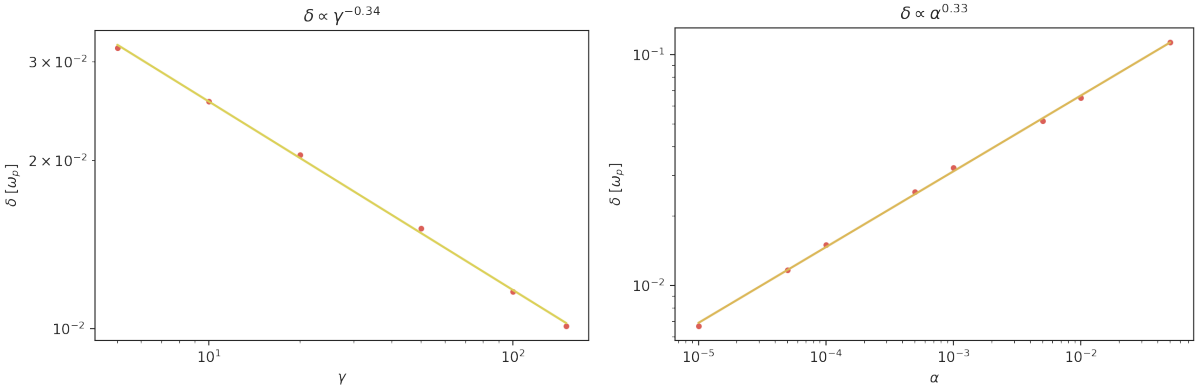

Figure 9.7: The constrained growth rate as a function of γ (**left**) and α (**right**). They both follow a power law that is fitted with a spectral index of $\frac{1}{3}$ as expected for the oblique instability.

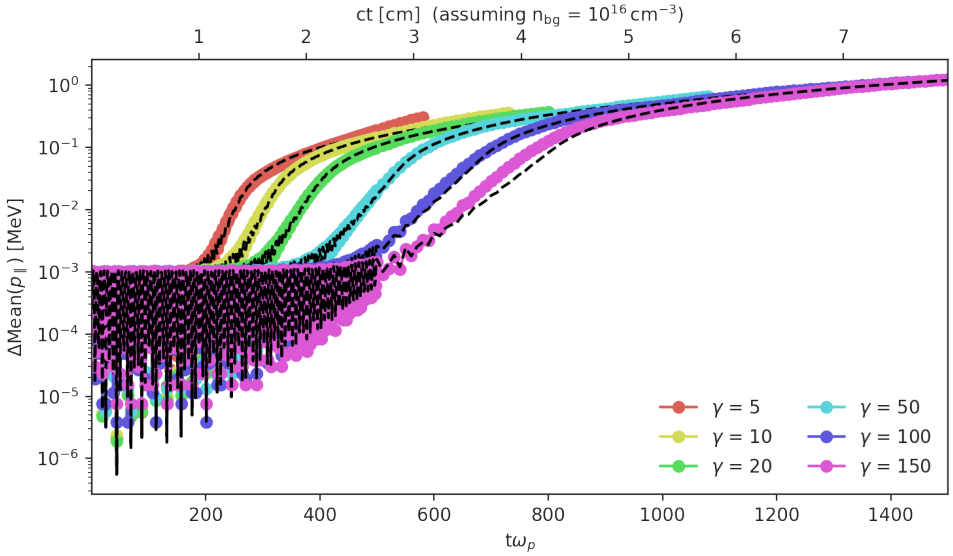

Figure 9.8: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ is shown for runs with varying Lorentz boost γ and density contrast $\alpha = 10^{-3}$. The mean value is calculated at regular intervals of $5\omega_p^{-1}$. However at late times, after $3t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (9.1) to the data is shown. The fit parameters are listed in table 9.1 and the scaling of growth rate and saturation level is shown in figure 9.10.

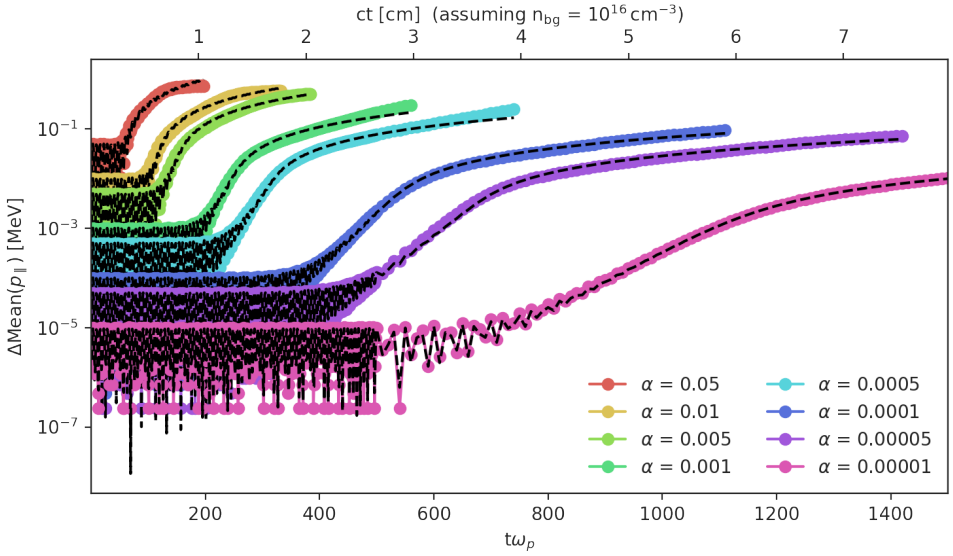

Figure 9.9: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ is shown for runs with varying density contrast α and Lorentz boost $\gamma = 5$. The mean value is calculated at regular intervals of $5\omega_p^{-1}$. However at late times, after $3t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (9.1) to the data is shown. The fit parameters are listed in table 9.1 and the scaling of growth rate and saturation level is shown in figure 9.10.

Table 9.1: The fit parameters (δ, V_1, A, t_{final}, κ) of (9.1) for $\Delta\text{Mean}(p_\parallel)$ from simulation runs with various values for α and γ. The derived quantity Λ is the value of $\Delta\text{Mean}(p_\parallel)$ at time t_{final}.

γ	α	δ [ω_p]	V_1 [MeV]	A [MeV]	t_{final} [ω_p^{-1}]	κ	Λ [MeV]
5	0.001	$2.06 \cdot 10^{-8}$	$2.61 \cdot 10^{-2}$	$1.01 \cdot 10^{-3}$	254	$1.48 \cdot 10^{3}$	$1.21 \cdot 10^{-2}$
10	0.001	$2.79 \cdot 10^{-8}$	$2.09 \cdot 10^{-2}$	$1.03 \cdot 10^{-3}$	317	$1.45 \cdot 10^{15}$	$1.78 \cdot 10^{-2}$
20	0.001	$2.48 \cdot 10^{-8}$	$1.75 \cdot 10^{-2}$	$1.04 \cdot 10^{-3}$	390	$1.20 \cdot 10^{3}$	$2.14 \cdot 10^{-2}$
50	0.001	$4.88 \cdot 10^{-8}$	$1.27 \cdot 10^{-2}$	$1.05 \cdot 10^{-3}$	537	$4.91 \cdot 10^{38}$	$4.42 \cdot 10^{-2}$
100	0.001	$6.91 \cdot 10^{-8}$	$1.01 \cdot 10^{-2}$	$1.03 \cdot 10^{-3}$	684	$6.33 \cdot 10^{38}$	$6.83 \cdot 10^{-2}$
150	0.001	$1.78 \cdot 10^{-7}$	$8.15 \cdot 10^{-3}$	$1.05 \cdot 10^{-3}$	811	$3.21 \cdot 10^{44}$	$9.84 \cdot 10^{-2}$
5	0.05	$5.68 \cdot 10^{-6}$	$6.83 \cdot 10^{-2}$	$4.98 \cdot 10^{-2}$	66.6	$8.63 \cdot 10^{36}$	$9.54 \cdot 10^{-2}$
5	0.01	$5.95 \cdot 10^{-7}$	$4.55 \cdot 10^{-2}$	$1.01 \cdot 10^{-2}$	120	$4.91 \cdot 10^{10}$	$3.40 \cdot 10^{-2}$
5	0.005	$2.43 \cdot 10^{-7}$	$3.82 \cdot 10^{-2}$	$5.01 \cdot 10^{-3}$	152	591	$2.69 \cdot 10^{-2}$
5	0.001	$2.40 \cdot 10^{-8}$	$2.62 \cdot 10^{-2}$	$1.01 \cdot 10^{-3}$	251	830	$1.19 \cdot 10^{-2}$
5	0.0005	$1.45 \cdot 10^{-8}$	$2.10 \cdot 10^{-2}$	$5.09 \cdot 10^{-4}$	317	$1.20 \cdot 10^{4}$	$9.20 \cdot 10^{-3}$
5	0.0001	$2.51 \cdot 10^{-9}$	$1.36 \cdot 10^{-2}$	$1.01 \cdot 10^{-4}$	532	73.4	$4.44 \cdot 10^{-3}$
5	0.00005	$2.00 \cdot 10^{-9}$	$1.06 \cdot 10^{-2}$	$5.05 \cdot 10^{-5}$	680	244	$3.52 \cdot 10^{-3}$
5	0.00001	$1.64 \cdot 10^{-10}$	$7.27 \cdot 10^{-3}$	$9.47 \cdot 10^{-6}$	$1.11 \cdot 10^{3}$	10.4	$1.27 \cdot 10^{-3}$

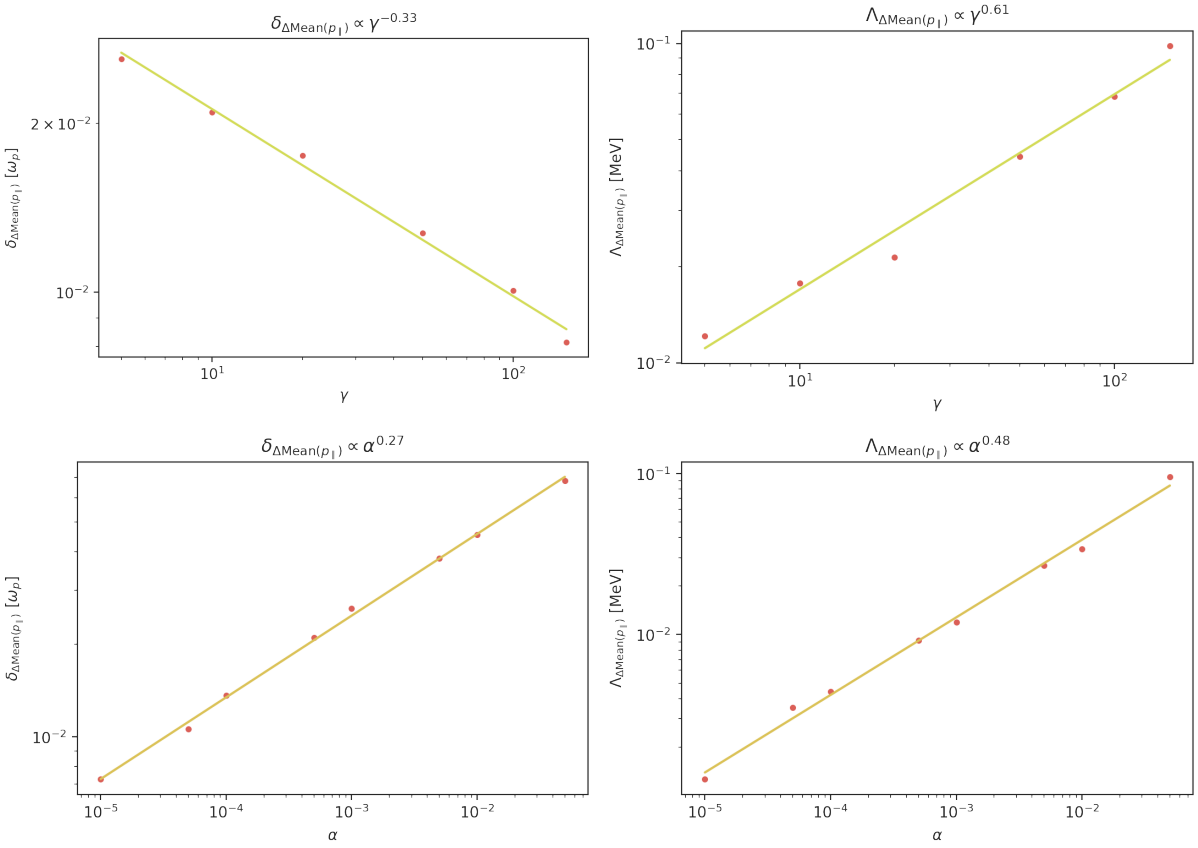

Figure 9.10: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the shift of mean parallel momentum $\Delta\text{Mean}(p_\parallel)$ with γ (**top**) and α (**bottom**). In all case the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (9.1).

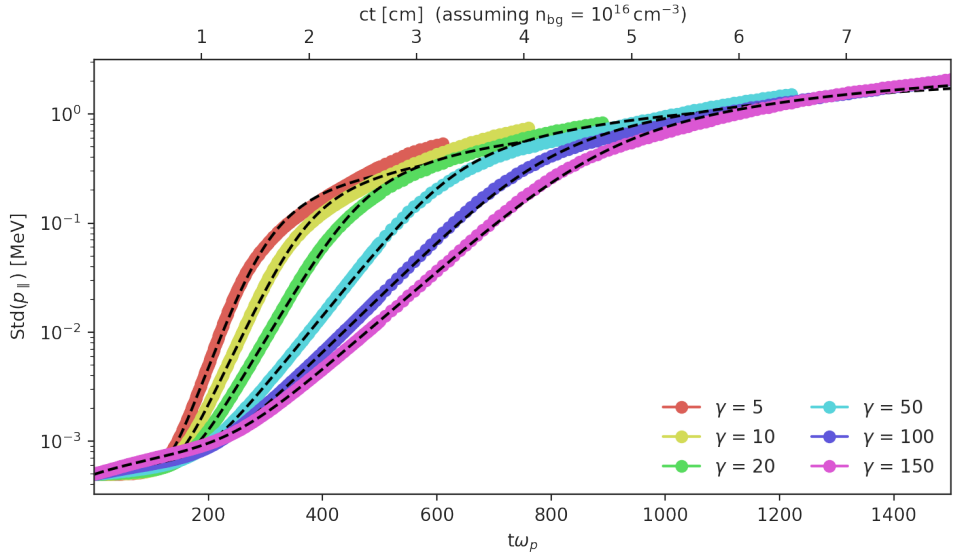

Figure 9.11: The evolution of the width of the longitudinal momentum distribution $\mathrm{Std}(p_\parallel)$ is shown for runs with varying Lorentz boost γ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However at late times, after $2t_{\mathrm{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 9.2 and the scaling of growth rate and saturation level is shown in figure 9.13.

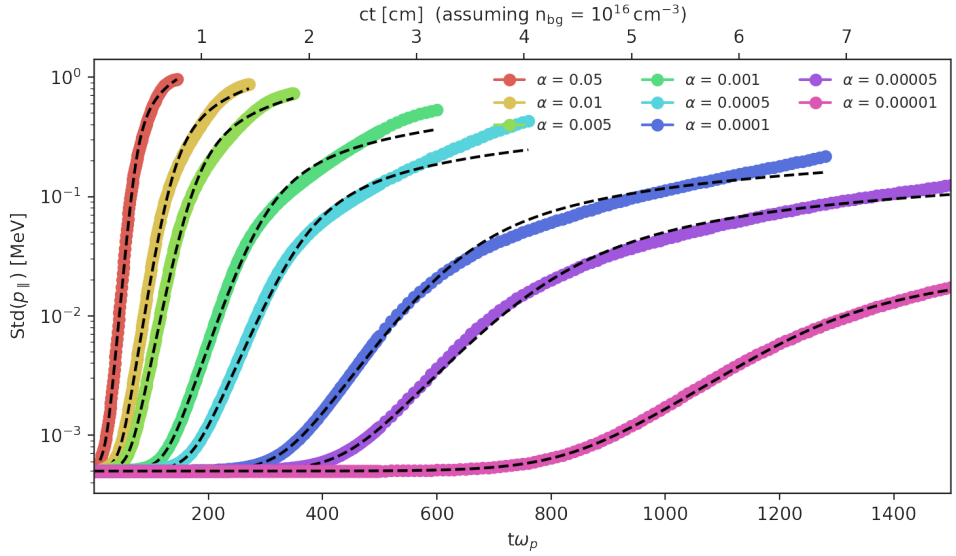

Figure 9.12: The evolution of the width of the longitudinal momentum distribution $\mathrm{Std}(p_\parallel)$ is shown for runs with varying density contrast α and Lorentz boost $\gamma = 5$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However at late times, after $2t_{\mathrm{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 9.2 and the scaling of growth rate and saturation level is shown in figure 9.13.

Table 9.2: The fit parameters (δ, D_0, D_1, t_{final} and κ) of (3.43) for $\text{Std}(p_\parallel)$ from simulation runs with various values for α and γ. The derived quantity Λ is the value of $\text{Std}(p_\perp)$ at time t_{final}.

γ	α	δ $[\omega_p]$	D_0 $[\text{MeV}^2\omega_p]$	D_1 $[\text{MeV}^2\omega_p]$	t_{final} $[\omega_p^{-1}]$	κ	Λ [MeV]
5	0.001	$3.22 \cdot 10^{-2}$	$6.62 \cdot 10^{-11}$	$1.99 \cdot 10^{-12}$	287	2.55	$4.90 \cdot 10^{-2}$
10	0.001	$2.54 \cdot 10^{-2}$	$9.02 \cdot 10^{-11}$	$4.39 \cdot 10^{-12}$	354	3.08	$7.23 \cdot 10^{-2}$
20	0.001	$2.02 \cdot 10^{-2}$	$1.32 \cdot 10^{-10}$	$8.31 \cdot 10^{-12}$	446	3.57	0.117
50	0.001	$1.48 \cdot 10^{-2}$	$2.89 \cdot 10^{-10}$	$2.09 \cdot 10^{-11}$	602	4.11	0.213
100	0.001	$1.18 \cdot 10^{-2}$	$4.44 \cdot 10^{-10}$	$3.95 \cdot 10^{-11}$	752	3.27	0.290
150	0.001	$1.05 \cdot 10^{-2}$	$9.41 \cdot 10^{-10}$	$4.88 \cdot 10^{-11}$	854	3.12	0.352
5	0.05	0.118	$5.70 \cdot 10^{-9}$	$3.88 \cdot 10^{-10}$	70.9	3.23	0.163
5	0.01	$6.91 \cdot 10^{-2}$	$8.89 \cdot 10^{-10}$	$4.11 \cdot 10^{-11}$	130	2.33	0.109
5	0.005	$5.47 \cdot 10^{-2}$	$3.52 \cdot 10^{-10}$	$1.72 \cdot 10^{-11}$	166	2.32	$8.90 \cdot 10^{-2}$
5	0.001	$3.23 \cdot 10^{-2}$	$5.60 \cdot 10^{-11}$	$2.74 \cdot 10^{-12}$	283	2.55	$5.10 \cdot 10^{-2}$
5	0.0005	$2.56 \cdot 10^{-2}$	$2.93 \cdot 10^{-11}$	$1.12 \cdot 10^{-12}$	353	2.85	$3.61 \cdot 10^{-2}$
5	0.0001	$1.51 \cdot 10^{-2}$	$4.76 \cdot 10^{-12}$	$1.89 \cdot 10^{-13}$	613	2.81	$2.37 \cdot 10^{-2}$
5	0.00005	$1.19 \cdot 10^{-2}$	$2.94 \cdot 10^{-12}$	$8.94 \cdot 10^{-14}$	770	2.60	$1.61 \cdot 10^{-2}$
5	0.00001	$6.93 \cdot 10^{-3}$	$2.39 \cdot 10^{-13}$	$1.70 \cdot 10^{-14}$	$1.24 \cdot 10^3$	8.41	$7.34 \cdot 10^{-3}$

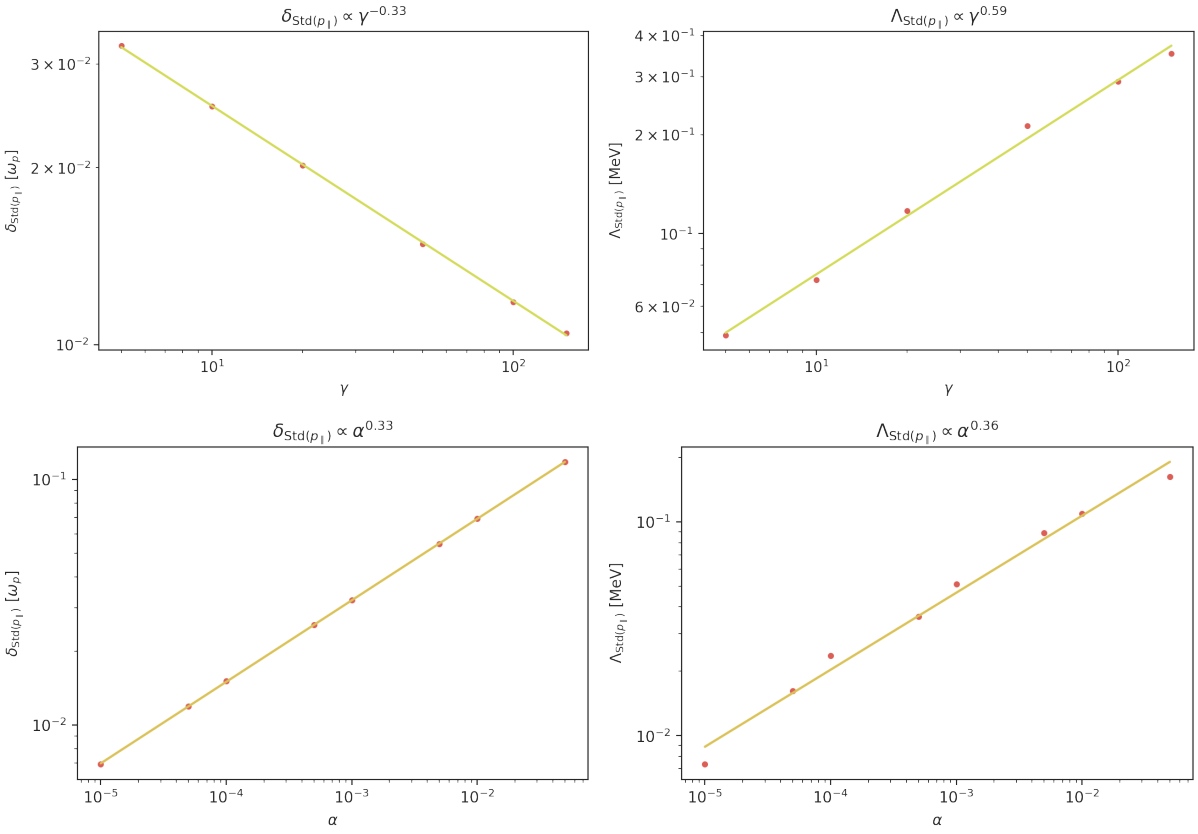

Figure 9.13: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the width of the parallel momentum $\mathrm{Std}(p_\parallel)$ with γ (**top**) and α (**bottom**). In all case the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.43).

