
Particle acceleration in magnetised mildly relativistic shocks

Gabriel Torralba Paz

Doctoral Thesis



Supervised by: prof. dr hab. Jacek Niemiec

Kraków, 2026

Abstract

In this thesis, fully kinetic particle-in-cell (PIC) simulations are used to study mildly relativistic collisionless shocks in strongly magnetised plasmas with oblique configurations of the ambient magnetic fields. Such shocks are expected to occur in relativistic plasma outflows, including jets of active galactic nuclei and microquasars, which are prominent sources of non-thermal radiation and potential acceleration sites of high-energy cosmic rays.

Previous studies have shown that while diffusive shock acceleration is efficient in weakly magnetised, subluminal shocks, and wakefield acceleration operates predominantly in the ultra-relativistic regime, both mechanisms are expected to be inefficient in mildly relativistic, strongly magnetised plasmas. This work therefore explores whether alternative mechanisms can support efficient particle heating and acceleration under these previously unexplored conditions.

Large-scale, high-resolution 2D3V PIC simulations are performed using the code THAT-MPI, which was extensively optimised and extended with new features for this study. The simulations probe shocks with Lorentz factor $\gamma_0 \approx 2.24$, propagating in strongly magnetised plasmas with $\sigma = 1$. Three magnetic field configurations are examined: a superluminal oblique shock ($\theta_{Bn} = 40^\circ$), a subluminal oblique shock ($\theta_{Bn} = 30^\circ$) close to the threshold angle between superluminal and subluminal configurations, and a quasi-parallel shock ($\theta_{Bn} = 10^\circ$).

The oblique shocks develop large-amplitude magnetosonic wave trains with sharp electrostatic potentials and decreasing amplitude downstream. Their structure is reminiscent of weak non-relativistic subcritical space shocks. The shocks are intrinsically non-stationary and develop strong transverse non-uniformities. This enables particle reflection and trapping within magnetosonic pulses, in the subluminal configuration close to the threshold angle giving rise to efficient heating and acceleration of both electrons and ions. The identified mechanism is analogous to the electron trapping and acceleration process previously reported in non-relativistic shocks, but in mildly relativistic shocks operates for both species and with significantly higher efficiency due to the assumed stronger magnetisation and higher shock speeds. In the subluminal case, electrons develop a high-energy power-law spectrum with index $p \approx 2$, limited by a cutoff set by the shock width, which scales with the Lorentz factor, wave amplitude, and ion-to-electron mass ratio. The ion acceleration through this mechanism is reduced by roughly two orders of magnitude. Some of the ions can overcome the shock potential and propagate upstream, where they undergo further energisation through the shock surfing acceleration process.

The quasi-parallel shock exhibits a markedly different structure. The shock is subcritical with respect to the whistler critical Mach number and forms an extended precursor populated by large-amplitude dispersive whistler waves. The interaction of ions and electrons with the precursor whistlers and additional wave modes associated with their nonlinear dynamics leads to strong heating of both species.

Overall, the results demonstrate that strongly magnetised, mildly relativistic oblique shocks can support an efficient particle energisation mechanism distinct from diffusive shock

acceleration and wakefield acceleration. While the resulting spectra are limited in extent, the process provides substantial energy gain and may act as an effective injection mechanism for further acceleration. These findings highlight the importance of shock non-stationarity and multidimensional effects and motivate future fully three-dimensional studies.

In addition to the PIC simulations, this thesis introduces a novel approach for analysing particle tracing data using neural networks. To assess its applicability, traced particle data from simulations studying electron acceleration in waves excited by the Buneman instability in high-Mach-number shocks are used as a test case. Three classes of algorithms are employed: classification, regression, and anomaly detection. The results demonstrate that neural networks have significant potential to enhance and streamline the analysis of particle trajectories in kinetic plasma simulations.

Streszczenie

W niniejszej rozprawie, z wykorzystaniem w pełni kinetycznych symulacji typu *particle-in-cell* (PIC), badane są średnio relatywistyczne fale uderzeniowe w silnie zmagetyzowanej, bezdenerżeniowej plazmie o skośnej konfiguracji zewnętrznego pola magnetycznego. Tego rodzaju fale uderzeniowe mogą występować w relatywistycznych wypływach plazmy, w tym w dżetach aktywnych jąder galaktyk oraz mikrokwarzarów, które są znanymi źródłami promieniowania nietermicznego i powszechnie uznawane są za potencjalne miejsca przyspieszania wysokoenergetycznych promieni kosmicznych.

Wcześniejsze badania wykazały, że o ile dyfuzyjne przyspieszanie na fali uderzeniowej jest efektywne w słabo zmagetyzowanych falach podświetlnych, a przyspieszanie typu *wakefield* zachodzi głównie w reżimie ultrarelatywistycznym, o tyle oba te mechanizmy wydają się nieefektywne w średnio relatywistycznych, silnie zmagetyzowanych plazmach. Niniejsza praca bada zatem, czy w tych dotąd nieeksplorowanych warunkach mogą działać alternatywne mechanizmy efektywnego grzania i przyspieszania cząstek.

W ramach niniejszej rozprawy przeprowadzono wysokiej rozdzielczości wielkoskalowe symulacje PIC w konfiguracji 2D3V z wykorzystaniem kodu THATMPI, który został na potrzeby pracy gruntownie zoptymalizowany i rozszerzony o nowe funkcjonalności. Symulacje obejmują fale uderzeniowe o współczynniku Lorentza $\gamma_0 \approx 2.24$, rozchodzące się w plazmie o dużej magnetyzacji, $\sigma = 1$. Analizowano trzy konfiguracje pola magnetycznego: skośną falę nadświetlną ($\theta_{Bn} = 40^\circ$), skośną falę podświetlną ($\theta_{Bn} = 30^\circ$) bliską kąta progowego rozdzielającego konfiguracje nadświetlne i podświetlne oraz falę quasi-równoległą ($\theta_{Bn} = 10^\circ$).

Skośne fale uderzeniowe mają strukturę impulsów magnetosonicznych o dużej amplitudzie, która stopniowo maleje wraz z odległością od frontu fali. Impulsy te są związane z potencjałami elektrostatycznymi o stromych gradientach. Ich struktura przypomina słabe, nierelatywistyczne, subkrytyczne fale uderzeniowe obserwowane w przestrzeni kosmicznej. Badane fale są z natury niestacjonarne i rozwijają silne poprzeczne niejednorodności, co umożliwia odbicie i pułapkowanie cząstek w obrębie impulsów magnetosonicznych. W konfiguracji podświetlnej bliskiej kąta progowego prowadzi to do efektywnego grzania i przyspieszania zarówno elektronów, jak i jonów. Mechanizm akceleracji jest analogiczny do procesu pułapkowania i przyspieszania elektronów opisanego wcześniej dla fal nierelatywistycznych, lecz w przypadku fal średnio relatywistycznych działa dla obu rodzajów cząstek i cechuje się znacznie wyższą wydajnością dzięki założonej silniejszej magnetyzacji i większej prędkości fali uderzeniowej. W przypadku podświetlnym elektrony o wysokich energiach tworzą widmo potęgowe o indeksie $p \approx 2$, ograniczone obcięciem wyznaczonym przez szerokość fali uderzeniowej, która skaluje się ze współczynnikiem Lorentza, amplitudą fali i stosunkiem masy jon-elektron. Efektywność przyspieszania jonów w ramach tego mechanizmu jest mniejsza o około dwa rzędy wielkości. Niektóre jony pokonują potencjał fali uderzeniowej i przemieszczają się w obszar przed frontem fali, gdzie ulegają dalszemu przyspieszaniu w procesie serfowania na fali (*shock surfing acceleration*).

Fala quasi-równoległa ma wyraźnie odmienną strukturę. Jest ona subkrytyczna względem krytycznej liczby Macha dla fal światła i tworzy rozległy prekursor wypełniony dyspersyjnymi

falami świstu (*whistlers*) o dużej amplitudzie. Oddziaływanie jonów i elektronów z tymi falami w prekursorze, a także z modami falowymi wynikającymi z ich nieliniowej dynamiki, prowadzi do silnego grzania obu rodzajów cząstek.

Ogólnie wyniki badań wykazują, że silnie zmagetyzowane, średnio relatywistyczne skośne fale uderzeniowe mogą efektywnie przyspieszać cząstki w procesie odmiennym od dyfuzyjnego przyspieszania na fali uderzeniowej oraz przyspieszania typu *wakefield*. Choć widma cząstek są formowane w ograniczonym zakresie energii, proces ten zapewnia znaczący przyrost energii i może pełnić rolę efektywnego mechanizmu wstępnego wrzucania cząstek do dalszych procesów przyspieszania. Wyniki te podkreślają znaczenie niestacjonarności fal uderzeniowych oraz efektów wielowymiarowych i stanowią motywację dla przyszłych badań z wykorzystaniem w pełni trójwymiarowych symulacji.

Oprócz symulacji PIC, niniejsza rozprawa wprowadza nowatorskie podejście oparte na sieciach neuronowych do analizy trajektorii wybranych cząstek śledzonych w symulacjach kinetycznych. W celu oceny przydatności tego podejścia wykorzystano dane śledzonych cząstek z symulacji badających przyspieszanie elektronów w falach wzbudzanych przez niestabilność Bunemana w szokach o dużej liczbie Macha. Zastosowano trzy klasy algorytmów: klasyfikację, regresję oraz detekcję anomalii. Wyniki pokazują, że sieci neuronowe mają duży potencjał do usprawnienia i przyspieszenia analizy trajektorii śledzonych cząstek w kinetycznych symulacjach plazmy.

Acknowledgements

I would like to express my gratitude towards my PhD supervisor prof. dr. hab. Jacek Niemiec, for his guidance during my PhD studies at the Institute for Nuclear Physics Polish Academy of Sciences, and for helping and supporting the start of my career as a scientist. I would like to also express my gratitude towards my collaborator and colleague dr. Artem Bohdan, who has been of great help during four long years of my studies.

I wish to thank profs. Masahiro Hoshino, Takanobu Amano and Shuichi Matsukiyo for their warm welcome during my internship and their continuous help in deepening my knowledge and improving my research and thesis.

I would also like to thank dr. Oleh Kobzar, prof. Martin Pohl, dr. Karol Fułat, Eloise Moore and dr. hab. Sabrina Casanova for their support and fruitful discussions during our international collaboration meetings.

Finally, I would like to express my sincere gratitude towards my family, who supported my career in pursuing my dream of becoming a doctor abroad.

This work has been supported by Narodowe Centrum Nauki through research Project No. 2019/33/B/ST9/02569.

We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grants no. PLG/2023/016861, PLG/2023/016378, PLG/2024/017881 and PLG/2025/019027.

We acknowledge Polish high-performance computing infrastructure PLGrid for awarding this project access to the LUMI supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CSC (Finland) and the LUMI consortium through PLL/2024/06/017077.

Contents

1	Introduction	5
2	Cosmic ray origin	7
2.1	Cosmic ray spectrum and sources	9
2.2	Jets in Galactic and extragalactic sources	12
2.3	Shocks in jets of AGN and MQs	15
3	Shocks in astrophysical plasmas	17
3.1	Collisionless plasmas	17
3.2	Classical description of a plasma shock wave	19
3.2.1	Rankine-Hugoniot jump conditions	20
3.2.2	MHD shock properties	22
3.2.3	De Hoffman-Teller frame and superluminality	26
3.2.4	Dissipation and the critical Mach number	28
3.2.5	The whistler critical Mach number	32
3.3	Relativistic shocks	33
3.3.1	Structure of a relativistic shock	33
3.3.2	Mildly relativistic shocks	35
4	Particle acceleration mechanisms in relativistic shocks	38
4.1	Diffusive shock acceleration (DSA) and its limitations	38
4.1.1	DSA in relativistic shocks	42
4.2	Synchrotron maser instability and wakefield acceleration	43
4.3	Electron trapping and acceleration in an oblique magnetosonic shock wave	44
4.3.1	Basic properties of oblique shocks	45
4.3.2	Electron reflection, trapping and acceleration	46
4.4	Ion acceleration by a magnetosonic pulse train	58
5	Particle-In-Cell method for plasma shock simulations	60
5.1	Kinetic description of a plasma	61
5.2	Particle-In-Cell (PIC) technique	62
5.2.1	Macroparticles	64
5.2.2	Integration of the relativistic equations of motion	64
5.2.3	Yee lattice	66
5.2.4	Current deposition and particle shape	68
5.2.5	Field weighting	70

5.3	Numerical limitations	70
5.3.1	CFL condition	71
5.3.2	Noise filtering and the Friedman filter	72
5.4	Generation of a shock	73
5.4.1	Boundary conditions	74
5.4.2	Initial tapering	75
5.5	New features of THATMPI	78
5.5.1	MPI subarrays	78
5.5.2	Parallel I/O	80
5.5.3	Relativistic particle injection	80
5.5.4	Improved particle tracing	82
5.6	Analysis tools of kinetic-based plasma simulations using neural networks	83
5.6.1	Analysis with momentum data	87
5.6.2	Analysis with E_x data	93
5.6.3	Conclusions on the traced data analysis with neural networks	95
6	PIC simulations of magnetised mildly relativistic shocks	97
6.1	Simulation setup	98
6.2	Superluminal shock (simulation $\theta 40$)	101
6.2.1	General structure of the shock transition	101
6.2.2	Ion and electron phase-space	101
6.2.3	Wave modes in the shock transition	104
6.2.4	Particle acceleration and heating in the magnetosonic pulse	109
6.2.5	Discussion on oblique superluminal shocks	111
6.3	Oblique subluminal shock (simulation $\theta 30$)	112
6.3.1	Electric potential in the magnetosonic wave	114
6.3.2	Late-stage evolution of the particle phase-space distributions	117
6.3.3	Electron trapping and acceleration in the magnetosonic shock	119
6.3.4	Electron spectra	123
6.3.5	Electron tracing	126
6.3.6	Ion trapping and acceleration	129
6.3.7	Estimates of the maximum particle energies	134
6.3.8	Discussion on oblique subluminal shocks	135
6.4	Quasi-parallel shock (simulation $\theta 10$)	136
6.4.1	Ion and electron phase-space	138
6.4.2	Upstream whistler waves	139
6.4.3	Nonlinear dynamics of large-amplitude precursor whistlers	149
6.4.4	Particle heating	157

6.4.5 Discussion on quasi-parallel shocks	167
7 Summary and discussion	169
A Comparison of the main simulations for different parameters	174
A.1 Simulation $\theta 40\text{-}\varphi_{B0}$	174
A.2 Simulation $\theta 40\text{-}1D$	177
A.3 Simulation $\theta 30\text{-}\varphi_{B0}$	179
A.4 Simulation $\theta 30\text{-}m400$	183
A.5 Simulation $\theta 10\text{-}\varphi_{B0}$	186
A.6 Simulation $\theta 10\text{-}1D$	189
B Wave polarisation convention	192
C Dispersion relation for a finite amplitude circularly polarised wave propagating along the magnetic field. Charge and current neutrality	193

CHAPTER 1

Introduction

COSMIC rays have been observed since the beginning of the twentieth century. They are highly energetic charged particles that travel across the Universe and originate from a wide variety of astrophysical sources. Some of the cosmic rays reach the Earth, where they are detected using telescopes and detectors. Their observed energies span an exceptionally broad range, from 1×10^9 eV up to 1×10^{21} eV (Bhattacharjee and Sigl 2000). Alongside the production of cosmic rays, high-energy photons and neutrinos are also produced. While cosmic rays do not point back to their sources due to deflection of their trajectories by the intergalactic and interstellar magnetic fields, photons and neutrinos propagate essentially undeflected and can therefore be used to constrain the physical processes responsible for production of such energetic particles.

A particular class of objects widely considered to be sources of high-energy cosmic rays are astrophysical jets. Jets are highly collimated outflows of matter propagating at relativistic speeds. They are observed in systems such as microquasars (MQs), active galactic nuclei (AGN), and gamma-ray bursts (GRB). Despite the large difference in physical scales between MQs, AGN, and GRBs, jet formation is thought to proceed through similar mechanisms in these systems. Intrinsic inhomogeneities in the jet plasma flow, as well as interactions between the jet and the surrounding medium, lead to the formation of shock waves within jets. These shocks are often associated with the production of high-energy cosmic rays.

Particle acceleration processes at shocks depend sensitively on the shock parameters, such as the shock Lorentz factor, γ_{sh} , the orientation of the background magnetic field with respect to the shock normal, θ_{Bn} , and the properties of the ambient plasma, including its magnetisation σ . The latter is defined as the ratio of the magnetic field energy density to the plasma bulk energy density (Romansky et al. 2024; Rassel et al. 2023; Khatoon et al. 2024). A widely considered particle energisation mechanism at shocks is *diffusive shock acceleration* (DSA), which is commonly invoked to explain the production of Galactic cosmic rays at non-relativistic shocks in supernova remnants. In the relativistic shock regime, DSA has been shown to be efficient only in unmagnetised or weakly magnetised shocks and is therefore unlikely to operate in astrophysical jets, which are expected to be at least moderately magnetised. Such shocks are typically mildly relativistic, with Lorentz factors $2 \lesssim \gamma_{sh} \lesssim 5$. Wakefield acceleration, which is efficient in magnetised ultra-relativistic shocks (Lyubarsky 2006; Hoshino 2008; Iwamoto et al. 2019), is ineffective in the mildly relativistic quasi-perpendicular regime (Ligorini et al. 2021a). This motivates the exploration of alternative acceleration processes. The goal of this thesis is to investigate the physics of mildly relativistic oblique shocks in strongly magnetised plasmas with $\sigma = 1$.

The main research method used in this study is fully kinetic particle-in-cell (PIC) numerical simulations. Large-scale 2D3V (two spatial dimensions and three velocity components) PIC simulations are employed to analyse mildly relativistic magnetised shocks for three configurations of an oblique background magnetic field:

- an oblique superluminal shock,

- an oblique subluminal shock with the magnetic-field inclination close to the threshold angle between the superluminal and subluminal configurations,
- and a quasi-parallel shock.

In addition to studying mildly relativistic magnetised shocks, this thesis introduces a novel approach to particle tracing analysis in kinetic plasma simulations based on neural networks. Two architectures – the *multi-layer perceptron* and the *convolutional neural network* – are applied to traced particle data using classification, regression, and anomaly detection techniques. The methods are demonstrated on a test case in which electron pre-acceleration at electrostatic waves produced by the Buneman instability is studied. The results show that neural networks have the potential to offer a fast and reliable tool for particle tracing analysis in kinetic plasma simulations.

The thesis is organised into seven chapters, including this [Introduction](#). The subsequent chapters are structured as follows:

- [Chapter 2](#) reviews the origin of cosmic rays and shocks in jets of AGN and MQs.
- [Chapter 3](#) introduces the concept and properties of shocks in astrophysical collisionless plasmas, with an emphasis on relativistic shocks.
- [Chapter 4](#) discusses particle acceleration mechanisms in relativistic shocks, emphasising a trapping and acceleration process relevant to the results of this thesis.
- [Chapter 5](#) outlines the PIC method, describes new features developed for the proprietary code THATMPI, and presents neural-network-based analysis tools for kinetic simulations.
- [Chapter 6](#) presents PIC simulation results for strongly magnetised, mildly relativistic shocks with the three magnetic-field obliquities.
- [Chapter 7](#) contains a summary and discussion of the results presented in [Chapter 6](#).

CHAPTER 2

Cosmic ray origin

COSMIC rays (CRs) were first discovered more than a century ago by Victor Hess (Hess 2018). It is now well established that CRs are high-energy charged particles arriving at Earth from a wide range of sources – from Galactic objects such as supernova remnants (SNRs), microquasars, and pulsar wind nebulae (PWNe), to extragalactic sources including active galactic nuclei (AGN), gamma-ray bursts (GRBs), and quasars. Nevertheless, despite decades of research since their discovery, the actual processes that accelerate CRs remain an open question. Since then, considerable effort is dedicated into understanding both the physical mechanisms capable of accelerating particles to energies beyond 10^{21} eV and the astrophysical environments in which these processes occur.

Victor Hess first detected cosmic rays during a balloon flight. Today, CRs are observed using the most advanced technologies in telescopes and particle detectors. In general, cosmic-ray (CR) detection methods can be divided into two categories: satellite-based and ground-based observations. Satellite-borne instruments are capable of directly detecting cosmic rays, but they are limited to relatively low energies. At higher energies, the required detector sizes become prohibitively large, leading to a dramatic increase in cost. Notable examples of satellite experiments include the Alpha Magnetic Spectrometer (AMS, Kounine 2012) and the Advanced Composition Explorer (ACE, Stone et al. 1998). Both operate in the low-energy CR regime and focus on specific particle populations: AMS is primarily dedicated to the study of antimatter cosmic rays, while ACE mainly investigates the solar wind and its associated energetic particles.

Ground-based telescopes, on the other hand, cannot directly detect cosmic rays, as most primary CRs interact with Earth's atmosphere before reaching the ground. The interaction of a primary CR with air molecules triggers a cascade of nuclear reactions that produce secondary high-energy particles. These, in turn, undergo further interactions, generating additional particles in a chain reaction known as an *extensive air shower* or *cosmic-ray shower*. The process continues until the produced secondary particles have sufficiently low energy for them to reach the ground, where the shower can be detected.

The characteristics of the shower depend strongly on the nature of the primary particle. If the primary particle is an electron or a γ -ray photon, the cascade consists mainly of electrons, positrons, and photons produced via bremsstrahlung and pair-creation processes (Gaisser 1990). In contrast, primary protons – being non-elementary particles – produce more complex and broader cascades containing both a hadronic and an electromagnetic component, the latter being similar to the purely electromagnetic showers described above. The hadronic component is dominated by pion production, followed by their decay into neutrons, muons, and other leptons. At higher energies, heavier mesons can also be produced, increasing the complexity of the shower. Similarly, heavier ions such as iron, yield even more complicated cascades, making the composition of the original CR more difficult to infer.

A wide range of ground-based detectors are currently in operation. For instance, the High Altitude Water Cherenkov (HAWC, Abeysekara et al. 2023) observatory, operating over the

energy range of 100 GeV to 50 TeV, utilises large water tanks with surrounding photomultiplier tubes to detect Cherenkov light produced by secondary particles in water.

Understanding the origin of cosmic rays does not necessarily rely on the direct or indirect detection of high-energy charged particles such as ions and electrons. It is known that the production of these charged particles is characterised by photon and neutrino emission. High-energy charged particles emitted from a source can interact with the surrounding environment, producing secondary particles – including photons, neutrinos, and pions – or exchanging energy through various processes.

The most common mechanisms responsible for high-energy photon emission include inverse Compton scattering, in which a highly energetic CR particle transfers part of its kinetic energy to a low-energy photon, and synchrotron radiation, where relativistic charged particles (typically electrons) emit radiation as they gyrate in the magnetic fields. Because photons and neutrinos have no charge, their trajectories are not deflected by magnetic fields, and they therefore point approximately back to their sources. Photons, however, can still interact with the medium, leading to energy losses and scattering, although the relevant cross-sections decrease with increasing photon energy. Neutrinos, instead, interact only very weakly with matter and thus propagate virtually unimpeded through space.

Protons and heavier nuclei can produce both photons and neutrinos when interacting with ambient material near the source. However, γ -rays can also originate from inverse Compton scattering of ultra-relativistic electrons. Consequently, the total γ -ray emission alone cannot fully constrain the CR content of the source. Neutrinos, on the other hand, are generated exclusively in hadronic interactions and are therefore directly linked to the hadronic component of the γ -ray emission. This connection provides valuable constraints on both the energy of the produced CRs and the nature of the sources that generated them.

The detection of γ -ray photons and neutrinos utilises the same principle as cosmic ray detections by using the Cherenkov light emitted when they interact with the medium. γ -ray observatories include, for instance, MAGIC (Baixeras et al. 2004), HESS (Hinton and HESS Collaboration 2004), VERITAS (Holder 2011) and SST-1M (Alispach et al. 2025), amongst others, which cover an energy bandwidth spanning GeV up to tens of TeV. Future observatories, such as the Cherenkov Telescope Array Observatory (CTAO, Wild and Ferrini 2022), will observe even higher energies, the former having an energy range from tens of GeV up to 300 TeV. Neutrino detectors must use instead a very dense target medium such as water or ice due to the extremely low cross-section for neutrinos. Examples of high-energy neutrino detectors are the IceCube Neutrino Observatory (Halzen and Klein 2010), that uses the Antarctic ice as its detection medium, and Super-Kamiokande (Kajita 1999), which relies on a large water tank (doped with gadolinium) to enhance neutrino interaction and detection.

A noticeable recent example is the detection of the flaring blazar TXS 0506+056, an active galactic nucleus whose jets point directly towards Earth. The event was first observed by IceCube (IceCube Collaboration et al. 2018) and subsequently confirmed in the γ -ray band by the Fermi-LAT satellite (Paiano et al. 2018). The coincident detection of neutrinos and γ -rays provides compelling evidence for the presence of ultra-high-energy protons within the blazar, likely accelerated by mechanisms such as shock acceleration.

2.1 Cosmic ray spectrum and sources

ONE of the most important clues for understanding the origin of CRs is their observed energy spectrum (Figure 2.1). Cosmic rays span an enormous range of energies, from about 1×10^9 eV up to nearly 1×10^{21} eV. CRs with energies below 1×10^{12} eV are affected by solar modulation, which varies with the Sun's ~ 22 -year activity cycle (Grieder 2001). Between 1×10^{12} eV and 1×10^{15} eV, the spectrum follows a power-law distribution, $\partial N/\partial E \propto E^{-\alpha}$, with a spectral index $\alpha \approx 2.7$. At around $E \sim 1 \times 10^{15}$ eV, the spectrum exhibits a smooth steepening, known as the "knee", where the index changes from $\alpha = 2.7$ to 3.0. A second steepening occurs near 5×10^{17} eV ($\alpha = 3.3$), followed by a hardening, the "ankle", around 5×10^{18} eV, where the index flattens again to $\alpha \approx 2.7$. At the highest energies, above 5×10^{19} eV, the flux drops sharply, forming an exponential cutoff known as the *GZK cutoff* (Greisen 1966; Zatsepin and Kuz'min 1966).

The changes in spectral slope are thought to be correlated with the composition of the particles and the astrophysical sources that produce cosmic rays. The current paradigm states that CRs below the knee ($E \lesssim 1 \times 10^{15}$ eV) are mainly composed of protons and originate predominantly from Galactic sources, primarily supernova remnants (SNRs) (Blasi 2013). Supernovae result either from the violent collapse of massive stars that have reached their fusion limit, or from the explosion of a white dwarf in a binary system due to the matter accretion from the companion star. Each explosion releases a vast amount of energy, with a total output of about 1×10^{63} eV (Bethe 1990). The ejected matter expands into the surrounding medium, forming a shell known as the *supernova remnant*.

The hypothesis that proposes SNRs as the main accelerators of Galactic CRs was first proposed by Baade and Zwicky (1934). The produced CRs can remain confined within the Galaxy by its magnetic field before they eventually escape into the intergalactic space. The average galactic CR energy production rate is estimated to be $\sim 1 \times 10^{52}$ eV s $^{-1}$ (Ginzburg and Syrovatsky 1961). From historical observations and models of observability, the rate of supernovae in the Milky Way is estimated at about 4.6 per century (Adams et al. 2013). This implies an average energy production rate for supernovae (SNe) of about 1×10^{51} eV s $^{-1}$, meaning that about 10% of the kinetic energy released by the SNe must be converted into CR energy.

The interaction between the supersonic expanding dense SNR shell and the surrounding interstellar medium generates non-relativistic shocks that typically propagate at ~ 3000 km s $^{-1}$ (about 1% of the speed of light) and can reach speeds of about 1×10^5 km s $^{-1}$. Particle acceleration is expected to occur at these shocks. In particular, diffusive shock acceleration (DSA) is expected to be the primary mechanism (see Section 4.1). Observations of SNRs show strong emission across the entire electromagnetic spectrum, from radio waves to X-rays, and several have even been detected in the TeV γ -ray band (Batzofin et al. 2024).

Besides SNRs, a range of other galactic objects can produce outflows strong enough to drive shocks. They include star-forming regions (Cao et al. 2025), pulsar wind nebulae (Abdo et al. 2013), the supermassive black hole Sagittarius A* at the Galactic Centre (HESS Collaboration et al. 2016), and microquasars. In particular, the latter are known to emit




Figure 2.1: The current cosmic ray spectrum. Shown are the measurements of the intensity dJ/dE of charged and neutral CRs, multiplied by the squared kinetic energy. The coloured measurements show the spectrum for different CR components while the black values represent the all-particle spectrum. The diagonal lines that divide the figure in shades of grey show the energy-integrated intensities. Figure reproduced from Navas et al. 2024, under Creative Commons license CC BY.

intense X-ray and γ -ray radiation, associated with particle acceleration in shocks formed within their jets (Fabrika 2004; Bromberg et al. 2011a).

The reason cosmic rays below the "knee" ($E \lesssim 1 \times 10^{15}$ eV) are considered to be of Galactic origin is the Milky Way's magnetic field. Cosmic rays produced by various Galactic sources propagate through the interstellar medium and some of them eventually arrive on Earth. CR propagation in the Galaxy is predominantly diffusive, and their interaction with the Galactic magnetic field acts to confine them within the Galaxy (Cesarsky 1980; Ptuskin et al. 1993; Strong et al. 2007). Given the average Galactic field strength of $B \approx 3 \mu\text{G}$ – $4 \mu\text{G}$ (Grasso and Rubinstein 2001), high-energy CRs with $E \lesssim 1 \times 10^{15}$ eV have enough time and distance to isotropise throughout the Galaxy. Consequently, the high-energy CR

spectrum observed from Earth appears isotropised. The simplest models of CR diffusion predict a spectral steepening of approximately $\alpha \simeq 0.7$ with respect to the spectral index at sources (Mollerach and Roulet 2018; Mueller et al. 1991). Assuming that CRs are primarily accelerated via DSA – which produces a power-law spectrum with $\alpha_{\text{DSA}} \approx -2$ (see Section 4.1) – the resulting observed spectrum becomes:

$$N(E)dE \propto E^{-(\alpha_{\text{DSA}} + \alpha_{\text{diff}})} = E^{-2.7}, \quad (2.1)$$

which is consistent with the observed spectral index, softening from $\alpha \approx 2.6$ to $\alpha \approx 2.85$ (Collaboration et al. 2019).

As previously mentioned, CRs below the knee consist mainly of protons, but a small fraction of heavier nuclei is also present in the total spectrum (see Figure 2.1). These heavy ion components cannot be fully explained by galactic confinement alone (Kampert and Unger 2012; Engelmann et al. 1990). For instance, if an iron nucleus (the heaviest stable CR component) were accelerated by DSA, it could reach energies up to $E \sim 1 \times 10^{17}$ eV. Hence, the cutoff energy due to confinement depends on the CR composition and the superposition of all contributions from different species, giving rise to the smooth "knee"-like feature observed in the spectrum (Hörandel 2004; Blasi and Amato 2012).

More recently, Alfaro et al. 2025 identified a subtle knee-like feature at $E \approx 40$ TeV, where the spectral index changes slightly from $\alpha \approx 2.5$ to $\alpha \approx 2.7$. The origin of this feature remains unknown, though it has been argued that either it represents the maximum energy achievable by a dominant nearby source (Liu et al. 2019) or is instead a cutoff that depends on the rigidity that only occurs in certain classes of sources, most likely SNRs (Zatsepin and Sokolskaya 2007; Yue et al. 2019).

Cosmic rays with even higher energies, $E \gtrsim 1 \times 10^{18}$ eV – commonly referred to as *ultra-high-energy cosmic rays* (UHECRs) – are too energetic to be confined by the Galactic magnetic fields and eventually escape the Milky Way. Consequently, the most energetic cosmic rays are believed to be of extragalactic origin (Axford 1994). A wide range of extragalactic sources have been proposed, including AGN (Netzer 2015), GRBs (Vietri 1995), and clusters of galaxies (Lindner et al. 2014; van Weeren et al. 2010). GRBs, in particular, encompass a variety of phenomena, from hypernovae (Hjorth et al. 2003; Paczyński 1998) to mergers between two neutron stars, or between a neutron star and a black hole (Abbott et al. 2017). Beyond the "ankle" of the cosmic ray spectrum, the number of detected particles decreases significantly, resulting in poor statistics. These cosmic rays are exclusively linked to extragalactic sources, although their exact origin remains uncertain.

Most of the sources discussed above exhibit powerful outflows associated with the observed cosmic ray emission and high-energy radiation. Of particular interest are jets, as their formation mechanisms can operate on both Galactic and extragalactic scales, despite the intrinsic differences between the underlying sources.

2.2 Jets in Galactic and extragalactic sources

FIRST evidence of jet-like structures dates back more than a century, when Curtis (1918) observed the optical jet from the elliptical galaxy M87, revealing what is now known as the AGN. Decades later, similar structures were identified on much smaller scales in rare Galactic sources known as microquasars, with SS433 being the first one ever discovered (Spencer 1979). Both AGN and microquasars contain compact objects capable of launching relativistic matter into highly collimated outflows (jets), ranging from mildly relativistic flows with $\gamma \gtrsim 1$ to ultra-relativistic jets with Lorentz factors of the order of 10 – 10^2 or even 10^3 .

AGN and microquasars, which exhibit relativistic outflows, are characterised by intense X-ray and γ -ray emission associated with synchrotron and inverse Compton radiation (Blandford et al. 2019; H.E.S.S. Collaboration* et al. 2024). These observations imply that such systems must be capable of efficiently accelerate electrons through mechanisms involving either plasma shocks, magnetic turbulence, or magnetic reconnection.

In the last decades, new Galactic sources have been found to produce collimated ejecta within stellar environments, such as very young massive stars (Reipurth et al. 1997) and accreting white dwarfs (Motch 1998). However, these sources generate non-relativistic outflows with broader opening angles that are predominantly associated with thermal emission.

A common feature among all jet-producing objects is the presence of an accretion disc surrounding the central compact object, which appears to play a crucial role in jet formation. The current understanding of jet production suggests that compact objects - either spinning black holes or neutron stars - accrete matter that settles into the equatorial plane, forming an *accretion disc* (Pringle 1981; Bardeen and Petterson 1975). The magnetic field of the compact object extracts angular momentum from the disc and drags part of the accreted material along the magnetic field lines, which are then ejected in the form of a jet perpendicular to the disc.

Essentially, there are two accretion modes typically associated with jet production around black holes: the standard and normal evolution (SANE) mode, which produces thin discs and narrow jets, and the magnetically arrested disc (MAD) mode, which generates thicker discs and broader, more powerful jets. The latter is generally favoured for AGN, as it can efficiently extract the rotational energy of the black hole to power strongly magnetised jets (McKinney et al. 2012; Tchekhovskoy et al. 2011).

Despite the similarities in their jet formation mechanisms, MQs and AGN differ fundamentally in both their physical nature and characteristic scales. *Microquasars* (MQs) are binary systems composed of a compact object – either a black hole or a neutron star – and a companion (donor) star. The compact object accretes matter from the donor star, forming an accretion disc. Simultaneously, a pair of anti-parallel, highly collimated jets are launched along the rotational axis of the binary system, as illustrated in Figure 2.2 (Escobar et al. 2021). All MQs are a subclass of *X-ray binaries*, but only those systems that exhibit relativistic jets are classified as microquasars.

The magnetic field structure in MQ jets can be inferred from polarimetric studies of their observed X-ray emission, which originates from high-energy electrons gyrating in a partially




Figure 2.2: Sketch of a microquasar. A compact object revolves around a massive star whose wind is accreted by the black hole, producing a pair of jets in opposite directions.

ordered magnetic field. Kaaret et al. 2024a found, for the γ -ray binary PSR B125963, that spectropolarimetric data reveals a strong toroidal magnetic field component at the location of the termination shock. This suggests that both particle acceleration and X-ray emission occur at a quasi-perpendicular shock. Similarly, Kaaret et al. 2024b reported that, in the case of SS433, the magnetic field lines are aligned with the jet bulk flow.