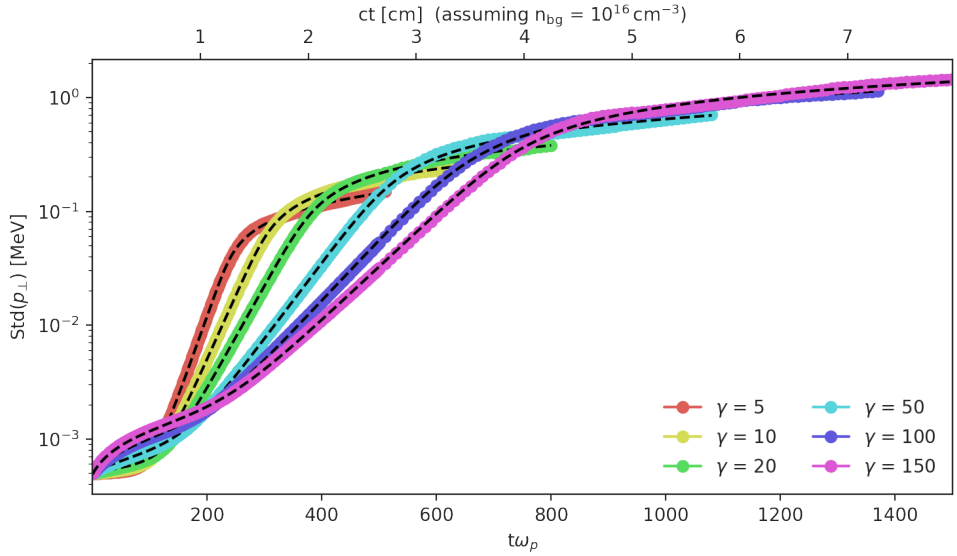

Figure 9.14: The evolution of the width of the transversal momentum distribution Std($p_\perp$) is shown for runs with varying Lorentz boost γ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However at late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 9.3 and the scaling of growth rate and saturation level is shown in figure 9.16.

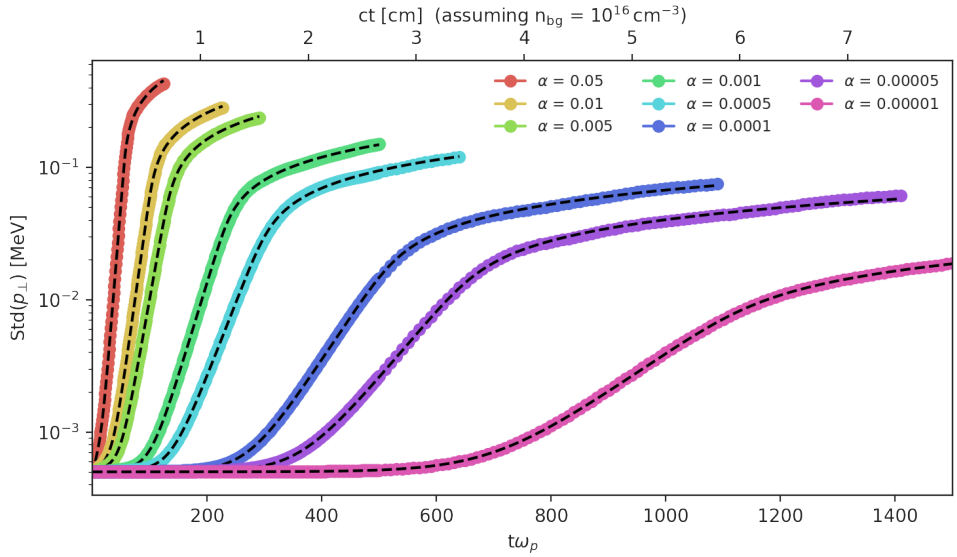

Figure 9.15: The evolution of the width of the transversal momentum distribution Std($p_\perp$) is shown for runs with varying density contrast α and Lorentz boost $\gamma = 5$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$. However at late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown. The fit parameters are listed in table 9.3 and the scaling of growth rate and saturation level is shown in figure 9.16.

Table 9.3: The fit parameters (δ, D_0, D_1, t_{final} and κ) of (3.43) for $\text{Std}(p_\perp)$ from simulation runs with various values for α and γ. The derived quantity Λ is the value of $\text{Std}(p_\perp)$ at time t_{final}.

γ	α	δ $[\omega_p]$	D_0 $[\text{MeV}^2\omega_p]$	D_1 $[\text{MeV}^2\omega_p]$	t_{final} $[\omega_p^{-1}]$	κ	Λ [MeV]
5	0.001	$3.25 \cdot 10^{-2}$	$9.97 \cdot 10^{-11}$	$1.06 \cdot 10^{-11}$	231	$1.62 \cdot 10^{-2}$	$3.19 \cdot 10^{-2}$
10	0.001	$2.59 \cdot 10^{-2}$	$2.20 \cdot 10^{-10}$	$2.06 \cdot 10^{-11}$	294	$5.04 \cdot 10^{-2}$	$5.28 \cdot 10^{-2}$
20	0.001	$2.09 \cdot 10^{-2}$	$4.77 \cdot 10^{-10}$	$3.72 \cdot 10^{-11}$	365	$8.74 \cdot 10^{-2}$	$7.72 \cdot 10^{-2}$
50	0.001	$1.55 \cdot 10^{-2}$	$1.37 \cdot 10^{-9}$	$8.27 \cdot 10^{-11}$	498	0.122	0.138
100	0.001	$1.24 \cdot 10^{-2}$	$3.14 \cdot 10^{-9}$	$1.66 \cdot 10^{-10}$	624	0.175	0.213
150	0.001	$1.09 \cdot 10^{-2}$	$4.91 \cdot 10^{-9}$	$2.16 \cdot 10^{-10}$	712	0.136	0.277
5	0.05	0.122	$1.03 \cdot 10^{-8}$	$1.76 \cdot 10^{-9}$	55.6	$6.46 \cdot 10^{40}$	0.108
5	0.01	$7.13 \cdot 10^{-2}$	$1.85 \cdot 10^{-9}$	$1.79 \cdot 10^{-10}$	101	144	$6.42 \cdot 10^{-2}$
5	0.005	$5.62 \cdot 10^{-2}$	$8.16 \cdot 10^{-10}$	$7.67 \cdot 10^{-11}$	130	32.7	$5.21 \cdot 10^{-2}$
5	0.001	$3.27 \cdot 10^{-2}$	$1.16 \cdot 10^{-10}$	$1.26 \cdot 10^{-11}$	228	82.7	$3.29 \cdot 10^{-2}$
5	0.0005	$2.57 \cdot 10^{-2}$	$4.70 \cdot 10^{-11}$	$6.17 \cdot 10^{-12}$	292	29.0	$2.60 \cdot 10^{-2}$
5	0.0001	$1.50 \cdot 10^{-2}$	$6.37 \cdot 10^{-12}$	$1.11 \cdot 10^{-12}$	505	28.1	$1.58 \cdot 10^{-2}$
5	0.00005	$1.18 \cdot 10^{-2}$	$3.12 \cdot 10^{-12}$	$5.84 \cdot 10^{-13}$	640	29.5	$1.24 \cdot 10^{-2}$
5	0.00001	$6.89 \cdot 10^{-3}$	$6.84 \cdot 10^{-13}$	$1.10 \cdot 10^{-13}$	$1.09 \cdot 10^{3}$	28.1	$6.95 \cdot 10^{-3}$

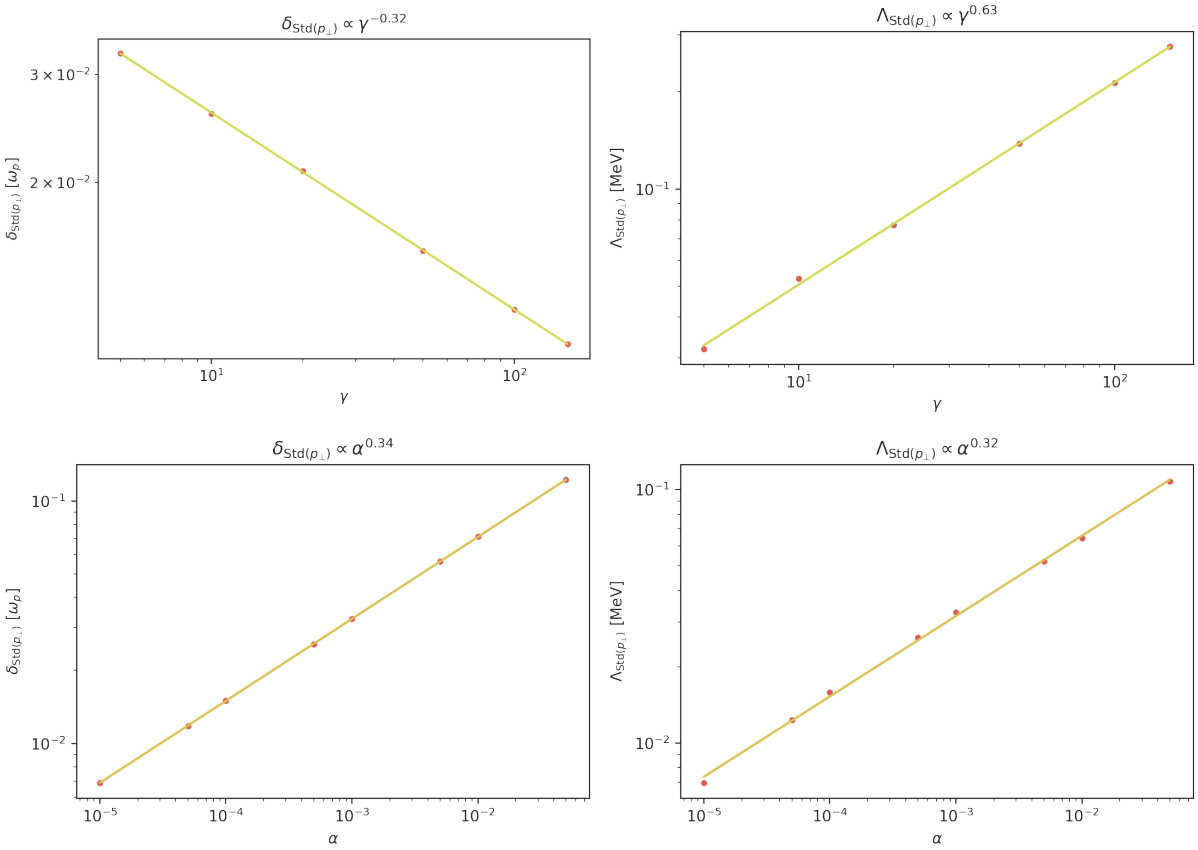

Figure 9.16: The scaling of the growth rate δ (**left**) and the saturation value Λ (**right**) of the width of the transversal momentum $\mathrm{Std}(p_\perp)$ with γ (**top**) and α (**bottom**). In all case the scaling is well described by a power-law. For each simulation the growth rate and the saturation level are extracted by fitting (3.43).

9.2 Finite Extent

The main difference between the neutral pair beam and a pure electron beam is that in the case of the non-neutral beam it carries non-zero charge and current densities. For an infinitely extended, homogeneous beam the produced fields cancel. This is not the case for a beam that has a finite length or width. The general effect of a finite extent have already been discussed in section 7.1, where we found that only a subset of the beam will participate in the full instability evolution leading to an overall reduced growth rate. Moreover in the case of a finite width, the increase of the transversal momentum spread due to the instability mechanism produces a strongly diverging beam that could make the further transport of the beam a challenge.
The effect of the electric and magnetic fields are in a sense opposing. The electric field produces a force that pushes particles away from the center of the beam (both in longitudinal and transversal direction), whereas the toroidal magnetic field produces a force that pushes particles towards the center in transversal direction (essential working as a focusing field).
We study the influence of the finite extent on a charged beam by considering separately a beam with a fixed finite length and a beam with a fixed finite width and scaling the charge ratio χ

from 0 (for a pure electron beam) to 1 (for a neutral beam) and noting the influence on beam evolution. The values for the length and width of the beam are chosen to be sufficiently large that the overall growth rate of the instability is roughly the same as in the infinitely extended scenario (compare section 7.1).

9.2.1 Finite Length

For the parameters of the simulation setup ($\gamma = 5$, $\alpha = 10^{-3}$) a bunch length of $50\,\frac{c}{\omega_p}$ is sufficiently long to achieve the regime where the overall instability growth rate approximates the growth rate of an infinitely long beam. We only qualitatively investigate how the evolution of the momentum distribution is affected by changing the charge ratio χ from a neutral to a single charge beam.
Figure 9.17 shows that for the evolution of $\Delta\text{Mean}(p_\parallel)$ increasing the beam charge leads to the early time oscillation we have already discussed. However the late time evolution is entirely unaffected and both the neutral and charged beams saturate at the same level. For the evolution of the longitudinal width $\text{Std}(p_\parallel)$ (figure 9.18) we now observe an early time oscillation that increases with larger beam charge as a new effect that only occurs when the beam has finite length. Nevertheless the late time evolution remains unaffected. Lastly the evolution of the transversal momentum spread $\text{Std}(p_\parallel)$, shown in figure 9.19, remains completely unaffected by the beam charge even for a finite length beam.

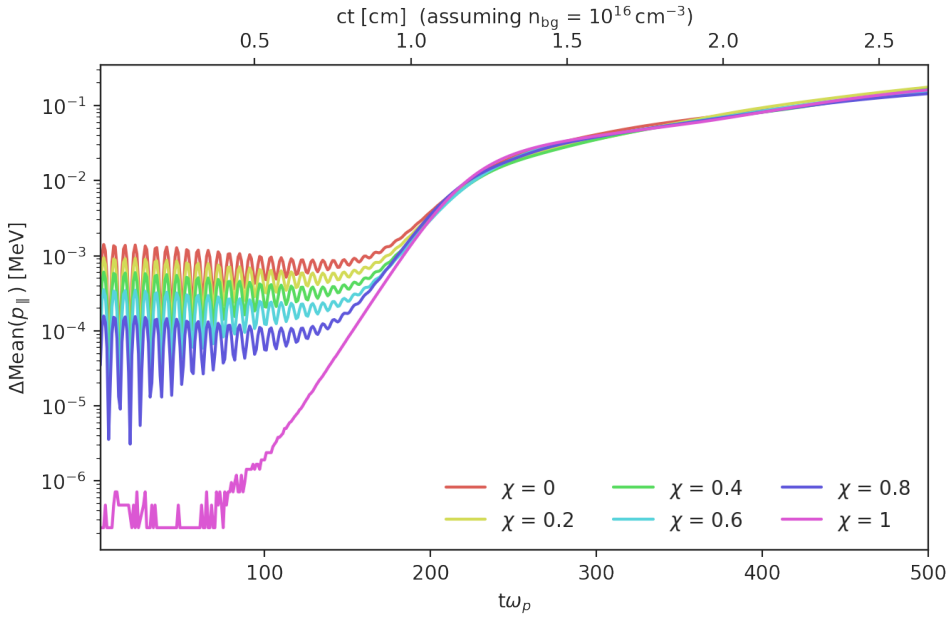

Figure 9.17: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\,(p_\parallel)$ is shown for runs with varying charge ratio χ, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

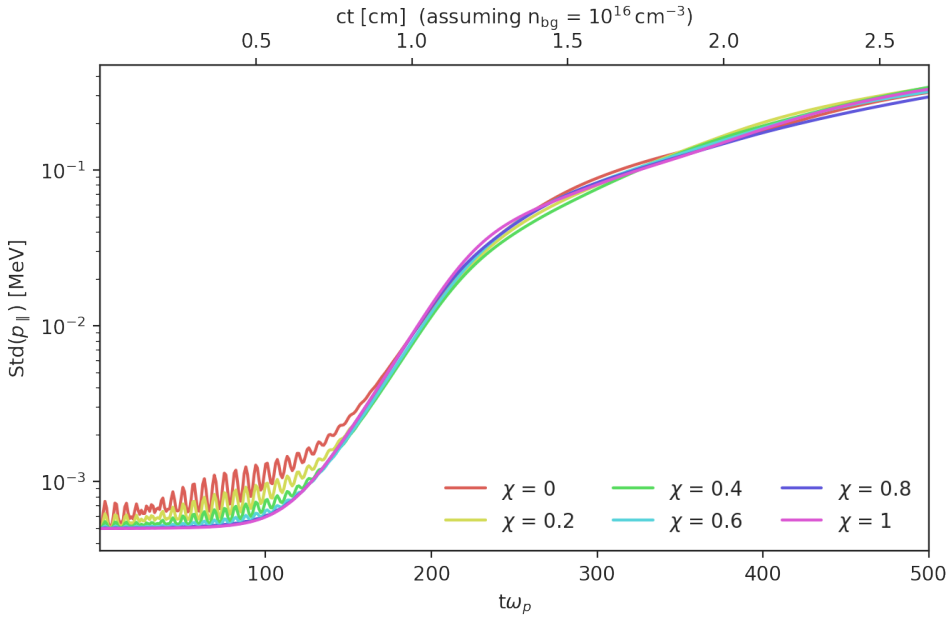

Figure 9.18: The evolution of the width of the longitudinal momentum distribution $\mathrm{Std}(p_\parallel)$ is shown for runs with varying charge ratio χ, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

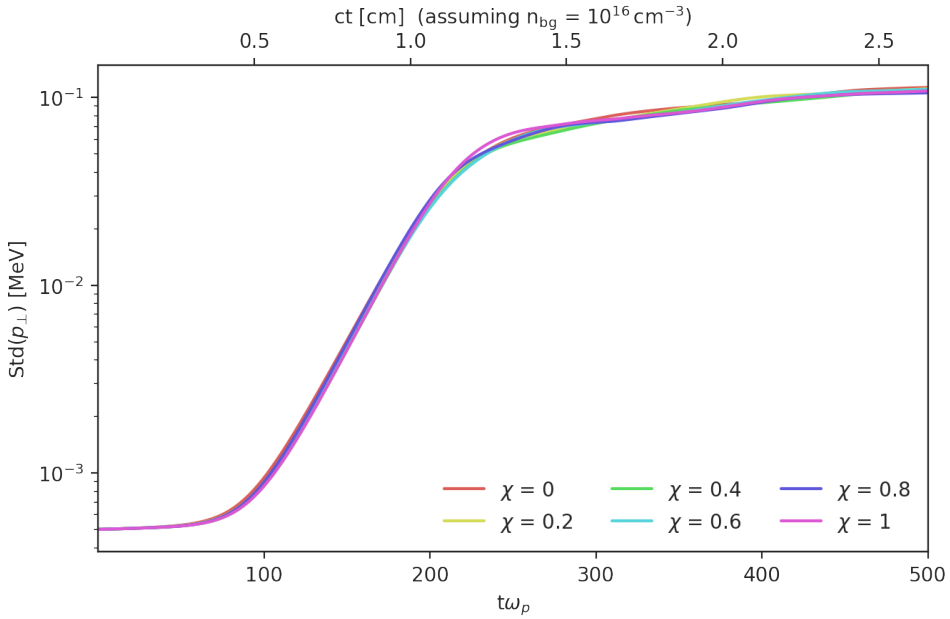

Figure 9.19: The evolution of the width of the transversal momentum distribution $\mathrm{Std}(p_\perp)$ is shown for runs with varying charge ratio χ, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

9.2.2 Finite Width

A bunch width of $2\,\frac{c}{\omega_p}$ is sufficient to achieve an overall instability growth rate that is approximately the same as for an infinitely wide beam. The presence of a current in the beam leads to a toroidal magnetic field, in the 2D PIC simulation that takes the form of a B_z component orthogonal to the plane of simulation (figure 9.20). The B_z component takes opposing signs on the left and right from the transversal center of the beam. On both sides the beam exerts a force towards the center on the beam particles. According to Biot-Savart's law the magnetic field is proportional to the current of the beam and thus proportional to $\alpha\,(1-\chi)$. Figure 9.21 and 9.22 show the evolution of the transversal beam width. Here we can clearly observe that before the increase of the transversal momentum leads to growth of the beam width, the beam is focused due to the toroidal field. For larger α and smaller χ the focusing is stronger. As a consequence of the focusing, the density of the beam is effectively increased, leading to a larger value for the density ratio α. As we have widely discussed the value of α affects the growth rate and the saturation level of the instability. We will see that the focusing leads to a slight modification of the late time evolution of the momentum distribution. Already in the evolution of the beam width we can see that despite the focusing at late times the pure electron beam overtakes the other beams which experienced less focusing.
For the shift of the mean longitudinal momentum $\Delta\mathrm{Mean}\big(p_\parallel\big)$ (figure 9.23) we can observe this slight enhancement at late times on top of the early time oscillation. Similarly for the width of the longitudinal momentum $\mathrm{Std}\big(p_\parallel\big)$ (figure 9.24) we also observe an enhancement of the widening from the time of saturation onwards. Additionally we see a very slight oscillation at very early times that has not been present for the infinitely wide beams.
For the transversal momentum spread $\mathrm{Std}(p_\perp)$ we find a severely altered evolution (figures 9.25 and 9.25), where the focusing due to the toroidal magnetic field leads to a strong early growth that is more pronounced for a stronger magnetic field (either due to smaller χ or larger α). Nevertheless at late times the effect of the instability dominates and the difference between a neutral pair beam and a pure electron beam resolves only to a slightly increases saturation level for the later. Comparing figures 9.25 ($\alpha = 10^{-3}$) and 9.26 ($\alpha = 10^{-4}$) we find that reducing the density contrast α (and thus the current of the beam) also reduces the influence of focusing effect.

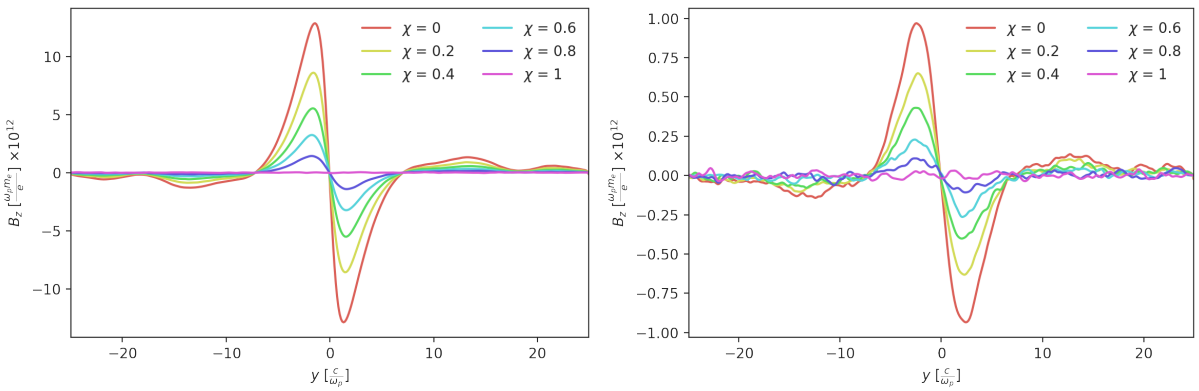

Figure 9.20: The toroidal Magnetic Field B_z (orthogonal to the plane of simulation) at time $100\omega_p^{-1}$ as a function of the beam charge ratio χ for beam density ratios $\alpha = 10^{-3}$ (**left**) and $\alpha = 10^{-4}$ (**right**).

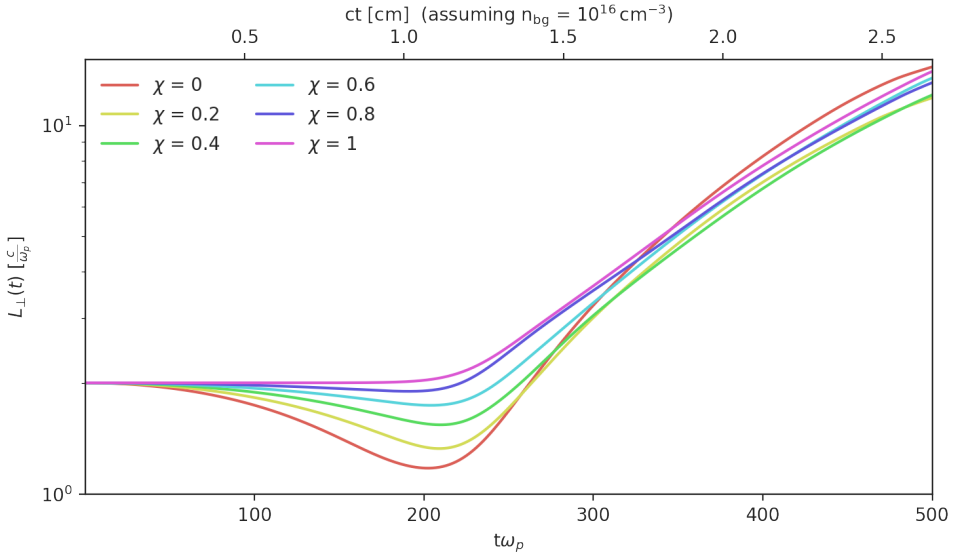

Figure 9.21: The evolution of the transversal beam width $L_\perp$ as a function of time for varying beam charge χ, Lorentz boost $\gamma = 5$ and density ratio $\alpha = 10^{-3}$.

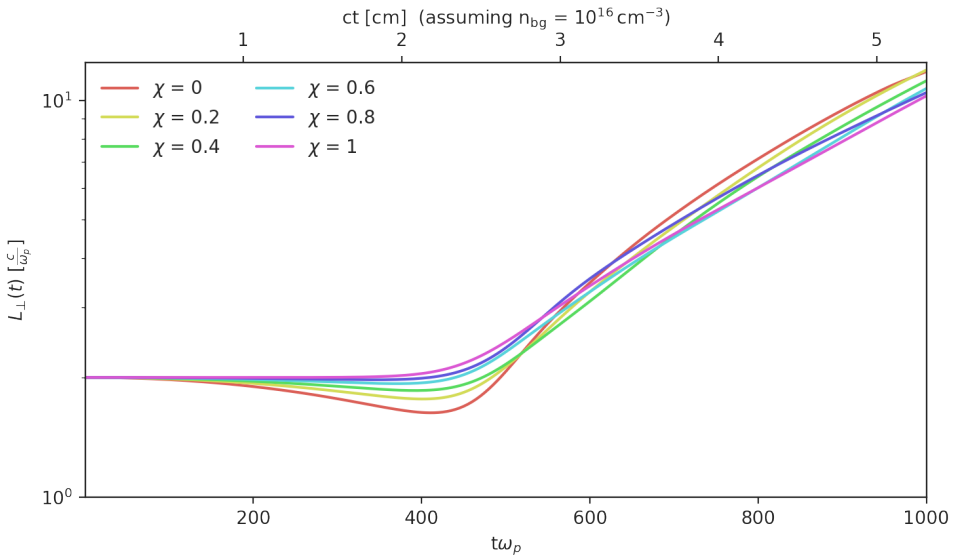

Figure 9.22: The evolution of the transversal beam width $L_\perp$ as a function of time for varying beam charge χ, Lorentz boost $\gamma = 5$ and density ratio $\alpha = 10^{-4}$.

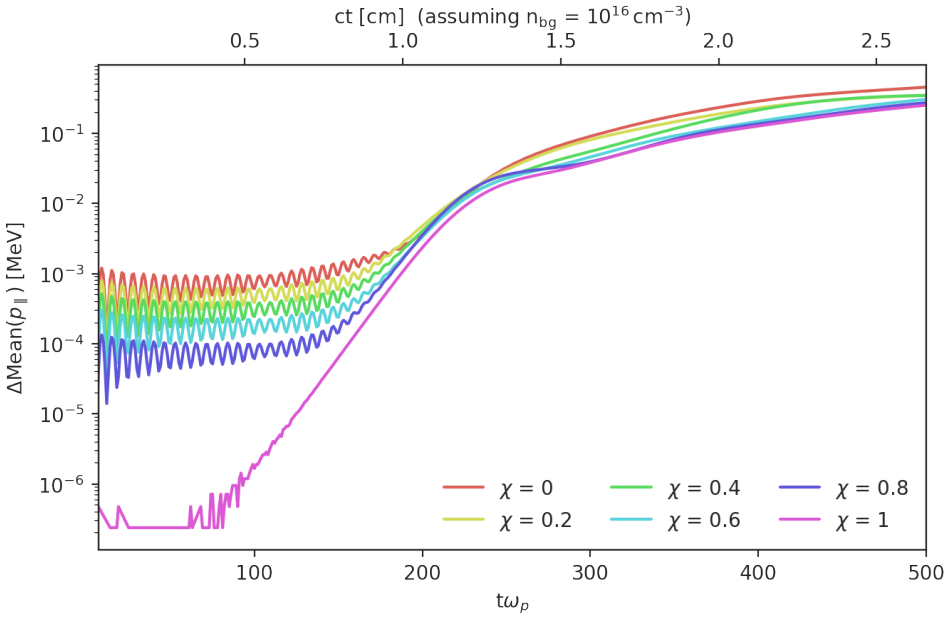

Figure 9.23: The evolution of the mean value shift of the longitudinal momentum $\Delta\text{Mean}\left(p_\parallel\right)$ is shown for runs with varying charge ratio χ, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

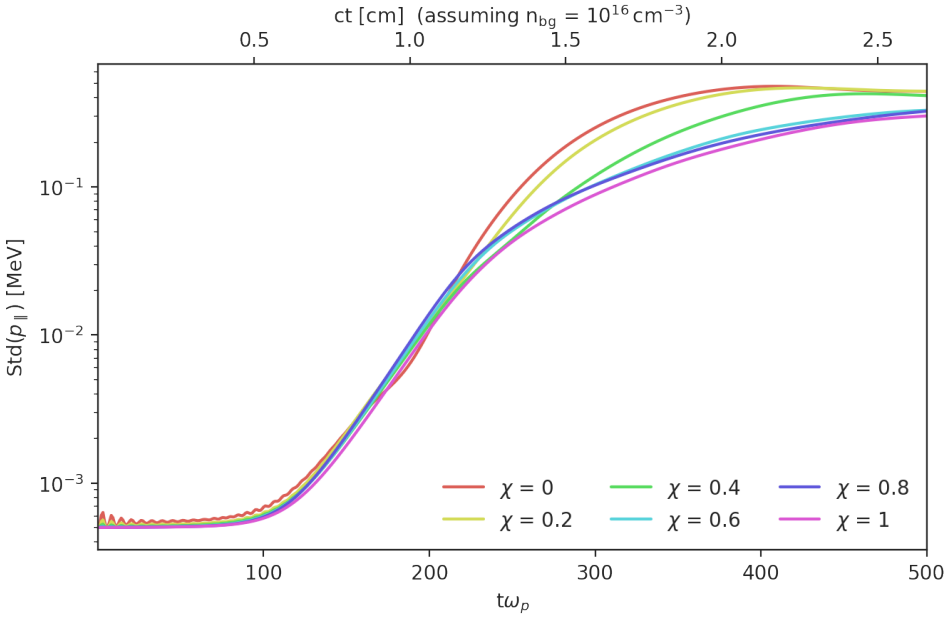

Figure 9.24: The evolution of the width of the longitudinal momentum distribution $\text{Std}\left(p_\parallel\right)$ is shown for runs with varying charge ratio χ, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

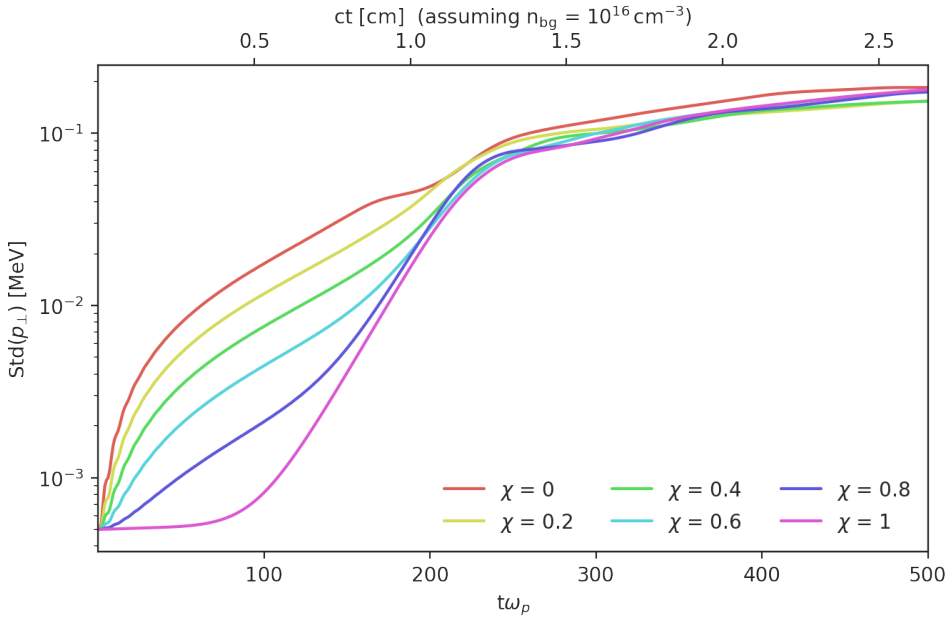

Figure 9.25: The evolution of the width of the transversal momentum distribution Std($p_\perp$) is shown for runs with varying charge ratio χ, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-3}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

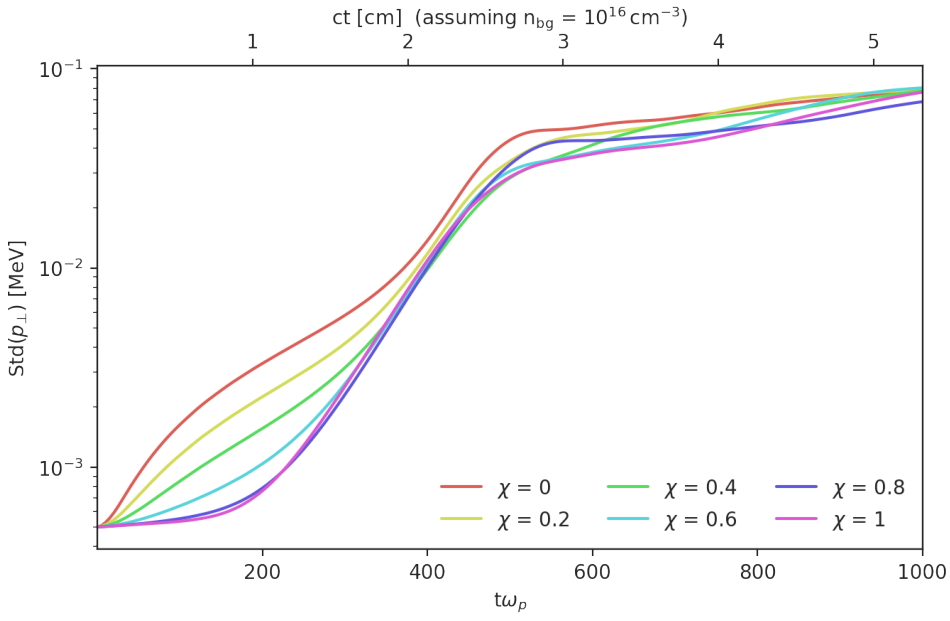

Figure 9.26: The evolution of the width of the transversal momentum distribution Std($p_\perp$) is shown for runs with varying charge ratio χ, Lorentz boost $\gamma = 5$ and density contrast $\alpha = 10^{-4}$. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$.

10 | Summary & Outlook

At the end of this thesis we will give a short summary of the results and make some suggestions on how to continue with these results in future investigations.

10.1 Summary

We started this thesis by investigating the unstable behavior of cold, neutral, infinitely extended pair beams propagating through a cold background plasma with 2D PIC simulations and found that the growth of electric fields is described by the oblique instability. The buildup of electric fields in turn leads to a two-fold back-reaction on the particles that make up the beam: The beam particles collectively lose energy and momentum that is transfered to the fields and the background plasma and by interacting with the electric fields the momentum of the beam particles can essentially diffuse, leading to a heating of the beam. A hotter beam is also more strongly diverging and in the case of the cascade contribution to the gamma ray spectrum could lead to similar phenomenology as IGMF: a halo of extended emission around the source and delayed arrival of the low energy component in the case of emission that is variable in time. We showed that the heating of the beam can be modeled by a Fokker-Planck equation with a diffusion coefficient that is proportional to the energy density of the electric field's dominating modes. The growth rates and diffusion constants can be extrapolated to astrophysical scales based on their scaling behavior at laboratory scales that neatly conforms to power laws. Furthermore we found that for a dilute beam the energy that is drained from the beam in equal parts builds up the electric fields and heats the background plasma isotropically.
Beyond the idealized conditions we investigated the influence of many corrections, namely a finite beam size, a warmer beam, a non-neutral beam, an external magnetic field and a warm or non-homogeneous background plasma and derived limits when these corrections become relevant. These corrections can have profound influence on the design of a laboratory experiment as well as the astrophysical phenomenology. For example we derived a relationship for the length of a beam and the integrated growth rate that sets a requirement for the length of a laboratory beam to produce detectable results. The same relationship also sets the time scale for which a blazar has to emit primary gamma rays for the instability to produce a significant effect on the secondary gamma ray contribution.
For a laboratory experiment that aims to investigate the unstable behavior of neutral pair beams, and hopefully corroborate the results of our PIC simulations, we identified two key challenges: The creation and handling of a neutral pair beam and the measurement of the instability. A pair beam consisting of electrons and positrons can be created from a pure electron beam by hitting a high density target where electrons and positrons are produced in a cascade process.