Microquasars have been observed in the γ -ray band up to tens – and in some cases hundreds – of TeV (H.E.S.S. Collaboration* et al. 2024; Neronov et al. 2025). The highest-energy emission appears closer to the MQ, where the outflow speed is greatest, while lower energies are found farther out from the binary as the outflow decelerates. Such emission is often modelled as synchrotron or inverse-Compton radiation from relativistic electrons, implying that some kind of mechanism, most likely associated with shocks, is capable of accelerating electrons efficiently.

Microquasars are, however, rare objects in the Galaxy compared to SNRs due to the specific conditions required for their formation. Fewer than twenty have been identified, and some remain unconfirmed. Among the confirmed MQs, all contain black holes as the compact object. Although X-ray binaries with neutron stars are well known (Cowley et al. 2003), none have yet been observed to produce jets. If neutron-star microquasars exist,


their magnetic fields must be relatively weak (Massi 2007). The most powerful MQs known are SS 433 (Abeysekara et al. 2018; LHAASO Collaboration 2024a; H.E.S.S. Collaboration* et al. 2024) and V4641 Sgr (Alfaro et al. 2024; LHAASO Collaboration 2024b), with the maximum energy observed of ~ 100 TeV and ~ 200 TeV, respectively. Such observations show emission from extended sources which correlates to a pair of anti-parallel jets while the X-ray binary is observed between this extended emission.

An *active galactic nucleus* (AGN) is a galaxy containing a supermassive accreting black hole. The AGN luminosity far exceeds that of its host galaxy. The central black hole is surrounded by a dense, thin accretion disc, that is itself surrounded by an optically thick, torus-shaped cloud of gas and dust that obscures visible light. On larger scales, ionised gas envelops the entire structure (Netzer 2015; Urry and Padovani 1995).

In radio-loud AGN, a distinctive feature is the presence of a pair of powerful relativistic jets that emerge from the vicinity of the black hole in approximately opposite direction. These jets can extend over hundreds of kiloparsecs, interacting with the intergalactic medium and forming large radio lobes. Figure 2.3 shows *Cygnus A*, an AGN at a redshift of approximately $z \sim 0.06$. The jets span several kiloparsecs, dwarfing the host galaxy, which appears as the bright blue point at the centre of the image.

Cosmic rays originating in AGN propagate through the intergalactic medium in a diffusive manner, analogous to Galactic cosmic rays, but under much weaker magnetic fields. The intergalactic magnetic field strength is estimated to be within the range $10^{-12} \text{ G} < B_{\text{IGM}} < 10^{-7} \text{ G}$ (Adams et al. 1997). Consequently, trajectories of any high-energy particles produced in the vicinity of a central engine or within jets are lightly deviated by the magnetic field and

Figure 2.3: Composite image of the AGN *Cygnus A* in optical (gold), X-ray (red) and radio (blue) wavelengths. The AGN looks like a regular galaxy when observed with an optical telescope, but the full structure appears in radio and X-rays. Radio emission traces most of the extended features, while X-ray emission is primarily associated with the central black hole and the hot spots. Figure reproduced from DePasquale et al. 2015, under Creative Commons license CC BY. The individual images are credited to: X-ray: NASA/CXC/SAO; Optical: NASA/STScI; Radio: NSF/NRAO/AUI/VLA



can travel vast distances – tens of millions of light years – before reaching the Milky Way. Upon entering the Galaxy, the stronger interstellar magnetic field dominates their motion.

AGN emit radiation across the entire electromagnetic spectrum. The accretion disc radiates predominantly in the ultraviolet and soft X-ray bands as a result of viscous dissipation and turbulence within the disc. In the immediate vicinity of the black hole, hard X-ray emission is also observed, although its precise origin remains uncertain. Strong optical and ultraviolet emission arises from the innermost gas clouds surrounding the black hole, but this radiation may be partially or fully obscured depending on the observer’s line of sight relative to the optically thick dusty torus (Ulrich et al. 1997).

The jets of AGN produce intense radiation from the radio to the γ -ray band. The plasma within these jets moves at relativistic speeds, beaming radiation along the direction of motion; as a result, the observed intensity is strongly angle-dependent. Due to the axisymmetric structure of AGN, one jet may appear brighter than the other, sometimes giving the impression of a single jet. In particular, *blazars* are AGN whose jets are oriented close to the line of sight, appearing extraordinarily bright due to relativistic beaming. The dominant mechanisms for the emitted radiation in AGN jets are synchrotron emission and inverse Compton scattering. Multiwavelength observations of numerous AGN reveal a stable synchrotron component, typically in the infrared or visible bands – redshifted from high-energy bands – which is observed to be highly polarised (Böttcher and Baring 2019; Thum et al. 2018), as well as a variable component generally attributed to inverse Compton processes during strong accretion episodes (Laing and Bridle 2013).

2.3 Shocks in jets of AGN and MQs

SHOCKS are among the most commonly invoked scenarios for explaining the acceleration mechanisms responsible for cosmic-ray production. Because astrophysical jets can attain very high velocities, shocks are expected to form naturally when the outflow interacts with the surrounding medium, whose properties vary significantly depending on the source.

In MQ jets, shocks are typically non-relativistic to mildly relativistic, and they can give rise to physical processes that are not present in SNR shocks. Under relativistic conditions, DSA becomes highly sensitive to the properties of the upstream plasma and may not be the dominant mechanism responsible for energising the electrons that produce the observed emission (Pohl et al. 2020, see also Section 4). It is therefore important to consider alternative acceleration processes that may operate in this regime, such as magnetic reconnection or magnetised shocks in mildly relativistic conditions – the latter being the primary focus of this thesis.

Shock formation in microquasar jets is considerably more complex than in SNRs, and multiple types of shocks can develop along the outflow. Initially, the jet expands freely at mildly relativistic speeds, pushing against the ambient medium. Once the jet ram pressure balances that of its surroundings, a recollimation shock forms (Bromberg et al. 2011b), which, as the name suggests, acts to refocus and collimate the jet. In addition, the interaction between the unshocked jet boundaries and the surrounding medium may generate a

shocked shear layer that encapsulates the jet. Farther downstream, the jet decelerates as the accumulated mass of the interstellar medium becomes comparable to the jet mass, leading to the formation of forward and reverse shocks propagating in opposite directions. Plasma processed by these shocks is heated and expands, inflating a cocoon around the jet. Finally, intrinsic inhomogeneities in the jet plasma can give rise to internal shocks (Miller-Jones et al. 2005), further contributing to the dynamical complexity of the system.

Shocks in AGN jets are expected to form at multiple locations along the outflow. Jets are not smooth, steady streams of plasma; instead, they are often highly turbulent. Such turbulence can produce dense plasma inhomogeneities or unsteady velocity fields within the jet, both of which can lead to shock formation (Blandford and Königl 1979). In addition, depending on the black hole accretion rate, the jet may consist of discrete plasma shells rather than a continuous flow, with collisions between these shells generating internal shocks.

At larger distances from the central engine, bright radio lobes form where the jet interacts with the intergalactic medium. In this region, a forward shock propagates into the ambient medium while a reverse shock travels back into the jet. These interactions give rise to intense emission regions commonly referred to as *hot spots*, which emit X-rays via inverse Compton scattering by particles accelerated at the shocks (Kataoka and Stawarz 2005; O’Dea and Saikia 2021; Pereira-Santaella et al. 2022). Despite the high bulk Lorentz factors of AGN jets – typically in the range $\gamma \sim 10$ to 100 – the shocks themselves generally propagate at mildly relativistic speeds.

Although the formation of a shock may appear conceptually simple, the physical processes governing its development are highly nonlinear. Interactions between charged particles and electromagnetic fields give rise to a rich variety of collective effects that determine the shock structure and drive particle heating and acceleration. Understanding these processes is therefore essential for interpreting the emission not only from AGN and micro-quasars, but from a wide range of astrophysical sources.

CHAPTER 3

Shocks in astrophysical plasmas

SHOCKS play a fundamental role in the study of high-energy processes and phenomena in the universe. They are found in a wide variety of space and astrophysical systems, ranging from interplanetary settings – such as the solar wind bow shock or the termination shock at the edge of the heliosphere – through Galactic environments including SNRs and relativistic plasma outflows in the vicinity of compact objects in binary systems, and extending to extragalactic systems such as jets of AGN and gamma-ray bursts, and shocks formed in the intragalactic medium during mergers of galaxy clusters. These shocks have long been associated with the generation of cosmic rays and high-energy emission, as evidenced by numerous observational studies (see Chapter 2). A common feature of most of space and astrophysical plasmas is their extremely low collisionality, due to very low particle densities and, in many cases, high temperatures. Particle-particle collisions are therefore insufficient to govern the plasma dynamics. Instead, collective electromagnetic interactions dominate, giving rise to a rich array of nonlinear processes. These processes not only mediate shock formation but also play a central role in particle heating and acceleration. The accelerated particles may eventually reach the Earth as CRs or contribute to the observed high-energy radiation through mechanisms such as synchrotron emission and inverse Compton scattering.

The purpose of this chapter is to introduce the concept of a collisionless plasma, describe the formation and dynamics of collisionless shocks, and discuss their physical consequences.

3.1 Collisionless plasmas

PLASMAS consist of charged particles which stay in global charge neutrality, so that

$$\sum_j q_j n_j = 0, \quad (3.1)$$

where q_j is the charge of species j and n_j its density. The most common type of plasma in the universe is the proton-electron plasma, as expected from the large abundance of hydrogen. However, plasmas of other compositions can also occur in some astrophysical environments. An example is pair plasma, consisting of electrons and positrons, which can be generated in the intense magnetic fields surrounding neutron stars (pulsars or magnetars) (Gaensler and Slane 2006) or during the operation of a central engine that leads to gamma-ray bursts. In astrophysical plasmas, particle densities are typically low enough that Coulomb collisions between particles are extremely rare. As a result, such plasmas are effectively *collisionless*. The degree of collisionality can be characterised by the mean free path of particles:

$$\lambda_c = \left(\frac{\bar{v}}{c}\right)^4 \frac{\epsilon_0}{n_e \sigma_T \ln \Lambda}, \quad (3.2)$$

where $\bar{v}$ is the mean speed of electrons, c is the speed of light, ϵ_0 the vacuum permittivity, n_e the electron number density, σ_T is the Thompson cross-section and $\ln \Lambda$ is the Coulomb logarithm, which accounts for the enhanced contribution of multiple small-angle deflections relative to rare large-angle collisions (Chandrasekhar 1943). As the density decreases, the mean free path increases and collisions become extremely rare. As an example, in the solar wind plasma, consisting of mostly protons and electrons, the electron mean free path is on the order of 10^8 km. Under such conditions, binary Coulomb collisions can be safely neglected, and the plasma dynamics are instead governed by collective interactions. These include particle-wave interactions, mediated through the Lorentz force, as well as wave-wave interactions. Both processes are inherently nonlinear.

To explain the effects of collective interactions, consider a test charged particle embedded in an isotropic plasma in charge equilibrium. Such a particle experiences the influence of the surrounding charges in the medium. If one considers a small sphere of radius r with the test particle located at its surface, the net charge enclosed within the sphere is not exactly zero. As a result, the test particle is effectively *screened*. If the radius of the sphere becomes larger, the probability of having the same number of positive and negative particles increases, and the net charge felt by the test particle correspondingly diminishes. This screening process defines a characteristic spatial scale, known as the *Debye length*, which is defined as the distance over which the screening effect decreases by a factor of natural exponential e :

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{n e^2}}, \quad (3.3)$$

where k_B denotes the Boltzmann's constant, T_e the electron temperature, n the plasma number density (assuming $n = n_e = n_i$) and e is the elementary charge. Since electrons respond more strongly to the electromagnetic fields due to their large charge-to-mass ratio (q/m) compared to ions, λ_D is most naturally expressed in terms of the electron temperature. A sphere of radius λ_D therefore encloses a number of electrons given by:

$$N_D = \frac{4\pi}{3} \lambda_D n_e. \quad (3.4)$$

In collisionless plasmas, there exists an inherent fundamental oscillation frequency that is a universal property of such systems. To illustrate this, consider an idealised, stationary plasma in which ions and electrons occupy identical positions, resulting in complete charge and current neutrality. If the position of an electron is perturbed by a small displacement δx , a restoring force appears due to the resulting charge imbalance. This perturbation generates an electric field that acts to return the electron towards its equilibrium position with a non-zero velocity. The electron thus starts to execute a periodic motion with a characteristic frequency

$$\omega_{p,e} = \sqrt{\frac{e^2 n_e}{\gamma_e m_e \epsilon_0}}, \quad (3.5)$$

called the *plasma frequency*, where γ_e is the Lorentz factor of the plasma's mean speed and m_e is the electron mass. The ion plasma frequency can be computed in the same way. Note that

the plasma frequency is Lorentz invariant and that $\omega_{p,e} \gg \omega_{p,i}$. The longitudinal oscillations arising from such periodic motion are referred to as *Langmuir waves*. Associated with the plasma frequency is a characteristic spatial scale known as the *skin depth*, which can be expressed as

$$\lambda_{s,j} = \frac{c}{\omega_{p,j}}, \quad (3.6)$$

where j represents a plasma species. This quantity represents the attenuation length of an electromagnetic wave in a plasma composed of the species j and is commonly used in kinetic-based simulations to characterise the kinetic scales associated with the particles.

An important parameter that characterises a plasma and helps determine its physical regime is the plasma beta, β_p , defined as the ratio of the thermal pressure to the magnetic pressure:

$$\beta_p = \frac{P_{T,i} + P_{T,e}}{P_B} = \frac{n_i k_B T_i + n_e k_B T_e}{B^2/2\mu_0} = \beta_i + \beta_e, \quad (3.7)$$

where μ_0 is the vacuum permeability. In equilibrium, $\beta_i = \beta_e$. A related quantity is the plasma *magnetisation* σ_j of a species j . It is defined here as the squared ratio of the cyclotron frequency Ω_j to the plasma frequency $\omega_{p,j}$. The total plasma magnetisation, σ , is then defined as the inverse sum of all particle species:

$$\sigma_j = \frac{\Omega_j^2}{\omega_{p,j}^2} = \frac{B^2/\mu_0}{\gamma_j n_j m_j c^2}, \quad \sigma^{-1} = \sum_j \sigma_j^{-1}. \quad (3.8)$$

It is important to note that the magnetisation parameter is not uniquely defined, and different authors adopt slightly different conventions. A commonly used alternative definition expresses the magnetisation as the ratio of magnetic pressure, $P_B = B^2/2\mu_0$, to the bulk plasma energy, $u_{p,j} = \gamma n_j m_j c^2$ (for a single species). With this choice, the magnetisation differs by a factor of 1/2 compared to the definition given in equation 3.8. Another definition, frequently employed in the context of ultra-relativistic shocks, replaces the magnetic pressure with the total electromagnetic energy density obtained from Poyntings theorem,

$$u_s = \frac{1}{2} \left(\epsilon_0 |\mathbf{E}|^2 + \frac{1}{\mu_0} |\mathbf{B}|^2 \right). \quad (3.9)$$

Since the motional electric field approaches the same order of magnitude as the magnetic field ($E \sim B$) in the limit $v \rightarrow c$, this formulation proves useful in simulations of ultra-relativistic shocks.

3.2 Classical description of a plasma shock wave

A *shock wave* (or simply a *shock*) is defined as an abrupt transition between two plasma states. This transition occurs when the shock propagates at a speed greater than the characteristic phase velocity of waves in the medium. The classical example is that of an acoustic shock in air, which occurs when the disturbance travels faster than the sound speed. In a

collisionless magnetised plasma, the relevant characteristic velocity may be the sound speed, the Alfvén speed, or the magnetosonic wave speed, depending on the physical conditions. The essential difference between a regular wave and a shock wave is that the formation of a shock is an irreversible process, and therefore entropy must increase across the shock transition. Nevertheless, there are certain quantities that must be conserved across the shock front. These conservation laws are expressed through the *Rankine-Hugoniot relations*, commonly referred to as the *jump conditions*.

3.2.1 Rankine-Hugoniot jump conditions

The jump conditions are the conservation laws of fluid dynamics (Landau and Lifshitz 1959) applied to discontinuities such as shocks. Assume two plasmas with different properties separated by a shock transition (see Figure 3.1, left). The *upstream plasma* (also called the unshocked plasma) is labelled with subindex 1 and the downstream plasma (also called the shocked plasma) is labelled with subindex 2. The plasma state on each side of the shock is characterised with several macroscopic parameters: the plasma pressure, P , density, ρ , temperature, T , and a background magnetic field, $\mathbf{B} = (B_x, B_y, 0)$, which is assumed here to lie in the xy -plane, without loss of generality. In the frame of reference of the shock ($v_{sh} = 0$), the upstream plasma flows with velocity v_1 , while the downstream plasma moves with velocity v_2 . By definition, $v_1 > v_2$. In a fluid, three fundamental quantities are conserved: mass, momentum, and energy. When a magnetic field is present, the conservation of magnetic flux must also be included. Together, these principles yield the following conservation equations (Mallick 2011):

$$n_1 u_1 = n_2 u_2 \quad (3.10)$$

$$w_1 u_1^2 + P_1 + \frac{B_1^2}{2\mu_0} - \frac{B_{1,x}^2}{\mu_0} = w_2 u_2^2 + P_2 + \frac{B_2^2}{2\mu_0} - \frac{B_{2,x}^2}{\mu_0} \quad (3.11)$$

$$w_1 u_{1,x} u_{1,y} + u_{1,x} \frac{B_{1,x} B_{1,y}}{\mu_0} = w_2 u_{2,x} u_{2,y} + u_{2,x} \frac{B_{2,x} B_{2,y}}{\mu_0} \quad (3.12)$$

$$w_1 \gamma_1 u_{1,x} + u_{1,x} \frac{B_1^2}{\mu_0} = w_2 \gamma_2 u_{2,x} + u_{2,x} \frac{B_2^2}{\mu_0} \quad (3.13)$$

$$- (\mathbf{B}_1 \cdot \mathbf{u}_1) \frac{B_{1,x}}{\mu_0} - (\mathbf{B}_2 \cdot \mathbf{u}_2) \frac{B_{2,x}}{\mu_0} \quad (3.14)$$

$$u_{1,x} B_{1,y} - u_{1,y} B_{1,x} = u_{2,x} B_{2,y} - u_{2,y} B_{2,x} \quad (3.15)$$

$$B_{1,x} = B_{2,x},$$

where $u_i = \gamma_i \beta_i$, $\beta_i = v_i/c$ is the normalised fluid four-velocity and w_i is the specific enthalpy defined as $w_i = n_i + p_i(\Gamma_{ad} - 1)/\Gamma_{ad}$, where Γ_{ad} is the adiabatic index of the plasma.

Equation 3.10 represents the mass flux conservation through the shock transition, $j_m = \rho u$, analogous to the current density. Note that the density is computed in the proper frame of reference of the fluid. When expressed in the laboratory frame, the density transforms as $n_l = \gamma_l n$, where γ_l is the Lorentz factor associated with the relative velocity between the proper frame and the laboratory frame. In this case, the continuity condition takes the form: $n_{l,1} v_1 = n_{l,2} v_2$.




Figure 3.1: Structure of an ideal shock wave in its rest frame (left). The magnetic field is compressed in the downstream region, perpendicular to the shock normal. After transformation to the de Hoffmann-Teller frame (right), the plasma velocity becomes aligned with the magnetic field, and the motional electric field vanishes.

Equation 3.11 is the momentum conservation in the direction parallel to the shock velocity. The first terms on each side of the equation correspond to the Reynolds stress and describe the change of fluid momentum. The second and third terms represent the fluid pressure and the magnetic pressure, respectively. The fourth terms correspond to the magnetic tension. While magnetic pressure acts to expand magnetic field lines and thereby reduce the fluid pressure, magnetic tension acts to straighten the magnetic field lines.

Equation 3.12 is analogous to Equation 3.11, but describes the conservation of momentum in the direction perpendicular to the shock velocity. Since there is no change of plasma parameters in this direction, only the off-diagonal terms involving $u_x u_y$ remain. This equation assumes the absence of shear, implying that the pressure tensor is diagonal, which is a valid approximation for most plasmas.

Equation 3.13 represents the conservation of energy. The first terms on both sides represent the energy flux carried by the particles, while the second and third terms are the expanded form of the electromagnetic energy flux: $-(\mathbf{u} \times \mathbf{B}) \times \mathbf{B} = B^2 \mathbf{u} - (\mathbf{B} \cdot \mathbf{u}) \mathbf{B}$. Only the x -component of the flow is considered, since the shock transition occurs in this direction.

Equation 3.14 describes the conservation of the magnetic flux ϕ_B . This equation can also be interpreted as the conservation of the transverse electric field, $\mathbf{E} = -\mathbf{u} \times \mathbf{B}$. Note that this is not the motional electric field, $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$, but rather a field derived from a Lorentz boost of the electric field:

$$E'_z = \gamma_i \left(E_z^0 + v_{i,x} B_y - v_{i,y} B_x \right) = ct, \quad (3.16)$$

where v_i , $i = 1, 2$ represents the bulk velocity of the upstream or downstream plasma, respectively.

Equation 3.15 directly follows from Gauss' law $\nabla \cdot \mathbf{B} = 0$.

The adiabatic index of the plasma, Γ_{ad} , which enters the definition of the enthalpy, depends in the fluid model on the dynamical regime under consideration and the number of available degrees of freedom, f , according to the general relation $\Gamma_{ad} = 1 + 2/f$. For a non-relativistic three-dimensional (3D) plasma, where each particle possesses three translational degrees of freedom, one obtains $\Gamma_{ad} = 5/3$. For an ultra-relativistic plasma the adiabatic index reduces to $\Gamma_{ad} = 4/3$, because relativistic particles store twice as much energy per degree of freedom as non-relativistic ones. Kinetic plasma simulations are often performed in two spatial dimensions, and the effective adiabatic index is modified accordingly because the number of available degrees of freedom differs from the fully three-dimensional case. One thus finds $\Gamma_{ad} = 2$ for the two-dimensional (2D) non-relativistic case and $\Gamma_{ad} = 3/2$ for the 2D ultra-relativistic case. However, the actual number of effective degrees of freedom in 2D simulations depends on the orientation of the large-scale magnetic field. For the field within the simulation plane, particle trajectories bend out of the plane, giving access to all three velocity components. Hence the relativistic adiabatic index in this case is $\Gamma_{ad} = 4/3$. For an out-of-plane field, particles are confined to the simulation plane and have two degrees of freedom.

It is important to note that the presence of accelerated particles at the shock can modify the adiabatic index, which is otherwise calculated under the assumption of a purely thermal particle distribution. Such modifications, in turn, affect the shock jump conditions, for example the density compression ratio (Stockem et al. 2012). In kinetic simulations, shocks may not be fully developed, or the shocked plasma may not be fully thermalised, so that the effective adiabatic index can take values between the non-relativistic and ultra-relativistic limits, with a considerable impact on the resulting jump conditions (Nishikawa et al. 2009).

3.2.2 MHD shock properties

Finding a solution for all six equations of the Rankine-Hugoniot (RH) jump conditions is difficult in the most general case. Therefore, it is convenient to consider limiting cases.

In the strong shock limit, the upstream thermal pressure is negligible compared to the bulk kinetic energy, and the enthalpy simplifies to $\omega_1 = n_1$. Both thermal and magnetic pressures are neglected in this approximation.

For an unmagnetised shock ($B_{1,2} = 0$), the shock speed in the downstream frame, $v_{sh} = -v_2$, can be computed by applying a Lorentz boost to Equations 3.10–3.13 (Stockem et al. 2012). This yields

$$\beta_{sh} = \frac{\Gamma_{ad,2} - 1}{\sqrt{\gamma_{1,2}^2 - 1}} \left(\gamma_{1,2} - \frac{n_1}{\omega_1} \right), \quad (3.17)$$

where $\beta_{sh} = v_{sh}/c$, $\Gamma_{ad,2}$ is the downstream adiabatic index, $\gamma_{1,2}$ is the upstream Lorentz factor as seen from the downstream, n_1 is the upstream density, and ω_1 its specific enthalpy. A higher upstream pressure, corresponding to a larger ω_1 , leads to a faster shock. In the strong shock limit, where the upstream thermal pressure is negligible (cold plasma, $\omega_1 = n_1$), the shock speed simplifies to $\beta_{sh} = (\Gamma_{ad,2} - 1) \sqrt{(\gamma_{1,2} - 1) / (\gamma_{1,2} + 1)}$ as derived in Spitkovsky 2008.

Shock Mach numbers

As mentioned above, a plasma shock must propagate faster than the sound speed, the Alfvén speed, or the fast magnetosonic speed. The simplest way to characterise the shock speed relative to a wave speed is through the *Mach number* M , defined as the ratio of the shock speed to the characteristic wave speed. By definition, a shock always has $M > 1$.

For a plasma, there are four different Mach numbers corresponding to four distinct wave speeds::

$$\text{Sonic Mach number: } M_s = \frac{v_{sh}}{c_s}, \quad c_s^2 = \frac{\Gamma_{ad,e} k_B T_e + \Gamma_{ad,i} k_B T_i}{m_i} \quad (3.18)$$

$$\text{Alfvénic Mach number: } M_A = \frac{v_{sh}}{v_A}, \quad v_A^2 = \frac{B^2}{\mu_0 \sum_j (n_j m_j)} \quad (3.19)$$

$$\text{Fast/Slow Mach number: } M_{f,s} = \frac{v_{sh}}{v_{\pm}}, \quad (3.20)$$

$$v_{\pm}^2 = \frac{1}{2} \left(v_A^2 + c_s^2 \pm \sqrt{(v_A^2 + c_s^2)^2 - 4v_A^2 c_s^2 \cos^2 \theta_{Bn}} \right)$$

where c_s is the sound speed, v_A is the Alfvén speed, $v_{\pm}$ is the fast or slow magnetosonic speed (depending on the sign used), θ_{Bn} is the obliquity angle of the magnetic field $\mathbf{B}$ with respect to the shock normal.

The sound speed in Equation 3.18 is given for a proton-electron ideal plasma. For a plasma in thermal equilibrium, $T_e = T_i = T$, the formula simplifies to $c_s^2 \simeq 2\Gamma_{ad} k_B T / m_i$. While the upstream plasma of a shock is typically in thermal equilibrium and this simplification is valid, the downstream plasma often has a much higher electron temperature than ion temperature, $T_e \gg T_i$, so that $c_s^2 = \Gamma_{ad,e} k_B T_e / m_i$. The Alfvén speed is derived from the dispersion relation in a cold plasma for a circularly polarised wave propagating parallel to the magnetic field in the limit of $\omega \ll \Omega_e$, where ω is the wave frequency and Ω_e is the electron cyclotron frequency (Kulsrud 2005, chapter 9). Finally, the magnetosonic speeds, both slow and fast, are obtained from the linearised MHD (Kulsrud 2005, chapter 5). Note that for a cold plasma, $c_s = 0$, thus the slow speed is zero while the fast speed equals the Alfvén speed; conversely, for an unmagnetised plasma ($B = 0$), the slow speed is also zero, but the fast speed equals the sound speed.

The above equations for the different speeds are computed in the limit of a non-relativistic shock speed. All four speeds differ if computed using a relativistic formulation. The relativistic Alfvén velocity is given by

$$v_{A,rel} = \frac{v_A}{\sqrt{1 + \frac{v_A^2}{c^2}}}, \quad (3.21)$$

and is derived in the same manner as the non-relativistic Alfvén speed. One can see that for a strong magnetic field, the non-relativistic Alfvén speed v_A could exceed the speed of light, whereas in the limit $v_A \ll c$, corresponding to a weakly magnetised or dense plasma, the relativistic Alfvén speed reduces to the classical non-relativistic value.

The sound speed can be computed from its definition

$$c_s^2 = \left(\frac{\partial P}{\partial \rho} \right)_s, \quad (3.22)$$

where the partial derivative of the pressure P with respect to the density ρ is evaluated at constant entropy density. However, in the relativistic regime, no general analytical expression exists (Coll et al. 2019; Faussurier 2024) and it can only be computed in specific limits. For instance, in the ultra-relativistic limit, where $\Gamma_{ad} = 4/3$, the sound speed approaches $c_s^2 \rightarrow c^2/3$.

The magnetosonic wave speed has a form analogous to the non-relativistic case, depending on both the sound speed c_s and the Alfvén speed v_A . Consequently, its computation in the relativistic regime carries the same challenges as for the sound speed (Komissarov 1999; Kirk and Duffy 1999a).

Compression ratio

A direct consequence of the jump relations is the increase in density downstream of the shock. In other words, the matter is highly *compressed* at the shock. This compression is also evident in the magnetic field, but only the tangential component is amplified. The ratio of downstream to upstream density, $r = n_2/n_1$ referred to as the *compression ratio*, can be derived from the jump conditions (Eqs. 3.10-3.15) in the downstream frame of reference. For an unmagnetised shock, the compression ratio is

$$r = \frac{n_2}{n_{1,2}} = 1 + \frac{\beta_{1,2}}{\beta_{sh}} = 1 + \frac{(\gamma_{1,2}^2 - 1) \omega_1}{\gamma_{1,2} (\Gamma_{ad,2} - 1) (\gamma_{1,2} \omega_1 - 1)}, \quad (3.23)$$

where $\beta_{1,2}$ is the upstream speed as seen from the downstream reference frame and ω_1 is the upstream specific enthalpy (Stockem et al. 2012). In the strong shock limit, this expression simplifies to

$$r = 1 + \frac{1 + \gamma_{1,2}^{-1}}{\Gamma_{ad,2} - 1}. \quad (3.24)$$

Table 3.1 presents the compression ratio in the limiting cases. It is useful to compute the compression ratio in both two and three dimensions because kinetic-scale plasma simulations often model shocks in reduced dimensionality due to computational constraints. The lower dimensionality affects the adiabatic index, which in turn modifies the compression ratio. Notably, for a given dimensionality, the compression ratio is identical in both the non-relativistic and ultra-relativistic limits.

Table 3.1: Compression ratio for 2D and 3D non-relativistic and ultra-relativistic plasmas.	Dims.	$\gamma_{1,2}$	$\Gamma_{ad,2}$	r
	3D	$\gtrsim 1$	$5/3$	4
	3D	$\rightarrow \infty$	$4/3$	4
	2D	$\gtrsim 1$	2	3
	2D	$\rightarrow \infty$	$3/2$	3

For weakly magnetised plasmas, the influence of the magnetic field on the shock dynamics is small, so that the compression ratio given by Equation 3.23 provides a good approximation. However, since the shock also compresses the magnetic field, part of the shock energy is expended on the field, causing the compression ratio to be lower than in the unmagnetised case. The compression ratio for the tangential magnetic field, r_B , depends on the density compression ratio, the Alfvén Mach number, and the magnetic field obliquity:

$$r_B = \frac{B_{2,t}}{B_{1,t}} = \frac{r (M_A^2 - \cos^2 \theta_{Bn})}{M_A^2 - r \cos^2 \theta_{Bn}}, \quad (3.25)$$

where B_t denotes the tangential component of the magnetic field. Note that $r_B > r$, and this remains true even for strong magnetic fields with $M_A \sim 1$ (Burgess and Scholer 2015, Chapter 2). The magnetic-field compression increases rapidly with M_A^2 in Equation 3.25, and in the strong shock limit ($M_A \gg 1$), $r_B \rightarrow r$.

There exists an analytical solution for the compression ratio in the non-relativistic regime, $v_{sh} \ll c$ (Burgess and Scholer 2015, Chapter 2), which accounts for the magnetic field; however, the full solution is cumbersome. Nonetheless, some interesting limits and general trends can be highlighted for a three-dimensional plasma:

- For $M_A \gg 1$, the compression ratio tends to $r \rightarrow 4$, corresponding to the first row in Table 3.1. Such high values are typically observed in supernova remnant shocks, where the Alfvénic Mach number can exceed 100. In contrast, the terrestrial bow shock, with $M_A \sim 5\text{--}10$, has a compression ratio slightly lower than $r = 4$.
- For $M_A \rightarrow 1$, the compression ratio scales approximately linearly with the Mach number, $r \propto M_A$. These low values occur at interplanetary shocks, where the magnetic field is dynamically important and the shocks are considered weak (Wilson et al. 2017).
- In quasi-parallel shocks, the tangential component of the magnetic field is small, so the energy spent on magnetic compression is negligible. These shocks can therefore be approximated hydrodynamically. Consequently, the compression ratio reaches higher values at low M_A and approaches $r \approx 4$ at high Mach numbers, as expected.
- As the obliquity increases, the compression ratio follows a similar trend as the quasi-parallel case, but the approach to $r \rightarrow 4$ with increasing M_A is slower. This is because the tangential magnetic-field component becomes more significant, increasing the energy required for field compression.
- The obliquity dependence described above applies strictly for cold plasmas, $\beta_p \rightarrow 0$. In hotter plasmas, this dependence becomes negligible, and all shocks exhibit similar trends regardless of obliquity. Analogous to the relativistic case, more energy is required to heat higher-energy particles, causing the compression ratio to increase more gradually towards $r = 4$ as the Mach number rises.

In the relativistic regime, an analytical solution for the compression ratio cannot be obtained (Kirk and Duffy 1999b). However, the ultra-relativistic limit is well-known:

- For $M_A \gg 1$ and $v_{sh} \rightarrow c$, the compression ratio tends to $r \rightarrow 4$, corresponding to the second row in Table 3.1.
- Since relativistic shocks require more energy to heat downstream relativistic particles, the compression ratio for intermediate values of M_A , but still ultra-relativistic, is slightly reduced compared to the ideal hydrodynamic expectation to conserve energy. Nevertheless, the dependence of the compression ratio on obliquity and plasma beta follows trends similar to the non-relativistic case.

The compression ratio is not Lorentz invariant, and therefore its value depends on the chosen frame of reference. In the upstream frame, it can be expressed as $r_{\text{up}} = r_{\text{dw}} \gamma_{\text{up}}^2$, where up and down refer to the upstream and downstream frames, respectively, and γ_{up} is the Lorentz factor of the upstream plasma as measured in the downstream frame. Furthermore, in high-Mach-number subluminal shocks, a small fraction of particles is reflected at the shock to prevent infinite steepening of the shock front (see next section). In this case, the compression ratio is evaluated using the values far upstream and far downstream, and its effective value is slightly reduced because some particles do not cross into the downstream region.

Shock potential

In an ion-electron plasma, an electrostatic potential, referred to as the *cross-shock potential*, develops due to charge separation at the shock front. The RH conservation relations (Section 3.2.1) describe the compression of the magnetic field as the shock propagates into the upstream medium. Upstream ions moving toward the shock interact with the compressed magnetic field. The sudden increase in magnetic field causes ions to gyrate with a relatively large gyroradius, while electrons, being lighter, gyrate with a much smaller radius. This difference in gyroradii leads to charge separation, which in turn generates an electrostatic potential across the shock which is governed by Gauss' law, $\nabla^2 \phi = -\rho/\epsilon_0$, assuming a stationary magnetic field, $\partial \mathbf{A}/\partial t = 0$, where $\mathbf{A}$ is the magnetic vector potential. For ions, this potential acts as a barrier, aiding their reflection upstream in combination with the compressed magnetic field. Electrons, on the other hand, are attracted by the electrostatic potential and gain energy. However, in strongly magnetised plasmas, $\sigma_e \gg 1$, the cross-shock potential may trap electrons within the shock, where they are accelerated. A more detailed analysis of this process is provided in Section 4.3.