The resulting beam will be unfocused, with a broad energy range and not neutral at first. The further handling would require non-conventional beam optics, because they need focus both charges simultaneously. This could for example be achieved by two triplets of quadrupole magnets right after each other or a single solenoid with a very high field strength. For the measurement of the instability we identified the change of the momentum distribution as an easily accessible observable. It would be enough to measure the width of the momentum distribution along one dimension after saturation to characterize the evolution of the system. Using the scaling relationships derived one could unambiguously identify the instability process as mechanism that drives this effect.
However an exciting alternative to using a complicated neutral pair beam setup would be to use a charged electron beam to drive the instability. We have shown that the unstable behavior of such a beam is the same as the neutral beam up to some minor additional effects that do not spoil the evolution of the momentum distribution as an observable.

10.2 Outlook

Using the results from this thesis it should be possible to perform a relatively simple laboratory experiment that probes the effect of the oblique plasma instability on an electron beam. The results can be trivially transfered to the behavior of a neutral pair beam. This allows to confirm or falsify our Fokker-Planck model for the evolution of the momentum distribution. Furthermore it allows to closely investigate the saturation regime of the instability where the continuously self-replenishing beam in the astrophysical scenario differs from the scenario studied in the laboratory. The Fokker-Planck model enables us to describe the effect of the collective behavior in the form of plasma instabilities on a single particle with just the drift and diffusion coefficient extracted from the PIC simulation. Scaling these coefficients to astrophysical parameters allows us to include the effects of plasma instabilities in state of the art cosmic ray propagation codes and study the phenomenology quantitatively.

References

[1] M. Aartsen et al. Neutrino emission from the direction of the blazar txs 0506+056 prior to the icecube-170922a alert. *Science*, 361(6398):147–151, 2018.

[2] A. A. Abdo et al. Fermi Large Area Telescope Observations of the Crab Pulsar And Nebula. *The Astrophysical Journal*, 708(2):1254–1267, Jan. 2010.

[3] A. A. Abdo et al. Fermi large area telescope observations of markarian 421: The missing piece of its spectral energy distribution. *The Astrophysical Journal*, 736(2):131, jul 2011.

[4] M. Ackermann et al. The Fermi Large Area Telescope on Orbit: Event Classification, Instrument Response Functions, and Calibration. *The Astrophysical Journal Supplement*, 203(1):4, Nov. 2012.

[5] M. Ackermann et al. The spectrum of isotropic diffuse gamma-ray emission between 100 mev and 820 gev. *The Astrophysical Journal*, 799(1):86, jan 2015.

[6] M. Ackermann et al. The search for spatial extension in high-latitude sources detected by the fermi large area telescope. *The Astrophysical Journal Supplement Series*, 237:32, 08 2018.

[7] S. Agostinelli et al. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[8] F. Aharonian et al. Observations of the Crab nebula with HESS. *Astronomy & Astrophysics*, 457(3):899–915, Oct. 2006.

[9] M. Alawashra and M. Pohl. Suppression of the TeV Pair-beam-Plasma Instability by a Tangled Weak Intergalactic Magnetic Field. *The Astrophysical Journal*, 929(1):67, Apr. 2022.

[10] J. Aleksić et al. The major upgrade of the MAGIC telescopes, Part II: A performance study using observations of the Crab Nebula. *Astroparticle Physics*, 72:76–94, Jan. 2016.

[11] J. Allison et al. Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, 2006.

[12] J. Allison et al. Recent developments in geant4. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, 2016.

[13] R. Alves Batista et al. CRPropa 3—a public astrophysical simulation framework for propagating extraterrestrial ultra-high energy particles. *Journal of Cosmology and Astroparticle Physics*, 2016(5):038, May 2016.

[14] R. Alves Batista et al. CRPropa 3.2 - an advanced framework for high-energy particle propagation in extragalactic and galactic spaces. *Journal of Cosmology and Astroparticle Physics*, 2022(9):035, Sept. 2022.

[15] R. Alves Batista and A. Saveliev. The Gamma-Ray Window to Intergalactic Magnetism. *Universe*, 7(7):223, July 2021.

[16] R. Alves Batista, A. Saveliev, and E. M. de Gouveia Dal Pino. The impact of plasma instabilities on the spectra of TeV blazars. *Monthly Notices of the Royal Astronomical Society*, 489(3):3836–3849, Nov. 2019.

[17] R. Alves Batista, A. Saveliev, G. Sigl, and T. Vachaspati. Probing intergalactic magnetic fields with simulations of electromagnetic cascades. *Phys. Rev. D*, 94(8):083005, Oct. 2016.

[18] T. D. Arber, K. Bennett, C. S. Brady, A. Lawrence-Douglas, M. G. Ramsay, N. J. Sircombe, P. Gillies, R. G. Evans, H. Schmitz, A. R. Bell, and C. P. Ridgers. Contemporary particle-in-cell approach to laser-plasma modelling. *Plasma Physics and Controlled Fusion*, 57(11):1–26, Nov. 2015.

[19] C. D. Arrowsmith, N. Shukla, N. Charitonidis, R. Boni, H. Chen, T. Davenne, A. Dyson, D. H. Froula, J. T. Gudmundsson, B. T. Huffman, Y. Kadi, B. Reville, S. Richardson, S. Sarkar, J. L. Shaw, L. O. Silva, P. Simon, R. M. G. M. Trines, R. Bingham, and G. Gregori. Generating ultradense pair beams using 400 GeV /c protons. *Physical Review Research*, 3(2):023103, May 2021.

[20] R. Banerjee and K. Jedamzik. Evolution of cosmic magnetic fields: From the very early Universe, to recombination, to the present. *Physical Review D*, 70(12), Dec. 2004.

[21] A. Beck, J. Derouillat, M. Lobet, A. Farjallah, F. Massimo, I. Zemzemi, F. Perez, T. Vinci, and M. Grech. Adaptive SIMD optimizations in particle-in-cell codes with fine-grain particle sorting. *Computer Physics Communications*, 244:246–263, Nov. 2019.

[22] V. Beckmann and C. SHRADER. The AGN phenomenon: open issues. In *Proceedings of An INTEGRAL view of the high-energy sky (the first 10 years) - 9th INTEGRAL Workshop and celebration of the 10th anniversary of the launch — PoS(INTEGRAL 2012)*, volume 176, page 069, 2013.

[23] V. Berezinsky and O. Kalashev. High-energy electromagnetic cascades in extragalactic space: Physics and features. *Phys. Rev. D*, 94:023007, Jul 2016.

[24] K. Bernlöhr et al. The optical system of the h.e.s.s. imaging atmospheric cherenkov telescopes. part i: layout and components of the system. *Astroparticle Physics*, 20(2):111–128, 2003.

[25] BIPM. *Le Système international d'unités / The International System of Units ('The SI Brochure')*. Bureau international des poids et mesures, ninth edition, 2019.

[26] A. Bogaerts, E. Neyts, R. Gijbels, and J. van der Mullen. Gas discharge plasmas and their applications. *Spectrochimica Acta Part B: Atomic Spectroscopy*, 57(4):609–658, 2002.

[27] J. P. Boris. Relativistic Plasma Simulation - Optimization of a Hybrid Code. In *Fourth Conference on Numerical Simulations of Plasmas*, 1970.

[28] D. Bose, V. R. Chitnis, P. Majumdar, and A. Shukla. Galactic and extragalactic sources of very high energy gamma rays. *European Physical Journal Special Topics*, 231(1):27–66, Jan. 2022.

[29] B. Brejzman and D. Ryutov. Powerful relativistic electron beams in a plasma and in a vacuum (theory). *Nuclear Fusion*, 14(6):012, dec 1974.

[30] A. Bret. Weibel, two-stream, filamentation, oblique, bell, buneman... which one grows faster ? *Astrophysical Journal*, 699, 03 2009.

[31] A. Bret. Hierarchy of instabilities for two counter-streaming magnetized pair beams. *Physics of Plasmas*, 23(6):062122, 2016.

[32] A. Bret and M. E. Dieckmann. Hierarchy of instabilities for two counter-streaming magnetized pair beams: Influence of field obliquity. *Physics of Plasmas*, 24(6):062105, 2017.

[33] A. Bret, L. Gremillet, and D. Bénisti. Exact relativistic kinetic theory of the full unstable spectrum of an electron-beam–plasma system with maxwell-jüttner distribution functions. *Phys. Rev. E*, 81:036402, Mar 2010.

[34] A. Bret, L. Gremillet, D. Bénisti, and E. Lefebvre. Exact relativistic kinetic theory of an electron-beam–plasma system: Hierarchy of the competing modes in the system-parameter space. *Phys. Rev. Lett.*, 100:205008, May 2008.

[35] A. Bret, L. Gremillet, and M. E. Dieckmann. Multidimensional electron beam-plasma instabilities in the relativistic regime. *Physics of Plasmas*, 17(12):120501, 2010.

[36] A. E. Broderick, P. Chang, and C. Pfrommer. The Cosmological Impact of Luminous TeV Blazars. I. Implications of Plasma Instabilities for the Intergalactic Magnetic Field and Extragalactic Gamma-Ray Background. *The Astrophysical Journal*, 752(1):22, June 2012.

[37] A. E. Broderick, C. Pfrommer, E. Puchwein, and P. Chang. IMPLICATIONS OF PLASMA BEAM INSTABILITIES FOR THE STATISTICS OF THE FERMI HARD GAMMA-RAY BLAZARS AND THE ORIGIN OF THE EXTRAGALACTIC GAMMA-RAY BACKGROUND. *The Astrophysical Journal*, 790(2):137, jul 2014.

[38] A. E. Broderick, P. Tiede, M. Shalaby, C. Pfrommer, E. Puchwein, P. Chang, and A. Lamberts. BOW TIES IN THE SKY. i. THE ANGULAR STRUCTURE OF INVERSE COMPTON GAMMA-RAY HALOS IN THE FERMI SKY. *The Astrophysical Journal*, 832(2):109, nov 2016.

[39] M. Bucher. Physics of the cosmic microwave background anisotropy. *International Journal of Modern Physics D*, 24(2):1530004–303, Jan. 2015.

[40] P. Chang, A. Broderick, C. Pfrommer, E. Puchwein, A. Lamberts, M. Shalaby, and G. Vasil. The linear instability of dilute ultrarelativistic $e^{\pm}$ pair beams. *The Astrophysical Journal*, 833, 10 2016.

[41] P. Chang, A. E. Broderick, and C. Pfrommer. The Cosmological Impact of Luminous TeV Blazars. II. Rewriting the Thermal History of the Intergalactic Medium. *The Astrophysical Journal*, 752(1):23, June 2012.

[42] P. Chang, A. E. Broderick, C. Pfrommer, E. Puchwein, A. Lamberts, and M. Shalaby. The Effect of Nonlinear Landau Damping on Ultrarelativistic Beam Plasma Instabilities. *The Astrophysical Journal*, 797(2):110, Dec. 2014.

[43] W. Chen, M. Errando, and J. Buckley. Novel search for TeV-initiated pair cascades in the intergalactic medium. In *42nd COSPAR Scientific Assembly*, volume 42, pages E1.14–26–18, July 2018.

[44] C. Chiuderi and M. Velli. *Basics of Plasma Astrophysics*. UNITEXT for Physics. Springer Milan, 2016.

[45] A. Cooray. Extragalactic background light measurements and applications. *Royal Society Open Science*, 3(3):150555, Mar. 2016.

[46] J. Cortina. Status and First Results of the Magic Telescope. *Astrophysics and Space Science*, 297(1-4):245–255, June 2005.

[47] J. Couperus Cabadağ, R. Pausch, A. Köhler, O. Zarini, J. Krämer, M. Garten, A. Huebl, R. Gebhardt, U. Helbig, S. Bock, K. Zeil, A. Debus, M. Bussmann, U. Schramm, and A. Irman. Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator. *Nature Communications*, 8:487, 09 2017.

[48] A. De Angelis, G. Galanti, and M. Roncadelli. Transparency of the Universe to gamma-rays. *Monthly Notices of the Royal Astronomical Society*, 432(4):3245–3249, 05 2013.

[49] A. De Angelis and M. Mallamaci. Gamma-ray astrophysics. *European Physical Journal Plus*, 133(8):324, Aug. 2018.

[50] R. de Maria. General method for final focus system design for circular colliders. *Phys. Rev. ST Accel. Beams*, 11:031001, Mar 2008.

[51] B. Degrange and G. Fontaine. Introduction to high-energy gamma-ray astronomy. *Comptes Rendus Physique*, 16(6-7):587–599, Aug. 2015.

[52] C. D. Dermer, M. Cavadini, S. Razzaque, J. D. Finke, J. Chiang, and B. Lott. Time Delay of Cascade Radiation for TeV Blazars and the Measurement of the Intergalactic Magnetic Field. *The Astrophysical Journal Letters*, 733(2):L21, June 2011.

[53] G. Dobler, D. P. Finkbeiner, I. Cholis, T. Slatyer, and N. Weiner. The fermi haze: A gamma-ray counterpart to the microwave haze. *The Astrophysical Journal*, 717(2):825, jun 2010.

[54] R. Durrer. The cosmic microwave background: the history of its experimental investigation and its significance for cosmology. *Classical and Quantum Gravity*, 32(12):124007, June 2015.

[55] R. Durrer and A. Neronov. Cosmological magnetic fields: Their generation, evolution and observation. *The Astronomy and Astrophysics Review*, 21, 03 2013.

[56] T. A. Dzhatdoev, E. I. Podlesnyi, and I. A. Vaiman. Can we constrain the extragalactic magnetic field from very high energy observations of GRB 190114C? *Phys. Rev. D*, 102(12):123017, Dec. 2020.

[57] I. Efthymiopoulos, C. Hessler, H. Gaillard, D. Grenier, M. Meddahi, P. Trilhe, A. Pardons, C. Theis, N. Charitonidis, S. Evrard, H. Vincke, and M. J. Lazzaroni. Hiradmat: A new irradiation facility for material testing at cern. 2011.

[58] A. Elyiv, A. Neronov, and D. V. Semikoz. Gamma-ray induced cascades and magnetic fields in the intergalactic medium. *Phys. Rev. D*, 80(2):023010, July 2009.

[59] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of laser-driven plasma-based electron accelerators. *Rev. Mod. Phys.*, 81:1229–1285, Aug 2009.

[60] T. Esirkepov. Exact charge conservation scheme for particle-in-cell simulation with an arbitrary form-factor. *Computer Physics Communications*, 135(2):144–153, 2001.

[61] J. D. Finke, S. Razzaque, and C. D. Dermer. Modeling the Extragalactic Background Light from Stars and Dust. *The Astrophysical Journal*, 712(1):238–249, Mar. 2010.

[62] K. Floettmann. Some basic features of the beam emittance. *Phys. Rev. ST Accel. Beams*, 6:034202, Mar 2003.

[63] A. Franceschini, G. Rodighiero, and M. Vaccari. Extragalactic optical-infrared background radiation, its time evolution and the cosmic photon-photon opacity. *Astronomy and Astrophysics*, 487(3):837–852, Sept. 2008.

[64] S. Funk. Ground- and space-based gamma-ray astronomy. *Annual Review of Nuclear and Particle Science*, 65(1):245–277, 2015.

[65] R. C. Gilmore, R. S. Somerville, J. R. Primack, and A. Domínguez. Semi-analytic modelling of the extragalactic background light and consequences for extragalactic gamma-ray spectra. *Monthly Notices of the Royal Astronomical Society*, 422(4):3189–3207, June 2012.

[66] E. Gschwendtner et al. Awake, the advanced proton driven plasma wakefield acceleration experiment at cern. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 829:76–82, 2016. 2nd European Advanced Accelerator Concepts Workshop - EAAC 2015.

[67] O. Gueta. The Cherenkov Telescope Array: layout, design and performance. In *37th International Cosmic Ray Conference*, page 885, Mar. 2022.

[68] A. V. Higuera and J. R. Cary. Structure-preserving second-order integration of relativistic charged particle trajectories in electromagnetic fields. *Physics of Plasmas*, 24(5):052104, 2017.

[69] I. Hofmann. Performance of solenoids versus quadrupoles in focusing and energy selection of laser accelerated protons. *Phys. Rev. ST Accel. Beams*, 16:041302, Apr 2013.

[70] J. Holder et al. The first VERITAS telescope. *Astroparticle Physics*, 25(6):391–401, July 2006.

[71] Y. Inoue and T. Totani. The Blazar Sequence and the Cosmic Gamma-ray Background Radiation in the Fermi Era. *The Astrophysical Journal*, 702(1):523–536, Sept. 2009.

[72] M. Kamionkowski and A. Kosowsky. The cosmic microwave background and particle physics. *Annual Review of Nuclear and Particle Science*, 49(1):77–123, 1999.

[73] Kneiske, T. M. and Mannheim, K. Bl lacertae contribution to the extragalactic gamma-ray background. *Astronomy & Astrophysics*, 479(1):41–47, 2008.

[74] M. Kono and M. Skoric. *Nonlinear Physics of Plasmas*. Springer Series on Atomic, Optical, and Plasma Physics. Springer Berlin Heidelberg, 2010.

[75] N. L. Kugland, D. D. Ryutov, C. Plechaty, J. S. Ross, and H. S. Park. Invited Article: Relation between electric and magnetic field structures and their proton-beam images. *Review of Scientific Instruments*, 83(10):101301–101301–26, Oct. 2012.

[76] C. K. Li, F. H. Séguin, J. A. Frenje, J. R. Rygg, R. D. Petrasso, R. P. J. Town, P. A. Amendt, S. P. Hatchett, O. L. Landen, A. J. Mackinnon, P. K. Patel, V. A. Smalyuk, T. C. Sangster, and J. P. Knauer. Measuring e and b fields in laser-produced plasmas with monoenergetic proton radiography. *Phys. Rev. Lett.*, 97:135003, Sep 2006.

[77] MAGIC Collaboration. Very-High-Energy gamma rays from a Distant Quasar: How Transparent Is the Universe? *Science*, 320(5884):1752, June 2008.

[78] MAGIC Collaboration, A. Neronov, D. Semikoz, and A. Korochkin. A lower bound on intergalactic magnetic fields from time variability of 1ES 0229+200 from MAGIC and Fermi/LAT observations. *arXiv e-prints*, page arXiv:2210.03321, Oct. 2022.

[79] S. P. D. Mangles et al. Laser-wakefield acceleration of monoenergetic electron beams in the first plasma-wave period. *Phys. Rev. Lett.*, 96:215001, May 2006.

[80] F. Miniati and A. Elyiv. Relaxation of Blazar-induced Pair Beams in Cosmic Voids. *The Astrophysical Journal*, 770(1):54, June 2013.

[81] R. Mirzoyan. Technological novelties of ground-based very high energy gamma-ray astrophysics with the imaging atmospheric cherenkov telescopes. *Universe*, 8(4), 2022.

[82] A. Neronov and D. V. Semikoz. A method of measurement of extragalactic magnetic fields by TeV gamma ray telescopes. *Soviet Journal of Experimental and Theoretical Physics Letters*, 85(10):473–477, July 2007.

[83] A. Neronov and D. V. Semikoz. Sensitivity of γ-ray telescopes for detection of magnetic fields in the intergalactic medium. *Phys. Rev. D*, 80(12):123012, Dec. 2009.

[84] A. Neronov and D. V. Semikoz. Sensitivity of γ-ray telescopes for detection of magnetic fields in the intergalactic medium. *Physical Review D*, 80(12):123012, Dec. 2009.

[85] A. Neronov and I. Vovk. Evidence for Strong Extragalactic Magnetic Fields from Fermi Observations of TeV Blazars. *Science*, 328(5974):73, Apr. 2010.

[86] S. Nguyen-Kuok. *Theory of Low-Temperature Plasma Physics*. Springer Series on Atomic, Optical, and Plasma Physics. Springer International Publishing, 2016.

[87] R. C. others. The optical system of the h.e.s.s. imaging atmospheric cherenkov telescopes. part ii: mirror alignment and point spread function. *Astroparticle Physics*, 20(2):129–143, 2003.

[88] A. Peratt. *Physics of the Plasma Universe*. Springer-Verlag, 1992.

[89] R. Perry and Y. Lyubarsky. The role of resonant plasma instabilities in the evolution of blazar-induced pair beams. *Monthly Notices of the Royal Astronomical Society*, 503(2):2215–2228, May 2021.

[90] A. Piel. *Plasma Physics: An Introduction to Laboratory, Space, and Fusion Plasmas*. Springer Berlin Heidelberg, 2010.

[91] C. V. Pieronek, A. J. Gonsalves, C. Benedetti, S. S. Bulanov, J. van Tilborg, J. H. Bin, K. K. Swanson, J. Daniels, G. A. Bagdasarov, N. A. Bobrova, V. A. Gasilov, G. Korn, P. V. Sasorov, C. G. R. Geddes, C. B. Schroeder, W. P. Leemans, and E. Esarey. Laser-heated capillary discharge waveguides as tunable structures for laser-plasma acceleration. *Physics of Plasmas*, 27(9):093101, 2020.

[92] Planck Collaboration et al. Planck 2015 results - i. overview of products and scientific results. *A&A*, 594:A1, 2016.

[93] A. Pukhov. Particle-in-Cell Codes for plasma-based particle acceleration. *arXiv e-prints*, 2015.

[94] I. Rafighi. *Plasma Instabilities from blazar-induced pair beams propagating through IGM*. doctoralthesis, Universität Potsdam, 2018.

[95] I. Rafighi, S. Vafin, M. Pohl, and J. Niemiec. Plasma effects on relativistic pair beams from TeV blazars: PIC simulations and analytical predictions. *Astronomy and Astrophysics*, 607:A112, Nov. 2017.

[96] D. Ryu, D. R. G. Schleicher, R. A. Treumann, C. G. Tsagas, and L. M. Widrow. Magnetic Fields in the Large-Scale Structure of the Universe. *Space Science Reviews*, 166(1-4):1–35, May 2012.

[97] P. San Miguel Claveria, X. Davoine, J. R. Peterson, M. Gilljohann, I. Andriyash, R. Ariniello, C. Clarke, H. Ekerfelt, C. Emma, J. Faure, S. Gessner, M. J. Hogan, C. Joshi, C. H. Keitel, A. Knetsch, O. Kononenko, M. Litos, Y. Mankovska, K. Marsh, A. Matheron, Z. Nie, B. O'Shea, D. Storey, N. Vafaei-Najafabadi, Y. Wu, X. Xu, J. Yan, C. Zhang, M. Tamburini, F. Fiuza, L. Gremillet, and S. Corde. Spatiotemporal dynamics of ultrarelativistic beam-plasma instabilities. *Phys. Rev. Research*, 4:023085, May 2022.

[98] G. Sarri, M. Dieckmann, I. Kourakis, A. Piazza, B. Reville, C. Keitel, and M. Zepf. Overview of laser-driven generation of electron–positron beams. *Journal of Plasma Physics*, 81:1–23, 05 2015.

[99] G. Sarri, K. Poder, J. Cole, W. Schumaker, A. Piazza, B. Reville, T. Dzelzainis, D. Doria, L. Gizzi, G. M. Grittani, S. Kar, C. Keitel, K. Krushelnick, S. Kuschel, S. Mangles, Z. Najmudin, N. Shukla, L. Silva, D. Symes, and M. Zepf. Generation of neutral and high-density electron–positron pair plasmas in the laboratory. *Nature communications*, 6:6747, 04 2015.

[100] G. Sarri, W. Schumaker, A. Di Piazza, M. Vargas, B. Dromey, M. E. Dieckmann, V. Chvykov, A. Maksimchuk, V. Yanovsky, Z. H. He, B. X. Hou, J. A. Nees, A. G. R. Thomas, C. H. Keitel, M. Zepf, and K. Krushelnick. Table-top laser-based source of femtosecond, collimated, ultrarelativistic positron beams. *Phys. Rev. Lett.*, 110:255002, Jun 2013.

[101] G. Sarri, J. Warwick, W. Schumaker, K. Poder, J. Cole, D. Doria, T. Dzelzainis, K. Krushelnick, S. Kuschel, S. Mangles, Z. Najmudin, L. Romagnani, G. Marrero Samarin, D. Symes, A. Thomas, M. Yeung, and M. Zepf. Spectral and spatial characterisation of laser-driven positron beams. *Plasma Physics and Controlled Fusion*, 59, 01 2017.

[102] A. Saveliev, C. Evoli, and G. Sigl. The Role of Plasma Instabilities in the Propagation of Gamma-Rays from Distant Blazars. *arXiv e-prints*, page arXiv:1311.6752, Nov. 2013.

[103] R. Schlickeiser, D. Ibscher, and M. Supsar. Plasma Effects on Fast Pair Beams in Cosmic Voids. *The Astrophysical Journal*, 758(2):102, Oct. 2012.

[104] R. Schlickeiser, S. Krakau, and M. Supsar. Plasma Effects on Fast Pair Beams. II. Reactive versus Kinetic Instability of Parallel Electrostatic Waves. *The Astrophysical Journal*, 777(1):49, Nov. 2013.

[105] C. B. Schroeder, C. Benedetti, E. Esarey, F. J. Grüner, and W. P. Leemans. Particle beam self-modulation instability in tapered and inhomogeneous plasma. *Physics of Plasmas*, 19(1):010703, 2012.

[106] M. Shalaby, A. E. Broderick, P. Chang, C. Pfrommer, A. Lamberts, and E. Puchwein. Growth of Beam-Plasma Instabilities in the Presence of Background Inhomogeneity. *The Astrophysical Journal*, 859(1):45, May 2018.

[107] M. Shalaby, A. E. Broderick, P. Chang, C. Pfrommer, E. Puchwein, and A. Lamberts. The growth of the longitudinal beam–plasma instability in the presence of an inhomogeneous background. *Journal of Plasma Physics*, 86(2):535860201, 2020.

[108] G. Sigl. *Astroparticle Physics: Theory and Phenomenology*. Atlantis Studies in Astroparticle Physics and Cosmology. Atlantis Press, 2016.

[109] L. Sironi and D. Giannios. Relativistic Pair Beams from TeV Blazars: A Source of Reprocessed GeV Emission rather than Intergalactic Heating. *The Astrophysical Journal*, 787(1):49, May 2014.

[110] J. Sitarek. TeV Instrumentation: Current and Future. *Galaxies*, 10(1):21, Jan. 2022.

[111] K. N. Sjobak, E. Adli, R. Corsini, W. Farabolini, G. Boyle, C. A. Lindstrøm, M. Meisel, J. Osterhoff, J.-H. Röckemann, L. Schaper, and A. E. Dyson. Strong focusing gradient in a linear active plasma lens. *Phys. Rev. Accel. Beams*, 24:121306, Dec 2021.

[112] R. Smets, N. Aunai, A. Ciardi, M. Drouin, M. Campos-Pinto, and P. Deegan. A new method to dispatch split particles in Particle-In-Cell codes. *Computer Physics Communications*, 261:107666, Apr. 2021.

[113] A. J. Smith and HAWC Collaboration. HAWC: Design, Operation, Reconstruction and Analysis. In *34th International Cosmic Ray Conference (ICRC2015)*, volume 34 of *International Cosmic Ray Conference*, page 966, July 2015.

[114] D. J. Spence and S. M. Hooker. Investigation of a hydrogen plasma waveguide. *Phys. Rev. E*, 63:015401, Dec 2000.

[115] M. Su and D. P. Finkbeiner. Evidence for gamma-ray jets in the milky way. *The Astrophysical Journal*, 753(1):61, jun 2012.

[116] M. Su, T. R. Slatyer, and D. P. Finkbeiner. Giant gamma-ray bubbles from fermi-lat: Active galactic nucleus activity or bipolar galactic wind? *The Astrophysical Journal*, 724(2):1044, nov 2010.

[117] M. Supsar and R. Schlickeiser. Plasma Effects on Fast Pair Beams. III. Oblique Electrostatic Growth Rates for Perpendicular Maxwellian Pair Beams. *The Astrophysical Journal*, 783(2):96, Mar. 2014.

[118] A. M. Taylor, I. Vovk, and A. Neronov. Extragalactic magnetic fields constraints from simultaneous GeV-TeV observations of blazars. *Astronomy and Astrophysics*, 529:A144, May 2011.

[119] P. Tiede, A. E. Broderick, M. Shalaby, C. Pfrommer, E. Puchwein, P. Chang, and A. Lamberts. Bow ties in the sky. II. searching for gamma-ray halos in the fermi sky using anisotropy. *The Astrophysical Journal*, 850(2):157, nov 2017.

[120] S. Vafin, P. J. Deka, M. Pohl, and A. Bohdan. Revisit of Nonlinear Landau Damping for Electrostatic Instability Driven by Blazar-induced Pair Beams. *The Astrophysical Journal*, 873(1):10, Mar. 2019.

[121] S. Vafin, I. Rafighi, M. Pohl, and J. Niemiec. The Electrostatic Instability for Realistic Pair Distributions in Blazar/EBL Cascades. *The Astrophysical Journal*, 857(1):43, Apr. 2018.

[122] J. van Tilborg, S. Steinke, C. G. R. Geddes, N. H. Matlis, B. H. Shaw, A. J. Gonsalves, J. V. Huijts, K. Nakamura, J. Daniels, C. B. Schroeder, C. Benedetti, E. Esarey, S. S. Bulanov, N. A. Bobrova, P. V. Sasorov, and W. P. Leemans. Active plasma lensing for relativistic laser-plasma-accelerated electron beams. *Phys. Rev. Lett.*, 115:184802, Oct 2015.

[123] J.-L. Vay and B. B. Godfrey. Modeling of relativistic plasmas with the particle-in-cell method. *Comptes Rendus Mécanique*, 342(10):610–618, 2014. Theoretical and numerical approaches for Vlasov-maxwell equations.

[124] T. M. Venters. Contribution to the extragalactic gamma-ray background from the cascades of very high energy gamma rays from blazars. *The Astrophysical Journal*, 710(2):1530, 2010.

[125] J. Villasenor and O. Buneman. Rigorous charge conservation for local electromagnetic field solvers. *Computer Physics Communications*, 69(2):306–316, 1992.

[126] I. Vovk, A. Taylor, D. Semikoz, and a. Neronov. Fermi/lat observations of 1es 0229+200: Implications for extragalactic magnetic fields and background light. *The Astrophysical Journal Letters*, 747:L14, 02 2012.

[127] J. Warwick, A. Alejo, T. Dzelzainis, W. Schumaker, D. Doria, L. Romagnani, K. Poder, J. Cole, M. Yeung, K. Krushelnick, S. Mangles, Z. Najmudin, G. Samarin, D. Symes, A. Thomas, M. Borghesi, and G. Sarri. General features of experiments on the dynamics of laser-driven electron–positron beams. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 909:95–101, 2018. 3rd European Advanced Accelerator Concepts workshop (EAAC2017).

[128] J. Warwick, T. Dzelzainis, M. E. Dieckmann, W. Schumaker, D. Doria, L. Romagnani, K. Poder, J. M. Cole, A. Alejo, M. Yeung, K. Krushelnick, S. P. D. Mangles, Z. Najmudin, B. Reville, G. M. Samarin, D. D. Symes, A. G. R. Thomas, M. Borghesi, and G. Sarri. Experimental observation of a current-driven instability in a neutral electron-positron beam. *Phys. Rev. Lett.*, 119:185002, Nov 2017.

[129] T. Weekes et al. Veritas: the very energetic radiation imaging telescope array system. *Astroparticle Physics*, 17(2):221–243, 2002.

[130] R. L. Workman et al. Review of Particle Physics. *PTEP*, 2022:083C01, 2022.

[131] D. Yan, J. Zhou, P. Zhang, Q. Zhu, and J. Wang. Impact of Plasma Instability on Constraint of the Intergalactic Magnetic Field. *The Astrophysical Journal*, 870(1):17, Jan. 2019.

[132] T. Yang, H. Cheng, Y. Yan, M. Wu, D. Li, Y. Li, Y. Xia, C. Lin, and X. Yan. Designing of active plasma lens for focusing laser-plasma-accelerated pulsed proton beams. *Phys. Rev. Accel. Beams*, 24:031301, Mar 2021.

[133] K. Yee. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, 1966.

A | Numerical Convergence

The Particle-in-Cell method discretizes the Vlasov equation (3.10) by discretizing the fields on a grid and sampling the distribution function of the particles with numerical particles. In the limit of infinitely small grid spacing and infinitely many macroparticles the PIC method solves the Vlasov equation exactly. For any realistic simulation using finite grid resolution and finite number of macroparticles great care has to be taken that the simulation converges, that means that difference between the simulation result and the exact solution is small. Since the exact solution is only known in a small number of simple cases we can not compare with the exact solution directly. We estimate the difference between a simulation of a given resolution and the exact solution by comparing the simulation result with a simulation of higher resolution. We increase the resolution until a regime is reached where the observables remain constant within a range of error that we deem acceptable.
Here we explicitly study the convergence behavior for the beam instability of a neutral pair beam with a given density ratio α and Lorentz booost γ with respect to the number of macroparticles used to resolve the particle distribution functions, the resolution of instable field modes and the resolution of the plasma wavelength. Furthermore we also study the conservation of energy in the PIC simulation.

A.1 Number of Macroparticles

The number of macroparticles used to sample the particle distribution function directly influences the amount of noise in the PIC simulation. Scanning the number of macroparticles per simulation cell N_p/N_{cell} from 1 to 50 shows that the qualitative behavior remains the same but all observables are less affected by noise in the simulations with a higher number of macroparticles. Quantitatively the growth rate of the instability is slightly reduced for a low number of macroparticles. For more than 10 macroparticles per cell the growth rate however stabilizes. We have to note that some observables like the diffusion parameters D_0 and D_1, which are **not** physical observables, are strongly affected by the number of macroparticles and still scale according to a power-law even in a regime that has converged. The physical behavior is not affected by this scaling.

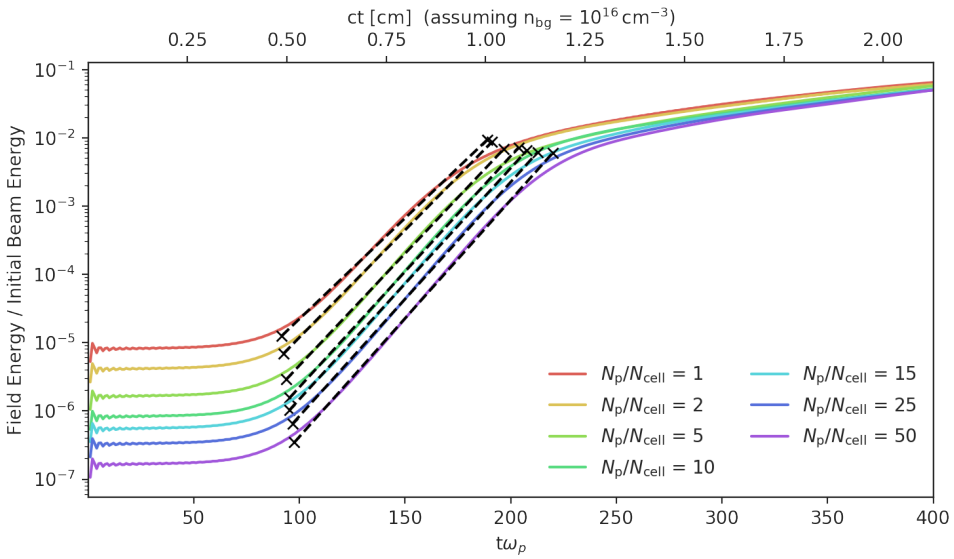

Figure A.1: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a density contrast of $\alpha = 10^{-3}$, Lorentz boost $\gamma = 5$ and varying number of macroparticles per grid cell N_p/N_{cell}. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

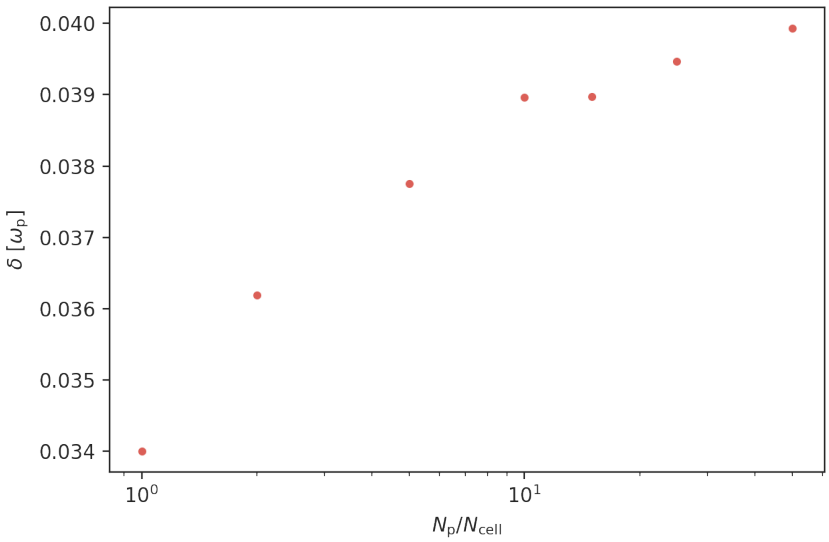

Figure A.2: The growth rate of the total field energy density as a function of the number of macroparticles per grid cell N_p/N_{cell} extracted from figure A.1. For more than 10 particle per grid cell the growth rate stabilizes.

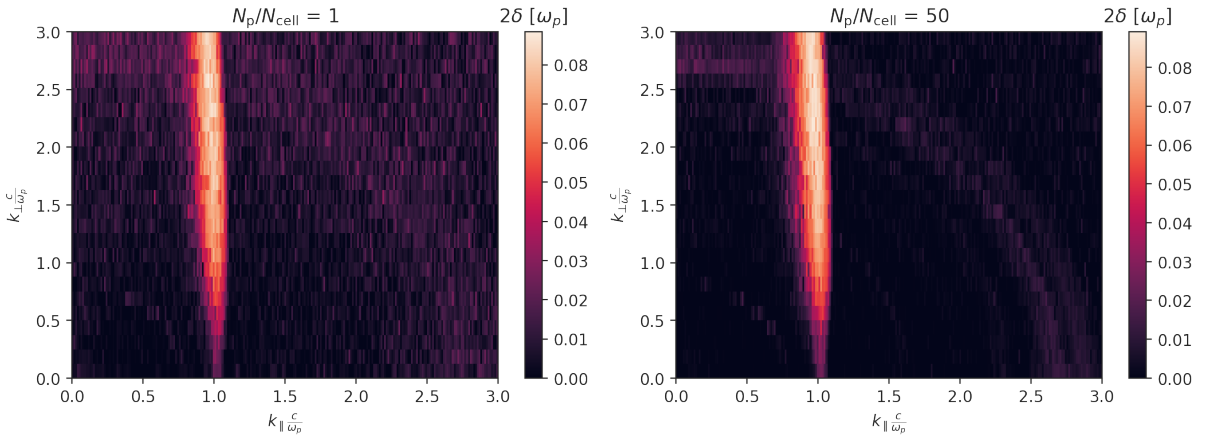

Figure A.3: Maps of the growth rate as a function of wavevector for each mode for two simulation runs with a different number of macroparticles. On the **left** the number of macroparticles per grid cell is small leading to a large amount noise. On the **right** the number of macroparticles is 50 times larger leading to a less noisy growth rate map.