3.2.3 De Hoffman-Teller frame and superluminality

In the presence of the background magnetic field, its obliquity, θ_{Bn} , with respect to the shock normal plays an important role. In general, if the magnetic field is nearly aligned with the shock normal, $\theta_{Bn} \sim 0^\circ$, the shocks are classified as *quasi-parallel*, whereas for a large inclination, $\theta_{Bn} \sim 90^\circ$, they are classified as *quasi-perpendicular*. Shocks are classified as *oblique* when their magnetic field orientation is neither purely parallel, $\theta_{Bn} = 0^\circ$, nor purely perpendicular, $\theta_{Bn} = 90^\circ$, to the shock normal. Since in a perpendicular and oblique shock the magnetic field forms an angle with respect to the direction of the upstream particle flow,

a motional electric field must appear. Because collisionless plasmas are perfectly conductive, i.e., they exhibit no resistivity, this motional electric field becomes

$$\mathbf{E} = -\mathbf{v} \times \mathbf{B}. \quad (3.26)$$

It is sometimes possible to move to a frame of reference in which the upstream velocity and magnetic field are aligned (see Figure 3.1, right). From the shock frame of reference $v_{\text{up}} = -v_{sh}$, a velocity transformation is performed with velocity $\mathbf{V}_{HT}$ so that the motional electric field disappears on both sides of the shock:

$$\mathbf{E} = -\mathbf{v}_{HT} \times \mathbf{B} = -(\mathbf{v} - \mathbf{V}_{HT}) \times \mathbf{B} = 0 \rightarrow \mathbf{v} \times \mathbf{B} = \mathbf{V}_{HT} \times \mathbf{B}, \quad (3.27)$$

where $\mathbf{v}$ is either the downstream or upstream velocity. This equation admits infinitely many solutions for $\mathbf{V}_{HT}$. To obtain a unique solution, an additional constraint must be imposed. A convenient choice is to perform a transformation that preserves the direction of the shock front, which is achieved by requiring $\hat{\mathbf{n}} \perp \mathbf{V}_{HT}$. In other words, the transformation is restricted to the shock plane. By taking the cross-product of $\hat{\mathbf{n}}$ on both sides of Equation 3.27 and using the upstream velocity, a unique solution for $\mathbf{V}_{HT}$ can be obtained:

$$\begin{aligned} \hat{\mathbf{n}} \times (\mathbf{v}_{\text{up}} \times \mathbf{B}) &= \hat{\mathbf{n}} \times (\mathbf{V}_{HT} \times \mathbf{B}) = \mathbf{V}_{HT} \cdot (\hat{\mathbf{n}} \cdot \mathbf{B}) - \mathbf{B} \cdot (\hat{\mathbf{n}} \cdot \mathbf{V}_{HT}) \xrightarrow{0} \\ \mathbf{V}_{HT} &= \frac{\hat{\mathbf{n}} \times (\mathbf{v}_{\text{up}} \times \mathbf{B})}{\hat{\mathbf{n}} \cdot \mathbf{B}} = -v_{\text{up}} \tan \theta_{Bn} \hat{\mathbf{b}}_{\perp} = v_{sh} \tan \theta_{Bn} \hat{\mathbf{b}}_{\perp}, \end{aligned} \quad (3.28)$$

where we used the fact that $\mathbf{a} \times (\mathbf{b} \times \mathbf{c}) = \mathbf{b}(\mathbf{a} \cdot \mathbf{c}) - \mathbf{c}(\mathbf{a} \cdot \mathbf{b})$ and $\hat{\mathbf{b}}_{\perp}$ is the perpendicular component of the magnetic field unit vector $\hat{\mathbf{b}} = \mathbf{B}/\|\mathbf{B}\|$. This frame of reference is called *de Hoffman-Teller frame* (dHT frame). The dHT frame can only exist for obliquity angles lower than a threshold angle, θ_{tr} , where $v_{sh} \tan \theta_{tr} = c$ (Krauss-Varban and Wu 1989). The threshold angle separates shocks into two categories depending on the upstream speed: if $v_{sh} \tan \theta_{Bn} < c$, the shock is referred to as *subluminal*, whereas if $v_{sh} \tan \theta_{Bn} \geq c$ the shock is *superluminal*. In a subluminal shock, particles can be reflected from the shock-compressed magnetic field and stream along the upstream magnetic field lines. If the shock obliquity is larger than the threshold angle, shock-reflected particles only half-gyrate in the shock and are ultimately advected downstream, as otherwise they would need to have speeds along the background field larger than c .

The distinction between subluminal and superluminal shocks can also be made in the upstream frame of reference by applying a Lorentz transformation to the obliquity angle. Particles can be reflected at the shock only if

$$v_{sh} < c \cos \theta_{Bn}, \quad (3.29)$$

and the threshold angle is

$$\theta_{tr} = \arccos\left(\frac{v_{sh}}{c}\right). \quad (3.30)$$

It is convenient to compute the threshold angle in the downstream frame of reference,

$$\cot \theta_{tr} = \gamma_{sh} (\beta_{sh} - \beta_{\text{up}}), \quad \beta_{sh} > 0, \beta_{\text{up}} < 0, \quad (3.31)$$

where $\beta_{sh} = v_{sh}/c$ and $\beta_{up}v_{up}/c$ are the normalised shock speed and upstream speed, respectively (Sironi and Spitkovsky 2011). This expression is derived from the definition of the threshold angle in the upstream rest frame, and applying a Lorentz boost β_{up} .

3.2.4 Dissipation and the critical Mach number

As explained at the beginning of Section 3.2, the formation of a shock is an irreversible process, leading to an increase in entropy. This results in heating of the affected medium, in this case the downstream plasma. In collisional plasmas, Coulomb collisions provide the primary mechanism for dissipation and heating. However, most astrophysical plasmas are collisionless, so an alternative dissipation mechanism must exist to convert the upstream plasma energy into downstream thermal energy.

For low values of the Alfvénic Mach number, the shock primarily dissipates energy through resistive processes and the intrinsic dispersive properties of the plasma:

- Resistive dissipation considers a resistive electric field given by Ohm's law $\mathbf{E} = \rho \mathbf{J}$, where ρ is the tensor resistivity. The energy exchange between particles and fields is governed by Poynting's theorem:

$$-\frac{\partial u}{\partial t} = \nabla \cdot \mathbf{S} + \mathbf{J} \cdot \mathbf{E} \quad (3.32)$$

where u is the system's internal energy density and $\mathbf{S}$ is the Poynting flux. The term $\mathbf{J} \cdot \mathbf{E}$ represents the work done by the electric field on charged particles, increasing their kinetic energy and heating the plasma. In steady state ($\partial/\partial t = 0$), the energy input from electromagnetic fields balances the heating: $-\nabla \cdot \mathbf{S} = \mathbf{J} \cdot \mathbf{E}$. This type of dissipation is also referred to as *Ohmic heating*.

- Plasma waves can exhibit dispersive behaviour, meaning that the phase and group velocities differ. In the MHD limit of long-wavelength waves, the waves are nearly dispersionless ($v_p = v_g$), and energy is not dissipated since the waves travel with the bulk plasma. However, below the ion kinetic scales (ion skin depth or gyroradius), dispersive behaviour becomes relevant and short wavelength waves can travel faster or slower than long-wavelength waves depending on their nature. If these short-wavelength waves, generated at the shock, are not fully damped, they may propagate upstream faster than the shock, effectively carrying energy away and dissipating it from the shock. Whistler waves are an example of such dispersive waves and can escape upstream if their group velocity exceeds the shock speed (see Section 3.2.5).

However, low Mach number shocks have a limit for which resistive dissipation is not enough to prevent the shock from steepening ad infinitum. Above this limit, shock have to dissipate energy by other means. This limit can be understood in terms of the Mach number and is called *critical Mach number*, M_{cr} .

Consider a fast magnetosonic wave located downstream and propagating toward the shock with a speed given by Equation 3.20 in the shock frame. Note that $v_+ > c_s$. Owing to the nonlinear nature of the wave, it gradually steepens and slows down. In principle,

the dissipation mechanisms discussed above prevent the wave from steepening indefinitely. As a result of the balance between nonlinear steepening and dissipation, the wave speed decreases until it approaches the sound speed c_s . At this stage, the subsequent evolution of the shock depends on the downstream flow speed. If the sound speed is smaller than the downstream speed (low Mach number), the steepened wave does not reach the shock front and therefore does not contribute to further shock steepening. If, instead, the sound speed exceeds the downstream speed (higher Mach number), the wave reaches the shock front and continues to steepen it. To avoid this unphysical situation, an additional dissipation mechanism – distinct from those described above – must operate to halt further steepening.

Shocks with Mach numbers below a critical value M_{cr} are termed *subcritical*, whereas shocks with Mach numbers above M_{cr} are termed *supercritical*. An equivalent distinction can be made based on the dominant dissipation mechanism that limits shock steepening: subcritical shocks are also referred to as *resistive shocks*, while supercritical shocks are termed *dispersive shocks*, at least in the case of weak shocks.

The dissipation processes at supercritical shocks are two-fold:

- In supercritical shocks, the majority of ions cross the shock and are transmitted downstream, contributing to the downstream population. However, a small fraction of ions can reach the shock and undergo (approximately) specular reflection due to the strong magnetic field gradient at the shock front, thereby gaining energy. After reflection, these ions propagate upstream along the magnetic field lines. This ion reflection mechanism limits further shock steepening, since part of the shock energy is transferred to the reflected ions and away from the shock and thus acts as a dissipative process. The reflected ions, which move faster than the shock may excite plasma instabilities in the shock upstream, which contribute to particle scattering and energy dissipation. Note that electrons are not involved in this process: electron kinetic scales are much smaller than ion kinetic scales, so electrons typically cross the shock without being reflected, and ion reflection alone is sufficient to provide dissipation.
- At kinetic scales, nonlinear interactions between particles and waves become essential for explaining the shock physics under different conditions. In a fluid description, these interactions can be interpreted as an effective viscous term. However, since fluid models cannot resolve kinetic scales, wave-particle interactions are instead described as an *anomalous viscosity*. The resulting plasma waves act as scattering centres, effectively dissipating energy. The details of wave-particle interactions depend sensitively on the shock parameters and are typically studied within the quasi-linear regime.

Derivation of the critical Mach number

The critical Mach number is derived here for a non-relativistic system, following Burgess and Scholer 2015. Consider a one-dimensional shock transition in which the magnetic field is perpendicular to the shock normal, $\mathbf{B} = B\hat{\mathbf{z}}$. Assume the presence of a resistive term $\rho\mathbf{J}$ arising from Ohmic heating or wave-particle scattering. For simplicity, the resistivity is taken to be isotropic, and the resistive term can be interpreted as an effective drag opposing the

plasma motion. For ions, the mass and momentum conservation equations are

$$\frac{d}{dx}(nV_{i,x}) = 0 \quad (3.33)$$

$$m_i n_i V_{i,x} \frac{\partial \mathbf{V}_i}{\partial x} = n_i e (\mathbf{E} + \mathbf{V}_i \times \mathbf{B} - \rho \mathbf{J}) - \nabla \mathcal{P}_i, \quad (3.34)$$

where n , V and $\mathcal{P}$ denote the number density, bulk velocity, and pressure tensor, respectively. The left-hand side of Equation 3.34 represents the inertial term in the stationary case, which follows from the definition of force, $f = d(mnv)/dt = mn \partial v / \partial t + mn \partial v / \partial x \, dx/dt = mn v \, dv/dx$. The first term on the right-hand side of Equation 3.34 is the Lorentz force acting on the ion fluid, including the resistive drag associated with $\rho \mathbf{J}$, while the last term corresponds to the force due to the ion pressure gradient. The contribution of magnetic pressure can be neglected provided that the magnetic field amplitude is sufficiently small.

Electrons can be treated as a massless fluid. Assuming quasi-neutrality, $n_i = n_e = n$, the mass conservation equation is identical to Equation 3.33, which implies $V_{i,x} = V_{e,x}$. The momentum and energy conservation equations for electrons are then given by

$$-n_e e (\mathbf{E} + \mathbf{V}_i \times \mathbf{B} - \rho \mathbf{J}) - \nabla p_e = 0, \quad (3.35)$$

$$V_{e,x} \frac{dp_e}{dx} = -\Gamma_{ad} p_e \frac{dV_{e,x}}{dx} + (\Gamma_{ad} - 1) \rho J^2, \quad (3.36)$$

where p_e is the electron scalar pressure. Because electrons have a low mass, they isotropise rapidly due to turbulence and resistive drag, and it is therefore customary to assume a scalar pressure. The first and second terms in Equation 3.36 follow from the polytropic equation of state $d(p\rho^{-\Gamma_{ad}})/dt = 0$ (Kulsrud 2005), while the third term is obtained from Poynting's theorem in the absence of electromagnetic energy flux, $\partial u_S / \partial t = \nabla \cdot \mathbf{S} + \mathbf{J} \cdot \mathbf{E} = \rho J^2$.

Given quasi-neutrality and stationarity, the Ampère-Maxwell law can be written as

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} = \mu_0 n e (\mathbf{V}_i - \mathbf{V}_e) = \mu_0 n e (V_{i,y} - V_{e,y}) \hat{\mathbf{j}}. \quad (3.37)$$

Using Equations 3.35 and 3.37 and evaluating $(\nabla \times \mathbf{B}) \times \mathbf{B}$, the electric field can be expressed as

$$\mathbf{E} = -\mathbf{V}_i \times \mathbf{B} - \frac{1}{ne} \nabla \left(p_e + \frac{B^2}{2\mu_0} \right) + \eta \mathbf{J}. \quad (3.38)$$

Similarly, Equation 3.34 becomes

$$m_i n_i V_{i,x} \frac{\partial \mathbf{V}_i}{\partial x} = -\frac{\partial}{\partial x} \left(p_e + \frac{B^2}{2\mu_0} \right) - \nabla \mathcal{P}_i. \quad (3.39)$$

It is convenient to introduce dimensionless variables normalised to their upstream values in order to obtain relations in which the Alfvénic Mach number explicitly appears. In addition, the length scale is normalised to the ion gyroradius $R_{g,i} = V_u / \Omega_i$:

$$\tilde{n} = \frac{n}{n_u}, \quad \tilde{\mathbf{V}} = \frac{\mathbf{V}}{V_u}, \quad \tilde{B} = \frac{B}{B_u}, \quad \tilde{p}_e = \frac{p_e}{B_u^2 / 2\mu_0}, \quad \tilde{\eta} = \frac{\eta \Omega_i}{\mu_0 V_u^2}, \quad M_A^2 = \frac{V_u^2 \mu_0 n m_i}{B_u^2}. \quad (3.40)$$

With these definitions, Equations 3.33, 3.36 and 3.39 can be rewritten as

$$nV_x = 1 \quad (3.41)$$

$$V_x \frac{dp_e}{dx} = -\Gamma_{ad} p_e \frac{dV_x}{dx} - 2(\Gamma_{ad} - 1)\eta \left(\frac{dB}{dx} \right)^2 \quad (3.42)$$

$$nV_x \frac{d\mathbf{V}}{dx} = \frac{-1}{2M_A^2} \frac{d}{dx} (p_e + B^2) \hat{\mathbf{x}} - \frac{1}{2M_A^2} \frac{d}{dx} \mathbf{p}_{i,x} \quad (3.43)$$

Here the tilde notation has been dropped for simplicity, without loss of generality. Note that in the second equation $\mu_0^2 J^2 = \mu_0^2 J_y^2 = (\nabla \times \mathbf{B})^2 = -(\frac{dB}{dx})^2$, and in the third equation a diagonal pressure tensor is assumed, such that $p_{i,x} = \mathcal{P}_i \hat{\mathbf{x}}$. Equation 3.43 can then be integrated to give

$$\left(\mathbf{V} + \frac{1}{2M_A^2} p_{i,x} \right) + \frac{1}{2M_A^2} (p_e + B^2) \hat{\mathbf{x}} = \text{ct}, \quad (3.44)$$

where Equation 3.41 has been used to eliminate the factor nV_x .

Now consider a cold upstream, such that $p_{e/u} = p_{i/u} = 0$, and a downstream region in which electrons are heated throughout the shock transition while ions remain cold, $p_{i,x} = 0$. The shock and upstream conditions can then be compared using Equation 3.44, which yields

$$V_x + \frac{1}{2M_A^2} (p_e + B^2) = 1 + \frac{1}{2M_A^2}. \quad (3.45)$$

Differentiating this expression gives

$$\frac{dp_e}{dx} = -2M_A^2 \frac{dV_x}{dx} - \frac{dB^2}{dx}, \quad (3.46)$$

which can be substituted into Equation 3.42 to obtain

$$n \frac{dV_x}{dx} = \frac{-1}{2M_A^2} \frac{V_x \frac{dB^2}{dx} + 2(\Gamma_{ad} - 1)\eta \left(\frac{dB}{dx} \right)^2}{V_x^2 - \Gamma_{ad} p_e / 2nM_A^2}. \quad (3.47)$$

The denominator on the right-hand side of Equation 3.47 can be written as $(V_x^2 - c_s^2)$, where $c_s^2 = \Gamma_{ad} p_e / n$ (note that c_s is still expressed in normalised units). A singularity in dV_x/dx therefore appears when $V_x = c_s$. As the plasma flows through the shock, the initially supersonic flow speed V_x decreases, while the electron temperature T_e – and hence the sound speed c_s – increase, as expected from the RH jump conditions (see Section 3.2.1). If the flow speed far downstream remains larger than the sound speed throughout the shock transition, Equation 3.42 admits a finite solution and the downstream flow remains supersonic. Even in this case, a shock still exists because the downstream flow speed remains larger than the fast magnetosonic speed v_+ (see Equation 3.20). This implies that downstream sound waves do not contribute to further shock steepening, and no additional dissipation mechanism is required to prevent indefinite steepening. In contrast, if the flow speed far downstream becomes smaller than the sound speed, a singularity develops within the shock transition when

$V_x = c_s$. In this case, an additional dissipation mechanism – such as particle reflection – is required to dissipate energy. The threshold at which this occurs defines the *critical Mach number*, $M_{cr} = V_{u,c}/v_+$. Shocks with $M > M_{cr}$ therefore require additional dissipation mechanisms to prevent infinite steepening.

The value of the critical Mach number depends on the magnetic field obliquity θ_{Bn} , the plasma beta β_p , and the adiabatic index Γ_{ad} . A parametric survey of the critical Mach number was performed by Edmiston and Kennel 1984. For a standard cold ($\beta_p = 0$) and adiabatic plasma with $\Gamma_{ad} = 5/3$, the critical Mach number is $M_{cr} = 2.76$ for a strictly perpendicular shock ($\theta_{Bn} = 90^\circ$) and $M_{cr} = 1.53$ for a strictly parallel shock ($\theta_{Bn} = 0^\circ$). For a hot plasma ($\beta_p \rightarrow \infty$) the critical Mach number decreases to $M_{cr} = 1.16$ for $\theta_{Bn} = 90^\circ$ and $M_{cr} = 1$ for $\theta_{Bn} = 0^\circ$. The critical Mach number is expected to be lower for parallel configurations because the fast magnetosonic speed is smaller, $v_+(\theta_{Bn} = 0^\circ) + c_s^2 = v_+(\theta_{Bn} = 90^\circ)$, making it easier for a shock to be subsonic. Hotter plasmas have a higher sound speed and thus the downstream flow is slower than c_s , leading to values of M_{cr} close to unity.

Although it is often stated that supercritical shocks must reflect ions at the shock as a dissipation mechanism to prevent infinite steepening, ion reflection is not exclusive to supercritical shocks. The boundary between subcritical and supercritical shocks is not a sharp transition but rather a smooth one (Omidi et al. 1990). Direct observations of the terrestrial bow shock by the International Sun-Earth Explorer missions ISEE-1 and ISEE-2 confirm the presence of reflected ions upstream of subcritical shocks, both for fast Mach numbers slightly above the critical value ($M_f/M_{cr} \simeq 1.1$) and even below it ($M_f/M_{cr} \simeq 0.6$) (Greenstadt and Mellott 1987).

3.2.5 The whistler critical Mach number

Whistler waves propagate parallel to the magnetic field and are the simplest type of dispersive waves: their phase and group velocities vary with the wavenumber, k , and frequency, ω . By comparing the shock speed with the whistler phase speed, one finds that in quasi-parallel magnetic-field configurations, whistlers emitted at the shock can run ahead of it, since their propagation speed exceeds the shock speed (Krasnoselskikh et al. 2002; Burgess and Scholer 2015). If the whistler group speed is lower than the shock speed, the wave cannot outrun the shock and dissipates downstream. If the group speed is higher but the phase speed is lower, the wave energy propagates upstream but the peaks move towards the shock and downstream. Thus, there exists a phase speed at which the whistler phase-stands with the shock, $v_{sh} = v_{ph}$. For such phase-standing whistlers, the group speed is lower than the phase speed, $v_g < v_{ph}$. This results in a train of waves in front of the shock whose length increases over time. The leading edge of this wave train is populated by high- k , high- ω whistlers that move faster, while waves of lower wavenumber and frequency remain closer to the shock front.

In the low β regime, one can derive an approximate condition under which whistlers are able to outrun the shock. Consider the dispersion relation of the whistler mode in the high-frequency limit, in which electron kinetic effects become non-negligible (Stringer 1963):

$$\omega^2 = \frac{v_A^2 k^2}{1 + k^2 c^2 / \omega_{p,e}^2} \left(1 + \frac{v_A^2 k^2}{(1 + k^2 c^2 / \omega_{p,e}^2) \Omega_i^2} \right). \quad (3.48)$$

For large frequencies, and assuming $\omega_{p,e} \gg \Omega_i$, the unity term can be neglected and the phase speed can be derived as:

$$\omega^2 = \frac{v_A^4 k^4}{(1 + k^2 c^2 / \omega_{p,e}^2)^2 \Omega_i^2} \rightarrow \frac{\omega}{k v_A} = \frac{v_{ph}}{v_A} = \frac{k c}{(1 + \kappa^2) \omega_{p,i}} \frac{\kappa}{1 + \kappa^2} \mu^{1/2}, \quad (3.49)$$

where $v_A/c = \frac{\Omega_i}{\omega_{p,i}}$, $\omega_{p,e}/\omega_{p,i} = \sqrt{m_i/m_e} = \mu^{1/2}$ and $\kappa = kc/\omega_{p,e}$. In quasi-parallel magnetic-field configurations, a factor of $\cos \theta_{Bn}$ is introduced. The critical whistler phase speed, i.e., the critical whistler Mach number, corresponds to the maximum value of the phase speed:

$$\frac{1}{v_A} \frac{dv_{ph}}{d\kappa} = \mu^{1/2} \cos \theta_{Bn} \frac{1 - \kappa^2}{(1 + \kappa^2)^2} = 0 \rightarrow \kappa = 1, \quad (3.50)$$

and therefore

$$M_w \equiv \frac{\max v_{ph}}{v_A} = \frac{\mu^{1/2}}{2} \cos \theta_{Bn}. \quad (3.51)$$

Note that the above expression is valid only for weakly magnetised plasmas. For higher magnetisation, both Alfvén speed and maximum phase speed of the whistler branch become relativistic. In this regime, an analytical expression for the whistler critical Mach number cannot be obtained, even for low- β plasmas, but it can be determined numerically from the full dispersion relation (see Section 6.4.2).

3.3 Relativistic shocks

RELATIVISTIC shocks, characterised by a propagation Lorentz factor $\gamma \gg 1$, are ubiquitous in the Universe and occur in both Galactic and extragalactic environments hosting relativistic plasma outflows. Within the Galaxy, such shocks are commonly found in pulsar wind nebulae (e.g., the Crab nebula, Abeysekara et al. 2017) and microquasars (Hess 2018), while in extragalactic systems they are present in the jets of GRBs and AGN (Sironi et al. 2015; Pohl et al. 2020). These sources are known to emit intense high-energy radiation, which is produced by ultra-relativistic particles. Relativistic shocks are widely recognised as key sites for the acceleration of these particles.

3.3.1 Structure of a relativistic shock

In relativistic shocks, the bulk plasma velocity is comparable to velocities of the accelerated particles. As a result, particle distributions near the shock are strongly anisotropic. The internal structure and dynamics of such shocks are governed by the properties of the upstream flow, including its Lorentz factor γ_{sh} , as well as the strength and orientation of the background magnetic field. In the ultra-relativistic regime, where $\gamma_{sh} \gg 1$, the dominant parameter controlling the shock physics is the plasma magnetisation σ .

In collisionless ultra-relativistic shocks, the deceleration and isotropisation of the incoming plasma are facilitated by the reflection of a fraction of incoming particles back upstream.


In moderately magnetised plasmas, with $10^{-2} \lesssim \sigma \lesssim 10^{-1}$, this reflection is primarily caused by the compression of the ambient magnetic field at the shock front. At lower magnetisations, the necessary magnetic turbulence is generated self-consistently through kinetic micro-instabilities in the shock precursor. These instabilities are driven by the interaction between the incoming upstream flow and the population of particles that are reflected and accelerated at the shock (see, e.g., Pohl et al. 2020).

In unmagnetised or weakly magnetised electronproton plasmas ($\sigma_e \lesssim 10^{-2}$), the dominant instabilities are the *Weibel (filamentation) instability* (Medvedev and Loeb 1999) and the *perpendicular current instability* (Lemoine et al. 2014b; Lemoine et al. 2014a). These instabilities arise from counterstreaming between the incoming upstream plasma and shock-reflected particles. Their interaction first generates current filaments at electron skin-depth scales, which subsequently evolve into ion-scale magnetic turbulence. Electrons interacting with the associated transverse electric fields are heated and isotropised in the shock precursor, enabling further energisation that is essential for the sustained generation of large-amplitude magnetic fields downstream (Takamoto et al. 2018).

Another instability observed in the weakly magnetised regime is the Bell instability, which is driven by a non-thermal ion beam (Reville et al. 2006). Although superluminal shocks cannot completely reflect ions at the shock front, this instability can nevertheless be excited as a secondary consequence of the Weibel instability (Lemoine and Pelletier 2011). The Bell instability has also been observed in mildly relativistic shocks, where the shock structure transitions from Weibel-mediated to Bell-mediated, owing to the slower growth rate of the Bell instability compared to that of the Weibel instability (Crumley et al. 2019).

As the magnetisation increases, the shock-compressed magnetic field becomes dynamically important. Upstream ions gyrate around this field, forming a ring-like distribution that is unstable to the *synchrotron maser instability* (SMI; Langdon et al. 1988). For a steady inflow of upstream plasma, this continuous gyration generates electromagnetic waves that propagate upstream. In electron-ion plasmas, the interaction of incoming electrons with precursor waves can produce electrostatic wakefields Lyubarsky 2006. The resulting wakefield

Figure 3.2: Schematic structure of a superluminal *moderately magnetised* relativistic shock in the downstream frame of reference. The downstream plasma is compressed and exhibits large-scale turbulence. The shock propagates to the right with speed v_{sh} , while the upstream plasma flows to the left with speed v_{up} . The upstream region is populated by small-scale perturbations arising from micro-instabilities between the inflowing plasma and reflected ions.



acceleration can energize electrons and ions to high energies (Hoshino 2008; Iwamoto et al. 2017a). The details of this process are discussed in Section 4.2.

It is important to note that ultrarelativistic magnetised shocks are typically quasi-perpendicular and superluminal. In the upstream frame of reference, the shock speed approaches the speed of light, $v_{sh} \rightarrow c$, and the threshold angle is close to zero, $\cos \theta_{tr} \approx 1 \rightarrow \theta_{tr} \approx 0^\circ$. Consequently, in the shock frame of reference, the threshold angle approaches 45° .

Figure 3.2 shows a schematic of a relativistic shock as seen from the downstream frame of reference. There are two very distinct media: downstream on the left and upstream on the right; separated by the shock itself. The downstream plasma exhibits higher density and higher temperature as opposed to the far upstream plasma. The plasma is heated by anomalous resistivity present at the shock transition, coming from the multiple instabilities that are produced as a consequence of upstream particles interacting with the shock. Moreover, the downstream magnetic field has higher amplitude as expected from the RH jump conditions (see section 3.2.1). Particle-wave coupling generates turbulent fluctuations that isotropise the downstream plasma via particle scattering.

Figure 3.2 presents a schematic structure of a relativistic *moderately magnetised* shock in the downstream frame of reference. The downstream region on the left and the upstream region on the right are separated by the shock transition. The downstream plasma is characterised by higher density and temperature than the far upstream region. Plasma heating occurs through anomalous resistivity within the shock transition, arising from instabilities generated by the interaction of upstream particles with the shock. The downstream magnetic field is compressed, as expected from the RankineHugoniot jump conditions (see Section 3.2.1). Particle-wave coupling produces turbulent fluctuations that isotropise the downstream plasma via particle scattering. Note that the structure of strongly magnetised, mildly relativistic shocks with $\sigma = 1$, which are studied in this thesis, differs substantially from the configuration described above, as presented in Chapter 6.

3.3.2 Mildly relativistic shocks

Shocks are referred to here as *mildly relativistic* if their Lorentz factors lie in the range $\gamma_{sh} \simeq 2\text{--}5$. Such shocks are most commonly considered in the context of AGN jets and microquasars, although the emission observed in some GRBs may also be explained by mildly relativistic outflows.

As discussed in Section 2.2, some AGN – in particular Fanaroff-Riley type II (FRII) sources (Fanaroff and Riley 1974) – produce giant radio lobes that terminate in bright hotspots. These hotspots mark the location of the jet termination shock, where the relativistic AGN jet impacts the intergalactic medium. They are associated with strong X-ray emission, commonly attributed to synchrotron radiation, and therefore require an efficient source of ultra-relativistic leptons.

Using MHD simulations, Matthews et al. 2019 proposed that mildly relativistic shocks can account for the observed electromagnetic emission from AGN hotspots and may also contribute to a hadronic component in the form of ultra-high-energy cosmic rays (UHECRs) under typical AGN conditions. Since ultra-relativistic shocks are generally inefficient at accelerating protons and heavier ions to UHECR energies, mildly relativistic shocks are favoured

in this context. However, Matthews et al. 2019 model assumes diffusive shock acceleration (DSA) as the dominant energisation mechanism, which may not always operate efficiently (see Section 4.1.1).

Observational evidence for mildly relativistic shocks has been reported by Laing and Bridle 2013, who studied the radio quasar 3C 279 and the blazar Mrk 501. By modelling the emission with a combined DSA and shock drift acceleration (SDA) framework, they inferred the presence of shocks propagating at speeds of $v_{sh} \simeq 0.7c$. Such shocks are expected to arise from collisions between ultra-relativistic MHD flows within blazar jets and may occur frequently, leading to time-dependent particle acceleration and variable emission spectra.

In the same work, Laing and Bridle 2013 also analysed a sample of eleven low-luminosity radio galaxies observed with the Very Large Array (VLA). In these sources, strong emission is detected closer to the base of the jet. Given the lower jet power in such systems, mildly relativistic shocks are again invoked, as the jets are not expected to reach ultra-relativistic speeds.

Microquasars are less powerful than AGN or GRBs due to their smaller physical scales. Consequently, the outflows launched from the accreting black hole or neutron star can be only mildly relativistic. Saavedra et al. 2023 investigated a high-mass γ -ray-emitting binary – more powerful than a typical X-ray binary – using observations from *NuSTAR* and *XMM-Newton*. They proposed a lepto-hadronic mildly relativistic jet as the origin of the observed non-thermal emission from the microquasar. Since not all microquasars exhibit extended structures, the presence of jets is often inferred indirectly through modelling of the observational data. Such studies have motivated the use of mildly relativistic outflows to model microquasar jets. In particular, de la Cita et al. 2017 employed relativistic hydrodynamical simulations to produce mildly relativistic jets capable of explaining the observed non-thermal radiation in microquasars.

The structure of mildly relativistic shocks and its dependence on upstream parameters are similar to those of ultrarelativistic shocks (see Sections 3.3.1 and 4.1.1). Despite extensive efforts to characterise the jet properties of AGN and microquasars, some key parameters remain poorly constrained observationally, most notably the magnetic field obliquity and the plasma magnetisation. Nevertheless, available observational evidence suggests that the plasma in these environments is likely at least moderately magnetised (see, e.g., Ghisellini et al. 2005; Ligorini et al. 2021a).

As noted above, typical magnetised ultrarelativistic shocks are quasi-perpendicular because the component of the magnetic field transverse to the shock normal is amplified by a factor γ_{sh} through a Lorentz transformation. Since such shocks are also superluminal, particle acceleration may be strongly suppressed, as many acceleration mechanisms require particles to return upstream in order to self-generate the turbulence necessary for scattering at the shock. Mildly relativistic shocks are therefore particularly attractive for particle acceleration studies. Their lower propagation speeds allow for a wider range of magnetic field obliquity angles for which the shock remains subluminal, enabling particle reflection far upstream and the subsequent excitation of instabilities that facilitate particle acceleration. Figure 3.3 illustrates this broader angular range in the upstream frame of reference. For shocks with Lorentz factors in the mildly relativistic range, $\gamma_{sh} = 2$ –5, the corresponding




Figure 3.3: Threshold angle θ_{tr} in the upstream frame of reference as a function of the shock Lorentz factor γ_{sh} . The two cross markers indicate example values at $\gamma_{sh} = 2.5$ and $\gamma_{sh} = 10$. For obliquity angles above θ_{tr} the shock is superluminal, whereas for angles below θ_{tr} the shock is subluminal.

threshold angle spans $\theta_{tr} = 30^\circ$ – 11.54° . As the shock speed increases, the threshold angle decreases accordingly.

CHAPTER 4

Particle acceleration mechanisms in relativistic shocks

PARTICLE acceleration mechanisms in relativistic shocks differ substantially from that in non-relativistic shocks. Several mechanisms that operate efficiently in non-relativistic regimes become strongly suppressed at relativistic shocks, while other processes arise only under relativistic conditions. This chapter first introduces the well-known diffusive shock acceleration mechanism and summarises the reasons why it cannot account for the emission observed in many environments hosting relativistic shocks. The wakefield acceleration is then discussed as an energisation mechanism that operates in relativistic magnetised shocks and can accelerate electrons to high energies, although it becomes inefficient in strongly magnetised relativistic plasmas. Finally, a less explored mechanism is presented that operates in the strongly magnetised regime, in which a magnetosonic shock wave traps electrons and accelerates them to very high energies through the motional electric field. The analysis of this process constitutes one of the main objectives of this thesis, and the corresponding results are presented in Sections 6.2 and 6.3.

4.1 Diffusive shock acceleration (DSA) and its limitations

DIFFUSIVE shock acceleration (DSA) is an energisation process proposed by Blandford and Ostriker 1978 and Bell 1978a, building on the original ideas of Enrico Fermi (Fermi 1954). In this framework, charged particles are scattered by magnetic irregularities or turbulence, gaining energy through repeated interactions. In Fermi's original formulation (Fermi 1949), a particle interacting with a moving magnetic cloud gains or loses energy depending on whether the encounter is head-on (opposite velocity directions) or tail-on (same velocity direction). For a random distribution of clouds (scattering centres), the average fractional energy gain is proportional to the squared ratio of the cloud speed and the particle speed:

$$\frac{\Delta\epsilon}{\epsilon} = \sum_i \left(\frac{u_i}{c} \right)^2, \quad (4.1)$$

where the sum runs over all interacting clouds with velocity u_i , and the particle is assumed to be ultra-relativistic, with $v_p \simeq c$. Because the energy gain is second order in the velocity ratio, this mechanism is commonly referred to as *second-order Fermi acceleration*. The efficiency of this process is significantly enhanced in the vicinity of a shock wave. In this case, the converging plasma flows upstream and downstream of the shock ensure that particle-scattering-centre interactions are predominantly head-on. As a result, the average energy gain becomes first order in the velocity ratio, leading to a much more efficient acceleration mechanism. This shock-associated version of Fermi acceleration, now known as DSA, can accelerate particles to very high energies and forms the basis of the standard theory of particle acceleration at astrophysical shocks (Axford et al. 1977; Krymskii 1977; Bell 1978a; Bell 1978b; Blandford and Ostriker 1978).

To elucidate the physics of this mechanism, it is convenient to adopt a microscopic description (Drury 1983). Consider a non-relativistic strong, planar shock that is infinitely extended and infinitely thin. As a shock, it satisfies the RankineHugoniot jump conditions. Consider relativistic particles with number density n approaching the shock from upstream, such that crossing the shock does not alter their momentum. In the upstream frame of reference, each particle has a velocity $\mathbf{v}$, which is generally distributed over angles rather than being strictly aligned with the shock normal. As a consequence, only a fraction of the particles cross the shock. For a nearly isotropic particle distribution, the flux of particles crossing the shock is given by

$$F = \int_0^1 n v_{\parallel} \frac{d\mu}{2} = n v \int_0^1 \mu \frac{d\mu}{2} = \frac{n v}{4}, \quad (4.2)$$

where $\mu = \cos \theta$ is the pitch-angle cosine. This result can be understood intuitively: half of the particles move toward the shock, with an average parallel velocity $v/2$, yielding the factor $1/4$. If the flux of particles escaping downstream after crossing the shock is $n V_{\text{dw}}$, the probability that a particle escapes downstream is

$$P(\text{escape}) = \frac{n V_{\text{dw}}}{n v / 4} = 4 \frac{V_{\text{dw}}}{v} = 4 \frac{V_{\text{up}}}{v r}, \quad (4.3)$$

where $r = V_{\text{up}}/V_{\text{dw}}$ is the compression ratio of the shock. The probability that a particle returns to the shock is therefore $P(\text{return}) = 1 - P(\text{escape})$.

Now consider the energy E of a single ultra-relativistic particle, for which $E \simeq pc$. Transforming to the shock frame modifies the particle energy according to $E_{\text{sh}} = \gamma_{\text{up}}(E + p_x V_{\text{up}}/c) \approx (E + p \cos \theta c V_{\text{up}}/c) = E(1 + V_{\text{up}} \cos \theta/c)$, where $V_{\text{up}} = -V_{\text{sh}}$ is the upstream flow speed, and we have assumed a non-relativistic shock such that $\gamma_{\text{up}} \rightarrow 1$. Since the particle is ultra-relativistic, its momentum magnitude is essentially unchanged by the shock. Therefore, the energy in the downstream frame is simply $E' = E(1 + (V_{\text{up}} - V_{\text{dw}}) \cos \theta/c)$. The relative energy gain per shock crossing cycle is therefore

$$\frac{\Delta E}{E} = \frac{E' - E}{E} = \frac{V_{\text{up}} - V_{\text{dw}}}{c} \cos \theta = \frac{V}{c} \cos \theta, \quad (4.4)$$

where $V = V_{\text{up}} - V_{\text{dw}} = V_{\text{up}}(1 - r^{-1})$ is the upstream velocity in the downstream frame.

As mentioned, not all particles cross the shock. Therefore, it is necessary to consider the probability distribution of particles crossing the shock to compute the average energy gain. From equation 4.2, this probability is

$$p(\theta) d\theta = A \mu d\mu = A \cos \theta d(\cos \theta), \quad (4.5)$$

where A is the normalisation constant such that $A \int_0^1 p(\theta) d\theta = 1$, which yields $A = 2$. Using this distribution, the average energy gain per crossing is

$$\left\langle \frac{\Delta E}{E} \right\rangle = 2 \int_0^{\pi/2} \frac{\Delta E}{E} p(\theta) d\theta = 2 \frac{V}{c} \int_0^{\pi/2} \cos^2(\theta) d(\cos \theta) = \frac{2}{3} \frac{V}{c} = \frac{2}{3} \frac{V_{\text{up}}}{c} (1 - r^{-1}). \quad (4.6)$$

For a strong shock, this simplifies to $\langle \Delta E/E \rangle = \frac{V_{\text{up}}}{2c}$. Unlike second-order Fermi acceleration, this energy gain is linearly proportional to the shock velocity, which is why it is often called *first-order Fermi acceleration*.

If a particle is able to return back to the shock from downstream, the process repeats with opposite velocity $-V$ and integration limits $\theta \in \left[\frac{\pi}{2}, 0\right]$, yielding the same average energy gain (Drury 1983). A full DSA cycle typically considers two crossings, giving $\langle \Delta E/E \rangle_{\text{cycle}} = \frac{4V}{3c}$. After N cycles, the total energy gain is

$$E_N = E_0 \left[1 + \frac{4}{3} \frac{V_{\text{up}}}{c} (q - r^{-1}) \right]^N, \quad (4.7)$$

where E_0 is the initial energy of the particle.

As particles gain energy through DSA, an increasing fraction escape either far downstream or upstream, so fewer particles reach higher energies. The probability of a particle completing at least N cycles is

$$P(n \geq N) = \prod_{i=1}^N \left(1 - \frac{4V_{\text{up}}}{vr} \right) \approx \left(1 - \frac{4V_{\text{up}}}{cr} \right)^N, \quad (4.8)$$

where a relativistic particle is assumed, $v \approx c$. Note that $P(n \geq N) = P(E \geq E_N)$, as both describe the same process. Taking the logarithm yields

$$\begin{aligned} \ln P(E \geq E_N) &= N \ln \left(1 - \frac{4V_{\text{up}}}{cr} \right) = \ln \left(\frac{E_N}{E_0} \right) \frac{\ln \left(1 - \frac{4V_{\text{up}}}{cr} \right)}{1 + \frac{4}{3} \frac{V_{\text{up}}}{c} (q - r^{-1})} \approx \\ &= - \ln \left(\frac{E_N}{E_0} \right) \frac{3}{r-1} \rightarrow P(E \geq E_N) \approx \left(\frac{E_N}{E_0} \right)^{\frac{-3}{r-1}}, \end{aligned} \quad (4.9)$$

where N was obtained from Equation 4.4 and the approximation $\ln(1+x) \approx \ln x$ for $x \rightarrow 0$ was used, which is valid for a non-relativistic shock. The particle energy distribution can then be derived from

$$\begin{aligned} P(E \geq E_N) &= A \int_0^E N(E) dE \rightarrow N(E) \propto \frac{dP(E \geq E_N)}{dE} = E^{-q}, \\ q &= \frac{3}{r-1} + 1 = \frac{r+2}{r-1}. \end{aligned} \quad (4.10)$$

For a strong shock with $r = 4$, this gives $q = 2$, which corresponds to the standard power-law index predicted by DSA.

DSA has been proposed as the primary acceleration mechanism responsible for CR production in a wide range of astrophysical sources. Recent developments in DSA theory take into account the back-reaction of the accelerated particles on the shock structure, which

can slightly modify the expected spectral index (Blasi 2013; Drury 1983). For energies below 10^{15} TeV, corresponding to CRs from Galactic sources, the observed spectrum – after accounting for galactic diffusion – closely matches the DSA-predicted power-law index, yielding $q \approx 2.21$ – 2.29 (Thoudam, S. et al. 2016). This agreement makes DSA a compelling candidate for explaining the bulk of the low-energy CR spectrum. Moreover, observations indicate magnetic field amplification due to self-generated turbulence near the shock, which further enhances the particle energisation rate and maximum energies gained in the process (Bell 2004; Ressler et al. 2014; Parizot et al. 2006; Morlino and Caprioli 2012).

Limitations to DSA

The theory of DSA assumes an infinitesimally thin shock. At kinetic scales, however, shocks have a finite thickness d_{sh} , which, in perpendicular shocks, is roughly of the order of the thermal ion gyroradius $d_{sh} \approx r_g$. For relativistic particles, as assumed in DSA theory, the shock thickness is negligible and particles are accelerated via standard DSA. However, many non-relativistic shocks with moderate plasma beta β_p , have non-relativistic upstream plasma distributions. The resulting spectra are softer, as the escape probability increases for slower particles (see Equation 4.3). Moreover, the shock thickness becomes comparable to the particle gyroradius. To ensure effective acceleration, the shock thickness must be smaller than the characteristic length scale of the accelerated particles (Amano et al. 2022). For electrons, this length scale is the gyroradius $r_{g,e} = m_e v_e / (eB)$, giving the condition

$$r_{g,e} \frac{m_e v_e}{eB} \gtrsim d_{sh} = v_{sh} \Omega_i = v_{sh} \frac{m_i}{eB} \rightarrow p_e \gtrsim \gamma_e m_i v_{sh} \rightarrow \frac{p_e}{\gamma_e m_e V_A} \gtrsim \frac{m_i}{m_e} M_A, \quad (4.11)$$

where M_A is the Alfvénic Mach number. This expression is convenient because it clearly shows the dependence on the mass ratio m_i/m_e and M_A . For ions, this condition is typically satisfied automatically due to their larger mass. Electrons, however, generally have lower momentum and require an additional *pre-acceleration* mechanism. The need to energise electrons prior to entering DSA is referred to as the *electron injection problem*.

At the same time, DSA requires sufficiently strong turbulence both upstream and downstream of the shock to scatter particles back towards the shock front. For electrons, this necessitates the presence of electron-scale waves capable of providing efficient pitch-angle scattering. Such waves are typically strongly damped, and sustaining them requires a continuous injection of turbulence. While pre-existing turbulence could in principle supply this energy, this scenario is unlikely, as MHD turbulence does not decay to electron scales on the timescales relevant for DSA. An alternative mechanism is the generation of turbulence through kinetic plasma instabilities. If such instabilities develop upstream, they can amplify the magnetic field and produce waves capable of scattering particles. A well-known example is Bell’s instability, which arises in the presence of a low-density population of relativistic CRs.

4.1.1 DSA in relativistic shocks

As noted in Section 3.3, DSA at relativistic shocks is highly sensitive to the physical conditions of the upstream plasma. This sensitivity arises because the bulk flow velocity becomes comparable to the velocities of the accelerated particles, leading to strongly anisotropic particle distributions near the shock. Key controlling factors for DSA efficiency include the Lorentz factor of the shock, the strength and obliquity of the background magnetic field, and the nature of the electromagnetic turbulence that governs particle scattering.

These effects were first identified in the test-particle regime using semi-analytical approaches and Monte Carlo simulations (see, e.g., Niemiec et al. 2006; Pohl et al. 2020, and references therein). In particular, studies of quasi-perpendicular, superluminal shocks – commonly encountered in ultrarelativistic flows – have shown that particle acceleration is generally inefficient, resulting in steep energy spectra. Efficient acceleration and the formation of extended power-law distributions occur only when intense magnetic turbulence is present in the shock vicinity, corresponding to extremely low effective magnetization. Under such conditions, ultrarelativistic shocks with $\gamma_{\text{sh}} \gg 1$ can generate broad power-law particle spectra with asymptotic indices $q \approx 2.3\text{--}2.6$ consistent with those inferred for synchrotron-emitting electrons in GRB afterglows (see, e.g., Pohl et al. 2020).

These results were later confirmed by particle-in-cell (PIC) simulations. In ultrarelativistic, unmagnetised shocks, a DSA-like process operates in which particles gain energy through multiple reflections between the upstream and downstream regions of the shock (Spitkovsky 2008). Particle scattering in this regime is provided by magnetic turbulence generated by the Weibel instability. However, for electron-proton plasmas with magnetisations exceeding $\sigma \sim 3 \times 10^{-5}$, the shock structure and the efficiency of particle acceleration become strongly dependent on the obliquity of the background magnetic field. In particular, superluminal shocks are found to be inefficient particle accelerators, whereas efficient acceleration can occur at subluminal shocks (Sironi and Spitkovsky 2011).

In subluminal ultrarelativistic configurations, ions are accelerated more efficiently than electrons for magnetisations $\sigma \gtrsim 10^{-2}$. Upstream particle scattering is mediated by turbulence self-consistently generated by shock-reflected particles, through either Weibel- or Bell-type instabilities, depending on the magnetisation. A DSA-like mechanism operates predominantly in quasi-parallel shocks, while for increasing magnetic obliquity toward the threshold angle θ_{tr} , shock drift acceleration (SDA, see Section 4.3.2) becomes the dominant energisation process (Sironi and Spitkovsky 2011; Pohl et al. 2020). Acceleration efficiency decreases for subluminal shocks at magnetisations $\sigma \gtrsim 0.1$, as the level of magnetic turbulence produced by upstream instabilities becomes insufficient to sustain efficient particle scattering. An analogous dependence on shock obliquity and magnetisation has been identified in mildly relativistic shocks (Sironi and Spitkovsky 2011; Crumley et al. 2019) – see also Rassel et al. 2023 for results on mildly relativistic shocks for high-beta plasma conditions typical in kilonova scenarios (Kulsrud and Zweibel 2008).

The high-energy radiation observed from blazar jets and the prompt emission of GRBs is generally interpreted as originating from electrons accelerated to non-thermal energy distributions at relativistic internal shocks. In addition, shocks with mildly relativistic speeds are expected to form in the terminal regions of AGN jets, where prominent hotspots are ob-

served. If the observed emission in these sources were to be explained by diffusive shock acceleration (DSA) operating at relativistic shocks, this would require nearly parallel and/or weakly magnetised shock configurations (Pohl et al. 2020).

However, relativistic shocks in astrophysical jets are most likely magnetised and quasi-perpendicular (i.e., superluminal, see Section 3.3). Under these conditions, DSA is expected to be highly inefficient and therefore unlikely to be the dominant mechanism responsible for producing the observed population of radiating electrons. This motivates the exploration of alternative acceleration processes. Wakefield acceleration is one such mechanism, but, as discussed below in Section 4.2, it is not efficient in a mildly relativistic regime. This thesis investigates another alternative process that can operate in strongly magnetised conditions. The foundations of this mechanism are presented in Sections 4.3 and 4.4.

4.2 Synchrotron maser instability and wakefield acceleration

RELATIVISTIC shocks with magnetizations $\sigma \gtrsim 10^{-2}$ are mediated by the magnetic reflection of the incoming flow off the shock-compressed magnetic field. Reflected particles gyrate at the shock front, forming a ring-like distribution that is unstable to the *synchrotron maser instability* (SMI, Langdon et al. 1988). For a steady inflow of upstream plasma, the electron-driven SMI produces a train of high-frequency X-mode waves that propagate upstream at the speed of light, filling the upstream region as *precursor waves* (Hoshino and Arons 1991). The SMI also emits O-mode precursor waves, although linear theory predicts their amplitude to be lower than that of the X-mode (Melrose et al. 1984). Recent PIC simulations have shown, however, that for very low electron magnetisations $\sigma_e \lesssim 10^{-2}$, the O-mode energy can exceed that of the X-mode due to mode conversion in turbulent shock conditions (Iwamoto et al. 2018). The amplitude of the emitted waves is strongly dependent on the electron magnetisation. Iwamoto et al. 2018 simulated a highly relativistic shock ($\gamma_0 = 40$) over a range of magnetisations from $\sigma_e = 10^{-3}$ to $\sigma_e = 1$. For $\sigma_e = 10^{-3}$, the precursor-wave amplitude relative to the upstream magnetic field is $\delta B/B_0 \approx 2.2$, and it decreases significantly with increasing magnetisation. At $\sigma_e = 1$, the precursor amplitude is $\delta B/B_0 \approx 0.2$, nearly an order of magnitude smaller. Ion gyration excites magnetosonic waves that propagate slower than the shock in the strong shock limit and therefore cannot travel upstream.

As the large-amplitude precursor waves propagate upstream in an ion-electron plasma, incoming electrons interact strongly with the waves and undergo intense transverse oscillations, whereas ions, owing to their much larger inertia, remain largely unaffected. This difference in response produces a time-dependent charge drift between electrons and ions, which gives rise to a longitudinal electrostatic field known as a *wakefield*. The resulting electrostatic wakefield can subsequently interact with electrons and ions and accelerate them to high energies through the process of *wakefield acceleration* (WFA).

The SMI and WFA have been extensively studied at ultra-relativistic shocks (Hoshino 2008; Lyubarsky 2006; Iwamoto et al. 2019; Iwamoto et al. 2018; Iwamoto et al. 2017b; Plotnikov et al. 2018). In ion-electron plasmas, the precursor wakefield carries sufficient energy to drive WFA when the electron magnetisation satisfies $\sigma_e \gtrsim 0.1$, and studies have

explored conditions up to $\sigma_e = 5$. At high magnetisations, the wakefield energy can greatly exceed that in pair plasmas, an enhancement attributed to a positive feedback loop between ion-electron coupling and the induced wakefields (Iwamoto et al. 2019; Iwamoto et al. 2022; Iwamoto et al. 2024). In ultra-relativistic shocks with $\gamma_{sh} \gg 1$, these turbulent wakefields can efficiently heat electrons, driving energy equipartition with ions prior to the shock encounter, while simultaneously enabling the acceleration of non-thermal particles.

Mildly relativistic shocks have also been found to emit precursor waves, as the ultra-relativistic shocks. Ligorini et al. 2021a performed 2D PIC simulations for a shock with Lorentz factor $\gamma_{sh} = 3.3$ and magnetisation $\sigma = 0.1$, considering both out-of-plane and in-plane strictly perpendicular magnetic field configurations. In the out-of-plane setup, X-mode waves are excited, with amplitude $\delta B/B_0 = 0.2$, much lower than in the ultra-relativistic case. In the in-plane configuration, X-mode waves are also produced, but O-mode waves appear as well due to magnetic turbulence at the shock. The ponderomotive force of these waves excites a longitudinal electric field, generating a wakefield with a small amplitude $E_x/(B_0 c) \approx 5 \times 10^{-3}$. As shown in Ligorini et al. 2021b, upstream electrons interact with the precursor waves and are energised by the wakefield. However, the resulting electron energy spectra indicate that WFA is inefficient in this regime, with electrons reaching at most about twice the average ion energy. Consequently, perpendicular mildly relativistic magnetised shocks are ineffective as particle accelerators.

4.3 Electron trapping and acceleration in an oblique magnetosonic shock wave

ION-ELECTRON plasma shocks develop a cross-shock potential due to charge separation, provided that the magnetic field is not very weak, as explained in Section 3.2.2. This potential reflects some ions back upstream and also facilitates the electron advection into the downstream region, since it accelerates electrons towards it. However, it is possible that magnetised shocks can also reflect electrons through a mechanism different from that responsible for ion reflection. Bessho and Ohsawa (1999) and Bessho and Ohsawa (2002) have demonstrated that a large-amplitude magnetosonic wave propagating obliquely to a strong magnetic field, such that $\sigma_e \gtrsim 1$, can reflect electrons near the end of the shock pulse. This occurs because the propagation of such large-amplitude wave is not perfectly stationary and fluctuations of the electromagnetic fields can alter the potential there. The reflected electrons become trapped within the shock potential and are subsequently accelerated to ultrarelativistic energies by a strong electric field in the shock region. It has been shown that electron energies can be extremely high at shocks with obliquity angles approaching the threshold angle, θ_{tr} , between subluminal and superluminal configurations (see Section 3.2.3). However, the electron reflection is difficult in superluminal shocks.

The electron trapping and acceleration have been shown to occur in non-relativistic and transrelativistic shocks with shock velocities up to $v_{sh} \sim 0.7c$, corresponding to Lorentz factors of up to $\gamma_{sh} \sim 1.4$. The maximum electron energy increases with v_{sh} (Bessho and Ohsawa 2002), suggesting that electron acceleration may become even more efficient in rel-

ativistic shocks. It will be demonstrated in this thesis that mildly relativistic shocks, which encompass a broad range of subluminal configurations (see Section 3.3), allow this mechanism to operate with remarkable efficiency.

In this section, the theory of electron trapping and acceleration in an oblique shock wave is analysed following the approach of Bessho and Ohsawa (1999), Bessho and Ohsawa (2002), and Zindo et al. (2005) (see also Ohsawa 2014). However, unlike those studies, the derivation presented here is fully relativistic.

4.3.1 Basic properties of oblique shocks

Let us assume a one-dimensional (1D) magnetosonic shock wave (or a magnetosonic pulse) propagating in the x -direction with velocity v_{sh} in the external oblique magnetic field in the $(x - z)$ plane, $\mathbf{B}_0 = (B_{0x}, 0, B_{0z})$, in an electron-proton plasma. One dimensionality means that the field quantities depend on the x -coordinate only, and so $\partial/\partial y = \partial/\partial z = 0$. In the wave frame of reference the shock is stationary and the upstream plasma moves at a velocity $v_{up} = -v_{sh}$. In this frame the background magnetic field forms an angle θ_{Bn} with the shock normal. The magnetic field components B_y and B_z in the pulse region can have finite values, whereas $B_x = B_{0x}$, due to $\nabla \cdot \mathbf{B} = 0$. The wave electric fields have finite E_x and $E_y = E_{0y}$ components, and $E_z = 0$, where $E_{0y} = -v_{sh}B_{0z}$ is the motional electric field. A sketch of the geometry is presented in Figure 4.1.

As a consequence of the sharp rise in the magnetic field strength in the front of the magnetosonic shock wave, an electrostatic potential is built up at the wave's location – marked in blue in the right panel of Figure 4.1. The potential is due to a strong electrostatic gradient of the longitudinal electric field (orange line in Fig. 4.1) caused by the charge separation, which in turn is caused by the difference in gyroradius of both ions and electrons, $r_{gi} \gg r_{ge}$. If the shock wave is of large amplitude, $B_M > B_0$, the wave potential has a similar profile to the transverse magnetic field B_z . In the strong Mach number regime, $M_A > 2$, and for a perpendicular magnetic field, $B_{0x} = 0$, the electrostatic potential depends linearly on the value of the magnetic field at the shock wave, B_M (Nakazawa and Ohsawa 1997):

$$\phi = \frac{B_0}{\mu_0 en_0} \left(\frac{B_M}{B_0} - 1 \right). \quad (4.12)$$

This expression holds also for an oblique magnetic field, $B_{0z} \gtrsim B_{0x}$, as long as $B_{Mz} \lesssim B_M$ and the wave amplitude is large (Miyahara et al. 2003). The plasma density and the transverse electric field, E_y , profiles are similar to those of B_z and ϕ , whereas B_y and E_x are nearly proportional to $\partial B_z / \partial x$. Hence, they become zero at positions at which B_z takes its maximum or minimum values. These relations were derived for small-amplitude waves and 1D simulations of nonrelativistic shocks demonstrated that they approximately hold also in large-amplitude magnetosonic waves (Bessho and Ohsawa 1999; Bessho and Ohsawa 2002; Zindo et al. 2005; Ohsawa 2014). It will be shown in Section 6.3 that these relations are similar also at magnetised mildly relativistic oblique shocks.




Figure 4.1: Sketch of a profile of the transverse magnetic field B_z in a stationary magnetosonic shock wave propagating at a speed v_{sh} in a background magnetic field $\mathbf{B}_0$ inclined at an angle θ_{Bn} to the x -axis (left). The value of B_z at the shock wave peak is indicated as B_M . Corresponding longitudinal electrostatic field, E_x , and the associated electrostatic potential, ϕ , in the wave (right). The horizontal dashed line marks the value of zero, while vertical dashed lines show the first maximum, $x = x_M$, and the first minimum, $x = x_m$, of the potential. Note that the $B_z - B_{0z}$ profile, marked with a thin green line, is similar to the electric potential profile, albeit with a larger amplitude.

4.3.2 Electron reflection, trapping and acceleration

If the background magnetic field B_0 is strong such that $\sigma_e \gtrsim 1$, electrons will likely follow magnetic field lines and the guiding-center theory becomes useful to analyse their motion (Grebogi and Littlejohn 1984; Cary and Brizard 2009). In the guiding-centre theory, a particle is affected by a magnetic field that makes it gyrate around the gyro-centre, denoted as $\mathbf{R}$, with a certain velocity $v_\perp$. The particle can also have a component of its velocity parallel to the magnetic field, $v_\parallel$, so the gyro-centre motion becomes a straight line and the particle's full trajectory is a helix. To understand the motion of a particle in the electromagnetic field, the most convenient approach is the Hamiltonian theory, in which different set of variables, named *phase-space coordinates* $(\mathbf{q}, \mathbf{p})$, satisfy the following relations:

$$\frac{d\mathbf{q}}{dt} = \frac{\partial \mathcal{H}}{\partial \mathbf{p}}, \quad \frac{d\mathbf{p}}{dt} = -\frac{\partial \mathcal{H}}{\partial \mathbf{q}}, \quad (4.13)$$

where $\mathcal{H}$ is the Hamiltonian. Phase-space coordinates can be any set of variables that depend on the usual coordinates, as long as they preserve phase-space. In the guiding-centre theory, and only for weakly time-dependent fields $\mathbf{B}$ and $\mathbf{E}$, the phase-space coordinates are $z = (\mathbf{R}, p_\parallel; \mu, \xi; t)$ where $p_\parallel = \gamma m_e v_\parallel$ is the parallel momentum, μ is the magnetic moment defined as:

$$\mu = \frac{m_e |u_\perp|^2}{2B(\mathbf{R}, t)}, \quad (4.14)$$

where $u_\perp$ is the perpendicular normalised four-velocity, $p_\perp = m_e u_\perp = \gamma m_e v_\perp$, ξ is the gyro-phase of the particle that defines the angular position of the gyro-motion, $d\xi/dt = \Omega$, and t is the time. For convenience, a non-covariant formulation of the Hamiltonian is used. While with this choice the derivations are not directly applicable to other frame of references, they are greatly simplified, as opposed to the covariant formulation. The Hamiltonian for an electron is:

$$\mathcal{H} = \gamma m_e c^2 + e\phi(\mathbf{R}, t), \quad (4.15)$$

where $\gamma = \sqrt{1 + p^2/(mc)^2} = \sqrt{1 + (2/mc^2)\mu B(\mathbf{R}, t) + p_\parallel^2/(mc)^2}$ is the Lorentz factor of the particle and ϕ is the electrostatic potential such that:

$$\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t}, \quad (4.16)$$

where $\mathbf{A}$ is the magnetic vector potential. From the Hamiltonian, it is possible to compute evolution of $\mathbf{R}$ and $p_\parallel$ (Cary and Brizard 2009, eqs. 6.24-25):

$$\mathbf{V} = \frac{d\mathbf{R}}{dt} = \frac{p_\parallel}{\gamma m_e} \frac{\mathbf{B}^*}{B_\parallel^*} + \mathbf{E}^* \times \frac{\hat{\mathbf{b}}}{B_\parallel^*}, \quad (4.17)$$

$$\mathbf{F}_{L,\parallel} = \frac{dp_\parallel}{dt} = e\mathbf{E}^* \cdot \frac{\mathbf{B}^*}{B_\parallel^*}, \quad (4.18)$$

where $\mathbf{E}^*$ and $\mathbf{B}^*$ are the effective fields, which are defined as the contribution of the background field and the particle's effect on the fields:

$$\mathbf{B}^* \equiv \mathbf{B} + \frac{cp_\parallel}{e} \nabla \times \hat{\mathbf{b}}, \quad (4.19)$$

$$B_\parallel^* \equiv \hat{\mathbf{b}} \cdot \mathbf{B}^*, \quad (4.20)$$

$$\mathbf{E}^* \equiv \mathbf{E} - \frac{1}{e} \left(mc^2 \nabla \gamma - p_\parallel \frac{\partial \hat{\mathbf{b}}}{\partial t} \right), \quad (4.21)$$

where $\hat{\mathbf{b}} = \mathbf{B}/|B|$. Since the shock wave is governed by the plasma bulk, it is assumed that a single particle does not significantly modify the electromagnetic field within the timescales of the considered process. Therefore, Equations 4.17 and 4.18 can be written as

$$\mathbf{V} = \frac{d\mathbf{R}}{dt} = \frac{p_\parallel}{\gamma m_e} \frac{\mathbf{B}}{B} + \frac{\mathbf{E} \times \mathbf{B}}{B^2}, \quad (4.22)$$

$$\mathbf{F}_{L,\parallel} = \frac{dp_\parallel}{dt} = -e \frac{\mathbf{E} \cdot \mathbf{B}}{B}. \quad (4.23)$$

Equation 4.22 represents a motion along the magnetic field and the well-known $\mathbf{E} \times \mathbf{B}$ drift. The evolution of the energy of the particle is calculated using Equation 4.18:

$$\begin{aligned} \frac{dp_{\parallel}}{dt} &= mc \frac{d}{dt} \left(\sqrt{\gamma^2 - 1 - \frac{2}{mc^2} \mu B} \right) = -e \frac{\mathbf{E} \cdot \mathbf{B}}{B} \rightarrow \\ &\frac{d}{dt} \left(\gamma^2 - \frac{2}{mc^2} \mu B \right) = -\frac{2e}{(mc)^2} \frac{\mathbf{E} \cdot \mathbf{B}}{B} p_{\parallel} \rightarrow \\ &\frac{d(\gamma m_e c^2)}{dt} - \frac{1}{\gamma} \frac{d(\mu B)}{dt} = -e \frac{\mathbf{E} \cdot \mathbf{B}}{B} \frac{p_{\parallel}}{\gamma m_e}. \end{aligned} \quad (4.24)$$

For convenience, let us take the x -coordinate of Equation 4.22, $\mathbf{V}_x \equiv v_x$:

$$v_x = \frac{p_{\parallel}}{\gamma m_e} \frac{B_x}{B} + \frac{(\mathbf{E} \times \mathbf{B})_x}{B^2} \rightarrow \frac{p_{\parallel}}{\gamma m_e} = \frac{B}{B_x} (v_x - v_{d,x}), \quad (4.25)$$

where $\mathbf{v}_d = (\mathbf{E} \times \mathbf{B})/B^2$ is the $\mathbf{E} \times \mathbf{B}$ drift velocity. Equation 4.25 can be substituted in equation 4.24 to obtain

$$\frac{d(\gamma m_e c^2)}{dt} - \frac{1}{\gamma} \frac{d(\mu B)}{dt} = -e \frac{\mathbf{E} \cdot \mathbf{B}}{B_x} (v_x - v_{d,x}). \quad (4.26)$$

We now introduce the following quantity:

$$F(s, t) = - \int_C \frac{\mathbf{E}(x, t) \cdot \mathbf{B}(x, t)}{B} ds = - \int_x \frac{\mathbf{E}(x, t) \cdot \mathbf{B}(x, t)}{B_x} dx, \quad (4.27)$$

where $ds = (B/B_x) dx$. Equation 4.27 represents the electric field potential in the direction parallel to the magnetic field if the latter is constant, $\partial \mathbf{A} / \partial t = 0$. Note, that $E_{\parallel}$ in Equation 4.27 is a mixture of longitudinal and transverse electric fields, so F is a pseudo-potential and not the ordinary electric field potential. Also note, that F attains its maximum value within the main pulse, near x_M , where the magnetic field B_z and the electric potential ϕ also reach their maxima (see Figs. 4.1b and 4.2).

By computing the evolution of F :

$$\frac{dF}{dt} = \frac{\partial F}{\partial t} + \frac{\partial F}{\partial x} \frac{dx}{dt} = \frac{\partial F}{\partial t} + -\frac{\mathbf{E} \cdot \mathbf{B}}{B_x} v_x, \quad (4.28)$$

we can now substitute it in Equation 4.26:

$$\begin{aligned} \frac{d(\gamma m_e c^2)}{dt} - \frac{1}{\gamma} \frac{d(\mu B)}{dt} &= e \frac{\partial F}{\partial t} \frac{dx}{dt} - e \frac{\partial F}{\partial t} v_{d,x} = e \frac{dF}{dt} - e \frac{\partial F}{\partial t} - \frac{\partial F}{\partial x} v_{d,x} \rightarrow \\ &\frac{d(\gamma m_e c^2)}{dt} - \frac{1}{\gamma} \frac{d(\mu B)}{dt} - e \frac{dF}{dt} = -e \left(\frac{\partial F}{\partial t} + \frac{\partial F}{\partial x} v_{d,x} \right), \end{aligned} \quad (4.29)$$

which resembles Equation 8 in Zindo et al. 2005, albeit here a relativistic formulation is used. The first term on the LHS of equation 4.29 is the the change in total energy of the

particle, the second term the change in energy in the perpendicular momentum, and the third term in the potential energy taken from, or given to the electron, related exclusively to the parallel electric field $E_{\parallel}$. Overall, the LHS shows the evolution of the electron's energy parallel to $\mathbf{B}$. The second term on the LHS may be neglected,

$$\frac{1}{\gamma} \frac{d(\mu B)}{dt} = \frac{1}{\gamma} \frac{d(m_e |u_{\perp}|^2/2)}{dt} = \frac{m_e |u_{\perp}|}{\gamma} \frac{d|u_{\perp}|}{dt} = m_e |v_{\perp}| \frac{d|u_{\perp}|}{dt} \approx 0, \quad (4.30)$$

since $v_{\perp} \approx 0$ in the strong magnetic field limit. Consequently, and given that the total particle kinetic energy is dominated by the parallel component of the velocity, the particle motion can be effectively regarded as one-dimensional. In this case Equation 4.29 simplifies to:

$$\frac{d}{dt} (\gamma m_e c^2 - eF) = -e \left(\frac{\partial F}{\partial t} + \frac{\partial F}{\partial x} v_{d,x} \right). \quad (4.31)$$

Note however, that as electrons become highly relativistic due to acceleration at the shock, their perpendicular four-velocities, $u_{\perp} = \gamma v_{\perp}$, increase significantly. The term in Equation 4.30 cannot then be neglected.

Upstream electrons that flow into the shock encounter a potential peak at $x = x_M$ (see right panel of Figure 4.1), and then a trough at $x = x_m$ right after. As it will be shown later in more details, an electron traversing the shock at its first encounter gains energy almost exclusively by the electrostatic potential. For the electron motion in the x -direction considered here, and under the assumption of stationary pulse propagation, the potential, ϕ , is positive in the wave and the electron is transmitted downstream. However, this is not always the case in a nonstationary pulse. Nonstationary waves can produce fluctuations in the electric potential, causing it to take locally small values (see references in, e.g., Bessho and Ohsawa 1999; Ohsawa 2014). It will be demonstrated below that if such nonstationary dips are created near the end of the large-amplitude pulse, at the trough of the potential at x_m , an electron can be reflected at this location, become subsequently trapped in the pulse and accelerated. Since the magnetic field is considered strong, the upstream electrons that reach the wave move almost exclusively along the magnetic field lines, rather than in the x -direction. Therefore, the relevant quantity is not the electric potential, ϕ , but rather the pseudo-potential, F . Analogously to the motion in the electric potential ϕ , the electron is reflected if $F_m < 0$ (i.e., lower than the upstream value $F_0 = 0$), and advected downstream if $F_m > 0$. A schematic illustration of these conditions is shown in Figure 4.2.

If the electron is reflected, the x -component of its velocity becomes zero at the trough, $v_x = 0$ (Equation 4.25), changing the value from negative to positive. From Equation 4.25, it is then possible to obtain the parallel velocity at the trough:

$$\frac{p_{\parallel}}{\gamma m_e} = v_{\parallel} = -\frac{E_{y,0}}{B_{0x}} \frac{B_{z,m}}{B_m} = v_{sh} \tan(\theta_{Bn}) \sin(\theta_m) > 0, \quad (4.32)$$

where B_m is the magnetic field at the reflection point x_m , θ_m is the field obliquity at x_m and $\theta_m > \theta_{Bn}$. At this location, the value of F is

$$F_m = - \int \frac{E_x B_{0x} + E_{0y} B_y}{B_{0x}} dx = - \int E_x dx + \frac{v_{sh} B_{0z}}{B_{0x}} \int B_y dx = \phi_m - v_{\parallel} \frac{B_m}{B_{mz}} A_{mz}, \quad (4.33)$$




Figure 4.2: Schematic profiles of the pseudo potential F differing by their values at the point x_m near the end of the main pulse. The profile in the upper panels shows a dip with $F_m < 0$, that provides conditions for the electron reflection near x_m . Such profiles are perturbations to F due to nonstationarity, after which F quickly recovers to the time-average profile shown in the lower panel with $F_m > 0$. In such a wave all electrons are advected downstream.

where $A_{mz} = -\int B_{my} dx > 0$ is the z -component of the magnetic vector potential and Equation 4.32 is used. Reflection occurs if $F_m < 0$. Then, there exists a threshold parallel velocity $v_{\parallel}^*$ such that

$$F_m(v_{\parallel}^*) = 0 \rightarrow v_{\parallel}^* = \phi_m \sin \theta_m A_{mz}^{-1} > 0. \quad (4.34)$$

Thus, electrons with parallel velocity larger than the threshold $v_{\parallel}^*$ are reflected at $x = x_m$ and go back towards the shock with positive v_x ; those with $v_{\parallel} < v_{\parallel}^*$ instead are advected downstream. This can be understood as follows: an electron that closely follows the magnetic field lines (with high $v \approx v_{\parallel}$ and low magnetic moment μ) will most probably be reflected, as it surpasses the threshold. On the other hand, an electron with large gyroradius – high μ and small $v_{\parallel}$ – will most likely be advected downstream, as it does not reach the threshold.

The threshold $v_{\parallel}^*$ can be modified in several ways. A lower potential ϕ_m makes reflection easier, as the threshold decreases. Note that $\phi_m > 0$ always, while F can still take negative values. Similarly, a larger A_{mz} also facilitates reflection. The most effective way to produce a strong variation in the electrostatic potential (i.e., large ϕ_M and small ϕ_m) is through a strong magnetosonic shock wave, which simultaneously leads to a significant increase in the vector potential A_{mz} .