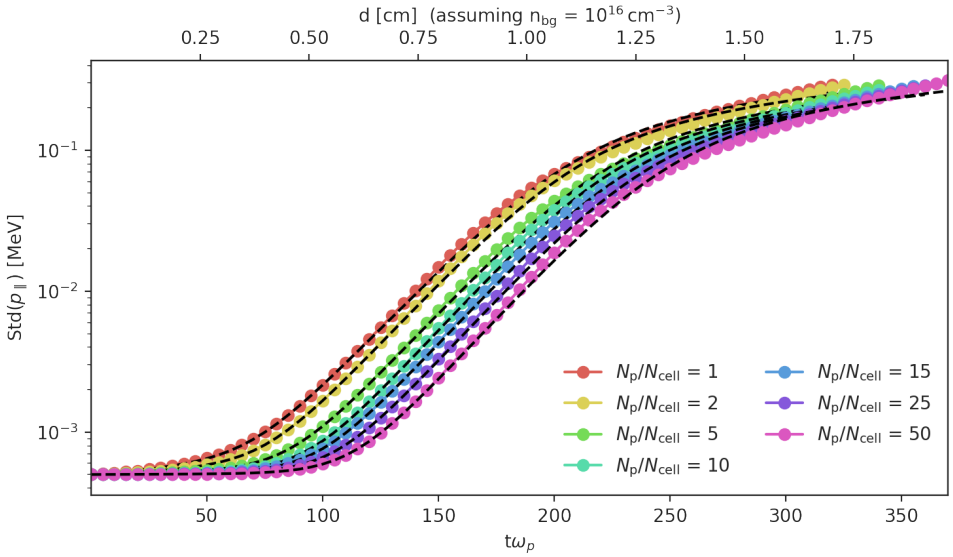

Figure A.4: The evolution of the width of the longitudinal momentum Std $(p_\parallel)$ is shown for runs with varying number of macroparticles per grid cell N_p/N_{cell}. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$, however for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

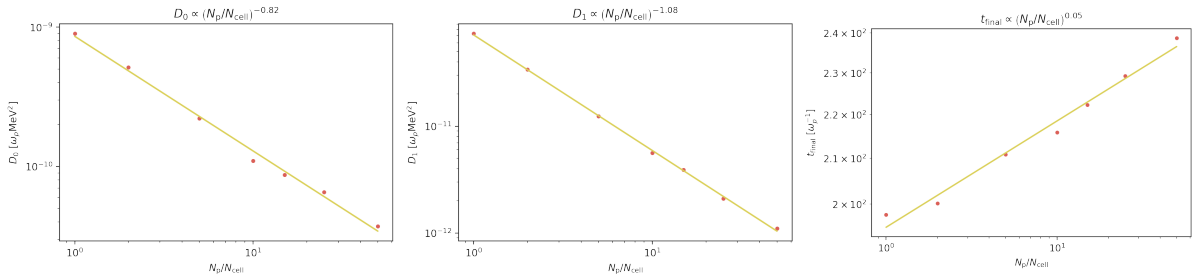

Figure A.5: The fitted values for the diffusion parameters D_0 (**left**), D_1 (**middle**) and the saturation time t_{final} (**right**) for the evolution of $\text{Std}(p_\parallel)$ extracted from figure A.4 as a function of the number of macroparticles N_p/N_{cell}. The diffusion parameters scale almost linearly with the number of macroparticles, whereas the saturation time only depends mildly on the number of cells.

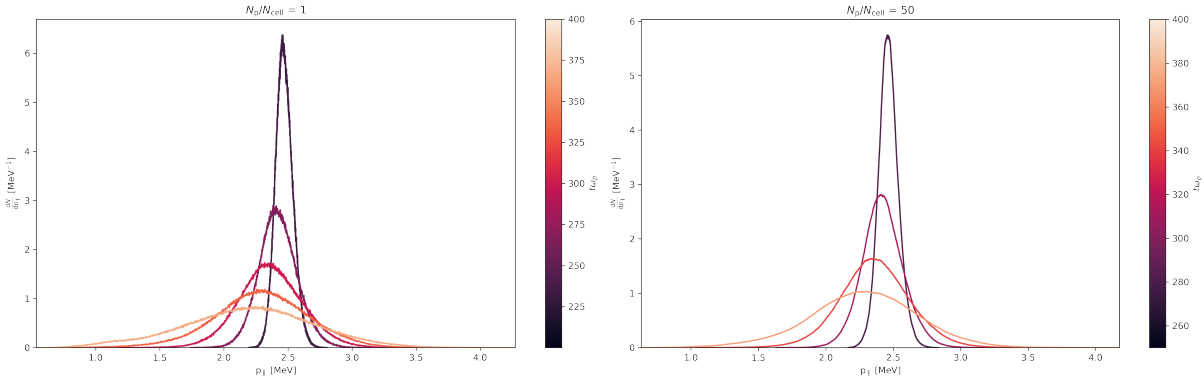

Figure A.6: The longitudinal beam momentum distribution for simulation runs with $N_p/N_{\text{cell}} = 1$ (**left**) and $N_p/N_{\text{cell}} = 50$ (**right**) at a number of time steps near the end of instability growth. In the simulation run with less macroparticles the momentum distribution shows stronger fluctuations.

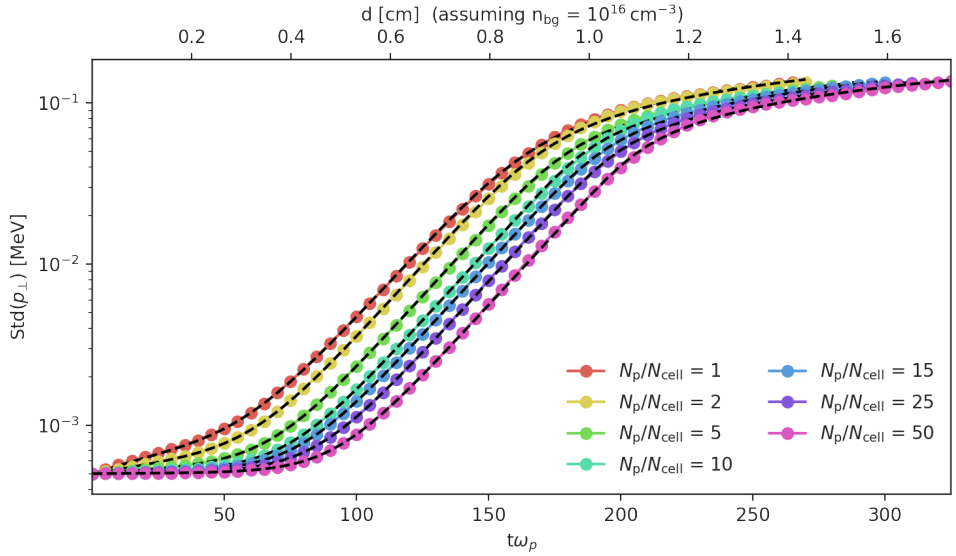

Figure A.7: The evolution of the width of the transversal momentum Std $(p_\perp)$ is shown for runs with varying number of macroparticles per grid cell N_p/N_{cell}. The standard deviation is calculated at regular intervals of $5\omega_p^{-1}$, however for late times, after $2t_{\text{final}}$ (fitted value), data points are not shown since the fit to the model does not describe the evolution very well anymore. In dashed lines a fit of (3.43) to the data is shown.

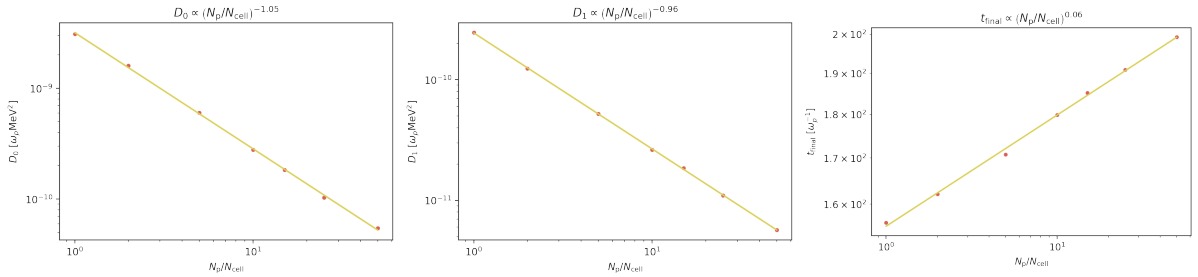

Figure A.8: The fitted values for the diffusion parameters D_0 (**left**), D_1 (**middle**) and the saturation time t_{final} (**right**) for the evolution of Std$(p_\perp)$ extracted from figure A.7 as a function of the number of macroparticles N_p/N_{cell}. The diffusion parameters scale almost linearly with the number of macroparticles, whereas the saturation time only depends mildly on the number of cells.

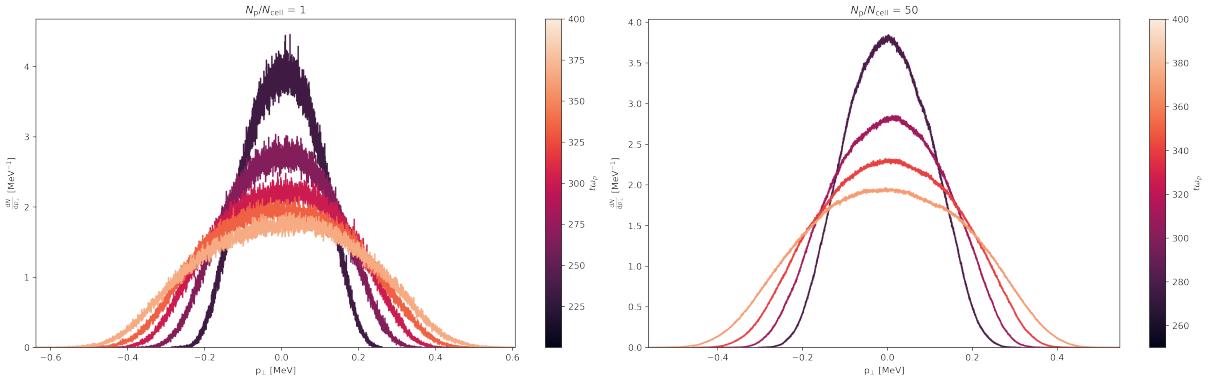

Figure A.9: The transversal beam momentum distribution for simulation runs with $N_p/N_{\text{cell}} = 1$ (**left**) and $N_p/N_{\text{cell}} = 50$ (**right**) at a number of time steps near the end of instability growth. In the simulation run with less macroparticles the momentum distribution shows stronger fluctuations.

A.2 Resolution of Modes

The maximum wavelength that can be resolved by a simulation domain with periodic boundary conditions and length L is twice the domain length $2L$. The second largest wavelength that can be resolved under these circumstances is L. Only a discrete number of wavevector that depend on the length of the simulation box can be resolved. The difference between neighboring modes for each dimension is Δk_i.

$$\Delta k_i = \frac{\pi}{L_i} \tag{A.1}$$

The resolution Δk is too small it can happen that the fastest growing modes are not resolved and the integrated growth rate is much smaller than the true physical growth rate. In principle the resolution has to be checked for each dimension separately, however we know for that oblique instability that the resonance is much narrow in $k_\parallel$ than $k_\perp$. For this reason we will scan both direction simultaneously under the condition that $\Delta k_\parallel = 0.1\Delta k_\perp$. We note that the width of the resonance depends on the chosen parameters of α and γ and thus for simulations with higher γ or lower α a higher resolution has to be chosen.
We find that if the resolution is too coarse the instability can be outright suppressed. For the chosen parameters of $\alpha = 10^{-3}$ and $\gamma = 50$ only for $\Delta k_\parallel = 0.1\Delta k_\perp < 0.1\frac{\omega_p}{c}$ does the integrated growth rate stabilize.

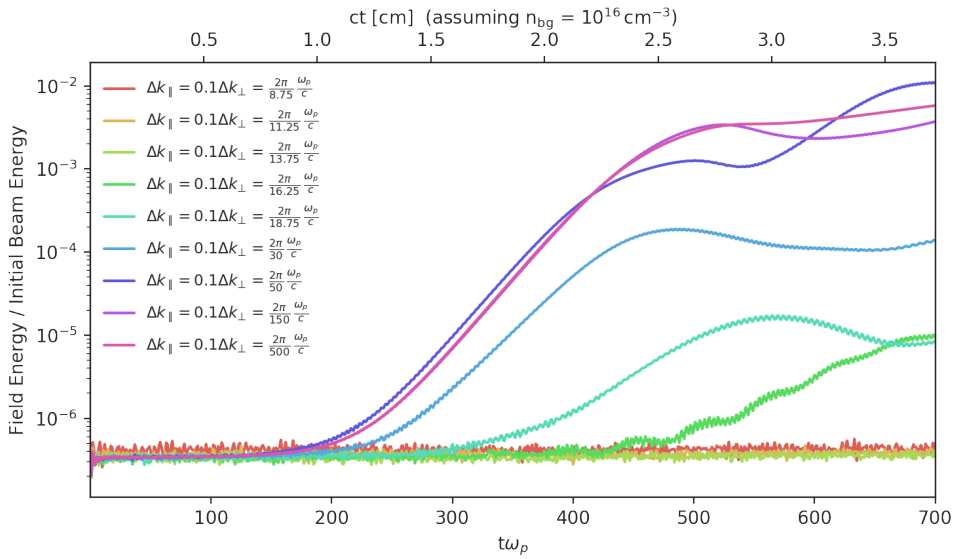

Figure A.10: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a density contrast of $\alpha = 10^{-3}$, Lorentz boost $\gamma = 50$ and varying resolution of field modes Δk. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

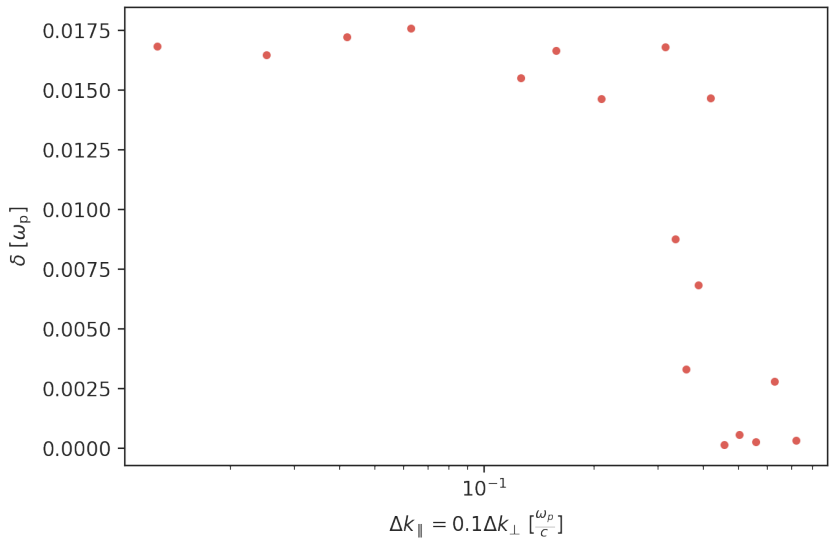

Figure A.11: The growth rate of the total field energy density as a function of resolution of field modes Δk extracted from figure A.10. For a resolution of $\Delta k_\parallel < 0.1\frac{\omega_p}{c}$ the growth rate stabilizes.

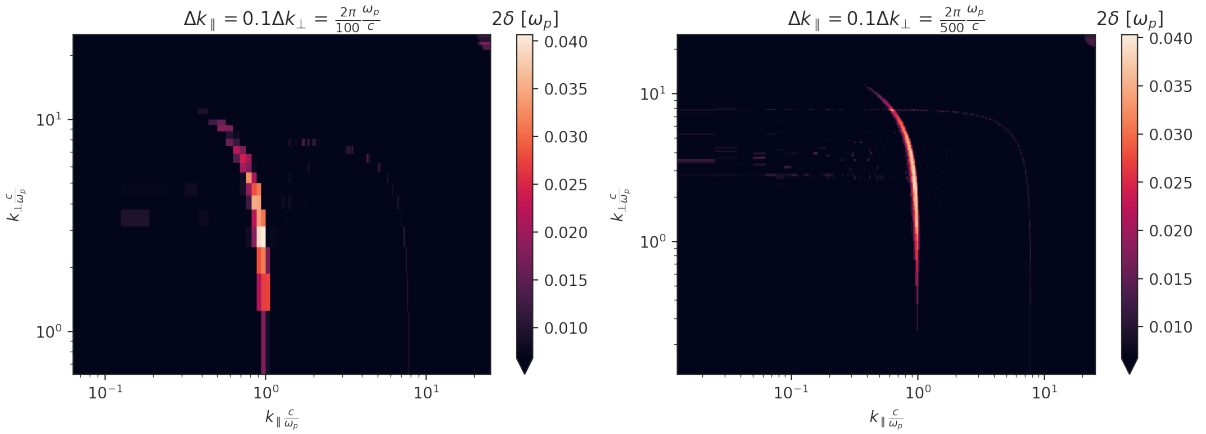

Figure A.12: Maps of the growth rate as a function of wavevector for each mode for two simulation runs with a different resolution of the field modes Δk. On the **left** the resolution is much more coarse than on the **right**.

A.3 Resolution of Plasmawavelength

The plasma wavelength is the fundamental length scale of the plasma system we want to investigate. It is possible to achieve a high resolution Δk as well as a maximum resolved mode $\max(k)$ that is larger than the plasma wave length while still only resolving the plasma wavelength in real space with only a few grid cells. More precisely $\max(k) > \frac{\omega_p}{c}$ if there are more than π^{-1} grid cells per plasma wavelength. This would of course be in violation of the requirement that the fundamental length scales have to be well resolved. We scan the number of grid cells per plasma wavelength $N_{\text{cell}}/L\lambda_p^{-1}$ and find that insufficient resolution does not only reduce the integrated growth rate but also shift the resonant modes to unphysical values $k_\parallel < \frac{\omega_p}{c}$. Only when the plasma wavelength is resolved with at least 8 cells do we observe convergence. This is equivalent to a maximum wavevector $\max = 8\pi\frac{\omega_p}{c}$, which is more than one order of magnitude larger than the resonant modes.

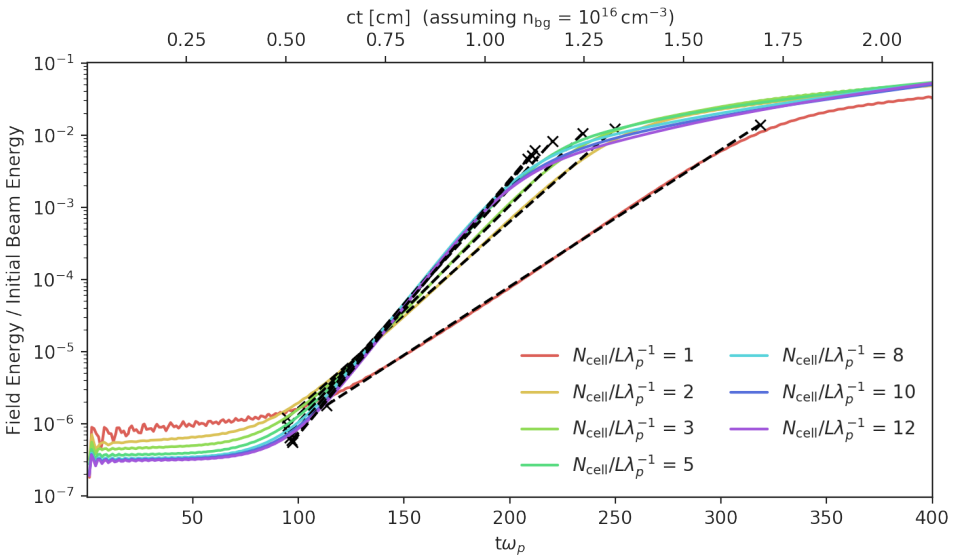

Figure A.13: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a density contrast of $\alpha = 10^{-3}$, Lorentz boost $\gamma = 5$ and varying resolution of the plasma wavelength $N_{\text{cell}}/L\lambda_p^{-1}$. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

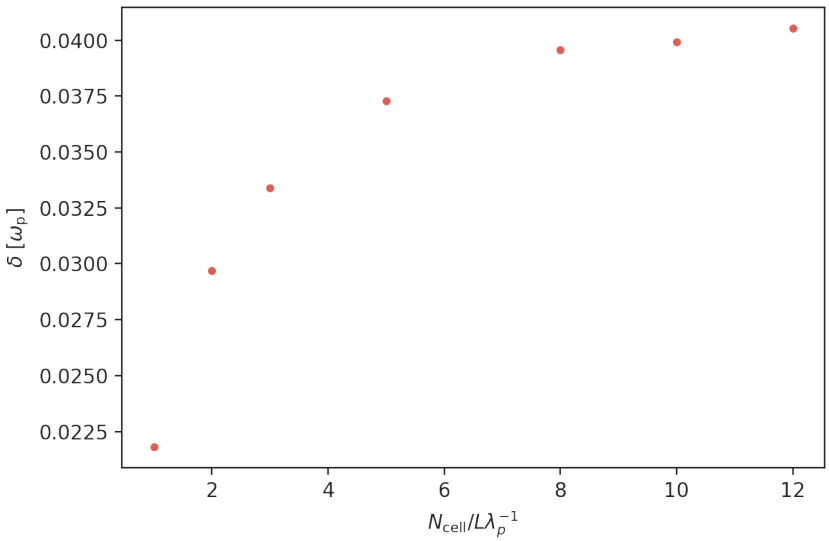

Figure A.14: The growth rate of the total field energy density as a function of the grid cells per plasma wavelength $N_{\text{cell}}/L\lambda_p^{-1}$ extracted from figure A.13. For 8 or more grid cells per plasma wavelength the growth rate stabilizes.