As is shown below, after reflection, the electrons become deeply trapped within the main pulse region, where the pseudo-potential F is large. They then travel towards the leading edge of the shock, where they experience another reflection in the opposite direction. Even when the values of F have returned to their typical, time-averaged positive levels by the time the electrons reach the rear of the main pulse, they are once again reflected forward (Bessho and Ohsawa 2002). Note that the first reflection off the trough is possible at both subluminal and superluminal shocks. However, the latter are faster than particles that follow magnetic field lines and so reflected electrons cannot move towards the front of the pulse and are eventually advected downstream.

It is possible to estimate the maximum electron kinetic energy obtained during the trap-

ping and acceleration process following the electron reflection. The energy evolution of an electron at the shock is

$$m_e c^2 \frac{d\gamma}{dt} = -e \mathbf{E} \cdot \mathbf{v} \rightarrow (\gamma - \gamma_{\text{up}}) m_e c^2 = e(\phi - \phi_0) - e E_{0y} \int v_y dt. \quad (4.35)$$

Using the guiding-center approximation (eq. 4.22) for the x and y component, the integral on the RHS gives

$$\int v_y dt = \int v_{\parallel} \frac{B_y}{B} dt - \int \frac{E_x B_z}{B^2} dt = \int v_{\parallel} \frac{B_y}{B} dt - \int \frac{E_x}{E_{0y} + v_{\parallel} \frac{B_{0x} B}{B_z}} dx, \quad (4.36)$$

where $dt = dx/v_x$. The first term on the RHS represents the potential difference of the longitudinal electric field E_x , while the second term is the energy gained by the electron from the motional electric field E_{y0} . The latter depends on the distance travelled by the particle; which in turn is determined by the $\mathbf{E} \times \mathbf{B}$ drift and the y -component of $v_{\parallel}$.




Figure 4.3: Schematic illustration of the path of an upstream electron crossing a strong magnetosonic shock wave and being reflected and trapped by the cross-shock potential. The black line with arrowheads indicates the electron trajectory. Along this trajectory, five major crosses mark key points labeled from A to E. The dashed line shows the pseudo potential F . The background shows a color gradient representing the longitudinal electrostatic field E_x : red corresponds to $E_x > 0$, blue to $E_x < 0$, and green to $E_x \approx 0$. The location of strong E_x gradient corresponds to the maximum value of F , while the trough at C is a minimum. In contrast to the stationary case illustrated in the bottom panel of Fig. 4.1, the pseudo potential F in this image has been perturbed and its value at the trough is lower than at point A, $F_C < F_A$, which allows for electron reflection. Once the electron reaches point E, it starts the cycle once again, where it will reflect, this time, at point C_2 .

Figure 4.3 shows a schematic trajectory of an upstream electron undergoing this process. It is convenient to compute the energy increases in steps, denoted by the marks A to E. Assume an electron starting at point A and moving towards the shock, where it is attracted by the potential as it moves towards point B where the potential, and therefore the pseudo-

potential, is maximum. The energy difference is:

$$\Delta E_{AB} = e(\phi_B - \phi_A) - eE_{0y} \int_A^B v_{\parallel} \frac{B_y}{B} dt + eE_{0y} \int_A^B \frac{E_x}{E_{0y} + v_{\parallel} \frac{B_{0x}B}{B_z}} dx. \quad (4.37)$$

The potential in the first term difference is large and positive, $\phi_B \gg \phi_A$. For the second and third term, one may assume that the upstream electron moves together with the bulk fluid, which is justified in a cold plasma. Far upstream of the shock, the fluid motion is in the x -direction, so in the guiding-center approximation, there is no motion in the z -direction, i.e., $\mathbf{R}_z = Z = 0$. Therefore, from Equation 4.22 one obtains

$$\frac{dZ}{dt} = v_{\parallel} \frac{B_z}{B} - \frac{E_{0y}B_{0x}}{B^2} = 0 \rightarrow v_{\parallel} = \frac{E_{0y}B_{0x}}{BB_z} < 0. \quad (4.38)$$

Using the previous equation and the fact that $B_y < 0$ at the front of the pulse $x_B < x < x_A$ due to Ampère-Maxwell law, the second term on the RHS of Equation 4.37 becomes positive

$$-eE_{0y} \int_A^B v_{\parallel} \frac{B_y}{B} dt > 0. \quad (4.39)$$

The parallel velocity of Equation 4.38 can be substituted in the third term of Equation 4.37 to obtain

$$e \int_A^B \frac{E_x}{1 + (B_{0x}/B_z)^2} dx < 0, \quad (4.40)$$

which is negative since the particle goes from A to B and $x_B < x_A$. The electron gains energy from the potential but loses energy by the $E_x \times B_z$ drift (third term) and the energy gain almost cancels out. The second term, albeit positive, becomes negligible since the magnetic field is compressed in z -direction ($B_z \sim B$) and thus the contribution by B_y is almost zero, $B_y/B \sim 0$.

The energy difference from B to C , ΔE_{BC} , yields similar results but with an opposite sign. Changing the limits in Equation 4.37 to B and C , one obtains a negative potential difference, $\phi_C < \phi_B$, the second term is nearly zero, and in the third term, the denominator is positive while the numerator is negative. Therefore, the energy gain in the first step, $A \rightarrow B$, is approximately canceled by the energy loss in the second step, $B \rightarrow C$.

At point C , the particle is reflected, as one can see from Equations 4.32 and 4.34. Then, electron energisation proceeds up to point D , where the electron once again reaches the location of the potential and pseduo-potential maximum:

$$\Delta E_{CD} = e(\phi_D - \phi_C) - eE_{0y} \int_C^D v_{\parallel} \frac{B_y}{B} dt + eE_{0y} \int_C^D \frac{E_x}{E_{0y} + v_{\parallel} \frac{B_{0x}B}{B_z}} dx. \quad (4.41)$$

As in the first step, $A \rightarrow B$, the potential difference is large and positive $\phi_D = \phi_B \gg \phi_C$. After reflection, v_x becomes positive,

$$v_x = \left(v_{\parallel} - v_{sh} \tan(\theta_{Bn}) \frac{B_z}{B} \right) \frac{B_{0x}}{B}. \quad (4.42)$$

Note, that due to the magnetic field compression, B_z around point C takes larger values than in the upstream, $B_z \sim B$, and $B_{0x}/B \sim 0$. Therefore, v_x can become quite small – especially for oblique subluminal shocks with obliquity $\theta_{Bn} \sim \theta_{tr}$. This corresponds to a small denominator in the third term on the RHS of Equation 4.41. Because the velocity of the guiding center position of the electron, v_x , is small, it takes a long time for the electron to reach the x –position of point D from point C. During this time, the electron travels a long distance in the y –direction due to the $\mathbf{E} \times \mathbf{B}$ drift. As a result, it gains a significant amount of energy from E_{y0} :

$$\Delta E_{CD} = e(\phi_D - \phi_C) - eE_{0y}(y_D - y_C) \gg 0. \quad (4.43)$$

Since $v_{\parallel}$ is also large, the energy increase due to the parallel electric field in second term is also significant. Therefore, the energy gain by this step is much larger than in the previous ones, $\Delta E_{CD} \gg \Delta E_{AB}, \Delta E_{BC}$.

In order to show the dependence of the energy gain ΔE_{CD} on the obliquity angle, let us assume an ultra-relativistic electron, $v_{\parallel} \approx c$, moving from C to D . If the magnetic field obliquity is not small, the magnetic field compression at the wave is significant, so that $B_{z,M} \sim B_M$, $B_y \sim 0$, allowing the denominator in the third term of Equation 4.41 to be simplified as:

$$E_{0y} + v_{\parallel} \frac{B_{0x}B}{B_z} \approx E_{0y} + cB_{0x}, \quad (4.44)$$

which allows the integral in the third term to be calculated. Noting that $\phi_C \sim 0$, the energy gain from $C \rightarrow D$ is

$$\Delta E_{CD} = e\phi_D + \frac{-e\phi_D}{1 + \frac{cB_{0x}}{E_{0y}}} = e\phi_D \frac{1}{1 + \frac{E_{0y}}{cB_{0x}}} = \frac{e\phi_D}{1 - \frac{v_{sh}}{c} \frac{B_{0z}}{B_{0x}}} = \frac{e\phi_D}{1 - \frac{\tan \theta_{Bn}}{\tan \theta_{tr}}}, \quad (4.45)$$

where $E_{0y} = -v_{sh}B_{0z}$ and $v_{sh}/c = \tan \theta_{tr}$. From this expression, one can see that the energy gain becomes extremely large for fast shock waves, $v_{sh} \rightarrow c$, and highly oblique magnetic fields close to the threshold angle, $\theta_{Bn} \rightarrow \theta_{tr}$. Note that if $\theta_{Bn} > \theta_{tr}$, the denominator becomes negative and $\Delta E_{CD} < 0$, which corresponds to a superluminal shock. In this case, electrons lose energy, which is transferred to the shock.

Once the electron reaches point D, the electrostatic field in the upstream-side begins to oppose its motion in the x –direction, removing energy from it until it reaches point E . The electron also undergoes an $\mathbf{E} \times \mathbf{B}$ drift in the negative y –direction, which further reduces its energy. It will be shown later that the electron is reflected once again, this time at point E . The energy difference is therefore:

$$\Delta E_{DE} = e(\phi_E - \phi_D) - eE_{0y} \int_D^E v_{\parallel} \frac{B_y}{B} dt + eE_{0y} \int_D^E \frac{E_x}{E_{0y} + v_{\parallel} \frac{B_{0x}B}{B_z}} dx. \quad (4.46)$$

The first term on the RHS is now negative, indicating that part of the kinetic energy is converted back into potential energy. In the second term, B_y becomes negative, while $v_{\parallel}$ remains

positive as the electron continues to follow the magnetic field lines; therefore, the second term is negative overall. In the third term, E_x is now positive. Although $v_{\parallel}$ is positive and large at point D, it rapidly decreases as the particle moves forward along x . This corresponds to a decreasing v_x up to point E in the upstream side of the potential peak. While B_z also decreases, if the motional electric field is not strong, $|E_{0y}| < v_{\parallel} B_{0x}$, the denominator of the third term is positive. Consequently, the third term is also negative and the overall energy gain is negative.

Since both second and third terms on the RHS of Equation 4.46 are negative, the electron performs a motion in the negative y -direction, $y_E - y_D < 0$, as can be seen in the sketch of Figure 4.3. However, if the motional electric field is large – either because the magnetic field is strong or the plasma is relativistic – the denominator of the third term can be negative and therefore the third term becomes positive. If this third term is large enough, it can cancel out with the second term and can even dominate over it, resulting in a positive motion along the y -direction. Therefore, the electron gains energy by E_{0y} , albeit small. It may also occur that, if the motion in x -direction is along a very short distance, the loss by the electrostatic potential becomes very small and the acceleration by E_{0y} may dominate. Therefore, in some special cases, the electron gains energy from D to E .

The motion from D to E , in contrast to the motion $C \rightarrow D$, is fast. The x -component of the parallel electron velocity decreases by the electrostatic potential since $\phi_D > \phi_E$, while the x -component of the $\mathbf{E} \times \mathbf{B}$ drift is large and negative $v_{dx} = E_{0y} B_z < 0$ – as is the case for $C \rightarrow D$. Both contributions causes a fast motion for the trajectory $D \rightarrow E$.

Comparing the total energy gain from C to E , one can see that the net gain is positive. From Equation 4.35, the energy gain associated with the electrostatic potential (first term) is positive since $\phi_E > \phi_C$ because point C is farther to the left with respect to the potential peak than point E is to the right $|x_D - x_M| > |x_E - x_M|$. This fact is also supported by simulations, which show that point E lies closer to the shock than point A and point C , which corresponds to $\phi_E > \phi_C > \phi_A = 0$. The second term is also positive since the particle moves overall in the positive y -direction

$$-eE_{0y} \int_C^E v_y dt = -eE_{0y}(y_E - y_C) > 0, \quad (4.47)$$

as shown in Figure 4.3. From C to D , the electron moves slowly in the x -direction towards the potential peak, allowing sufficient time for motion in y , whereas from D to E it moves rapidly to point E and thus has little time to drift in the negative y -direction.

Considering the entire trajectory $A \rightarrow E$, the energy gain along $A \rightarrow B$ and the energy loss along $B \rightarrow C$ approximately cancel each other. As a result, the net energy increase arises from the remaining segment $C \rightarrow E$, which is positive, as shown above. Consequently, the electron gains energy from A to E .

Once the electron reaches point E , it reverses its motion in the x -direction and propagates back toward the potential maximum. The condition for reflection at point E is analogous to Equation 4.32,

$$v_{\parallel} = v_{sh} \tan \theta_{Bn} \frac{B_z}{B} \rightarrow \beta_{\parallel} = \frac{\tan \theta_{Bn}}{\tan \theta_{tr}} \frac{B_z}{B}, \quad (4.48)$$

where the definition of the threshold angle is used (see Section 3.2.3). In contrast to the first reflection downstream of the shock, this reflection involves a change in the sign of the x -component of the electron velocity, from positive to negative.

Let us consider an obliquity angle slightly below the threshold, $\theta_{Bn} \lesssim \theta_{tr}$. In this case, the reflection condition reduces to $\beta_{\parallel} = B_z/B$. On the downstream side, this condition is not satisfied while the electron is being energised after the first reflection, since it has a large parallel velocity, $\beta_{\parallel} > B_z/B$. However, upon reaching the upstream side, the electron loses a small fraction of its parallel velocity, allowing the condition to be satisfied when $B_z \sim B$. Including the dependence on the obliquity angle through the factor $\tan \theta_{Bn} / \tan \theta_{tr}$, one finds that for smaller obliquity angles the reflection occurs at lower parallel velocities. This implies a longer deceleration phase and therefore a reflection point located farther upstream. Conversely, when the obliquity angle approaches the threshold angle, the reflection occurs very close to the potential peak at $x = x_M$.

If the obliquity angle exceeds the threshold angle, i.e. the shock is superluminal, then $\tan \theta_{Bn} / \tan \theta_{tr} > 1$ and the reflection condition is satisfied at higher, rather than lower, parallel velocities. In this case, the electron is reflected on the downstream side after the first reflection, while it is still gaining energy. As a result, electrons in superluminal shocks do not cross into the upstream region. This behaviour is also consistent with Equation 4.45, which shows that reflected electrons in superluminal shocks lose energy instead of gaining it.

The electron reflected on the upstream side subsequently crosses the potential maximum and moves back toward the downstream region, where it is reflected again at the potential minimum located at $x = x_m$. In contrast to upstream electrons reflected at point C , electrons that have already undergone acceleration are more easily reflected at a similar location, denoted C_2 , with $x_{C_2} \sim x_C$. For upstream electrons, reflection requires $F_C < F_A = 0$, implying that F_C must be negative. However, electrons reflected at point E originate from a region where $F_E > F_A$, so the reflection condition $F_E > F_{C_2}$ is satisfied more readily, since F_{C_2} can now be zero or even positive. As a result, these electrons can undergo repeated reflections between the potential extrema.

This process repeats multiple times, leading to successive energy gains. As the parallel velocity continues to increase, the reflection condition remains satisfied, causing the electron to become trapped. Once trapped, the electron is unlikely to escape and continues to gain energy efficiently.

Particle trapping may also occur for particles moving in the opposite direction, i.e., downstream particles streaming upstream. However, the likelihood of trapping for electrons moving from the downstream towards the shock is lower. The conditions are now less ideal and electrons may have a very low parallel velocity coming from downstream, which impairs electron reflection in the upstream-side even if $F_{dw} > F_{up}$. This probability will be even lower due to magnetic field turbulence that may be generated at the shock. The turbulence would disturb the electron trajectories and decrease the number of trapped electrons.

The mechanism derived in this section is based on the gyrokinetic approximation. Within this framework, the magnetic moment μ is, by definition, conserved in time. As a result, the process is adiabatic, at least to first order. Once electrons become trapped, they continue to gain energy but have a very limited probability of escape, implying that electron diffusion is

weak. As time progresses, an increasing number of electrons are trapped and accelerated, while those trapped earlier continue to gain energy. Consequently, electrons injected at earlier times reach higher energies, whereas those trapped later populate lower energies. This leads, ultimately, to the formation of a hard power-law energy spectrum. As will be shown in Section 6.3, the resulting distribution approaches a power law of the form $dN/dE \propto E^{-p}$ with index $p = 1$.

This mechanism shares certain similarities with other electron acceleration processes, in particular shock drift acceleration (SDA) and shock surfing acceleration (SSA):

- *Shock drift acceleration* (SDA) is a particle energisation mechanism operating in collisionless shocks. It relies on the presence of a magnetic field gradient across the shock transition and can accelerate both ions and electrons (Wu 1984; Leroy and Mangeney 1984; Ball and Melrose 2001). Consider an upstream electron approaching a shock with a quasi-perpendicular magnetic field. As it encounters the shock front, the electron experiences a rapid increase in magnetic field strength. This strong gradient generates a grad- B drift (Krauss-Varban and Wu 1989), causing the particle to drift along the shock surface. At the same time, the motional electric field constant in the shock frame acts opposite to the grad- B drift. The work done by this electric field on the drifting particle results in an increase in its kinetic energy. Ions also experience a grad- B drift; in their case, the drift direction aligns with the motional electric field, allowing them to be energised through the same mechanism.

In both SDA and the mechanism discussed in this section, electron energisation is ultimately driven by the motional electric field. The crucial difference between the two mechanisms lies in the role of the cross-shock electric potential. In classical SDA, the cross-shock potential is typically neglected, and particle energisation arises from the grad- B drift along the shock front. Although electrons can efficiently gain energy through this process, they are eventually either advected downstream or reflected upstream, limiting the duration of their interaction with the shock.

By contrast, the mechanism presented here requires a cross-shock potential that is sufficiently strong to reflect and trap electrons. In this case, the dominant drift is the $\mathbf{E} \times \mathbf{B}$ drift. As a result of this process, electrons remain trapped while they are accelerated for extended periods of time and do not readily escape, in contrast to the transient energisation characteristic of SDA.

- *Shock surfing acceleration* (SSA) also energises electrons through the motional electric field, but the underlying mechanism significantly differs. In high-Mach-number shocks, the Buneman instability can be excited – a type of two-stream instability driven by the interaction between reflected ions and upstream electrons. This instability generates electrostatic waves (often referred to as *Buneman waves*) that can trap electrons within their potential wells and accelerate them perpendicular to the magnetic field via the motional electric field (Matsumoto et al. 2012b).

While the trapping and acceleration of electrons in SSA bear some resemblance to the mechanism described here, a crucial distinction is that in SSA the electrostatic potential

is associated with a regular plasma wave rather than with a shock structure. Buneman waves typically have relatively small amplitudes, which limits both the number of electrons that can participate in SSA and the maximum energy they can attain. As a result, SSA is generally regarded as an injection mechanism rather than an efficient accelerator, particularly in non-relativistic shocks such as those in supernova remnants (SNRs) (Bohdan et al. 2019b; Bohdan et al. 2019a). By contrast, the mechanism presented here operates at lower Mach numbers, corresponding to stronger magnetic fields, and involves a shock wave with substantially larger amplitude.

When a particle is accelerated, its energy is extracted from the electromagnetic field. This is evident from Eq. 4.31, where the left-hand side represents the parallel energy of an electron, while the right-hand side involves the quantity F . If a non-negligible fraction of particles becomes trapped, the value of F is modified accordingly. This behaviour can be clarified by again considering Eq. 4.31. Consider perturbations of F between x_m and x_M whose characteristic timescale is much shorter than an ion gyroperiod, which determines the width of the electrostatic potential. Under these conditions, Eq. 4.31 can be written as

$$\frac{d}{dt} (\gamma m_e c^2 - eF) = \frac{d\epsilon}{dt} = -e \frac{\partial F}{\partial t}, \quad (4.49)$$

where $\epsilon = \gamma m_e c^2 - eF$ denotes the electron energy parallel to $\mathbf{B}$, including both kinetic and potential contributions. If an electron gains parallel energy, $d\epsilon/dt > 0$, then F must decrease, $\partial F/\partial t < 0$. Conversely, when the electron loses parallel energy, F increases. For an upstream electron, once reflection occurs, the particle travels from x_m to x_M and loses parallel energy, which is transferred to F . As a result, the profile of F increases slightly over the region $x = x_m - x_M$. This increase may raise the value of F above zero, making electron trapping more difficult. However, as discussed earlier, a trapped electron does not reach the upstream value of F but instead experiences a slightly larger value. Consequently, even if $F_m > 0$, trapped electrons remain confined and do not escape. Newly arriving upstream electrons, on the other hand, can escape more easily. As time evolves, F_m relaxes back to a lower value, allowing additional upstream electrons to become trapped. This process therefore proceeds cyclically.

The mechanism described in this section has been previously investigated by Bessho and Ohsawa 1999, Bessho and Ohsawa 2002, and Zindo et al. 2005 using one-dimensional PIC simulations. In particular, Bessho and Ohsawa 2002 examined a non-relativistic magnetosonic shock wave in a magnetic field with obliquity $\theta_{Bn} = 45^\circ$ and field components along x and z directions, for two Alfvénic Mach numbers, $M_A = 2.2$ and $M_A = 1.8$. In their setup, the electrons are magnetised ($\sigma_e = 9$), while, given the reduced mass ratio $m_i/m_e = 100$, the ions are only weakly magnetised ($\sigma_i = 0.09$). Their results show that a fraction of upstream electrons crossing the shock are reflected by the pseudo-potential, and subsequently become trapped within the shock region. This behaviour is evident in the phase-space distributions of both $p_{e,x}$ and $p_{e,z}$, corresponding to electrons moving along the magnetic field. The trapped electrons undergo acceleration due to the motional electric field and $\mathbf{E} \times \mathbf{B}$ drift, leading to a significant increase in energy. For the stronger shock case ($M_A = 2.2$), the maximum electron Lorentz factor reaches $\gamma_{e,\max} \approx 130$ at a simulation time of $t\omega_{pe} = 850$,

whereas for the weaker shock ($M_A = 1.8$) the maximum Lorentz factor is lower $\gamma_{e,\max} \approx 40$. The number of trapped electrons within the shock grows continuously over time, with trapping found to be more efficient in the weaker-shock case. In addition, shorter simulations spanning a range of shock velocities indicate that faster shocks generate stronger electrostatic potentials, which in turn allow trapped electrons to reach higher maximum energies.

Stationary, wave-like solutions – such as those assumed for the operation of this mechanism – can exist for perpendicular magnetic field configurations, even in the relativistic regime (Ohsawa 1986). Relativistic effects become important when $\sigma_e \gtrsim 1$. For oblique magnetic field geometries, the wave solution remains approximately stable. However, at sufficiently large magnetic field strengths ($\sigma_e \gg 1$) or for large wave amplitudes, non-stationary behaviour becomes significant and can affect the efficiency of the acceleration process (Lembege and Dawson 1989). One notable outcome of these effects is the emergence of electron wavelets at the leading edge of the shock. Several interpretations of this phenomenon were proposed by Lembege and Dawson (1989), although no definitive explanation was established. These wavelets may nonetheless influence electron injection into the acceleration process by perturbing the initial conditions, potentially suppressing or enhancing electron trapping.

Temperature effects also influence the efficiency of electron trapping. If the upstream plasma has a finite temperature, a fraction of the upstream electrons have higher velocities. Consequently, when these hotter electrons enter the shock region and begin to follow the magnetic field lines, they are more likely to be reflected at the end of the magnetosonic pulse, as the reflection condition (Equation 4.34) is more readily satisfied. In addition, the upstream plasma may be further heated by other processes occurring ahead of the shock. For instance, mildly relativistic shocks can generate precursor waves in both perpendicular and oblique magnetic field configurations (see Section 4.2). Although the wakefield acceleration (WFA) efficiency is relatively low in mildly relativistic shocks, these precursor waves can slightly pre-heat the electrons before they encounter the shock front. This pre-heating increases the probability of electron reflection, thereby increasing the fraction of electrons that become trapped.

4.4 Ion acceleration by a magnetosonic pulse train

IN a manner analogous to electrons, ions can also become trapped when the magnetisation is sufficiently high ($\sigma_i \gtrsim 1$). In this case, the ion-electron plasma is strongly magnetised, and the total magnetisation is approximately equal to the ion magnetisation ($\sigma \sim \sigma_i$). Consequently, electrons are highly magnetised, with $\sigma_e = \frac{m_i}{m_e} \sigma_i$, which, for a realistic mass ratio $m_i/m_e = 1836$, becomes about three orders of magnitude larger. However, the conditions required for ion trapping differ fundamentally from those for electrons. Instead of an electrostatic potential peak, ion trapping requires the presence of a potential trough. Such a structure cannot be formed by a single shock wave, as this only generates a potential peak. Instead, a sequence of magnetosonic pulses forming the full shock transition can create the troughs necessary to trap ions.

For such a structure to form, two main conditions must be satisfied: firstly, the presence of a strong magnetic field, determined by the ion magnetisation σ_i ; and secondly, the existence of a train of large-amplitude waves. Moreover, since the magnetic field is strong, the Alfvén speed is high, and the shock must therefore be at least mildly relativistic in order to have $M_A = v_{sh}/v_A > 1$. The right panel of Figure 4.1 illustrates such a trough, which is formed by two wave pulses and coincides approximately with the reflection point for electrons. Similar structures are observed in the simulations presented in this thesis (see Sections 6.2 and 6.3).

Upstream ions, like electrons, are guided along the magnetic field lines when interacting with the shock. However, unlike electrons, ions lose energy while moving toward the potential peak and are re-accelerated as the potential decreases beyond the peak. Using the pseudo-potential F is convenient for describing this process. After crossing the shock, ions may either be reflected and trapped or advected downstream, depending on the values of F_M and F_{M2} , which correspond to the first and second pseudo-potential peaks, respectively. Here, F_M is the same quantity used for electrons. If $F_M > F_{M2}$, ions are advected downstream, as the energy gained from the shock potential dominates. Conversely, if $F_M < F_{M2}$, the ion becomes trapped. Trapped ions undergo a similar energisation mechanism as electrons, gaining energy from the motional electric field. A key difference is that ions have a much larger gyroradius, typically comparable to the trough width. As they gain energy, their gyroradius increases and may eventually exceed the trough size, allowing them to escape either downstream or upstream.

Because ion trapping requires large-amplitude waves which are nonstationary, the non-stationarity effects and additional perturbations to the shock structure can perturb ion motion during acceleration. Scattering off these perturbations may redirect ions opposite to the motional electric field, causing energy loss. Consequently, the maximum ion energy is controlled by the scattering strength: ions may either be detrapped after reaching very high energies or scatter before attaining them.

CHAPTER 5

Particle-In-Cell method for plasma shock simulations

PLASMAS can be studied through both experimental and theoretical approaches. Experimental techniques have advanced significantly, allowing plasmas to be generated in laboratory settings. However, the physical conditions achievable in the laboratory differ substantially from those in astrophysical environments, and replicating such extreme conditions remains infeasible. Consequently, theoretical approaches are essential for studying astrophysical plasmas. Yet, plasmas are highly nonlinear systems with many degrees of freedom, making purely analytical treatments impractical. Numerical methods thus provide the most effective means to investigate plasma dynamics. A variety of numerical techniques have been developed for this purpose.

Magnetohydrodynamics (MHD) treats the plasma as a fluid and is typically applied to large-scale phenomena spanning tens of astronomical units. In MHD simulations, the plasma evolves as a conducting fluid under the influence of electromagnetic fields. However, because the plasma is treated as a fluid, kinetic particle effects are not captured.

Kinetic methods are used to investigate phenomena that MHD cannot resolve. These simulations generally cover domains of only a few thousand ion gyroradii –many orders of magnitude smaller than typical MHD domains. In kinetic simulations, the plasma is represented by particles interacting with electromagnetic fields defined on a grid: particles deposit currents onto the grid and the Lorentz force resulting from electromagnetic fields derived from Maxwell’s equations accelerates the particles. Although fully kinetic simulations could, in principle, be applied to larger systems, they are computationally very demanding.

Hybrid methods provide an intermediate approach, combining elements of both MHD and kinetic descriptions. In these models, lighter species such as electrons or positrons are treated as a fluid, while heavier species such as protons or ions are treated kinetically. Hybrid simulations span larger domains than fully kinetic simulations but remain smaller than those accessible with MHD. While electron-scale physics is not resolved, ion kinetic effects are captured, allowing hybrid models to bridge the gap between MHD and fully kinetic approaches if the electron kinetic effects do not considerably influence the large-scale ion or fluid dynamics.

In this chapter, the Particle-In-Cell (PIC) technique – a fully kinetic simulation method – is introduced. The code description is based on a modified version of THATMPI (two-and-a-half-dimensional astrophysical Stanford code with MPI), itself an adaptation of the original TRISTAN code (three-dimensional Stanford code) developed by Buneman (1993) and later parallelised using MPI by Niemiec et al. (2008). In addition, new features implemented in THATMPI are presented. This chapter also presents results obtained using a novel machine-learning-based method to analyse PIC simulation data.

5.1 Kinetic description of a plasma

CONSERVATION laws of mass, momentum, and energy can be used to describe a plasma as a fluid. However, this approach does not account for the discrete nature of the particles that compose the plasma. Therefore, a kinetic description is necessary. Since plasmas consist of charged particles, Maxwell's equations are also required to describe the electromagnetic fields.

Consider a six-dimensional phase-space in which a system of discrete charged particles evolves in a self-consistent electromagnetic field. For a given particle species k , the exact *probability density function* in phase-space is:

$$N_k(\mathbf{r}, \mathbf{p}, t) = \sum_i^N \delta(\mathbf{r} - \mathbf{r}_i(t)) \delta(\mathbf{p} - \mathbf{p}_i(t)), \quad (5.1)$$

where N is the total number of particles and the position $\mathbf{x}_i$ and momentum $\mathbf{p}_i$ of the i -th particle evolves according to particle equations of motion: $d\mathbf{x}/dt = \mathbf{v}$ and $d\mathbf{p}/dt = \mathbf{F}_L$. Equation 5.1 is known as the *Klimontovich equation* (Heber 1968). Together with Maxwell's equations, it provides an exact description of the discrete nature of a plasma. However this approach is not practical, as the number of particles can easily reach $N = 10^{15}$. In such the case, it is convenient to adopt a statistical description of the plasma. To this end, a probability density function (PDF) $f(\mathbf{r}, \mathbf{p}, t)$ is introduced, such that the number of particles, dN , contained in an infinitesimal phase-space volume is described as:

$$dN = f(\mathbf{r}, \mathbf{p}, t) d^3r d^3p. \quad (5.2)$$

Thus, the PDF corresponds to the ensemble average of the Klimontovich distribution:

$$f_k(\mathbf{r}, \mathbf{p}, t) = \langle N_k(\mathbf{r}, \mathbf{p}, t) \rangle. \quad (5.3)$$

The PDF function is commonly referred to as the *particle distribution function*.

In the absence of particle collisions, the phase-space volume $d^3r d^3p$ is conserved in time and so is $f(\mathbf{r}, \mathbf{p}, t)$. This can be expressed as:

$$\frac{df}{dt} = \frac{\partial f}{\partial t} + \mathbf{p} \cdot \nabla f + \mathbf{F}_L \cdot \frac{\partial f}{\partial \mathbf{p}} = 0. \quad (5.4)$$

Equation 5.4 is known as the *Vlasov equation*. A generalisation of this equation is the *Boltzmann equation* (Kulsrud 2005, Chapter 8), in which the PDF is no longer conserved and an additional collisional term appears on the RHS of Equation 5.4.

The force term $\mathbf{F}_L$ that appears in Equation 5.4 corresponds to the Lorentz force, $\mathbf{F}_L = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$, where $\mathbf{E}$ and $\mathbf{B}$ denote the electromagnetic fields. These fields are generated self-consistently by the charged particles, but may also include externally imposed

components. The fields are described by Maxwell's equations:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon_0}, \quad (5.5)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (5.6)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (5.7)$$

$$\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{j} + \varepsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right), \quad (5.8)$$

where ρ and $\mathbf{j}$ are the charge density and current density, respectively, defined as the zeroth and first order moment of the PDF $f(\mathbf{r}, \mathbf{p}, t)$:

$$\begin{aligned} \rho &= \frac{q}{m} \int f(\mathbf{r}, \mathbf{p}, t) d^3p \\ \mathbf{j} &= \frac{q}{m} \int f(\mathbf{r}, \mathbf{p}, t) \mathbf{p} d^3p. \end{aligned} \quad (5.9)$$

The particles evolve under the influence of the electromagnetic fields according to the equations of motion:

$$\begin{aligned} \frac{d\mathbf{p}}{dt} &= \frac{d(\gamma m \mathbf{v})}{dt} = \mathbf{F}_L = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \\ \frac{d\mathbf{r}}{dt} &= \mathbf{v}, \end{aligned} \quad (5.10)$$

where $\gamma = (1 - (v/c)^2)^{-1/2}$ is the particle's Lorentz factor. By solving Equations 5.10 for each particle in the system, the full evolution of the plasma can be determined. This approach to solving the Vlasov - Maxwell system of equations is known as the method of characteristics.

Once the basic equations describing the behavior of charged particles and electromagnetic fields are defined, the next step is to simulate their evolution numerically. This involves discretizing both the equations and the physical quantities so that they can be solved computationally. The numerical technique that models plasma dynamics in this discrete, particle-based way is called the particle-in-cell (PIC) method.

5.2 Particle-In-Cell (PIC) technique

PARTICLE-IN-CELL methods were first developed in the late 1950s, with significant contributions from Oscar Buneman (Buneman 1959) and John Dawson (Dawson 1962). Since its introduction, the method has been extensively refined and reviewed over the following decades (Hockney and Eastwood 1988; Birdsall and Langdon 1991; Fehske et al. 2008). The premise is simple: charged particles, move discretely within a computational domain, while the electromagnetic field is discretised on a spatial grid that covers the same domain.

The particles and fields interact self-consistently: the particles produce currents and charge accumulation that drive the fields, and the fields act upon the charge particles by the Lorentz force, updating the particles' motion.

The plasma is evolved within a computational domain where each particle is characterised by its position and momentum and roam freely, only perturbed by the Lorentz force. On the other hand, the electric field, magnetic field, current density and charge density are discretised on a spatial grid defined over the box. Using the kinetic formalism introduced in Section 5.1, particles and fields advance discretely in time while interacting at each time step. A single PIC cycle consists of four main stages, as illustrated in Figure 5.1. First, the relativistic equations of motion are integrated for each particle to update positions and momenta under the Lorentz force. Second, the currents generated by the particles are computed and deposited onto the spatial grid. In certain implementations of the PIC scheme, including the one explained here, the explicit computation of the charge density is not required. Third, Maxwell's equations (specifically Faraday's law and the Ampère-Maxwell law) are solved on the grid to update the electromagnetic fields using the deposited currents. Finally, the fields are interpolated to particle positions to compute the Lorentz force for the next step, closing the loop. The following sections explain each of these steps in detail and describe how charged particles and fields are treated within the PIC framework.




Figure 5.1: Schematic of the four main stages of a PIC loop. i indicates the particle while n represents the current computational step. Electromagnetic fields are computed on the grid.

5.2.1 Macroparticles

Typical astrophysical plasma shocks have number densities on the order of 10 particles per cm^3 , while the characteristic spatial scales of interest are around $10^5 \lambda_{\text{si}}^3$. This corresponds to a total number of particles of approximately 10^{20} . Simulating such an enormous number of particles is computationally unfeasible, even for state-of-the-art supercomputers capable of reaching performances of up to 400 PFLOPS. For this reason, the PIC technique relies on the use of macroparticles.

A macroparticle represents an ensemble of a large number (typically thousands) of real particles – such as ions or electrons – that collectively interact with the electromagnetic field in the same way as an individual point-like particle. The total electric charge of this ensemble is distributed over a finite spatial region rather than being localised at a single point. Each macroparticle is defined by an effective mass m and charge q that are several orders of magnitude larger than those of a single real particle. However, since particles interact with the electromagnetic field through their charge and current densities (Equation 5.9), the interaction of a macroparticle is identical to that for a real particle. Therefore, using the charge-to-mass ratio, q/m , the evolution of a macroparticle can be computed using the same equations of motion (Equation 5.10). The only difference is that the charge is now distributed over a finite size.

5.2.2 Integration of the relativistic equations of motion

The most commonly used integration scheme for solving the equations of motion in PIC simulations is the leapfrog method, a second-order, time-centred algorithm that advances the particle quantities using staggered time steps. That is, positions and forces are evaluated at integer time steps n , while the particle momenta are computed at half-integer time steps $n \pm 1/2$. A schematic representation of the leapfrog method is shown in Figure 5.2. By discretising the equations of motion (Equation 5.10) the evolution of momentum and positions is obtained:

$$\begin{aligned} \frac{\mathbf{p}^{n+1/2} - \mathbf{p}^{n-1/2}}{\Delta t} &= \mathbf{F}_L^n = q(\mathbf{E}^n + \mathbf{v}^n \times \mathbf{B}^n) \\ \frac{\mathbf{x}^{n+1} - \mathbf{x}^n}{\Delta t} &= \mathbf{v}^{n+1/2}, \end{aligned} \quad (5.11)$$

where Δt is the duration of a computational time step. The techniques for integrating this equation are called **pushers** since they *push* (apply a force) the particles.

The final objective is to obtain $\mathbf{p}^{n+1/2} = \mathbf{p}^{n-1/2} + \Delta \mathbf{p}$ in order to advance the particle to its next position. To achieve this, an assumption must be made for the value of $\mathbf{v}^n$, since it cannot be computed directly from Equations 5.11, as doing so would require knowledge of the future-step value $\mathbf{v}^{n+1}$. Boris 1970 proposed evaluating $\mathbf{v}^n$ as the time average of the four-velocity $\mathbf{u} = \gamma \mathbf{v}$, such that

$$\mathbf{u}^n = \frac{\mathbf{u}^{n+1/2} + \mathbf{u}^{n-1/2}}{2}. \quad (5.12)$$




Figure 5.2: Schematic of the leapfrog scheme. The solid lines represent the particle pushing step while the dotted lines show the electromagnetic push. Particles move in time every step $x^{n \rightarrow n+1}$ by applying the Lorentz force $F_{L,n}$ while their velocity is computed every half-step $v^{n \pm 1/2}$. Similarly, the electric field is computed every full step E_n while the magnetic field is computed every half-step $B_{n \pm 1/2}$.

However, this approach does not compute the $\mathbf{E} \times \mathbf{B}$ drift velocity accurately and can therefore introduce errors in certain specific cases.

Vay 2008 addressed this issue by computing instead the time average of the velocity:

$$\mathbf{v}^n = \frac{\mathbf{v}^{n+1/2} + \mathbf{v}^{n-1/2}}{2}, \quad (5.13)$$

which correctly reproduces the $\mathbf{E} \times \mathbf{B}$ drift velocity. Nevertheless, this pusher does not preserve phase-space volume over time, leading to unphysical particle trajectories and a larger cumulative error in the total system energy.

Finally, Higuera and Cary 2017 proposed a method that resolves the limitations of the previous pushers. The choice is as follows:

$$\mathbf{v}^n = \frac{\mathbf{u}^{n+1/2} + \mathbf{u}^{n-1/2}}{2\bar{\gamma}}, \quad \bar{\gamma} = \sqrt{1 + \left(\frac{\mathbf{u}^{n+1/2} + \mathbf{u}^{n-1/2}}{2c} \right)^2}, \quad (5.14)$$

where $\bar{\gamma}$ corresponds to the Lorentz factor associated with the four-velocity $\mathbf{u}^n$ if it were computed using the Boris pusher in Equation 5.12, such that $\bar{\gamma}^2 = 1 + (u_{\text{Boris}}^n/c)^2$. Since this last pusher is employed in the simulations presented in this thesis, the HC (Higuera-Cary) pusher is reviewed in further detail below.

The HC pusher is divided into three stages: an initial electrostatic half-push ε , which modifies the magnitude of the particle velocity; a magnetic rotation β , which rotates the particle velocity vector; and a final electrostatic half-push, which yields the updated four-velocity $\mathbf{u}^{n+1/2}$.

The electrostatic half-push is derived from the discrete form of the Lorentz force. Using

the electric field E^n , the velocity is advanced by half a time-step according to

$$\frac{m\Delta\mathbf{u}^n}{\Delta t} = m \frac{\mathbf{u}_- - \mathbf{u}^{n-1/2}}{\Delta t} = q \frac{\mathbf{E}^n}{2} \rightarrow \mathbf{u}_- = \mathbf{u}^{n-1/2} + \frac{q\Delta t}{2m} \mathbf{E} = \mathbf{u}^{n-1/2} + \varepsilon \quad (5.15)$$

where $\mathbf{u}_-$ represents the half-advanced velocity and $\varepsilon = (q\Delta t/2m)\mathbf{E}$ is the electrostatic half-push. The second stage applies the magnetic rotation to $\mathbf{u}_-$:

$$\mathbf{u}_+ \left(1 + \frac{\beta^2}{\gamma_+}\right) = \mathbf{u}_- - \frac{\beta \times \mathbf{u}_-}{\gamma_+} + \frac{(\beta \cdot \mathbf{u}_-)\beta}{\gamma_+^2}, \quad \beta = \frac{q\Delta t}{2m} \mathbf{B}^n, \quad (5.16)$$

where

$$\gamma_+^2 = \frac{1}{2} \left(\gamma_-^2 - \beta^2 + \sqrt{(\gamma_-^2 - \beta^2)^2 + 4(\beta^2 + (\beta \cdot \mathbf{u}_-)^2)} \right), \quad \gamma_- \equiv \gamma(\mathbf{u}_-). \quad (5.17)$$

Finally, the last electrostatic half-push is performed:

$$\mathbf{u}^{n+1} = \mathbf{u}_+ + \varepsilon + \frac{\mathbf{u}_+}{\gamma_+} \times \beta. \quad (5.18)$$

Once $\mathbf{u}^{n+1} = \gamma \mathbf{v}^{n+1}$ is obtained, the particle position is updated as described in Equation 5.11. Despite its apparent complexity, this method performs as efficiently as the Vay pusher while preserving phase-space volume.

Although higher-order integration schemes exist, the leapfrog method remains as the most used scheme, as it requires only the current position and half-step velocity, effectively minimising memory usage. Moreover, it is computationally fast and exhibits excellent energy conservation on the long term that allows the method to be time-reversible.

5.2.3 Yee lattice

The electromagnetic field, current density, and charge density must be defined discretely on a spatial grid. The most commonly used approach is the *Finite-Difference Time-Domain* (FDTD) method, first introduced by Yee (1966). This relatively simple technique solves time-dependent partial differential equations by employing centred finite-difference approximations for both spatial and temporal derivatives.

The vector field components are defined on a lattice composed of cells. In Cartesian coordinates, each cell has the shape of a parallelepiped, although other coordinate systems may produce cells of different geometries. The code employed in this work uses a uniform cubic grid, where $\Delta x = \Delta y = \Delta z = \Delta = 1$, corresponding to the length of each cube edge. A schematic representation of a single cubic cell is shown in Figure 5.3. In this figure, the positions of the Cartesian components of $\mathbf{E}$, $\mathbf{B}$, and $\mathbf{j}$, as well as the charge density ρ , are distinct from one another. This arrangement is essential to ensure a consistent discrete formulation of Maxwell's equations. Specifically, the charge density is defined at the cubes vertices, the electric field and current density components are located along the cubes edges,




Figure 5.3: Schematic of a cubic cell in a PIC simulation. The charge density is defined at the cell vertices, while the electric and magnetic fields are defined at the midpoints of the edges and faces, respectively, to ensure proper satisfaction of Maxwells equations.

and the magnetic field components are positioned at the centres of the cubes faces, oriented inward towards the cell.

Let us consider a cell whose lower-left-front corner has coordinates (i, j, k) , as illustrated in Figure 5.3. The vector field components are defined such that

$$\begin{aligned}
 E_x(i, j, k) &\rightarrow E_x(i + \frac{1}{2}, j, k) & j_x(i, j, k) &\rightarrow j_x(i + \frac{1}{2}, j, k) \\
 E_y(i, j, k) &\rightarrow E_y(i, j + \frac{1}{2}, k) & j_y(i, j, k) &\rightarrow j_y(i, j + \frac{1}{2}, k) \\
 E_z(i, j, k) &\rightarrow E_z(i, j, k + \frac{1}{2}) & j_z(i, j, k) &\rightarrow j_z(i, j, k + \frac{1}{2})
 \end{aligned}
 \tag{5.19}$$

$$\begin{aligned}
 B_x(i, j, k) &\rightarrow B_x(i, j + \frac{1}{2}, k + \frac{1}{2}) \\
 B_y(i, j, k) &\rightarrow B_y(i + \frac{1}{2}, j, k + \frac{1}{2}) \\
 B_z(i, j, k) &\rightarrow B_z(i + \frac{1}{2}, j + \frac{1}{2}, k).
 \end{aligned}
 \tag{5.21}$$

This arrangement of vector field components ensures that a variation in the magnetic field $\mathbf{B}$ induces a counter-clockwise electric field defined on the cell surface where the $\mathbf{B}$ component is located, as given by Equation 5.7. This location coincides with that of $\mathbf{E}$ on

the Yee lattice. Conversely, a variation in the electric field $\mathbf{E}$ around a given surface generates a magnetic field $\mathbf{B}$ normal to that surface, as well as a current density circulating along its boundary, as described by Equation 5.8.

Only the last two of Maxwell's equations – the Maxwell-Faraday and Ampère-Maxwell equations – are required to determine the complete electromagnetic field. The system of Maxwell's equations is overdetermined; therefore, these two equations are sufficient to ensure that the remaining ones (Gauss's laws) are satisfied, provided that the initial conditions of the system are also fulfilled. The discrete forms of the Maxwell-Faraday and Ampère-Maxwell equations are:

$$\left(\frac{\partial \mathbf{E}}{\partial t}\right)^{n+1/2} = \frac{\mathbf{E}^{n+1} - \mathbf{E}^n}{\Delta t} = c^2 (\nabla \times \mathbf{B})^{n+1/2} - \mathbf{j}^{n+1/2}, \quad (5.22)$$

$$\left(\frac{\partial \mathbf{B}}{\partial t}\right)^n = \frac{\mathbf{B}^{n+1/2} - \mathbf{B}^{n-1/2}}{\Delta t} = -(\nabla \times \mathbf{E})^n. \quad (5.23)$$

As with the equations of motion, the leapfrog scheme is used to integrate the electromagnetic field in time. For a two-dimensional simulation, Equations 5.22 and 5.23 can be simplified by taking $\nabla = (\partial/\partial x, \partial/\partial y, 0)$.

Particle motion requires knowledge of the electromagnetic fields at integer time steps (see Equation 5.11). However, only the electric field is defined at integer steps, while the magnetic field is computed at half-integer steps. To resolve this, the magnetic field is advanced by half a time-step, $\Delta t/2$, yielding $\mathbf{B}^n$. Once the particles have been moved, $\mathbf{B}^n$ is advanced by another half-step to obtain $\mathbf{B}^{n+1/2}$.

5.2.4 Current deposition and particle shape

Current deposition refers to the computation of the current density produced by a macroparticle of finite size on the spatial grid. Depending on the macroparticle shape and position, the corresponding current density – an additive quantity – is distributed across the Yee lattice.

Macroparticles can have different shapes. The two most common are illustrated in Figure 5.4: the cloud-in-cell (CIC) and the triangular-shaped cloud (TSC), corresponding to first-order and second-order shape functions, respectively. For a given particle shape, the electric charge q_p is distributed amongst neighbouring grid cells q_i according to the portion of the particle shape that overlaps each cell,

$$q_i = q_p S(\mathbf{r}_i), \quad (5.24)$$

where $S(\mathbf{r}_i)$ is the shape factor of the macroparticle at cell i . The shape factor thus determines the fraction of the particle's charge deposited in each grid cell. It is computed as the area (or volume, in three dimensions) corresponding to the overlap between the particle shape function and the grid cell, as illustrated in Figure 5.4.

To illustrate charge deposition more clearly, let us consider a one-dimensional grid. For a CIC shape, the macroparticle has an effective width of Δx , such that its charge is distributed




Figure 5.4: Electric charge distribution (top figures) of a macro-particle and its corresponding particle shape function (bottom figures) for the cloud-in-cell (CIC) approximation (left figures) and the triangular-shape-cloud (TSC) approximation (right figures).

between two adjacent cells. The corresponding shape function is given by

$$S_1(x_i) = \begin{cases} 1 - \frac{|x_i - x_p|}{\Delta x}, & \frac{|x_i - x_p|}{\Delta x} \leq 1 \\ 0, & \frac{|x_i - x_p|}{\Delta x} > 1, \end{cases} \quad (5.25)$$

where x_i denotes the discretised grid position and x_p is the position of the particle. The term $|x_i - x_p|$ accounts for both neighbouring cells to the left and right of the particle.

For a TSC shape, the macroparticle shape extends over $2\Delta x$, distributing its charge among three neighbouring cells. The corresponding shape function is

$$S_2(x_i) = \begin{cases} \frac{3}{4} - \left(\frac{|x_i - x_p|}{\Delta x}\right)^2, & 0 \leq \frac{|x_i - x_p|}{\Delta x} \leq \frac{1}{2} \\ \frac{1}{2} \left(\frac{3}{2} - \frac{|x_i - x_p|}{\Delta x}\right)^2, & \frac{1}{2} \leq \frac{|x_i - x_p|}{\Delta x} \leq \frac{3}{2} \\ 0, & \frac{|x_i - x_p|}{\Delta x} > \frac{3}{2} \end{cases}. \quad (5.26)$$

Although this higher-order shape function is computationally more expensive, it significantly reduces numerical noise by acting as a low-pass filter, effectively suppressing high-frequency modes of the discretised quantities. The presence of numerical noise can be understood as sharp variations in magnitude arising from limited spatial resolution (as in the CIC shape),

which produce a high-frequency response in Fourier space. For this reason, the TSC model is widely employed in modern PIC simulations.

Note that macroparticles can pass freely through one another. It is assumed that the collision rate remains sufficiently low for the plasma to be considered effectively collisionless, such that two ensembles of macroparticles may cross without interacting directly.

Once the charge distribution has been computed, the current density is computed as

$$\mathbf{j}^{n+1/2} = \rho \mathbf{v}^{n+1/2}, \quad (5.27)$$

where ρ is the charge density and $\mathbf{v}^{n+1/2}$ is obtained from Equation 5.11. While the current density is discretised on the grid, the particle velocity has no fixed position in space. Consequently, the velocity must be appropriately split according to the grid cells the particle crossed through. To achieve this, Umeda's method of current deposition is employed (Umeda et al. 2003).

In this approach, the particle velocity is decomposed into two components, $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$, which may be located in the same cell or separate cells, as illustrated in Figure 5.5. This decomposition is performed by introducing a relay point, whose position depends on the particle positions at $t^{n+1} = t^n + \Delta t$:

$$x_r = \begin{cases} \frac{x^n + x^{n+1}}{2}, & i^n = i^{n+1} \\ \frac{i^n + i^{n+1}}{2}, & i^n \neq i^{n+1} \end{cases}, \quad (5.28)$$

where i^n is the coordinate of the cell i where the particle is located at step n . The other spatial coordinates are treated analogously. The corresponding velocities are given by $\mathbf{v}_1 = (\mathbf{x}_r - \mathbf{x}_1)/\Delta t$ and $\mathbf{v}_2 = (\mathbf{x}_2 - \mathbf{x}_r)/\Delta t$, which are used to compute the total current density:

$$\mathbf{j} = \mathbf{j}_1 + \mathbf{j}_2 = q_1 \mathbf{v}_1 + q_2 \mathbf{v}_2. \quad (5.29)$$

5.2.5 Field weighting

Once the electromagnetic fields have been computed on the grid, the Lorentz force acting on each particle is obtained by interpolating the fields to the particle positions. As with charge deposition, the interpolation weights are determined by the particle shape factor. It is essential to use the same shape factor consistently for both charge deposition and field interpolation to prevent spurious forces (Dawson 1983).

5.3 Numerical limitations

The previous section described the algorithms that constitute the PIC loop. However, all numerical simulations – not only those based on the PIC method – are subject to errors and discrepancies inherent to any discrete algorithm. Consequently, specialised numerical techniques are required to mitigate such errors and prevent unphysical results. These techniques include numerical filtering, algorithms specifically designed for shock simulations, and various other computational strategies.




Figure 5.5: Schematic of the current deposition algorithm by Umeda et al. 2003. Panel a) shows the case of a particle moving within the same cell with coordinates (i, j) , panel b) the particle that moves diagonally from cell (i, j) to cell $(i + 1, j + 1)$, which corresponds to the bottom expression in Equation 5.28 for both coordinates x and y .

5.3.1 CFL condition

Due to the discretisation of space and time on the Yee lattice, there exists a maximum frequency and wavenumber that can be resolved, determined by the Nyquist frequency. Larger grid spacing (i.e., lower resolution) reduces the maximum wave speed that can be correctly represented. Consequently, the lattice imposes a maximum velocity for which a plasma simulation remains physically valid. Any wave propagating faster than this limit – such as a light wave, the fastest type of wave – results in unphysical propagation and may invalidate the simulation.

Consider a wave-like quantity $A(\mathbf{r}, t) = A_0 e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t)}$ travelling at the speed of light, $\omega = kc$. The spatial derivative of a single component using the *symmetric difference quotient* method is

$$\begin{aligned} \frac{\partial A}{\partial x} &\approx \frac{A(x + \Delta x/2, t) - A(x - \Delta x/2, t)}{2(\Delta x/2)} = \\ &A(x, t) \frac{e^{ik\Delta x/2} - e^{-ik\Delta x/2}}{\Delta x} = i \frac{\sin(k\Delta x/2)}{\Delta x/2} A(x, t) = iK A(x, t), \\ K &= \frac{\sin(k\Delta x/2)}{\Delta x/2}. \end{aligned} \quad (5.30)$$

The other two spatial components are obtained analogously. The time derivative is treated similarly, yielding $\partial A / \partial t = i\Omega A$, with $\Omega = \sin(\omega \Delta t/2) / (\Delta t/2)$.

The dispersion relation $\omega(k)$ can be obtained from the wave equation:

$$\frac{\partial^2 A}{\partial t^2} = c^2 \nabla^2 A \rightarrow \Omega^2 = c^2 \sum K_j^2 \rightarrow \left(\frac{\sin(\omega \Delta t / 2)}{\Delta t / 2} \right)^2 = c^2 \sum_{j=x,y,z} \left(\frac{\sin(k_j \Delta j / 2)}{\Delta j / 2} \right)^2. \quad (5.31)$$

To prevent numerical instabilities – and therefore wave growth – the wave frequency must remain real, $\omega \in \mathbb{R}$, which corresponds to the condition $\sin^2 x \leq 1$ for any wave-like quantity x . In the limiting case of a cubic cell of size Δ , the maximum allowable speed of light is

$$\frac{1}{(\Delta t)^2} = c_{\max}^2 \frac{1}{3\Delta^2} \rightarrow c_{\max} = \frac{1}{\sqrt{3}} \frac{\Delta}{\Delta t}, \quad (5.32)$$

known as the Courant-Friedrichs-Lewy (CFL) condition (Courant et al. 1928). The factor $\sqrt{3}$ arises from the three spatial dimensions; for D dimensions, it can be generalised by replacing $\sqrt{3}$ with $\sqrt{D}$, provided the cell sides are equal.

For any light speed lower than $c_{\max}$, the electromagnetic field remains stable, with no spurious phase or magnitude errors. If the speed exceeds this limit, the wave frequency becomes complex, leading to exponential growth and numerical instability. In THATMPI, the speed of light is set to $c = 0.5$ for $\Delta = \Delta t = 1$.

5.3.2 Noise filtering and the Friedman filter

The CFL condition introduced in the previous subsection imposes a maximum wave speed for a stable simulation. Any wave propagating faster than this limit may give rise to numerical noise. In particular, relativistic particles interacting with such waves can trigger the *Cherenkov numerical instability*, in which particles effectively move faster than the corresponding light waves in the grid, generating unphysical radiation. This spurious radiation perturbs the simulation and can render its results invalid. The instability is named after the analogous physical phenomenon in which charged particles exceed the speed of light in a medium, producing Cherenkov radiation.

Several methods can be employed to reduce numerical noise:

- Increasing the number of particles per cell, N_{ppc} , improves resolution. In PIC codes, higher resolution corresponds to representing physical lengths such as the skin depth λ_s with more grid cells. However, increasing N_{ppc} substantially, raises the number of computations required, extending simulation time and increasing memory usage, which is limited by the available RAM.
- Higher-order algorithms or schemes can provide more accurate results by adding more points in space or time. Examples include using higher-order particle shape functions (such as TSC) or higher-order time derivatives in Maxwells equations. While using these methods the number of calculations increases, it is generally preferred over simply increasing N_{ppc} , as they require less memory and overall computational power.

- In Fourier-based codes, a low-pass filter can remove unphysical high-frequency waves that may trigger numerical instabilities. Although this technique could also be applied in codes that operate in real space, it would require performing a Fourier transform of the electromagnetic field at every time step, significantly increasing computation time.
- Some numerical models can mitigate the growth of the Cherenkov instability by carefully choosing the speed of light, ensuring it satisfies the CFL condition while remaining below the critical threshold discussed previously.
- Digital filtering of specific quantities can further reduce noise. Examples include the *binomial filter*, the *Godfrey filter*, and the *Friedman filter*. The latter is employed in THATMPI (see below).

The Friedman filter, developed by Friedman 1990, modifies the update of the magnetic field $\mathbf{B}^{n+1/2}$ by introducing a parameter θ_F that adjusts the contribution of past electric fields and thereby controls the effectiveness of the filter. By discretising space and time and incorporating this parameter, Equation 5.23 becomes

$$\mathbf{B}^{n+1/2} = \mathbf{B}^{n-1/2} - \Delta t \nabla \times \left(\left(1 + \frac{\theta_F}{2}\right) \mathbf{E}^n - \theta_F \left(1 - \frac{\theta_F}{2}\right) \mathbf{E}^{n-1} + \frac{1}{2}(1 - \theta_F)^2 \theta_F \bar{\mathbf{E}}^{n-2} \right), \quad (5.33)$$


where $\bar{\mathbf{E}}^{n-2} = \mathbf{E}^{n-2} + \theta_F \bar{\mathbf{E}}^{n-3}$ and $0 \leq \theta_F \leq 1$. For $\theta_F = 0$, the standard update equation (Equation 5.23) is recovered. Greenwood et al. 2004 extensively reviewed the Friedman filter and found $\theta_F = 0.05$ as the optimal value for effectively suppressing Cherenkov emission. This value is employed in the THATMPI code.

5.4 Generation of a shock

To generate a plasma shock, two regions satisfying the jump conditions described in Section 3.2.1 are required. The shocks in this thesis are produced using the "reflecting wall" setup. This method employs an artificial conducting wall that reflects particles back into the simulation domain. Consider a simulation box with a plasma beam moving toward the left, where a conducting wall is positioned, as illustrated in Figure 5.6. An initial background magnetic field $\mathbf{B}_n$ can be initialised: if it lies in the simulation plane ($x - y$ plane), the field is *in-plane*; if it lies in the $x - z$ plane, it is *out-of-plane*.

When the upstream beam reaches the wall, particles are reflected, reversing their velocity component normal to the wall (assumed here to be along the x -direction) and updating their positions accordingly. The interaction between the reflected particles and the incoming beam leads to formation of the shock. To reduce numerical artifacts, the conducting wall must satisfy Maxwell's equations (see Subsection 5.4.1). Additionally, initial tapering of the plasma, detailed in Subsection 5.4.2, can be applied to suppress transient effects at the beginning of the simulation.

Figure 5.6: Sketch of a computational domain and the reflecting wall, located in the $x - y$ plane. The particle (shown as a point) is reflected on the left side of the box instead of crossing it. The interaction of reflected and incoming particles forms a shock. The background magnetic field, $\mathbf{B}_0$, lies in the simulation plane $x - y$ (an *in-plane* setup). An *out-of-plane* magnetic field would lie in the $x - z$ plane.



5.4.1 Boundary conditions

Since particles are generated within a computational box, they may interact with the boundaries in different ways. The shock simulations presented in this thesis utilise two different boundaries: a *reflecting boundary* where particles are reflected; and a *periodic boundary*, where particles reappear on the opposite side of the box upon crossing.

The periodic boundary is the simplest to implement. Particles crossing the boundary re-enter the box on the opposite side with the same velocity they had prior to crossing, $v_{i,-} = v_{i,+}$, where the subscripts $-$ and $+$ represent the velocity immediately before and after crossing the boundary, respectively. Electromagnetic fluctuations behave analogously $\delta E_- = \delta E_+$ and $\delta B_- = \delta B_+$.

A reflecting wall, also referred to as a conducting wall, reflects particles off the wall, modifying their velocities and positions. Particle reflection inverts the component of the velocity normal to the wall, analogous to a mirror. For a wall normal to the x -direction, the x -component of the velocity is negated $v_{x,p}^n = -v_{x,p}^{n+1}$.

The treatment of the electromagnetic field differs from that of particles. The reflecting wall can be modelled as an idealised charged sheet with surface charge density σ_c . Let us consider an infinite conducting wall located at $x = 0$. The electric field at the wall is only characterised by a normal component, E_x , while the tangential components are cancelled by the charges on the sheet $E_y = E_z = 0$. Outside the wall, the electric field is not necessarily zero but must remain continuous. In particular, the tangential components are antisymmetric with respect to the wall:

$$E_{y/z,-} = -E_{y/z,+}, \quad E_{y/z,0} = 0. \quad (5.34)$$

The normal component E_x is derived from Gauss' law (Equation 5.5). At the wall, since the tangential fields are zero,

$$\frac{\partial E_x}{\partial x} = \frac{\rho}{\epsilon_0} \rightarrow E_{x,+} - E_{x,-} = \frac{\sigma_c}{\epsilon_0}, \quad (5.35)$$

where $d\sigma_c \equiv \rho dx$ represents an infinitesimal surface charge. Since the wall is a numerical construct intended solely to reflect particles, it is typically assumed to carry no surface

charge, $\sigma_c = 0$. This yields a symmetric normal field component with respect to the wall:

$$E_{x,-} = E_{x,+}. \quad (5.36)$$

The tangential components of the magnetic field at a conducting wall are determined from the Ampère-Maxwell law (Equation 5.8). For a stationary solution, the time derivative is set to zero. Additionally, the current density at the wall is zero, $\mathbf{j} = \rho \mathbf{v} = 0$, since there is no charge present. Therefore,

$$\nabla \times \mathbf{B} = 0 \rightarrow \frac{\partial B_z}{\partial x} = \frac{\partial B_x}{\partial z}, \quad \frac{\partial B_y}{\partial x} = \frac{\partial B_x}{\partial y}. \quad (5.37)$$

For an infinite conducting wall, there can be no variation along the y or z directions with respect to B_x . Consequently, the tangential components satisfy

$$\frac{\partial B_y}{\partial x} = 0 \rightarrow B_{y,-} = B_{y,+} \quad \frac{\partial B_z}{\partial x} = 0 \rightarrow B_{z,-} = B_{z,+}. \quad (5.38)$$

Thus, the tangential components of the magnetic field are symmetric with respect to the wall.

The normal component B_x is obtained from Gauss' law (Equation 5.6). Again, for an infinite wall, there can be no variation along y or z , yielding

$$\frac{\partial B_x}{\partial x} = 0 \rightarrow B_{x,-} = B_{x,+}. \quad (5.39)$$

Therefore, the normal component of the magnetic field is also symmetric with respect to the wall.

In the simulations presented in this work, the computational domain employs a reflecting wall at the left boundary and periodic boundary conditions along the y -direction (top and bottom boundaries).

5.4.2 Initial tapering

Due to the various numerical approximations used to setup a shock in a PIC simulation initial transients may arise and escape back to the simulation domain, disrupting the upstream plasma beam and generating unphysical effects. Several initial tapering algorithms can be employed to significantly reduce these unwanted transients:

- In the method developed in Wieland et al. 2016 the background magnetic field $\mathbf{B}_0$ is smoothed over a distance Δ_x from the transition zone at x_{wall} :

$$\mathbf{B}(x) = \frac{B_0}{2} \left[\cos \left(\frac{x - x_{\text{wall}}}{\Delta_x} \right) + 1 \right] (\cos(\phi)\hat{\mathbf{y}} + \sin(\phi)\hat{\mathbf{z}}) \quad (5.40)$$

where ϕ is the angle between the y and z components. This correction smoothly transitions the magnetic field from 0 at $x = x_{\text{wall}} - \Delta_x$ to B_0 at $x = x_{\text{wall}}$. The corresponding motional electric field $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$ is corrected accordingly:

$$\mathbf{E}(x) = \frac{vB_0}{2} \left[\cos \left(\frac{x - x_{\text{wall}}}{\Delta_x} \right) + 1 \right] (\sin(\phi)\hat{\mathbf{y}} - \cos(\phi)\hat{\mathbf{z}}). \quad (5.41)$$

From Equation 5.8, the stationary configuration produces a finite current density:

$$\frac{\partial \mathbf{E}}{\partial t} = \nabla \times \mathbf{B} - \mu_0 \mathbf{J} = 0 \rightarrow \mathbf{J} = \frac{1}{\mu_0} \nabla \times \mathbf{B}. \quad (5.42)$$

This current can be implemented as an ion drift $\mathbf{v}_d$ relative to the electron velocity, $\mathbf{J} = ne\mathbf{v}_d$. Using Equation 5.40, the drift velocity is

$$\mathbf{v}_d(x) = \frac{B_0 c \pi}{2ne\Delta_x} \sin\left(\frac{x - x_{wall}}{\Delta_x}\right) (\sin(\phi)\hat{\mathbf{y}} - \cos(\phi)\hat{\mathbf{z}}). \quad (5.43)$$

It is generally recommended to set the tapering width Δ_x on the order of the shock width, approximately equal to the ion gyroradius, $\Delta_x \sim R_{g,i}$.

- A simpler tapering procedure can be applied that does not require corrections to the magnetic field. The initial upstream velocity is smoothed over a short distance Δ_x from the wall position x_{wall} using a cosine-like profile:

$$v_{x,i} = \begin{cases} v_{up} \left[1 - \sin^2\left(\frac{\pi}{2} \frac{x_i - x_{wall}}{\Delta_x}\right) \right], & x_{wall} < x_i < \Delta_x \\ v_{up}, & x_i > \Delta_x \end{cases}, \quad (5.44)$$

where v_{up} is the upstream bulk velocity. To prevent spurious forces, the motional electric field is modified accordingly. As in the previous tapering method, it is recommended to set the tapering width Δ_x to be of the order of the shock width.

- In oblique shocks, the upstream plasma with bulk velocity $\mathbf{V} = V\hat{\mathbf{x}}$ acquires a transverse drift V_d after crossing the shock along the direction of the magnetic field (Tran and Sironi 2024). This drift is particularly pronounced in strongly magnetised plasmas, where particles tend to follow the magnetic field lines even in the downstream region, despite turbulence and heating. To prevent this effect when particles reach the wall, an additional tapering can be applied.

On the left-most side of the reflecting wall, a constant electric field corresponding to the motional electric field associated with the velocity drift and upstream magnetic field is set over a distance $\Delta_x = 10$ to:

$$E_y = -V_z B_0 \cos \theta_{Bn}, \quad \phi = 90^\circ \quad (5.45)$$

$$E_z = V_y B_0 \cos \theta_{Bn}, \quad \phi = 0^\circ, \quad (5.46)$$

where V_y and V_z represent the transverse drift V_d , whose orientation depends on whether the magnetic field is out-of-plane (top) or in-plane (bottom). This boundary mimics a reflecting wall moving in the transverse direction with speed V_d , hence it is referred to as a *sliding wall*. Particles are reflected five cells beyond the tapering region Δ_x .

Particle reflection is performed in the sliding wall frame of reference. Given a particle velocity $\mathbf{v} = (v_x, v_y, v_z)$ and an out-of-plane magnetic field producing a drift in the

z -direction, $V_d = V_z$, the particle is first transformed to the sliding wall frame, reflected in the x and z directions, and then transformed back to the simulation frame. The resulting velocity components are:

$$v_{x,refl} = \frac{-v_x/\gamma^2}{D} \quad (5.47)$$

$$v_{y,refl} = \frac{v_y/\gamma^2}{D} \quad (5.48)$$

$$v_{z,refl} = \frac{2V_z - v_z \left(1 + \frac{V_z^2}{c^2}\right)}{D} \quad (5.49)$$

$$D = 1 + \frac{V_z^2}{c^2} - 2\frac{V_z v_z}{c^2}.$$

The expressions for a drift in the y -direction $V_d = V_y$ are obtained analogously. This drift correction is only necessary for oblique shocks; for parallel shocks, the drift vanishes, and for perpendicular shocks, no electric field correction is needed since there is no x -component of the magnetic field.

The sliding wall method is particularly useful for studies of oblique shocks, as considered in this thesis. However, extended test simulations indicate that tapering is unnecessary for the setup used here, since the effects of the reflecting wall do not influence the magnetosonic wave pulse that forms the shock transition.

Figure 5.7 compares simulations with and without a sliding wall, using the same parameters as the simulation described in Section 6.3. In a simulation with a sliding wall, the downstream conditions are artificially modified, resulting in a higher shock speed and compression ratio. Near the left boundary in the test run without the wall, the magnetic field decreases more rapidly ($x/\lambda_{si} \approx 0-32$) compared to the simulation with a sliding wall, where the decrease is more gradual. Despite these differences, the shock transition itself is generated properly in both cases, and any physical processes occurring at the shock are unaffected by the wall conditions. Consequently, the sliding wall is not used in the simulations presented in Chapter 6.




Figure 5.7: Comparison of the normalised compressed magnetic field B_z/B_{0z} for one of the main simulations, shown with and without a sliding wall. Section 6.1 provides a detailed description of the simulation parameters.

5.5 New features of THATMPI

NEW features have been added to the THATMPI code to improve computational efficiency, readability, and usability, as well as to implement new algorithms. A key addition is the Higuera-Cary (HC) particle pusher, discussed previously in Section 5.2.2. The following subsections describe four major upgrades to the code.

5.5.1 MPI subarrays

A typical computational box may contain more than five million cells. Using a single computational process to simulate such a large domain would be impractical, as the required runtime would be unfeasibly long. To overcome this, the computational box is divided into rectangular subdomains, each handled by a separate CPU-core process, thereby distributing the workload across all available processes.

Particles and electromagnetic field perturbations can traverse the entire computational box, and their information must be communicated between processes whenever they reach the boundary of a subdomain. This communication is handled using the Message Passing Interface (MPI), a portable and widely used standard for parallel computing. Each subdomain includes an additional layer of cells surrounding its physical region, known as *ghost cells*. Figure 5.8 illustrates how this process works: the ghost cells of neighbouring domains overlap in the physical box, so that a particle crossing a ghost cell in one domain is correctly transferred to the corresponding ghost cell of the neighbouring domain.




Figure 5.8: Schematic of a particle passing from one domain (blue) to the neighbouring domain (green).

Particle communication requires only a single layer of ghost cells. In contrast, electromagnetic fields require multiple layers of ghost cells, since neighbouring cells are needed to accurately compute spatial derivatives (see Section 5.2.3). Depending on the field updating scheme, the ghost layers are typically two to three cells wide.

Passing information between domains requires multiple rows or columns that account for ghost cells. Standard MPI routines, such as `MPI_Send` and `MPI_Recv`, can only transmit contiguous memory blocks (effectively 1D arrays). Higher-dimensional arrays are flattened in memory, so sending multi-row ghost boundaries with multiple MPI calls introduces significant overhead. Using a temporary 1D buffer reduces MPI calls but adds extra computation for copying data.

A more efficient solution is to use *MPI subarrays*. An MPI subarray defines a specific subregion $\bar{a}$ within a larger array $\bar{A}$, allowing MPI to directly transfer data between corresponding subarrays of neighbouring domains. For example, the *right* ghost boundary of domain A can be sent directly to the *left* ghost boundary of domain B (see Figure 5.9).