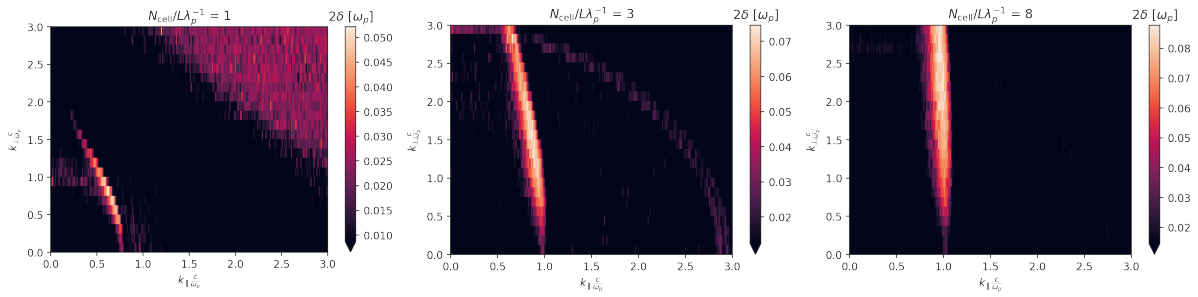

Figure A.15: Maps of the growth rate as a function of wavevector for each mode for three simulation runs that resolve the plasma wavelength if a different number of grid cells. On the **left** the plasma wavelength is only resolved with a single grid cell. As a result the fastest growing modes are shifted to smaller $k_\parallel$ and curve leftwards with increasing $k_\perp$. In the **middle** the plasma wavelength is resolved with 3 grid cells which reduces the shift and tilt. Only on the **right**, where 8 grid cells resolve each plasma wavelength, does the tilt become negligible and the fastest growing modes are found at $k_\parallel \sim \frac{\omega_p}{c}$.

A.4 Energy Conservation

The PIC scheme implemented in EPOCH does not strictly conserve energy. Thus it is possible that the energy of the system unphysically changes during a simulation run, an effect that is often called "numerical heating" when the total energy of the system increases. Here we want to check that the numerical heating effect is not out of control. We find that even under the worst circumstances the relative change of the internal energy of the system remains confined to the one percent level. We also find that the numerical heating at the beginning of the simulation run, before and during instability growth, can be reduced by increasing the number of particles per grid cell and that numerical heating can be reduced after instability saturation by resolving the plasma wavelength with more grid cells.
The conservation of energy can be improved by increasing the number of macroparticles or increasing the number of grid cells.

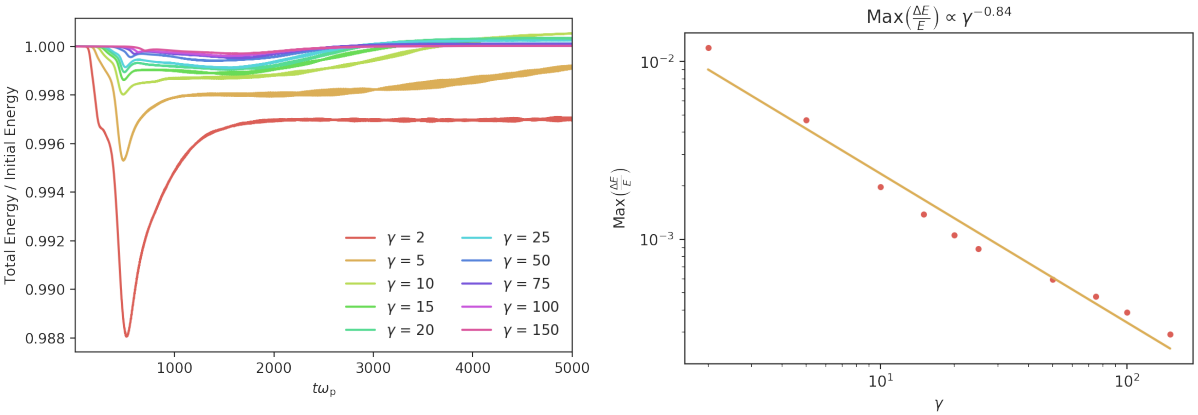

Figure A.16: On the **left** we show the evolution of the total energy of the system (normalized to the initial energy) for simulation runs with varying Lorentz boost γ. On the **right** we show the maximum deviation as a function of γ. The relative deviation decreases almost linearly with γ

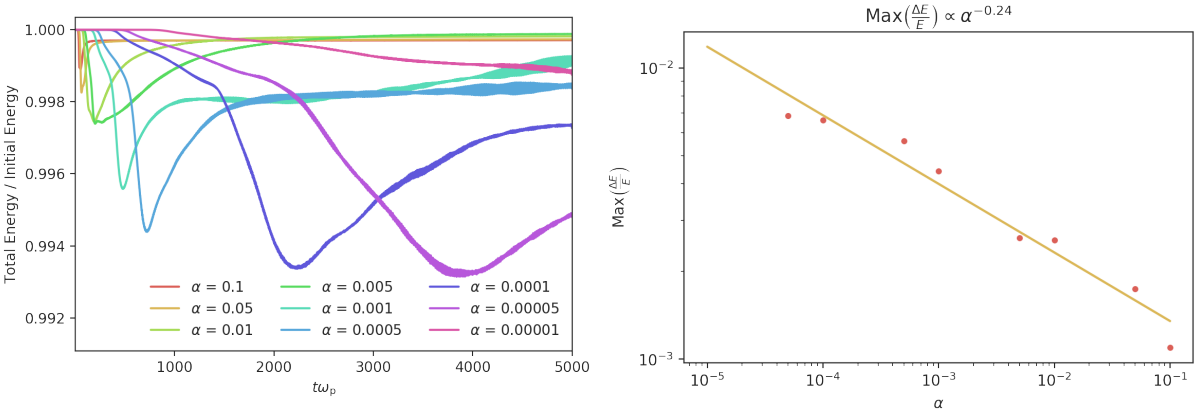

Figure A.17: On the **left** we show the evolution of the total energy of the system (normalized to the initial energy) for simulation runs with varying density ratio α. On the **right** we show the maximum deviation as a function of α. The numerical heating increases with smaller α.

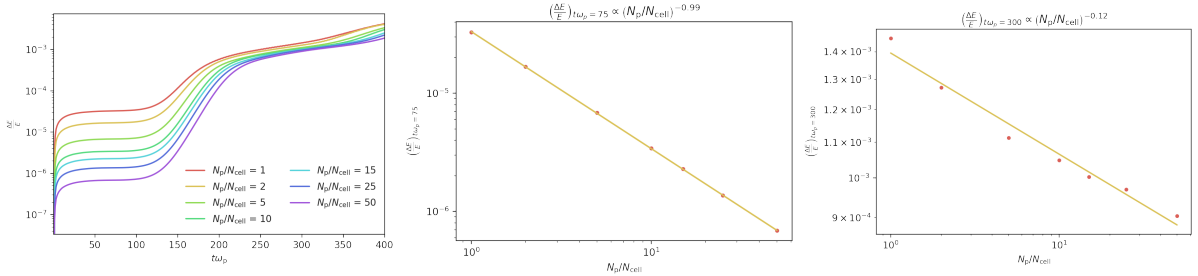

Figure A.18: On the **left** we show the evolution of the total energy of the system (normalized to the initial energy) for simulation runs with varying number of macroparticles N_p/N_{cell}. In the **center** we show the relative change of energy at time $t\omega_p = 75$, which is during the initial noise regime. In this regime the numerical heating scales linearly with the number of macroparticles. On the **right** we show the relative change of energy at time $t\omega_p = 300$, which is after the saturation takes place. In this regime the numerical heating is only mildly affected by the number of macroparticles.

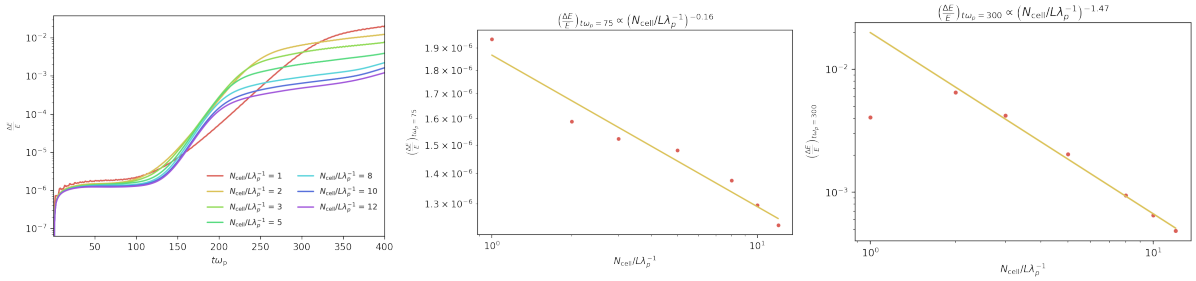

Figure A.19: On the **left** we show the evolution of the total energy of the system (normalized to the initial energy) for simulation runs with varying grid cells per plasma wavelength $N_{\text{cell}}/L\lambda_p^{-1}$. In the **center** we show the relative change of energy at time $t\omega_p = 75$, which is during the initial noise regime. In this regime the numerical heating is only mildly affected by the number of grid cells. On the **right** we show the relative change of energy at time $t\omega_p = 300$, which is after the saturation takes place. In this regime the numerical heating is very strongly affected by the number of grid cells.

B | Cold Beam in 3D

To save computational time we reduced the physical 3D problem of the beam instability to a system that resolve two spatial dimensions, one parallel to the beam direction and one transverse, and three momentum dimensions. This system should in principle capture most of the relevant physics, however some unphysical phenomena can be observed. For example the momentum distribution in the direction that is not resolved in space does not evolve. However for symmetry reasons the momentum distribution of both transversal dimensions should behave the same. For this reason we perform 3D simulation with reduced grid resolution and compare them with the 2D results to make sure that the insights we gained from the 2D system translate to the full 3D system. The settings of the 3D simulations can be found in table B.1. We note that the reduced grid resolution and the reduced number of particles per grid cell does not reduce the total number of macroparticles leading to similar computational effort for the simulation (per simulation time) as in the 2D higher resolution setup. This however also means that observables that can be averaged over the entire simulation domain suffer from far less noise than a 2D simulation with similar parameters.
We performed a scan of the instability growth for α and γ respectively and found that the scaling of the cold oblique instability can also be observed in 3D. Furthermore the evolution of the momentum distribution in all three dimensions is qualitatively the same as observed in the 2D PIC simulation for the two dimensions that are spatial resolved.

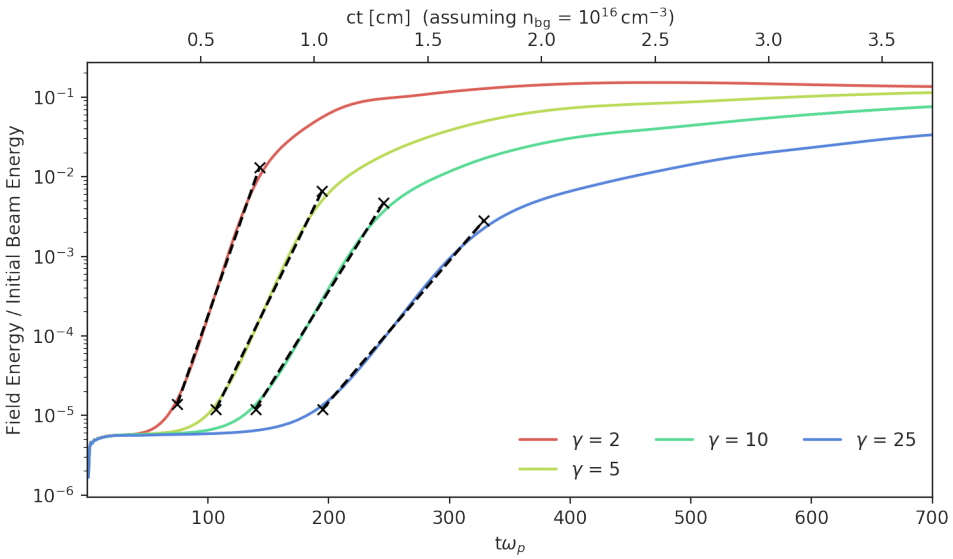

Figure B.1: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a density contrast of $\alpha = 10^{-3}$ and varying Lorentz boost γ. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

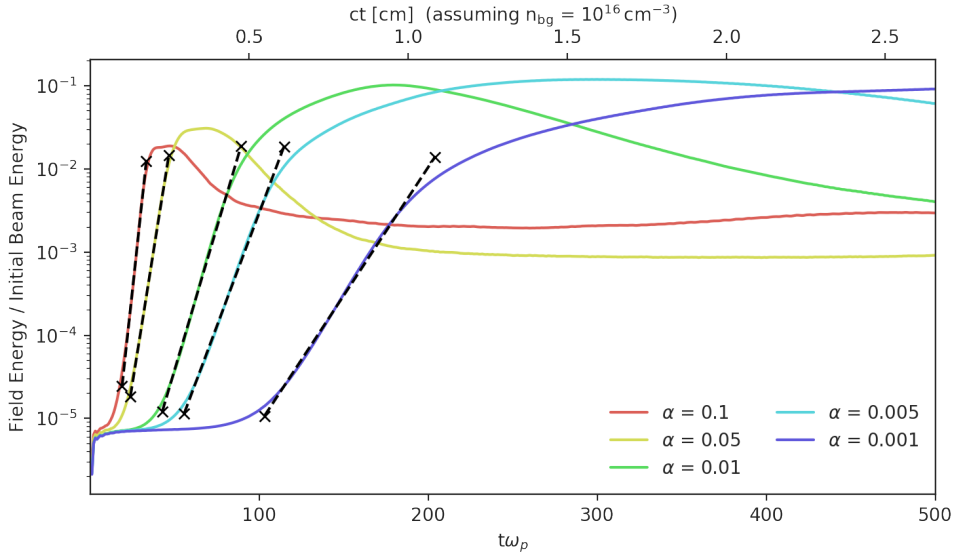

Figure B.2: The energy density of the combined electric and magnetic fields across all modes and all directional components normalized to the initial energy density of the pair beam as a function of time (bottom axis) and propagation distance (top axis). Each simulation run has a Lorentz boost of $\gamma = 5$ and varying density contrast α. For each simulation run the start and end point of the linear growth phase, as determined by fit to (5.5), is indicated by cross and the linear growth rate is indicated by a dashed line.

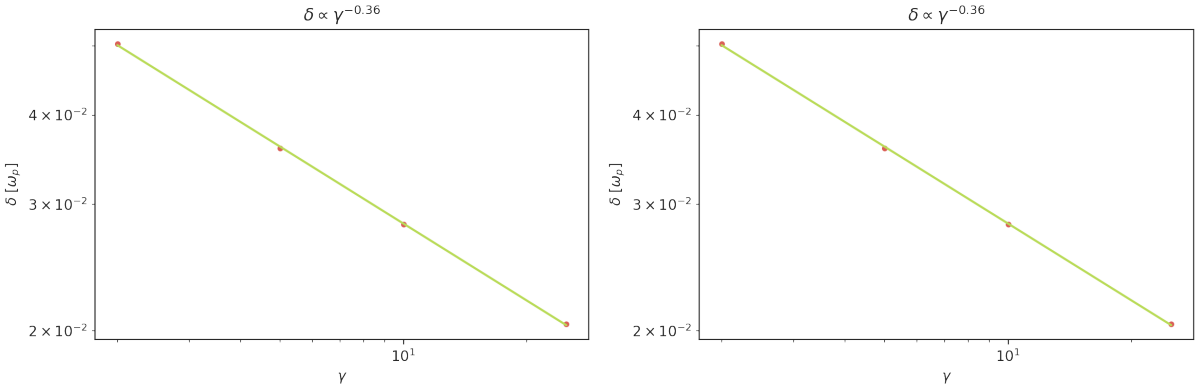

Figure B.3: The growth rate of the total field energy density as a function of γ (**left**, extracted from figure B.1) and α (**right**, extracted from figure B.2).

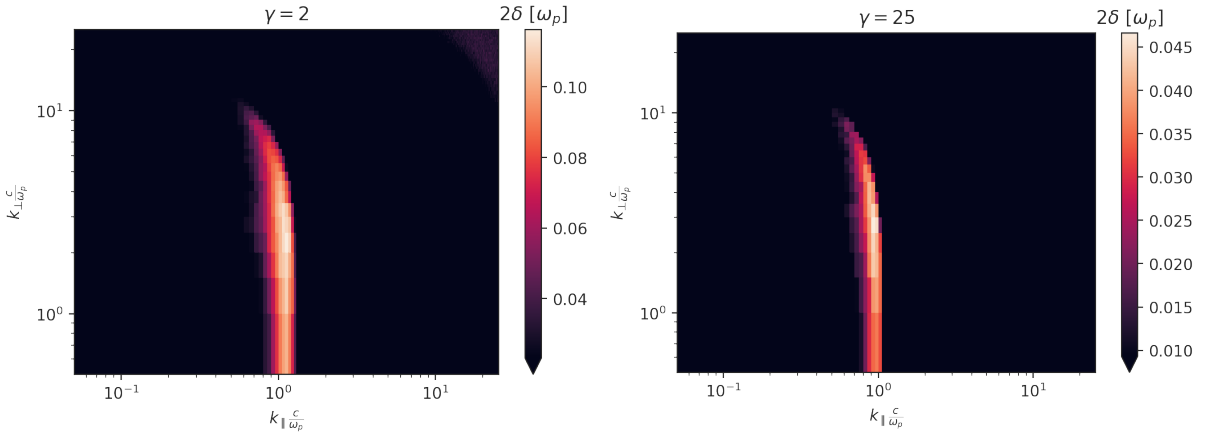

Figure B.4: Map of the growth rate as a function of wavevector for each mode for two simulation runs. On the **left** the Lorentz boost $\gamma = 2$ is lower than on the **right** where $\gamma = 25$. In both cases the density contrast $\alpha = 10^{-3}$ is the same. For a higher Lorentz boost the maximum growth rate is lower and also the width of the resonance is much smaller.

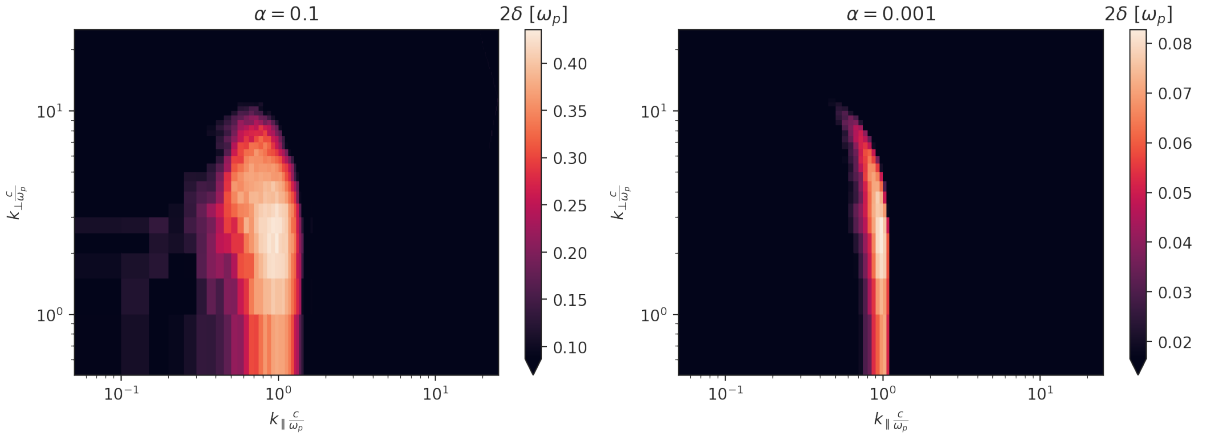

Figure B.5: Map of the growth rate as a function of wavevector for each mode for two simulation runs. On the **left** the density contrast $\alpha = 10^{-1}$ is lower than on the **right** where $\alpha = 10^{-3}$. In both cases the Lorentz boost $\gamma = 5$ is the same. For a lower density contrast the maximum growth rate is lower and also the width of the resonance is much smaller.