This approach minimises both MPI overhead and array copying. In 2D simulations, the gain is modest since ghost boundaries are only a few cells wide. In 3D, where ghost boundaries can span ~ 15 cells in the z -direction, MPI subarrays provide a substantial reduction in communication overhead and significantly improve simulation efficiency.




Figure 5.9: Schematic of domain communication using subarrays. Each subarray consists of multiple cells along the x and y directions. Ghost boundaries are exchanged between neighboring subarrays: a subarray sends its boundary data to the adjacent subarray, which simultaneously sends its own boundary data back.

5.5.2 Parallel I/O

In THATMPI, simulation data are periodically output using the HDF5 library. Originally, each process wrote its own domain to a separate file, resulting in as many files as processes. Post-processing then required a single process to combine all these files into a complete dataset for analysis. While manageable for few processes, modern PIC simulations often use thousands of processes, producing an overwhelming number of files that slow post-processing and are inefficient on HPC file systems like LUSTRE.

To address this, THATMPI implements parallel I/O using *Parallel HDF5*. Each process writes its domain data, such as electromagnetic fields, charge density, etc., directly into the appropriate region of a single HDF5 file, computing the size and offset based on its domain's position in the computational box. Ghost cells are excluded, but their contributions are communicated to neighbouring real cells before writing to ensure a consistent full array. Although parallel I/O slightly increases simulation runtime, it drastically reduces total time by eliminating the need for slow post-processing over thousands of files.

5.5.3 Relativistic particle injection

In the non-relativistic regime, particle velocities are distributed according to the Maxwell-Boltzmann (MB) probability density function (PDF), also known as a *Maxwellian*:

$$f(v) = \left(\frac{m}{2\pi k_B T} \right)^{3/2} 4\pi v^2 e^{-\frac{mv^2}{2k_B T}}. \quad (5.50)$$

In 1D, this reduces to a Gaussian with mean zero and standard deviation $SD = \sqrt{k_B T/m}$, while in 3D, the distribution factorises as $f(v_x, v_y, v_z) = f(v_x)f(v_y)f(v_z)$ because the velocity components are independent. The PDF for the magnitude of the velocity can then be obtained by transforming the 3D distribution into spherical coordinates, assuming spherical symmetry.

For relativistic plasmas, special relativity significantly alters the velocity distribution. The Maxwell-Jüttner (MJ) distribution describes this case (Jüttner 1911):

$$f(\gamma) = \frac{\gamma^2 \beta(\gamma)}{\Theta K_2\left(\frac{1}{\Theta}\right)} e^{-\gamma/\Theta} \quad (5.51)$$

where $\gamma = 1/\sqrt{1 - v^2/c^2}$ is the Lorentz factor, $\beta = v/c$, $\Theta = k_B T/mc^2$ is the normalised temperature, and K_2 is the modified Bessel function of the second kind. In the non-relativistic limit ($v \ll c$ and $\Theta \ll 1$), the MJ distribution reduces to the classical Maxwellian.

In PIC simulations, ions and electrons are injected with initial bulk velocities and temperatures. For non-relativistic simulations, each velocity component can be drawn from a Gaussian distribution using standard algorithms such as Box-Müller or Marsaglia polar methods (Box and Muller 1958; Marsaglia and Bray 1964). The resulting 1D Gaussian can then be combined to form a 3D MB distribution, and then apply a boost to obtain a moving




Figure 5.10: Comparison for the Maxwell-Boltzmann (MB) and Maxwell-Jüttner (MJ) probability distribution functions for two different thermal velocities ($v_{th}/c = \sqrt{21}/11$ and $v_{th}/c = \sqrt{3}/2$) and Lorentz boosts ($\Gamma_{boost} = 1.0$ and $\Gamma_{boost} = 1.5$). The black line shows the marginal probability density function $f(u_x)$ for a boosted Maxwell-Jüttner distribution.

upstream beam. However, for high temperatures or large boosts, this procedure yields an incorrect particle distribution and particles must be drawn directly from the MJ distribution.

A new particle injection method was implemented in THATMPI following the method described in Zenitani and Nakano 2022, based on the modification of Canfield's approach (Canfield et al. 1987). This algorithm takes advantage of the similarity between the Maxwell-Jüttner and gamma distributions $\Gamma(x; a, b)$, where a and b are the parameters of the gamma distribution. Figure 5.10 illustrates the difference between MB and MJ distributions: the top row shows non-boosted distributions with thermal velocities $v_{th}/c = \sqrt{21}/11$ and $v_{th}/c = \sqrt{3}/2$, while the bottom row depicts a boosted distribution with Lorentz factor $\Gamma_{boost} = 1.5$.

The MJ distribution provides wider wings for the unboosted PDF, while in the boosted case the MB distribution shows a tail distributed over a longer range of velocities.

It is important to note that temperature definitions differ between the non-relativistic and relativistic cases:

$$kT_{nr} = \frac{1}{2}mv^2, \quad kT_r = (\gamma(v) - 1)mc^2, \quad (5.52)$$

which coincide in the limit $v \ll c$. This new injection method ensures that high-temperature or relativistically boosted particle distributions are correctly initialised, preventing unphysical artifacts in the simulation and allowing accurate modeling of relativistic plasma flows.

5.5.4 Improved particle tracing

Kinetic plasma simulations allow to follow (or trace) individual particles. While it would be technically possible to save every particle's full history to disk, doing so would be computationally prohibitive and require an impractically large amount of storage. Therefore, only a subset of particles – typically a few thousand – are traced. Particle tracing records data such as the particle's position, momentum, and the electromagnetic field at its location. This is extremely useful for understanding particle dynamics, including heating and acceleration mechanisms (Ardaneh et al. 2015; Bohdan et al. 2019a; Kobzar et al. 2021).

THATMPI includes a particle tracing feature. For each traced particle, a file is created, and every n -th step (usually 10–20 steps), the particle data and local electromagnetic fields are saved to disk. Each particle has a unique tracing ID to ensure the correct mapping between particles and their saved data. Since particles can cross domain boundaries in parallel simulations, these IDs must be communicated between processes as well.

The previous implementation of particle tracing, however, had limitations: files were simple text files, difficult to read and analyse, and data was appended line by line, leading to inefficiencies. To address this, I upgraded the particle tracing system in two major ways using the HDF5 library:

- Each particles data is now stored in an HDF5 file containing datasets for magnetic field, electric field, position, momentum, Lorentz factor, and time. Datasets are pre-allocated at the beginning of the simulation as *extendable datasets*, a particular type of dataset that can be extended in memory. When the simulation is resumed, the datasets are extended by the number of expected writes. For example, if the simulation is resumed for 10,000 steps and data is saved every 10 steps, the datasets are extended by 1,000. This approach optimises writing performance and makes the data easier to read.
- Instead of generating separate files for each particle, all traced particles can now be saved in a single HDF5 file using Parallel HDF5. All processes participate collectively in file creation. Within the file, an HDF5 group is created for each particle, containing the corresponding datasets. Although this approach requires parallelisation, it significantly reduces the file system load, which is a common issue in high-performance computing environments using Lustre or similar systems.

5.6 Analysis tools of kinetic-based plasma simulations using neural networks

PIC simulations follow individual particles and therefore can be traced. However, the cost of I/O operations is large and becomes unfeasible to write data to disk each step or even every 10–20 steps for all particles. Currently, thousands of particles can be traced without heavily affecting the performance of the simulation, as opposed to millions of particles that are simulated. Analysing particle tracing data requires inspecting one by one all the traced particles, which lengthens the analysis and introduces a human bias into the results. Therefore, I developed a novel method for analysing traced particles more efficiently and reliably by using neural networks (NNs). A more detailed description can be found in Torralba Paz et al. 2025.

For this study, we used traced particle data obtained in a recent work on non-relativistic strictly perpendicular shocks of young SNRs performed by Bohdan et al. 2019a. Such shocks are characterised by a high sonic and Alfvénic Mach numbers and propagate in a low temperature medium. Under these conditions, the Buneman instability is excited in the shock foot (Amano and Hoshino 2009; Matsumoto et al. 2012a; Bohdan et al. 2017). This kinetic instability arises from the relative drift between shock-reflected ions and incoming upstream electrons, generating electrostatic Buneman waves that can accelerate a small fraction of upstream electrons to high energies through shock surfing acceleration (SSA, see Section 4.3.2).

The dataset analysed in this section is taken from run E2 in Bohdan et al. 2019a. In this simulation, $N = 252\,000$ particles were traced for a total duration of $t = 4500\omega_{pe}^{-1}$ for which a segment of $300\omega_{pe}^{-1}$ is saved as a dataset and used for the analysis. During this short period, particles are energised by the Buneman waves. This run in particular features Buneman waves with wavenumber almost parallel to the x -direction. Figure 5.11 shows two electrons: one of them crosses the waves without any significant interaction (blue curve) and the other one is accelerated by the electrostatic waves (red curve) as clearly seen in panel (b) at $t\omega_{pe} \approx 6270-6300$. The acceleration pattern is evident in momentum (panels d-f) and in E_x (panel c); in the latter the acceleration episode is associated with an average negative E_x field values. However, most pre-accelerated electrons do not show clear acceleration patterns as the one shown in Figure 5.11, which complicates the analysis. Thus, we utilise neural networks that are able to distinguish even the most intricate patterns, as long as they are trained with sufficient data. This analysis focuses on the characteristic features of SSA seen in E_x and particle momenta.

For the analysis the label y is computed for each electron, corresponding to their maximum normalised kinetic energy:

$$y_i = \max \frac{E_{k,i}}{m_e c^2} = \max (\gamma_i - 1), \quad (5.53)$$

where $E_{k,i}$ is the kinetic energy of the i -th electron. The distribution of y_i for the utilised dataset is presented in Figure 5.12. The histogram shows two distinct populations: a thermal




Figure 5.11: Interaction of two distinct sample electrons with the Buneman waves in the shock foot. The electron for which the data are shown in red is accelerated via SSA, whereas the dark blue electron passes the Buneman instability region effectively unaffected. Panel (a): the map of the E_x electric field component at time $t\omega_{pe} = 6375$. Overlaid are the positions of the electrons at the same moment (black dots), and their trajectories over the past $200\omega_{pe}^{-1}$. The two plus signs indicate the period of trapping of the red-lined electron. Panel (b): evolution of the kinetic energy of electrons, $\gamma - 1 = E_k/m_e c^2$. Panel (c): normalised amplitude of E_x electric field at electron positions in the simulation frame. Panels (d-f): evolution of electron momenta components.

population at low energies that contains most of the particles; and a non-thermal population containing $\sim 2\%$ of the electron population.

We employ two different types of neural networks for the analysis to test on the dataset:

- **Convolutional Neural Network (CNN):** the core components of this NN are the convolutional layers, which, as the name suggests, perform a convolution operation using a convolutional kernel with a given size over the input. CNNs are extremely useful in extracting patterns from continuous data such as images (2D) or time series (1D), the latter being our case.
- **Multilayer Perceptron (MLP):** one of the simplest NNs that are commonly used for general purpose due to their simplicity. They consist of a fully connected network of layers with varying size. Although they are usually outperformed by CNN in many types of problems, MLPs are faster and can be used as baseline when comparing to other NNs, as is done in this study.




Figure 5.12: Histogram of the maximum kinetic energy of the electrons for the dataset obtained from run E2 in Bohdan et al. 2019a. There are clearly two distinct populations: a thermal electron population at low energies and a non-thermal population (around 2% of the total population) that closely resembles a power-law at higher energies. The dashed curve represents a Maxwellian. The vertical dotted lines split the histogram in the thermal, supra-thermal, and high-energy populations that will be used for classification as explained later in the text. Note this histogram does not correctly portray the efficiency of SSA (and thus the power-law index of the non-thermal population) as this shows the maximum value of the kinetic energy for a certain period of time. However, the energy of the particles in the thermal population remains almost constant during this period and therefore the Maxwellian fits well.

The use of the NNs depends on the analysis algorithms that are performed on the data. We employ three different algorithms:

- **Classification:** a supervised machine learning algorithm whose goal is to categorise input data into predefined discrete classes or labels. The output of the algorithm consists of k neurons where k is the number of classes defined in the dataset. Each neuron contains a softmax activation function (Goodfellow et al. 2016). The value of each neuron has a number between 0 and 1, representing the likelihood that the input data falls into that class. If the classes are mutually exclusive, the sum of output values equals 1. To use this algorithm in our analysis, we split the dataset in three classes given by the values 0, 1 and 2 given by the dotted lines in Figure 5.12, corresponding to low-energy, high-energy and very-high-energy particles, respectively. The details can be found in Table 5.1.
- **Regression:** a supervised machine learning algorithm employed to make predictions

Table 5.1: Maximum particle kinetic energy ranges for which we define classes in our dataset. Also shown are the numbers of particles in each class for our dataset. Note the imbalance in the number of particles per class.

Maximum Kinetic Energy	Class	No. particles
$\gamma - 1 \leq 0.1$	0	247722
$0.1 < \gamma - 1 \leq 0.2$	1	3788
$\gamma - 1 > 0.2$	2	490

about a continuous-valued feature in the input data, compared to classification, which is discrete. The output layer of the algorithm is an array of neurons with liner activation function $y = x$ such that each predicts the value of a feature in the dataset.

- Anomaly Detection (AD): an unsupervised machine learning algorithm designed to identify input values that differ from the norm. This algorithm is suitable for imbalanced datasets in which infrequent values are categorised as outliers or anomalies as the method's name suggests. It uses an autoencoder as the neural network. This NN comprises a series of layers that, from the input layer, gradually decrease the number of neurons, effectively encoding the data. At the middle of the NN, a bottleneck is created where the most crucial information is compressed into a very small number of neurons. Subsequently, the autoencoder performs the reverse operation, effectively decoding the data, increasing the amount of neurons per layer. Finally, the output layer consists of a layer with the same number of neurons as the input data. Each neuron contains the reconstruction of the input data. If the reconstruction closely resembles the original input data, this data is tagged as non-anomalous because it is a common occurrence in the dataset. Conversely, if the reconstruction differs, the data is tagged as anomalous, a rare occurrence in the dataset.

The results obtained using an MLP and CNN for classification, regression and anomaly detection are shown for two different input datasets: a dataset containing all three components of momentum (p_x, p_y, p_z) and another one that only contains the x -component of the electric field E_x computed at particle's position. Particle momentum is directly related to the kinetic energy of a particle and the corresponding time series are quite smooth (see Figure 5.11d-5.11f). These well-behaved characteristics result in more precise and sensitive predictions. Conversely, the electric field data poses a challenge to NNs since the E_x time series obtained from the simulation present a high level of noise (see Figure 5.11c). Furthermore, the electric field along a particle's trajectory is not intrinsically linked to the kinetic energy of the particle, albeit both momentum and electric field show patterns that reveal whether the electron was accelerated or not. These characteristics of the electric field dataset result in a worse, but expected, performance. Nevertheless, both momentum and electric field datasets yield promising results. The results using the dataset with the electric field will not be presented here and thus I refer to my article Torralba Paz et al. 2025.

5.6.1 Analysis with momentum data

Classification

Figures 5.13 and 5.14 show the confusion matrix and the progression of loss and accuracy across training and validation epochs for NNs, along with a table presenting the precision, recall, and F1-score of the three classes defined in Table 5.1. These metrics are provided for CNN and MLP models, respectively. The confusion matrices are used here to evaluate the performance of our classification models, as they compare predicted classes (columns) with respect to the true classes (rows). The best performance is achieved with a mostly diagonal matrix. The performance metrics – precision P , recall R , and $F1$ -score – are calculated from the confusion matrix, M , as follows:

$$P_i = \frac{M_{ii}}{\sum_j M_{ij}}, \quad R_j = \frac{M_{jj}}{\sum_i M_{ij}}, \quad F1_i = 2 \frac{P_i \cdot R_i}{P_i + R_i}, \quad (5.54)$$

where i and j represent predicted (columns) and true (rows) values, respectively. The precision metric quantifies the accuracy of positive predictions made by the model, the recall measures the ratio of true positive predictions to all actual positive instances, and the $F1$ -score combines both P and R . The use of a harmonic mean in the definition of the $F1$ -score is a particularly useful metric for our imbalanced dataset, as it can be high only if both P and R are high.

One can see in the bottom panels of Figures 5.13 and 5.14 that it takes approximately 60–70 epochs for the loss to stabilise during the training and validation phases of the CNN and MLP networks with the classification algorithm. Note the spikes visible in the loss and accuracy figures. The training dataset is divided into small batches, each containing a random subset of the training data. This subset changes every epoch. Occasionally, batches may contain unfavorable values, such as batches consisting entirely of class 0 particles, which can slightly affect the results and lead to deviations in the weights of the NN, resulting in a higher loss than expected. However, the NN optimiser recognises this temporary increase in loss and returns to a low loss after a small number of epochs. This behavior can be clearly observed in the bottom panels Figure 5.14, e.g., between epochs ~ 70 and ~ 75 , and epochs ~ 110 and ~ 115 .

As shown in the tables of Figures 5.13 and 5.14, class 0, being the most abundant, achieves a very high F1-score, as expected. Both class 1 and class 2 also achieve high F1-scores. Interestingly, class 1 has a lower F1-score than class 2, despite containing fewer particles. This can be attributed to the distinctive time series shapes of the very high-energy particles in class 2 (characterised by large-amplitude momentum variations, see Fig. 5.11d-f), which aid the NN in their recognition. For both NNs, the results are notably robust, with F1-scores exceeding 0.7 for all classes. Overall, both the CNN and MLP demonstrate remarkably good performance.

Regression

Figures 5.15 and 5.16 show for the CNN and MLP, respectively, the linear regression applied to the true (x -axis) and predicted (y -axis) values of the maximum kinetic energy in panels




Figure 5.13: Results for classification with CNN. Top panel: the confusion matrix. The numbers show the distribution of data across all possible combinations of true and predicted labels for the 3 considered classes. Middle panel: precision P , recall, R , and $F1$ -score derived from the confusion matrix. Bottom panel: the loss and accuracy over the epochs required to train and validate the NN.

Figure 5.14: Results for classification with MLP (compare Fig. 5.13).

(a), the loss obtained during training in panels (b), and the histogram of both true and predicted values in panels (c). The linear regression fits are shown in panels (a) with orange lines, whereas the black lines $y = x$ represent a perfect score. In our calculations, we use the Huber loss function, which yields the lowest loss among the loss functions we have tested.

For both the CNN and MLP, the regression lines lie very close to the black line, as desired. To qualitatively assess the results, the coefficient of determination R^2 is computed. We obtain $R^2 = 0.9886$ for the CNN and $R^2 = 0.9782$ for the MLP, indicating near-perfect fits. In both cases we observe that high-energy particles are slightly overpredicted, showing higher energies than they should. These effects are visible in the histograms of true (blue) and predicted (orange) values. For the CNN, there is a slight overprediction at high energies,

but overall, the predicted histogram closely matches the true one. The MLP also shows a similar slight overprediction on the high-energy side. One might expect the opposite, with particles being underpredicted due to the large number of low-energy particles on which the NNs are trained. However, it is possible that the strong effect of sample weighting on high-energy particles biases the regression towards these particles. In conclusion, both the CNN and MLP yield very similar goodness-of-fit, indicating that either could be used for particle data analysis. We note however, that the MLP performs slightly faster.

Anomaly detection

The results for anomaly detection were obtained using an autoencoder based on convolutional layers. While an autoencoder based on fully-connected layer was tested, its results were significantly inferior and thus only the CNN-based autoencoder results are presented.

As the first approach, the autoencoder was trained exclusively with low-energy particles, with $\max(\gamma - 1) < 0.07$, to prevent exposure to any uncommon data that could bias the results towards high-energy particles, thereby potentially deteriorating the outcome significantly. This approach to anomaly detection is called semi-supervised anomaly detection, as we can confidently identify a portion of our dataset composed of common particles only, i.e. low-energy thermal population. Higher-energy particles are excluded from this sample, including those within the energy range $\max(\gamma - 1) \approx 0.07-0.1$, in which the thermal and non-thermal populations blend together. As the second approach, the autoencoder underwent training using the entire training dataset, including the non-thermal population, to investigate the effect of uncommon input data during the training phase.

Top left panel of Figure 5.17 shows the loss obtained after training the NN with low-energy particles only. We choose the thresholds for each momentum component to be the tenth highest loss calculated, as indicated with the dashed lines. Although we could have chosen the highest loss as the threshold, the NN training is not perfect, and some particles may produce excessively high losses, which could bias the testing results and lead to worse overall predictions. Therefore, we introduce some flexibility by selecting a slightly lower value. Top right panel of Figure 5.17b illustrates the testing loss. If the loss for an input variable exceeds the threshold (indicated again with dashed lines), the variable is considered an anomaly. Since particles are characterised by three momentum components, we consider a particle an anomaly if the loss in any of the momentum components is higher than the defined threshold.

Middle panels of Figure 5.17 summarises our results. In each panel, the horizontal dashed lines represent the thresholds computed for each feature, while the vertical lines at $\max(\gamma - 1) < 0.1$ indicate our separation between energetic and non-energetic (thermal) particles, identical to the one used in classification to distinguish class 0 from class 1 and 2 particles (see Table 5.1). Ideally, all energetic particles would lie in the top-right quadrant. However, time series of energetic particles located close to the crossing point where both dashed lines intersect become similar to those of non-energetic particles. An energised particle may therefore be mistaken for a thermal one, leading the autoencoder to fail in differentiating them, often labeling these ambiguous time series as non-anomalies. The results from the middle panel are summarised in the table showing the confusion matrix, bottom




Figure 5.15: Results for regression with CNN. Figure 5.16: Results for regression with MLP (compare Fig. 5.15). Top panel: regression of the true value of the energy and the energy predicted by the NN. Middle panel: evolution of loss during training. Bottom panel: comparison of the training values and the predicted values during the testing phase.

of Figure 5.17d. Out of 1073 energetic particles (actual anomalies) in our testing sample of 63,000 particles, the autoencoder correctly predicts 1069 particles, yielding a superb sensitivity (recall) of 0.996. However, at the same time the autoencoder predicts more anomalies than it should have, with 313 non-energetic particles classified as anomalies, resulting in

the precision of 0.774, which is still remarkable. These misclassified non-energetic particles likely reside very close to the vertical dashed line (energy separation), meaning that their time series are somewhat similar to those of energetic particles. Overall, AD with CNN-based autoencoder performs remarkably well with the accuracy (the ratio of correctly predicted observations to the total observations) of 99.6% and total F1-score of 87.9%.

This result can be elucidated with an example: if we trace new particles and feed them to the autoencoder, we can be confident that more than 3/4 of the anomalies (approximately 77%) are pre-accelerated particles, and we miss less than 0.5% of energetic particles.




Figure 5.17: Results for semi-supervised anomaly detection with CNN. Top panels: the training and testing loss for p_x , p_y and p_z , respectively. The dotted vertical lines indicate the thresholds used for AD, which during testing, split the histogram into anomalies and non-anomalies. Middle panel: the maximum kinetic energy vs. the testing loss. The horizontal dashed lines represents thresholds for each momentum component and the vertical lines depict our division between energetic and non-energetic particles. Table: the confusion matrix obtained from results in the middle panel.

Figure 5.18 shows results for anomaly detection in which the NN was trained using

particles selected randomly from the entire dataset, i.e., including supra-thermal and high-energy populations. The thresholds have been chosen to correspond to 1% of the dataset. One can see in the top left panels, that the thresholds (dashed lines) are now positioned in the middle of the histogram rather than on the right, as in the previous case of AD. We have checked that higher or lower thresholds result in worse performance of the NN. As expected, including the energetic populations in the training dataset degrades the overall performance. The middle panels clearly shows a wider spread in all momentum components compared to the Figure 5.17, which is largest for the p_y component. To address this, we redefined the anomaly criterion to include two variable losses above the threshold instead of just one, as in the semi-AD case. The final result is shown in the confusion matrix in the bottom table of Figure 5.18d. Although the precision is high, $P = 97.9\%$, as we obtain a small number of false anomalies (13), we are missing nearly 42% of energetic particles (448), which results in low recall of 57.5% and F1-score of 72.5%. The AD performance is thus clearly better when only low-energy particles are used for NN training.




Figure 5.18: Results for fully unsupervised anomaly detection with CNN (compare Fig. 5.17).

5.6.2 Analysis with E_x data

Classification

Figures 5.19 and 5.20 present classification results using electric field data with CNN and MLP models, respectively, and are compared with the momentum-based analysis in Section 5.6.1. Training converges faster than for momentum data, although the increasing validation loss indicates overfitting. Both NNs classify class 0 particles with high accuracy, achieving F1-scores comparable to those obtained with momentum data. For class 1 particles, the use of the electric field data yields high precision but reduced recall, leading to an F1-score comparable to the momentum-based CNN but significantly lower for the MLP due to many class 1 particles classified as class 0. Classification of class 2 particles remains poor for both NNs, with F1-scores around 0.5, primarily due to confusion with class 1. Overall, classification based on electric field data performs worse than momentum-based classification, reflecting the higher noise level and limited information content of a single electric field component compared to the three-component momentum data.

True \ Pred.	Pred.		
	0	1	2
0	61685	242	0
1	121	800	34
2	0	56	62

Class	Precision	Recall	F1-score
0	0.998	0.996	0.997
1	0.729	0.838	0.779
2	0.646	0.525	0.579




Figure 5.19: Results for classification with CNN using only the electric field data. See Figure 5.13 for more details.

True \ Pred.	Pred.		
	0	1	2
0	61808	119	0
1	271	659	25
2	6	61	51

Class	Precision	Recall	F1-score
0	0.996	0.998	0.997
1	0.786	0.690	0.735
2	0.6711	0.432	0.526




Figure 5.20: Results for classification with MLP using electric field data (compare Fig. 5.19).

Regression

Figures 5.21 and 5.22 show the regression results for the CNN and MLP models, respectively. Consistent with the classification results, the CNN outperforms the MLP, achieving $R^2 = 0.9011$ compared to $R^2 = 0.7407$. In both cases, the regression is biased towards lower energies due to the large thermal particle population, leading to a lower value of the R^2 scores. Although the histograms in Figures 5.21c and 5.22c appear similar, the linear regressions feature a significant spread in the predicted values that is not fully captured by the histogram analysis.




Figure 5.21: Results for regression with CNN. See Figure 5.15 for more details.

Figure 5.22: Results for regression with MLP using electric field data (compare Fig. 5.15).

Anomaly detection

For anomaly detection using electric field data, the model was trained exclusively on low-energy particles, as this approach performed better than training on the full dataset. The results are shown in Figure 5.23, with the anomaly threshold set to 1% of the full dataset, consistent with the momentum-based analysis. The training and testing loss distributions are nearly identical, as evidenced by the left plot, indicating that the network fails to distinguish energetic particles from the thermal population. This behaviour arises from the intrinsically noisy nature of the electric field time series, which obscures clear non-thermal signatures. As illustrated in the right plot, low-energy particles can exhibit large reconstruction losses, while most energetic particles are misclassified as non-anomalies. This is reflected in the confusion matrix (see the table in Fig. 5.23), which shows a poor classification of anomalies. Overall, anomaly detection based solely on electric field data proves ineffective for analysing particle trajectories in PIC simulations.




Figure 5.23: Results for anomaly detection with CNN for the electric field data (compare Fig. 5.15).

5.6.3 Conclusions on the traced data analysis with neural networks

Neural networks are employed to predict the maximum kinetic energy of particles passing through the Buneman instability region, where some of these particles undergo SSA. Using the three momentum components, supervised learning algorithms – both classification and regression – achieve good performance despite the highly non-linear nature of the dataset.

Anomaly detection is also explored as an unsupervised alternative using both fully unsupervised and semi-supervised strategies. While the semi-supervised approach outperforms the fully unsupervised one, both perform significantly worse than supervised methods, owing to the absence of labels.

Alternatively, a single component of the electric field, E_x , is considered. Despite the high noise level, both classification and regression produce remarkable results, although the performance is inferior to the momentum-based analysis. Overall, the CNN outperforms the MLP. Anomaly detection, however, yields poor results and its use is discouraged for such noisy data.

These results demonstrate that the energies of particles interacting with the Buneman instability can be predicted from either momentum time series or even a single electric field component. Performance depends strongly on the chosen input data and learning strategy, providing guidance for future applications on other acceleration processes. As numerical simulations continue to increase in scale, neural-networkbased techniques are expected to enable more efficient and accurate analysis of simulation outputs, particularly for particle-tracking data in fully kinetic PIC and hybrid kinetic simulations.

This work, published in Torralba Paz et al. [2025](#), is the first application of neural networks to particle tracing data from PIC simulations. Recent work by Swierc et al. [2025](#) applies similar methods to study particle injection into DSA and SDA in hybrid simulations, reaching similar conclusions to those presented here and further highlighting the potential of neural networks for kinetic plasma analysis.

CHAPTER 6

PIC simulations of magnetised mildly relativistic shocks

UNDERSTANDING particle acceleration processes in astrophysical sources that emit high-energy radiation, such as microquasars and AGN, is fundamental to unravelling the origin of cosmic rays. One of the most compelling scenarios for particle acceleration in these systems involves plasma shocks. Recent observations of the microquasar SS433 indicate that shock-driven particle acceleration is the most plausible mechanism responsible for the intense emission observed in its jets (H.E.S.S. Collaboration* et al. 2024). Likewise, shock-acceleration scenarios are widely invoked to explain the radiation produced in AGN jets (Zech and Lemoine 2021; Thimmappa et al. 2022; Spada et al. 2001b; Spada et al. 2001a).

Two well-established mechanisms are frequently invoked to explain particle acceleration in plasma shocks: DSA and WFA. As discussed in Section 4.1.1, DSA operates efficiently in unmagnetised plasmas or weakly magnetised plasmas under subluminal shock conditions. However, the shock conditions relevant to many astrophysical environments with relativistic outflows – in particular AGN and microquasars – are typically moderately or strongly magnetised. In the latter case, the magnetic field energy is comparable to the bulk plasma energy. Moreover, shocks in these sources are expected to be mildly relativistic. Under such conditions, DSA is not expected to operate efficiently. Similarly, while WFA is highly effective at accelerating electrons and ions in ultra-relativistic shocks, its efficiency decreases significantly in the mildly relativistic regime and is therefore insufficient to account for the observed emission (see Section 4.2). This indicates that an alternative acceleration mechanism(s) must operate in these magnetised environments to produce the observed cosmic rays and radiation.

In this chapter, the results of high-resolution, large-scale 2D3V PIC simulations of mildly relativistic shocks in a strongly magnetised plasma are presented. This parameter regime has not been studied before. The magnetic field obliquity is found to play a crucial role in determining both the shock dynamics and the efficiency of particle energisation. Accordingly, three representative shock configurations are investigated:

1. **A superluminal oblique shock**, which demonstrates that magnetic field obliquities exceeding the threshold angle are unfavourable for efficient particle acceleration;
2. **A subluminal shock**, with an obliquity slightly below the threshold angle, which exhibits highly efficient ion and electron acceleration via the mechanism first proposed by Bessho and Ohsawa (1999) for nonrelativistic shocks;
3. **A quasi-parallel shock**, which generates large-amplitude whistler waves that predominantly lead to particle heating rather than non-thermal acceleration.

For all shock configurations considered, the resulting shock structures are shown to differ substantially from those known in weakly magnetised plasmas, highlighting the qualitative impact of strong magnetisation on collisionless shock dynamics.

6.1 Simulation setup

THIS chapter presents three different PIC simulations, with their parameters summarized in Table 6.1. The simulations are configured in 2D3V, meaning two spatial dimensions (x and y) and all three components of the vector fields: electromagnetic fields, momentum, and current. Figure 6.1 shows a sketch of the simulation setup. An ion-electron plasma beam lying in the $x - y$ plane flows in the negative x -direction with velocity $v_{\text{up}} < 0$. The shock is generated using the *reflecting wall* setup as explained in Section 5.4. A background magnetic field is initialised on a Yee lattice, forming an angle $\theta_{Bn} = \{10^\circ, 30^\circ, 40^\circ\}$ with the x -axis, and an angle $\varphi_B = 90^\circ$ with the $y - z$ plane, corresponding to an out-of-plane magnetic field configuration. The plasma is initialised in a current-neutral configuration, $\mathbf{j} = \sigma (\mathbf{E} + \mathbf{v}_{\text{up}} \times \mathbf{B}) = 0$. Therefore, a motional electric field $\mathbf{E}_0 = -\mathbf{v}_{\text{up}} \times \mathbf{B}$ is introduced, which for an out-of-plane magnetic field points in the negative y -direction, $E_{0y} = v_{\text{up}} B_{0z} < 0$.




Figure 6.1: Sketch of the shock setup used for the simulations. The plasma beam moves in the negative x -direction with speed v_{up} . The plasma reflects off a conducting wall on the left and the interaction of the reflected and incoming particles forms a shock wave moving in the positive x -direction with speed v_{sh} . A background magnetic field $\mathbf{B}_0$ is initialised, forming an angle θ_{Bn} with the x -axis and an angle φ_B with the $y - z$ plane.

The simulations use a reduced mass ratio $R = m_i/m_e = 100$, which is sufficient to clearly separate ion and electron dynamics. Using a reduced mass ratio also increases the ion cyclotron and plasma frequencies, allowing for longer simulations in physical time scales Ω_i^{-1} , ω_{pi}^{-1} . For comparison, using the real mass ratio $R = 1836$ would increase the simulation time by approximately a factor of eighteen to cover one ion gyroperiod, and by roughly a factor of four to cover the inverse of the plasma frequency.

The number of particles per cell and per species (ions and electrons) is set to $n = 10$ ppc (particles per cell). The simulation resolution is determined by the electron skin depth in units of the grid cell size, Δ , which is set to $\lambda_{se} = 40\Delta$. This ensures that electron-scale

Sim. #	$\theta_{Bn}(^{\circ})$	$\theta_{Bn,up}(^{\circ})$	$L_y(\lambda_{si})$	$T_{\max}(\Omega_i^{-1})$	β_{sh}	$\beta_{sh,up}$	$M_{A,up}$	$\theta_{tr}(^{\circ})$
$\theta 40$	40	20.6	7.2	113.1	0.58	0.971	1.106	29.0
$\theta 30$	30	14.5	4.8	314.1	0.49	0.963	1.077	32.1
$\theta 10$	10	4.51	4.8	314.1	0.33	0.945	1.035	36.1

Table 6.1: Simulation parameters for the runs presented in this chapter. The vertical line separates the initial setup parameters (left) from those derived from the simulation (right). From left to right: magnetic field obliquity in the simulation frame, θ_{Bn} , and in the upstream frame, $\theta_{Bn,up}$; transverse box size, L_y ; maximum simulated time, $T_{\max}$; shock speed in the simulation frame, β_{sh} , and in the upstream frame, $\beta_{sh,up}$; Alfvénic Mach number in the upstream frame, $M_{A,up}$; and the threshold angle in the simulation frame, θ_{tr} .

physics is resolved over 40 cells, providing high spatial resolution. The corresponding ion skin depth is $\lambda_{si} = \sqrt{R}\lambda_{se} = 400\Delta$. The plasma frequencies are set automatically as $\omega_{pi} = 1.25 \times 10^{-3} \Delta t$ for ions and $\omega_{pe} = 1.25 \times 10^{-2} \Delta t$ for electrons.

To suppress numerical artifacts, we employ the techniques described in Chapter 5, including the Friedman filter for the electric field, a fourth-order field pusher, and the newly implemented Higuera-Cary pusher. The CFL condition is satisfied by setting the grid cell size and time step to $\Delta = \Delta t = 1$, and by choosing a speed of light $c = 0.5 < c_{\max} = 2^{-1/2}$.

The Lorentz factor of the upstream beam is $\gamma_{up} = \sqrt{5} \approx 2.24$, while the shock Lorentz factor depends on the specific configuration, as indicated in Table 6.1. The beam is initially cold ($v_{i/e,th} = 0$), corresponding to a plasma beta $\beta_p = 0$ and a sound speed $c_s = 0$, which implies a virtually infinite sonic Mach number, $M_s \rightarrow \infty$. With zero temperature, the Debye length cannot be resolved ($\lambda_D = 0$). In principle, an unresolved Debye length may lead to numerical artifacts in the plasma after a short simulation time. However, no such artifacts are observed within the duration of our simulations. Since the plasma is completely cold, there are no initial thermal fluctuations, and the plasma naturally thermalises downstream after crossing the shock or becomes heated through physical acceleration processes. It should be noted that if the temperature were slightly higher while still unresolved, numerical artifacts could develop on shorter timescales and potentially compromise the simulation results.