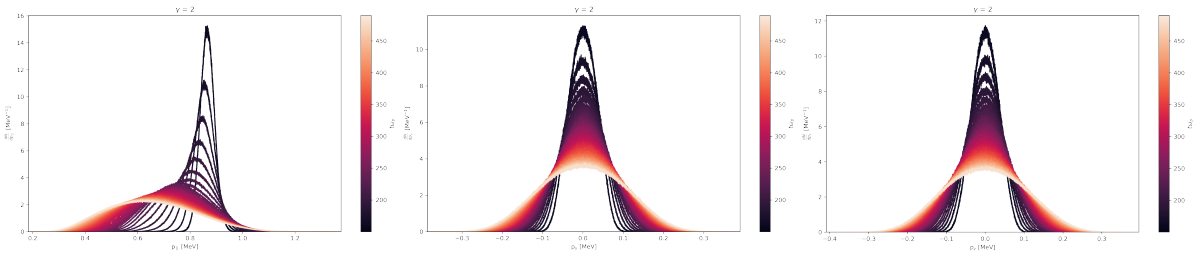

Figure B.6: The beam momentum distribution for a simulation run with $\gamma = 2$ and $\alpha = 10^{-3}$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **middle** and the **right** plots show the two transversal momentum distributions

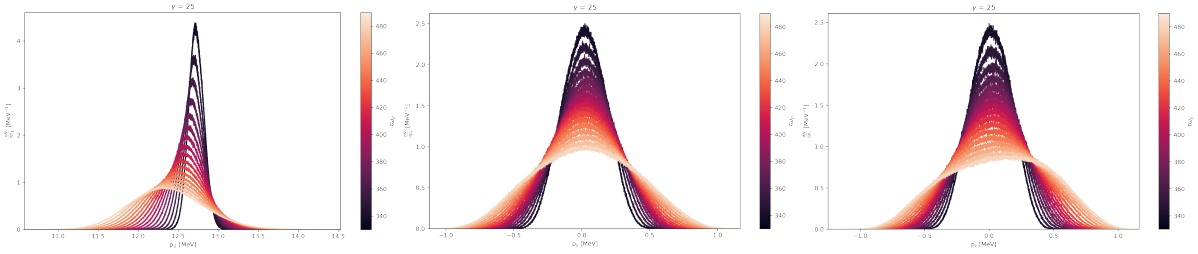

Figure B.7: The beam momentum distribution for a simulation run with $\gamma = 25$ and $\alpha = 10^{-3}$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **middle** and the **right** plots show the two transversal momentum distributions

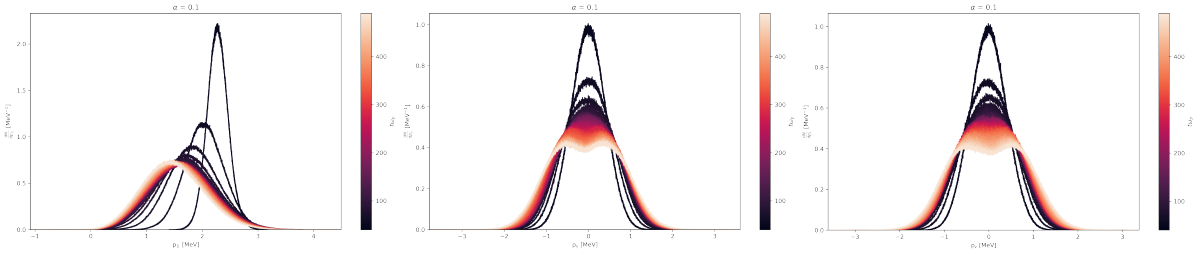

Figure B.8: The beam momentum distribution for a simulation run with $\alpha = 0.1$ and $\gamma = 5$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **middle** and the **right** plots show the two transversal momentum distributions

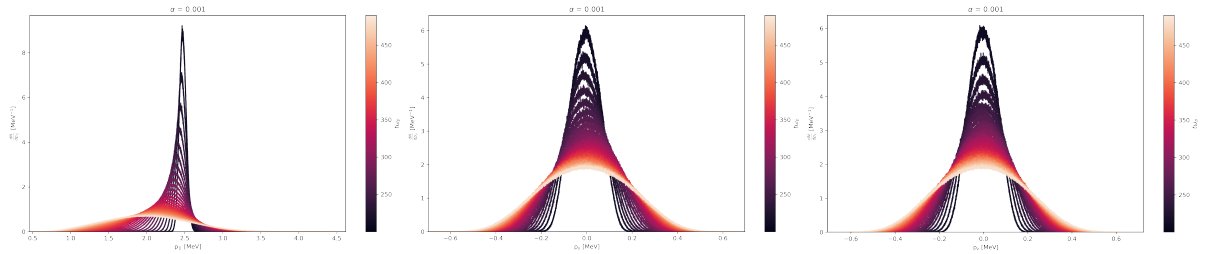

Figure B.9: The beam momentum distribution for a simulation run with $\alpha = 10^{-3}$ and $\gamma = 5$ at a number of time steps near the end of instability growth. The **left** plot shows the longitudinal momentum distribution. The **middle** and the **right** plots show the two transversal momentum distributions

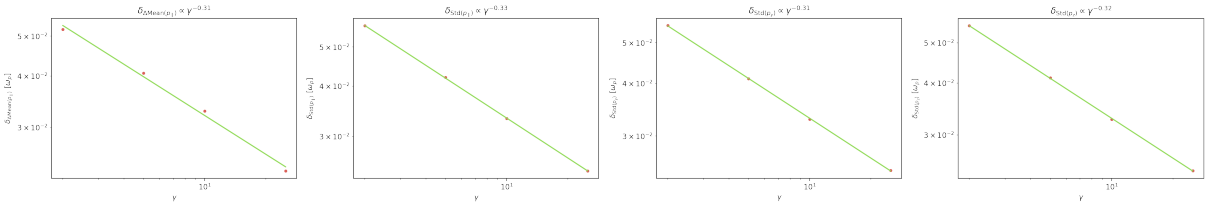

Figure B.10: The growth rate determined by a fit to the evolution of $\Delta\text{Mean}(p_\parallel)$ (**far left**), $\text{Std}(p_\parallel)$ (**center left**), $\text{Std}(p_y)$ (**center right**) and $\text{Std}(p_z)$ (**far right**) as a function of the Lorentz boost γ.

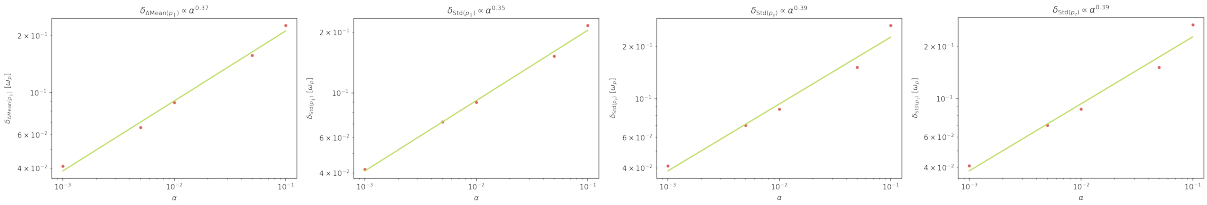

Figure B.11: The growth rate determined by a fit to the evolution of $\Delta\text{Mean}(p_\parallel)$ (**far left**), $\text{Std}(p_\parallel)$ (**center left**), $\text{Std}(p_y)$ (**center right**) and $\text{Std}(p_z)$ (**far right**) as a function of the density ratio α.

Table B.1: Overview of the 3D PIC Simulation settings for the simulation runs.

Settings of the PIC Simulation	
Number of Dimensions	3 (x beam direction, y & z transverse)
Boundary conditions in x	periodic
Boundary conditions in y	periodic
Boundary conditions in z	periodic
Box Length L_x	$125\,\frac{c}{\omega_p}$
Box Width L_y	$12.5\,\frac{c}{\omega_p}$
Box Width L_z	$12.5\,\frac{c}{\omega_p}$
Number of Cells N_x	1000
Number of Cells N_y	100
Number of Cells N_z	100
Number of particles per cell N_p (per species)	1
Timestep Δt	0.95 CFL-Criterion[a] $\approx 0.088\omega_p^{-1}$
Maxwell Solver	Yee
Field Order	6
Particle Pusher	Higuera & Cray
Particle Shape Function	Third Order B-Spline
Current Smoothing	5-fold (1-2-3-4 steps)
Background particles	Electrons & immobile Protons
Initial Background temperature T	$(\gamma - 1)\,\alpha \cdot$ 200 eV
Beam particles	Electrons & Positrons
Distribution Function	$f(\mathbf{p}; \mu, \sigma_x, \sigma_y) \propto \exp\left(-\frac{(p_x-\mu)^2}{2\sigma_x^2} + \frac{p_y^2+p_z^2}{2\sigma_y^2}\right)$
μ [MeV]	$0.511\sqrt{\gamma^2-1}$
σ_x [keV] (alias Δp_x)	0.5
σ_y [keV] (alias Δp_y)	0.5
α	0.001 for γ scan
γ	5 for α scan
Total time T	$700\,\omega_p^{-1}$ for γ scan & $500\,\omega_p^{-1}$ for α scan

[a] The CFL criterion depends on the grid spacing and the field order. For sixth order field interpolation it is $\frac{120}{149}\frac{1}{\sqrt{\Delta x^{-2}+\Delta y^{-2}+\Delta z^{-2}}}c^{-1}$

C | EPOCH Input Deck Example

Here we give a commented example input deck for the full simulation of a pair beam in 2D as shown in chapter 5 for the parameters $\gamma = 5$ and $\alpha = 10^{-3}$. All other simulation results in that chapter can be reproduced by adjusting the values of `alpha` and `gamma` in the `constant` block of the input deck.

```
1  begin:constant
2
3    # Define Constants
4    # Constants can be written to a file for easy acces later
5
6    ncell_x = 4000
7    ncell_y = 400
8    dens = 1e16*1e6
9    alpha = 0.001
10   omega_p =  sqrt(dens*qe^2/(me*epsilon0))
11   Length_x = 500*c/omega_p
12   Length_y = 50*c/omega_p
13   ppercell = 25
14   gamma_peak = 5.00
15
16   p_ref = sqrt(gamma_peak^2 - 1)*me*c
17   sigma_par = 0.0005/0.511
18   sigma_trans = 0.0005/0.511
19   sigma_max_load = 0.0025/0.511*me*c
20   sigma_max_plot_1d = (10+gamma_peak/5)*me*c
21
22   p0 = p_ref
23   p1 = sigma_par*me*c
24   p2 = sigma_trans*me*c
25 end:constant
26
27 begin:control
28
29   # Number of Grid Cells
30   nx = ncell_x
31   ny = ncell_y
32
33   # Final time of simulation
34   t_end = 5000/omega_p
35
36   # Size of domain
37   x_min = 0
38   x_max = Length_x
```

```
39    y_min = 0
40    y_max = Length_y
41
42    # custom parameters
43    dt_multiplier = 0.95 # multiplier for timestep relative to courant condition (
          default)
44    dlb_threshold = 0.1
45    field_order = 6 # Filed Order Interpolation, can be 2, 4, 6 for Yee Grid
46    stdout_frequency = 2000
47    use_random_seed = T # determine random seed from system clock
48    smooth_currents = T # Apply Current Smoothing
49    smooth_iterations = 5 # Number of Current Smoothing Steps
50    use_accurate_n_zeros = F # more expensive way to calculate output file names
51    print_constants = T # Write constant block to file
52  end:control
53
54
55  begin:boundaries
56
57    # Define Boundary conditions for each simulation
58    # box boundary separately
59
60    bc_x_min = periodic
61    bc_x_max = periodic
62    bc_y_min = periodic
63    bc_y_max = periodic
64  end:boundaries
65
66  begin:species
67    # ElectronBeam
68    name = ElectronBeam
69    charge = -1
70    mass = 1.0
71    number_density = alpha*dens
72    nparticles = ppercell * nx * ny
73
74    # Distribution function of Beam is shifted Gaussian
75    # Has to be noramlized so that the maximum is 1
76    dist_fn = exp(-(px - p0)^2/(2*p1^2) - py^2/(2*p2^2) - pz^2/(2*p2^2))
77    # The range of each coordinate that the momentum distribution
78    # is sampled with an acceptance rejection method
79    dist_fn_px_range = (p_ref-sigma_max_load, p_ref+sigma_max_load)
80    dist_fn_py_range = (-sigma_max_load, sigma_max_load)
81    dist_fn_pz_range = (-sigma_max_load, sigma_max_load)
82  end:species
83
84
85  begin:species
86    # PositronBeam
87    name = PositronBeam
88    charge = 1
89    mass = 1.0
90    number_density = alpha*dens
91    nparticles = ppercell * nx * ny
92
93    dist_fn = exp(-(px - p0)^2/(2*p1^2) - py^2/(2*p2^2) - pz^2/(2*p2^2))
```

```
94    dist_fn_px_range = (p_ref-sigma_max_load, p_ref+sigma_max_load)
95    dist_fn_py_range = (-sigma_max_load, sigma_max_load)
96    dist_fn_pz_range = (-sigma_max_load, sigma_max_load)
97  end:species
98
99  begin:species
100   # Background Electrons
101   name:ElectronBG
102   charge = -1
103   mass = 1.0
104   temperature_ev = 0.2*(gamma_peak-1)
105   number_density = dens
106   nparticles = ppercell * nx * ny
107 end:species
108
109 begin:species
110   # Background Ions
111   name:IonBG
112   charge = 1
113   mass = 1836.0
114   temperature_ev = 0.2*(gamma_peak-1)
115   number_density = dens
116   nparticles = ppercell * nx * ny
117
118   # Protons are set to immobile due to their much larger mass
119   # Particle Push is not applied to immobile particles
120   immobile = T
121 end:species
122
123 begin:dist_fn
124
125   # Define how momentum distribution histogram of beam electrons
126   # should be written to output files
127
128   name = px_beam
129   ndims = 1
130
131   direction1 = dir_px
132   range1 = (0, 0) # This automatically adjustes the range as the distribution
          widens
133
134   resolution1 = nint(3*sigma_max_plot_1d/(me*c)*50)
135
136   restrict_px = (p_ref-2*sigma_max_plot_1d,p_ref+sigma_max_plot_1d) # discard
          individual particles with extrem momentum change for histogram
137
138   include_species:ElectronBeam
139 end:dist_fn
140
141 begin:dist_fn
142   name = py_beam
143   ndims = 1
144
145   direction1 = dir_py
146   range1 = (0, 0)
147
```

```
148    resolution1 = 5000
149
150    include_species:ElectronBeam
151  end:dist_fn
152
153  begin:dist_fn
154    name = px_bg
155    ndims = 1
156
157    direction1 = dir_px
158    range1 = (0, 0)
159
160    resolution1 = 5000
161
162    restrict_px = (-me*c, me*c)
163
164    include_species:ElectronBG
165  end:dist_fn
166
167  begin:dist_fn
168    name = py_bg
169    ndims = 1
170
171    direction1 = dir_py
172    range1 = (0, 0)
173
174    resolution1 = 5000
175
176    restrict_px = (-me*c, me*c)
177
178    include_species:ElectronBG
179  end:dist_fn
180
181  begin:output
182
183    # write the energy every plasma period
184
185    dt_snapshot = 1/omega_p
186    dump_last = T
187    name = o1
188
189    total_energy_sum = always + single + species
190    distribution_functions = never
191  end:output
192
193  begin:output
194
195    # fields and distribution functions are written every
196    # 5 plasma periods until time = 500/omega_p
197
198    dt_snapshot = 5/omega_p
199    dump_last = F
200    name = o2
201
202    ex = always + single
203    ey = always + single
```

```
204     bz = always + single
205     distribution_functions = always + single
206
207     time_stop = 500/omega_p
208   end:output
209
210   begin:output
211
212     # fields and distributions are written every 50 plasma periods
213     # even after time = 500/omega_p
214
215     dt_snapshot = 50/omega_p
216     dump_last = T
217     name = o3
218
219     ex = always + single
220     ey = always + single
221     bz = always + single
222     distribution_functions = always + single
223
224   end:output
```

D | Explanation of Key Parameters

Table D.1: Overview of key parameters and quantities that define the behavior of the beam plasma instability

Parameter	Description
γ	The mean Lorentz boost of the e^+e^- beam.
$\alpha = \frac{n_\mathrm{b}}{n_\mathrm{bg}}$	The ratio of the beam plasma number density to the background plasma number density. Should be smaller than 1. Also called density contrast.
$\epsilon = \frac{n_\mathrm{b}(\gamma-1)m_e}{\frac{3}{2}n_\mathrm{bg}k_bT}$	The ratio of the beam kinetic energy to the background plasma kinetic energy.
δ	The growth rate of the (exponentially growing) amplitude the of an unstable mode. The energy density of a mode is proportional to the square of the amplitude and thus the energy density grows with the rate 2δ. The growth rate can also be derived for other quantities that grow exponentially.
Λ	The saturation value that a quantity reaches at end of a linear growth phase.
$\sigma_\parallel/\sigma_\perp$	Initial spread of beam momentum distribution distribution parallel/transversal to the beam direction.
$L_\parallel/L_\perp$	The extension of the beam longitudinal/transversal to the beam direction. In an idealized scenario this extension is infinite due to periodic boundary conditions of the simulation setup.
$\chi = \frac{n_{e^+}}{n_{e^-}}$	The ratio of positrons to electrons in the beam. In a neutral beam $\chi = 1$.

$j \propto n_{e^-} - n_{e^+} = \frac{1-\chi}{1+\chi}$

The current induced by the e^+e^- beam. For a neutral beam ($\chi = 1$) this is current vanishes.

$\dot{n} = \frac{1}{n_{\mathrm{bg}}} \frac{\partial n_{\mathrm{bg}}}{\partial x}$

Rate of change of the background plasma number density relative to the background plasma density. In the case of short coherence length this is a local quantity. In the case of long coherence length this a global quantity that gives the slope of the background plasma density experienced by the beam.

Acknowledgements

First and foremost I need to thank my advisor Prof. Dr. Günter Sigl for allowing me to embark on the academic journey that culminated in this thesis and making me part of his research group. Without his guidance this thesis would not be possible. I am also grateful for the freedom he allowed me in the way I pursed my research. Furthermore I want to thank my co-advisor Prof. Dr. Florian Grüner for his important input on experimental aspects, Dr. Carl B. Schroeder for his patient advice and Prof. Dr. Martin Pohl for sharing his vast expertise on the entire domain of plasma instabilities and PIC simulations.
The research presented in this thesis was only possible in close collaboration with many competent researchers. My closest collaborator was Dr. Oindrila Ghosh who's work on the theory of plasma instabilities provided important input for this thesis on many occasions. Dr. Benno Zeitler and Dr. Ryan David Stark contributed with important work on the transport and focusing of particle beam that contains electrons and positrons.
When I arrived in Hamburg in the fall of 2019 our small research group additionally consisted of Dr. Pranjal Trivedi, Dr. Peter Niksa and Dr. Julien Manshanden who warmly welcomed me and have by now moved on to the next stage of their careers. Thanks to the arrival of Dennis Maseizik and Simone Rossoni our group did not feel too lonely. Unfortunately in the spring of 2020 the world was turned upside down by the Covid-19 pandemic. In the following years we had to make due with mostly interacting from a distance. But even this hardship passed eventually and we are back to bonding over shared lunches and climbing sessions. Beyond our little research group I want to thank all members of the II. Insitut for theoretical Physics, the Desy Theory Group and the Quantum Universe cluster of excellence for all academic and social events over the last years.
This thesis is of course not the final word on plasma instabilities and the gamma ray portal. The research must go on and I am very glad that Suman Dey will continue our research. Thank you and good luck!
This work would not have been possible with the excellent computing resources provided by the Maxwell HPC and the people who run the cluster.

Last but certainly not least I need to extend my utmost gratitude to my parents for supporting me in more ways than I can count or put into words. Nothing in my life would have been possible without you.