The initial plasma is strongly magnetised, with a total magnetisation of $\sigma = 1$. The magnetisation for each species, defined such that $\sigma^{-1} = \sigma_i^{-1} + \sigma_e^{-1}$, is $\sigma_i \approx 1$ for ions and $\sigma_e \approx 100$ for electrons. The corresponding cyclotron frequencies are therefore set to $\Omega_i \approx \omega_{pi} \approx 1.25 \times 10^{-3} \Delta t$ for ions and $\Omega_e \approx 1.25 \times 10^{-1} \Delta t$ for electrons.

Due to the strong magnetic field, the non-relativistic Alfvén speed exceeds the speed of light, $v_A \approx B/\sqrt{\mu_0 n_i m_i} \approx 1.5 > c$. Therefore, the relativistic Alfvén speed must be used (see Equation 3.21), which is $\beta_{A,rel} \approx 0.83$ in the simulation rest frame, and $\beta_{A,rel,up} \approx 0.89$ in the upstream frame of reference. This corresponds to an Alfvénic Mach number close to $M_{A,up} = v_{sh,up}/v_{A,rel,up} \sim 1.08$, with the value depending on the shock speed. In this regime, the shock structure differs significantly from that of the commonly studied supercritical shocks ($M_A \gg 1$, *strong shocks*).

The computational box is periodic in the y -direction, with a fixed transverse size L_y , while the longitudinal size L_x increases over time, starting from $L_x = 11\,000 \Delta \approx 25 \lambda_{si}$ and

extending up to $L_x = 100\,000\,\Delta \approx 275\,\lambda_{\text{si}}$. Given the planarity of the shock, this transverse size is sufficient to capture the relevant physics while conserving computational resources, which allows for longer simulation times.

The results are presented in the order listed in Table 6.1: first, the superluminal oblique shock, followed by the subluminal oblique shock near the threshold angle, and finally the quasi-parallel shock.

6.2 Superluminal shock (simulation $\theta 40$)

6.2.1 General structure of the shock transition

Figure 6.2 shows the late-stage ($t\Omega_i = 113.08$) structure of a superluminal shock (simulation $\theta 40$, see Table 6.1), whose obliquity, $\theta_{Bn} = 40^\circ$, considerably exceeds the threshold angle, $\theta_{tr} \approx 29^\circ$, in the simulation frame. The shock transition features a series of sharp peaks in both density (top panel) and the magnetic field (bottom panel), starting at $x \approx 65.2 \lambda_{si}$. Each of these peaks is reminiscent of a *wave pulse* that act together as the shock transition. The compression of the B_z field component at the location of the first pulse is $B_z/B_{0z} \approx 3.9$ ($\delta B_z/B_0 \approx 1.83$) and of the ion density $n_i/n_0 \approx 18.7$. Associated with this pulse, a strong electrostatic gradient develops with an amplitude of $\delta E_x/E_0 \approx 0.8$, where $E_0 = |-\mathbf{v}_{up} \times \mathbf{B}_0| = v_{up} B_{0z}$. Here and throughout, a fluctuation amplitude is defined with respect to the mean field of the respective component, e.g., $\delta B_z/B_0 = (B_z - B_{0z})/B_0$.

Following the first pulse, which is here referred to as the *shock front*, several additional pulses appear at $x/\lambda_{si} \approx 56.4, 58.6, 60.6$, and 62.6 , each exhibiting gradually lower amplitudes in the magnetic field, density, and electrostatic gradient. These six pulses propagate with a velocity $\beta_{sh} \approx 0.58$, corresponding to a Lorentz factor of $\gamma_{sh} \approx 1.23$. In the upstream frame, this translates to a Lorentz factor $\gamma_{sh,up} \approx 4.2$ and an Alfvénic Mach number of $M_{A,up} \approx 1.1$. Note that the general structure of the shock is consistent with that of a large-amplitude magnetosonic wave discussed in Section 4.3.1, as further demonstrated in Section 6.2.3.

The downstream compression ratio of the shock, measured within the region $x/\lambda_{si} = 30$ – 50 , is $r = n_{i,dw}/n_0 \approx 2$, with a corresponding magnetic field compression ratio of $r_B = B_{z,dw}/B_{0z} \approx 2.6$. In contrast to the compression ratio of a strong relativistic shock ($r \sim 4$), both r and r_B are lower, as expected. Although no analytical expression exists for the compression ratio in the relativistic regime, the same general trends for the compression ratio observed in non-relativistic shocks remain applicable in mildly relativistic shocks (see Section 3.2.2). Given that the Alfvénic Mach number is close to unity, the dependence of the compression ratio on M_A is approximately linear; for an oblique magnetic field, the increase in compression ratio for a fixed M_A occurs slightly more rapidly than in the perpendicular case. Moreover, as anticipated, $r_B > r$ (Burgess and Scholer 2015).

6.2.2 Ion and electron phase-space

Figure 6.3 presents the ion and electron $\mathbf{p} - x$ phase-space distributions at the same time as in Figure 6.2. The shock front is visible in $p_{x,i}$ as a peak at $x \approx 65\lambda_{si}$. Ions crossing the shock front are abruptly decelerated by the cross-shock potential, while the compressed magnetic field deflects their trajectories in both the x and y -directions. In the case analysed here, since the ions are cold and magnetised, they collectively gyrate in the magnetic field, performing a cycloidal motion in $p_{x,i}(x)$. In $p_{y,i}(x)$, ions reach a maximum momentum before rapidly decreasing to negative values and eventually approaching zero farther downstream after a series of gyrations (Nakazawa and Ohsawa 1997).

Electrons, in contrast, are accelerated by the cross-shock potential as they cross the shock




Figure 6.2: Maps of the ion density, n_i/n_0 (top panel), the longitudinal electric field, E_x/E_0 (middle panel), and out-of-plane component of the magnetic field, B_z/B_0 , for simulation $\theta 40$ at time $t\Omega_i = 113.08$. All quantities are normalised to their upstream values, denoted by the subscript 0. The plot has been stretched in y -direction to show the structures in detail. Although the maximum value of ion density is $n_i/n_0 \approx 18.7$, the colourbar maximum value is $n_i/n_0 = 8$ that allows to see lower density features.

front, leading to a decrease in $p_{x,e}$ towards more negative values. The z -component of the electron momentum, $p_{z,e}$, exhibits strong modulation due to the abrupt increase of the magnetic field amplitude. Because electrons are highly magnetised, they predominantly follow magnetic field lines, which are oriented mainly along the z -direction, as the magnetic field itself is nearly aligned with the z -axis. Consequently, $p_{z,e}$ decreases even further relative to $p_{x,e}$. As electrons are advected downstream, the magnetic field amplitude fluctuates within the magnetosonic pulse that constitutes the shock transition, causing quasi-periodic variations in $p_{z,e}$. Note, that electrons traversing the shock front are already slightly heated in the extended upstream region, $x/\lambda_{si} \approx 65$ – 90 . This is due to wakefield excitation by precursor waves. At the same time ions remain largely unaffected by the wakefield (see Section 6.2.3).

As particles cross the magnetosonic pulse, they scatter off the electromagnetic fluctuations, which results in their heating and a broadening of their momentum distributions farther downstream. Ions isotropise slowly, as evidenced by the gradual broadening of their distribution in p_i that becomes wider between $x/\lambda_{si} \approx 35$ and 60 , whereas electrons isotropise more rapidly and reach a semi-isotropic state on the scale of approximately $15\lambda_{si}$.

The downstream region does not appear to be fully thermalised nor stationary even at the latest simulation time. This is evidenced in Figure 6.3 by finite-velocity motion right downstream of the magnetosonic pulse at $x/\lambda_{si} \approx 30$ – 50 (see also discussion in Section 6.2.3) and also by the persistent two-beam structure at $x \lesssim 30\lambda_{si}$, composed of incoming particles and those reflected from the left-boundary wall. The latter feature is visible in the x - and




Figure 6.3: Phase-space for ions (left) and electrons (right) for all components of momentum: p_x (top), p_y (middle) and p_z (bottom) at $t\Omega_i = 113.08$. The colour-scale shows the logarithm of the number of particles with a normalised position, x/λ_{si} , and momentum, p/mc , where $m = m_i, m_e$. Compare Fig. 6.2.

z -components of the ion and electron momenta. The y -components remain symmetric with respect to $p_y = 0$, consistent with the magnetic field lying mainly in the $x - z$ plane. In addition, the z -component of both ion and electron momenta exhibits a finite drift (see, e.g., $p_{z,i}$ at $x/\lambda_{si} \approx 40$ – 50), which is characteristic of low-Mach-number shocks (Treumann 2009, Chapter 2). Because the plasma is highly magnetised, ions and electrons do not experience strong diffusion or significant scattering by waves, which results in slow thermalisation of both populations. Note, however, that although the downstream region has not yet reached a steady state, the magnetosonic wave pulse at the forward edge of the shock transition is a well established physical feature that propagates with a velocity β_{sh} much larger than the reflected beam speed, $v_x \approx 0.1c$. Similar behaviour would also occur in a two-beam setup, in which two counter-streaming plasma flows generate shocks propagating in opposite directions. In that configuration, both plasmas would form shocks while the downstream particles would continue to stream along magnetic field lines with little obstruction for a long time, resulting in a momentum distribution qualitatively similar to that shown in Figure 6.3 (see also discussion in Section 5.4).




Figure 6.4: Transversely averaged profiles of the normalised magnetic field, δB_x , B_y and δB_z (a), and electric field components, E_x , δE_y and E_z (b), around the shock transition at $x/\lambda_{si} = 55\text{--}70$ and time $t\Omega_i = 113.08$. Black vertical lines mark the positions of the B_z and E_y peaks observed in Fig. 6.2. The distance between the peaks is approximately $\Delta x \approx 2\lambda_{si}$.

6.2.3 Wave modes in the shock transition

Structure of the magnetosonic wave pulse

Since the average shock structure is one-dimensional along the x -direction (compare Fig. 6.2), it is convenient to analyse spatial profiles. Figure 6.4 shows the profiles of the magnetic and electric field components, in panels (a) and (b), respectively. The black vertical lines indicate the positions of the B_z peaks observed in Figure 6.2. The amplitude of these peaks is progressively lower the farther downstream they are, with $\delta B_z/B_0 = [1.83, 1.54, 1.42, 1.34, 1.20]$, counting from the shock front at $x/\lambda_{si} \approx 65.2$. Note that the E_y electric field profile closely follows that of δB_z . Fluctuations in B_y as well as in E_x and E_z are finite, and both B_y and E_x become zero at positions of the B_z extrema, indicating that these components are proportional to gradient of the B_z magnetic field, $\partial B_z/\partial x$. This demonstrates that the average structure of this oblique superluminal mildly relativistic shock is similar to that of small- and large-amplitude oblique magnetosonic waves described earlier in the nonrelativistic shock regime (see Section 4.3.1). Note, however, that the shock is non-stationary and exhibits significant transverse non-uniformity. The latter is evident in Figure 6.2, which shows clear modulations in the E_x and B_z amplitudes along the crests of the magnetosonic pulses. These modulations are present in all pulse peaks, though they become more pronounced in the second and subsequent ones, and their amplitudes can reach values comparable to the local mean fields. They occur on spatial scales similar to the ion gyroradius, yet they do not form a coherent wave pattern with a defined propagation direction. Their morphology also varies substantially on timescales comparable to the ion gyroperiod. In the averaged profiles

shown in Figure 6.4, the amplitude of the B_z fluctuations in the first peak oscillates within the range $\delta B_z/B_0 \approx 1.4\text{--}1.8$, and the first peak can occasionally be weaker than the second one. These characteristics make it difficult to attribute the fluctuations to any known instability. They may instead result from the strong local non-stationarity of the large-amplitude waves, which alters the instantaneous conditions for ion motion and locally perturbs the ion gyration. These perturbations then propagate along the wave crests.

The observed average shock structure is reminiscent of that of a weak subcritical shock, in which the transition region is dominated by dispersive wave modes in the magnetosonic-whistler branch with dissipation through anomalous collisions. Weak shocks have been observed in Earth's Magnetosheath and in the interplanetary space (see Balogh and Treumann 2013 and Treumann 2009 and references therein). A weak subcritical shock forms when the fast magnetosonic Mach number is lower than the critical Mach number, $M_f < M_{cr}$ (see Section 3.2.4). In a cold plasma ($\beta_p \rightarrow 0$), the fast magnetosonic Mach number approaches the Alfvén Mach number, $M_f \approx M_A$. For a stable shock-like solution in the subcritical regime, anomalous dissipation mechanisms and wave dispersion are sufficient to dissipate energy and prevent infinite wave steepening. The structure of weak shocks has been investigated by Sagdeev (1966) in the two-fluid model and comprehensively summarised in Treumann (2009) and Balogh and Treumann (2013). As noted above, in the case of a cold magnetised shock with $M_A \gtrsim 1$ and $M_s \rightarrow \infty$, the transition is characterised by a steep ramp followed by short-wavelength oscillations in the downstream region, which are gradually damped by anomalous dissipation. This behaviour thus closely resembles that of B_z in Figure 6.4.

Figure 6.4 shows that the approximate distance between the peaks in the magnetosonic wave pulse is $\Delta x \approx 2\lambda_{gi}$, which is on the order of the ion gyroradius in the background magnetic field. The shock structure can therefore be interpreted as being set by the ion gyration in the wave's magnetic field with the $\mathbf{E} \times \mathbf{B}$ drift. Ions and electrons approaching the wave from the upstream perform a cycloid motion in the electromagnetic field of the wave (Nakazawa and Ohsawa 1997). The oscillations in phase-space at the forward part of the magnetosonic shock, visible in Figure 6.3, are consistent with the expected amplitudes of the particle momenta. For example, the estimated cycloid amplitude in the $p_{y,i}$ component is $\Delta p_{y,i}/m_i c \simeq e\Delta E_y/m_i \omega \approx 2$, which agrees well with fluctuation amplitude at the first phase-space peak. This estimate was made using $\omega/\Omega_i \approx 1.5$ and $\Delta E_y/E_0 \approx 2$.

Downstream entropy mode waves

Figure 6.5 shows the temporal evolution of the ion density. The shock is moving to the right at a speed $v_{sh} = 0.58c$ and no strong waves appear propagating upstream. On the other hand, downstream of the magnetosonic pulse waves are visible moving with a phase velocity of $v_{ph} \approx -0.25c$ in the simulation frame, which corresponds to $v_{ph} \approx -0.72c$ in the shock rest frame. As these waves advect further downstream, they gradually damp and lose coherence. The same structures are also evident in the electron density (not shown), albeit with smoother profiles.

It is important to note that, although the downstream waves move away from the shock in the simulation frame of reference, since thermalisation of the flow is very slow, the downstream still has a finite speed in the simulation frame. This can be observed in $p_{x,i}$ and $p_{z,i}$




Figure 6.5: Ion density average profile evolution in time interval $t\Omega_i = 50.26\text{--}113.08$. The profile has been smoothed using a Butterworth filter, acting as a low-pass filter. The downstream-directed waves have a phase speed of $v \approx -0.25c$ as indicated by the black line. The right black line indicates the shock speed, $v_{sh} = 0.58c$.

in Figure 6.3. The x -component of the mean ion speed downstream of the magnetosonic wave pulse is $v_x/c \approx -0.22$, which approximately coincides with the speed of the density waves. Thus, these waves move with the downstream plasma flow and appear stationary in the downstream frame of reference.

Figure 6.6 shows a zoom-in of the shock transition structure downstream of the strong magnetosonic pulse, in the region $x/\lambda_{si} = 35\text{--}55$ (compare Figs. 6.2 and 6.4). One can see that density waves are not associated with any significant oscillations in the magnetic field. This, together with the observed stationarity of the waves excludes the possibility that they are magnetosonic waves. The waves should rather be identified with the *entropy waves*. Entropy mode waves appear as a trivial, non-propagating, $\omega = 0$, solution to the dispersion relation for ideal MHD waves (Kulsrud 2005). They are characterised by a constant entropy (as opposed to other MHD waves), move with the plasma flow and carry no magnetic field energy, unlike other ideal MHD waves. This is confirmed in Figure 6.6, which shows the map




Figure 6.6: Top: map of ion density fluctuations $\delta n_i/n_0 = (n_i - n_0)/n_0$ for $x/\lambda_{si} = 35$ –55, which corresponds to downstream of the strong magnetosonic pulse in Figs. 6.2 and 6.4. Middle: profiles of ion density (blue line) and the normalised temperature, $\Theta = k_B T/m_i c^2$ (orange line). The temperature is obtained by fitting a Maxwellian for each component of the ion momentum (see Fig. 6.3) at different locations in x -direction. Bottom: profiles of the normalised magnetic field components, B_x/B_0 , B_y/B_0 , B_z/B_0 . Density and temperature are anti-correlated while the magnetic field shows very weak or no fluctuations.

of the ion density fluctuations, $\delta n_i/n_0$ (top panel), the profile of the ion density fluctuations and the normalised temperature, $\Theta = k_B T/m_i c^2$ in the middle panel, and the components of the magnetic field, B_x/B_0 , B_y/B_0 , B_z/B_0 (bottom panel). To obtain the temperature, Θ , the one-dimensional distribution for each component of ion momentum, $f(p_{i,j})$ has been computed and fitted to Maxwell-Jüttner marginal distribution function for each ion momentum component by averaging phase-space (see Fig. 6.3) in the x -direction over a short distance, $\delta x = 0.1\lambda_{si}$. The total temperature is obtained as the sum of temperatures of each component obtained from the fit, $\Theta = \Theta_x + \Theta_y + \Theta_z$. One can see that the ion density and temperature fluctuations are anti-correlated far from the shock, at $x/\lambda_{si} \approx 37$ –49, while the magnetic field shows no fluctuations in this region. Therefore, the pressure remains approximately constant, no entropy is generated, and the pressure balance is provided by the plasma pressure. This is consistent with entropy-mode-type waves. The entropy waves are most probably generated in the nonlinear structure of the magnetosonic pulse. These waves carry energy away from the shock, thereby providing additional means of dissipation. Note that the rightmost waves in Figure 6.6 at $x/\lambda_{si} \approx 50$ –55) show correlated ion and temperature fluctuations and some magnetic field fluctuations. These waves are part of the shock pulse and are therefore magnetosonic in nature, albeit highly damped, as indicated by the weak magnetic field fluctuations.


The entropy-mode waves contribute moderately to ion and electron heating, as evidenced by the modulations in the phase-space distributions shown in Figure 6.3. Note that downstream of the entropy-wave region, at $x/\lambda_{si} \approx 20$ –30, additional phase-space density modulations can also be observed. Given the persistence of the two-beam structure in this region (see Section 6.2.2), a two-stream ion-beam instability is likely triggered there, generating waves that are apparent in E_x (Figure 6.2, middle panel), as well as in B_y and E_z (not shown). Ions and electrons may scatter off these waves, thereby enhancing particle thermalisation.

Upstream precursor waves

The upstream region is populated by high-frequency, low-amplitude waves, visible in Figure 6.4 for $x \gtrsim 70 \lambda_{si}$. Such waves are a common feature of relativistic magnetised shocks. Figure 6.7 provides a close-up view of the region $x/\lambda_{si} = 80\text{--}90$, showing the transverse components of the electromagnetic field. The top panel displays phase-correlated oscillations in δB_z and δE_y , whereas the bottom panel the oscillations in B_y and E_z with a phase difference of π . These oscillations are identified here with the precursor waves generated through SMI, and correspond to the X-mode waves with δB fluctuations polarised along $\mathbf{B}_0$, and O-mode wave with the magnetic field components, δB_y , perpendicular to $\mathbf{B}_0$. The zero phase difference between δB_z and δE_y in the X-mode wave, and anti-phase-correlation of the O-mode wave components, results in a net Poynting flux in the positive x -direction.

The linear theory of the SMI predicts a generation of mainly X-mode waves (Hoshino 2008), but in simulation run $\theta 40$ both the X-mode and O-mode waves have similar amplitudes. Similar precursor wave characteristics were earlier observed in 2D simulations of ultrarelativistic and mildly-relativistic weakly and moderately magnetised shocks with the background magnetic field in the simulation plane (Iwamoto et al. 2018; Ligorini et al. 2021a). There, the emergence of the O-mode was attributed to the mode conversion from the X-mode within the shock structure, where substantial turbulence in the out-of-plane magnetic field – comparable to B_0 – locally alters the direction of the ambient magnetic field. As a result, X-modes acquire both δB_y and δB_z oscillations. Upon transmission upstream without a change in polarization, X-modes carrying δB_z fluctuations are converted into O-modes. A similar process occurs in run $\theta 40$, which can be attributed to the finite δB_y component within the magnetosonic pulse and to the presence of significant fluctuations in both B_y and B_z along the transverse y -direction (see above).

Figure 6.7: Profile of the transverse electric and magnetic field for the region $x/\lambda_{si} = 80\text{--}90$ in Figure 6.2. The top panel shows the X-mode waves shown as a phase correlation in $\delta E_y = E_y - E_{0y}$ and $\delta B_z = B_z - B_{0z}$ while the bottom panel shows O-mode waves with a phase anti-correlation in E_z and B_y .



Associated with precursor waves, a longitudinal wakefield electric field δE_x is generated through the *parametric decay instability* (PDI) (Kruer 1988; Hoshino 2008). As observed in Ligorini et al. (2021b), strictly perpendicular, mildly relativistic moderately magnetised shocks are inefficient particle accelerators through WFA mechanism, due to the small amplitude precursor waves and the resulting weak amplification of the wakefield. Similarly,

oblique strongly magnetised superluminal shocks produce only a weak wakefield, that again renders the WFA inefficient and leads to only slight electron heating in the region upstream of the magnetosonic pulse (see Section 6.2.2). Note that this result is not expected to change significantly even if the resolution used in this simulation ($\lambda_{se} = 40\Delta$) were increased to $\lambda_{se} = 80\Delta$ as in Ligorini et al. (2021a) and Ligorini et al. (2021b) to fully resolve the SMI dynamics, and consequently, the precursor wave amplitudes. Note also that similar precursor wave characteristics are observed in the simulation of a superluminal shock with the in-plane background magnetic field, simulation $\theta 40\text{-}\varphi_B 0$ (see Appendix A.1).




Figure 6.8: Electron phase-space, $p_{x,e}$ (top) and normalised potential, $\tilde{\phi} = e\phi/(\gamma - 1)m_i c^2$ and pseudo-potential, $\tilde{F} = eF/(\gamma - 1)m_i c^2$ (bottom), at two different time steps and locations: $t\Omega_i = 64.08$, $x/\lambda_{si} = 0\text{--}40$ (left) and $t\Omega_i = 113.08$, $x/\lambda_{si} = 50\text{--}70$ (right), centred around the shock transition. The dashed vertical lines correspond to the position of the strongest peaks in $\tilde{F}$. The right panels are shown at the same time step as in Fig. 6.3.

6.2.4 Particle acceleration and heating in the magnetosonic pulse

According to the mechanism described in Section 4.3.2, upstream electrons can be reflected and briefly trapped and accelerated in a superluminal magnetosonic wave, although they cannot remain trapped for extended periods of time and the resulting acceleration is not very efficient. However, due to the inherent non-stationarity of the shock – expected in




Figure 6.9: Ion (a) and electron (b) spectra at three different time steps, $t\Omega_i = \{50.3, 75.4, 100.5\}$ in the region around the shock transition, $\Delta x_{sh}/\lambda_{si} = -15$ – 1 . The dashed line in panel (b) shows an approximate fit of the spectrum at time $t\Omega_i = 100.5$ to a Maxwellian distribution.

large-amplitude shock waves – the efficiency of trapping and acceleration varies significantly with time and location, and transient episodes of enhanced acceleration efficiency may occur intermittently. Figure 6.8 compares results at two different time steps, showing $p_{x,e}$, $\tilde{\phi}$, and $\tilde{F}$, where the pseudo-potential, $\tilde{F}$, is computed using Equation 4.27 and its peaks are marked with dashed vertical lines. One can see that behind some of the pseudo-potential peaks the phase-space peaks towards large positive $p_{x,e}$ values are developed. The most pronounced energisation occurs at $x/\lambda_{si} \approx 32.2$ in the left panel of Figure 6.8, with the electrons reaching very high momenta in excess of $p_{x,e} \approx 300$, and at $x/\lambda_{si} \approx 58.5$ in the right panel. At these locations $\tilde{F}$ develops deep dips with $\tilde{F} < 0$, which is a condition for the electron reflection. At the location of other phase-space peaks within the magnetosonic pulse the average pseudo-potential dips are only slightly below zero, and in some cases even rise above it. The electron reflection at these locations must therefore have occurred either at positions along the wave crests where local minima of $\tilde{F} < 0$ were produced by the transverse wave modulations, or at an earlier time. Nevertheless, although the reflected electrons can at times gain a significant amount of energy from both the electrostatic potential, ϕ , and the motional electric field, E_{0y} , they can never cross a potential peak towards the upstream and must ultimately lose their energy while being advected downstream. In some instances, electrons may also become temporarily reflected, trapped and further energised at the pseudo-potential dips behind the magnetosonic pulse, in the entropy mode region, at $x/\lambda_{si} \approx 22$ – 32 in the left panel of Figure 6.8 and at $x/\lambda_{si} \approx 50$ – 58 in the right panel (see Section 6.2.3).

Thus the overall interaction of electrons within the superluminal shock results in their heating. In contrast, as shown in the following section, subluminal shocks can effectively trap

electrons and accelerate them to higher energies, leading to the development of a distinct non-thermal electron population.

Ions undergo similar mechanism, being reflected at potential peaks and trapped within potential troughs (see Section 4.4). Corresponding features to those observed in the electron phase-space are visible in both $p_{x,i}$ and $p_{z,i}$ (compare Fig. 6.3). Similar to electrons, ions are not efficiently accelerated, and their interaction with the shock primarily results in their heating.

Temporal evolution of particle spectra in the magnetosonic wave pulse region is shown in Figure 6.9a for ions and in Figure 6.9b for electrons. The ion spectrum is non-Maxwellian, which evidences the very slow ion thermalisation. One can see that enhanced ion energisation can occasionally occur, e.g., at time $t\Omega_i = 50.3$, which produces a small bump at the highest energies (see discussion above). However, as the shock evolves, these high-energy ions gradually lose energy, and the spectrum steeply falls at the high-energy end.

The electron spectrum, in contrast, displays a low-energy component that resembles a Maxwellian. In addition, a distinct high-energy tail is formed, that does not evolve significantly with time. The formation of this tail is consistent with the acceleration driven by the cross-shock potential, in which all upstream particles participate as they traverse the shock.

6.2.5 Discussion on oblique superluminal shocks

The results presented above for an oblique superluminal magnetised shock show that acceleration via trapping in the shock structure is inefficient for both electrons and ions. Particles that are reflected by the wave potential structures do not become fully trapped and ultimately escape, resulting primarily in particle heating with only a minor contribution to non-thermal particle populations.

Additional simulations have been performed to verify the dependence of the results on the 2D simulation setup and dimensionality. For this purpose, a 2D simulation, $\theta 40\text{-}\varphi_B 0$, was conducted with the background magnetic field lying in the simulation plane ($\varphi_B = 0^\circ$), along with a one-dimensional simulation, run $\theta 40\text{-}1\text{D}$. A detailed comparison with simulation $\theta 40$ is presented in Appendices A.1 and A.2.

Comparison with $\theta 40\text{-}\varphi_B 0$ shows that the large-amplitude magnetosonic wave structure is very similar for both magnetic field orientations. In both cases, the shocks are nonstationary and perturbed along the transverse direction, and the particle distributions are largely comparable. This demonstrates that the shock properties are insensitive to the choice of simulation-plane orientation. It is therefore expected that the shock structure would remain similar in fully three-dimensional (3D) simulations, and that runs $\theta 40$ and $\theta 40\text{-}\varphi_B 0$ reliably capture the physics of a superluminal magnetised shock.

In contrast, while the 1D simulation reproduces a shock structure nearly identical to the 2D runs, the phase-space distributions show clear differences, with significantly less efficient electron and ion trapping and acceleration. This highlights the importance of 2D effects for accurately capturing particle energisation.

6.3 Oblique subluminal shock (simulation $\theta 30$)

FIGURE 6.10 shows the structure of a subluminal mildly relativistic shock with an obliquity of $\theta_{Bn} = 30^\circ$. The shock speed is lower than in simulation $\theta 40$, $\beta_{sh} \approx 0.49$, with simulation frame Lorentz factor $\gamma_{sh} \approx 1.15$, and corresponding upstream-frame Lorentz factor of $\gamma_{sh,up} \approx 3.7$. The threshold angle is larger than in the superluminal shock, $\theta_{th} \approx 32.1$, and so the shock in simulation $\theta 30$ is subluminal. The average amplitude of the electrostatic gradient at the shock front, $\delta E_x/E_0 \approx 1$, is 20% larger than in the superluminal shock. The compression ratio downstream of the magnetosonic wave pulse for the magnetic field is slightly higher than in the superluminal case, $r_B \approx 3$, but the density compression ratio remains the same.




Figure 6.10: Maps of the normalised ion density, n_i/n_0 (top panel), the longitudinal electric field, E_x/E_0 (middle panel), and out-of-plane component of the magnetic field, B_z/B_0 , for simulation $\theta 30$ at $t\Omega_i = 314.1$. Compare Figure 6.2.

The overall structure closely resembles that of the superluminal shock, but the amplitudes of the peaks that constitute the magnetosonic wave pulse are, on average, lower. At $t\Omega_i = 314.1$ the amplitudes of the first five peaks are, $\delta B_z/B_0 = [1.19, 1.35, 1.26, 1.14, 1.08]$. Note that in this case, substantial fluctuations in the transverse wave structure also occur, exhibiting properties similar to those observed in the superluminal shock. The average B_z amplitude in the first peak varies within the range $\delta B_z/B_0 \approx 1$ –1.5. It also shows significant variation relative to the amplitude of the second peak, with differences between the two reaching up to $\pm 0.3B_0$. Similar variability is also observed in the other field quantities (see also Section 6.3.1). Overall, the magnetic field amplitude difference between simulation $\theta 40$ and $\theta 30$ can be attributed to the difference in shock speed. Since the shock is subluminal, ions and electrons are now greatly energised. The shock transfers part of its energy to both species, resulting in a considerably slower shock ($\beta_{sh} \approx 0.49$ for simulation $\theta 30$ compared to $\beta_{sh} \approx 0.58$ for simulation $\theta 40$) despite a difference in magnetic field obliquity of only 10° .




Figure 6.11: Phase-space for ions (left) and electrons (right) for simulation 030. The results are shown at $t\Omega_i = 113.08$ to enable direct comparison with Figure 6.3

Precursor waves and corresponding wakefield are observed in the shock upstream, $x/\lambda_{si} \gtrsim 120$, and are linked to the upstream electron heating, as in the superluminal shock. However, precursor waves have smaller amplitudes than in the superluminal shock, $\max(\delta B_y/B_0) \approx 0.02$, which is expected since the transverse background magnetic field, B_{z0} , is lower. Entropy mode waves are also present downstream of the magnetosonic wave, at $x/\lambda_{si} \approx 80\text{--}100$, but have slightly lower amplitude and larger speed than in simulation 040.

Figure 6.11 shows the ion and electron phase-space at time $t\Omega_i = 113.08$, to enable direct comparison with Figure 6.3 for the superluminal case. Unlike the superluminal shock, the phase-space in simulation 030 shows a very large number of high-energy ions and electrons in the region of the magnetosonic wave, with particle momenta already reaching $p_i/m_i c \approx 30$ and $p_e/m_e c \approx 2.5 \cdot 10^3$. As discussed in detail in Section 6.3.3, this acceleration is due to

the mechanism described in Section 4.3. Since the shock is subluminal – and, in particular, very close to the threshold angle – this mechanism is very efficient.

Since the shock transition and magnetosonic wave structure in simulation $\theta 30$ is similar to that in simulation $\theta 40$, the following sections focus on identifying and describing the particle acceleration processes occurring at the shock.

6.3.1 Electric potential in the magnetosonic wave

Figure 6.12 presents the structure of the shock, showing the transversely-averaged profiles of the E_x electric field and the normalised electric potential, $\tilde{\phi} = e\phi/(\gamma_{\text{up}} - 1)m_i c^2$, together with the profile of δB_z magnetic field. The wave structure can be then directly compared with a theoretical sketch shown in Figure 4.1b. Note that the profiles are computed in the shock frame of reference, indicated with the subscript sh and including the normalisation, $(\gamma_{sh} - 1)m_i c^2$, $E_{0,sh}$ and $B_{0,sh}$. Thus the field amplitude differ from results in the simulation frame by a factor of ~ 2.34 . The electrostatic potential is computed under the assumption that the electric field varies only in the x -direction, $\partial\phi/\partial x = -E_x$. This approximation is justified by the fact that the shock remains predominantly planar, even though some transverse perturbations are present. However, computing the potential directly from the exact relation $\mathbf{E} = -\nabla\phi$ does not qualitatively alter either the potential level or its variability. One can see that the strong electrostatic potential peaks present in the superluminal shock are also seen in the subluminal case. The first peak at $t\Omega_i = 314.1$ has an amplitude of $\tilde{\phi} \approx 0.47$. Due to wave nonstationarities and fluctuations in the transverse wave structure, the y -averaged E_x amplitude varies within $\delta E_x/E_0 \approx 0.8\text{--}1.1$, and consequently the elec-




Figure 6.12: Transversely-averaged profile of the normalised longitudinal electric field, E_x/E_0 , the electrostatic potential energy, $e\phi$, normalised to the upstream ion kinetic energy $(\gamma - 1)m_i c^2$, and the compressed magnetic field fluctuations, $\delta B_z/B_0$, at time $t\Omega_i = 314.1$ (compare Figure 6.10). All quantities are computed in the shock frame of reference, indicated by the subscript sh . The horizontal dashed line at zero shows the upstream values, and the upper horizontal dashed line shows the value of the electrostatic potential at the first peak of the wave, $\tilde{\phi} = 0.47$.

trostatic potential fluctuates in the range $\tilde{\phi} \approx 0.3\text{--}0.5$, with an average value of $\tilde{\phi} \approx 0.38$. The typical change in the average amplitude is $\Delta\tilde{\phi} \approx 0.05$ on a timescale of $1/\Omega_i$, although variations as large as $\Delta\tilde{\phi} \approx 0.2$ can occur over the same interval. The second and subsequent peaks exhibit similarly strong variability. Part of the variability in the y -averaged profiles originates from transverse wave modulations on spatial scales of $\sim \lambda_{si}$, which introduce variations along the wave crests and troughs of up to $\Delta\tilde{\phi} \approx 0.06$ in the first peak and $\Delta\tilde{\phi} \approx 0.09\text{--}0.15$ in the second peak, corresponding to approximately 15% and up to 38% of the respective average potentials. A small fraction of the potential fluctuations arises from inhomogeneous upstream wakefield electric fields.

The average value of $\tilde{\phi} \approx 0.38$ indicates that nearly 40% of the upstream ion kinetic energy is expended in shaping the shock potential. For comparison, strong shocks typically use only about 20% of the energy for this purpose. The amplitude of the potential can be attempted to be understood on the ground of theoretical modeling. If the magnetic field is *strictly* perpendicular to the shock normal and electrons are magnetised, the wave potential arises from the $\mathbf{E} \times \mathbf{B}$ drift acting on the incoming electrons. The shock potential at its maximum (which coincides with the location of the maximum magnetic field)) can then be determined by the compression ratio B/B_0 as

$$e\phi_{sh} = \frac{B_{0,sh}^2}{\mu n_{0,sh}} \left(\frac{B_{M,sh}}{B_{0,sh}} - 1 \right) \rightarrow \frac{e\phi_{sh}}{\gamma_i m_i c^2} = \sigma_i \left(\frac{B_{M,sh}}{B_{0,sh}} - 1 \right), \quad (6.1)$$

where the subscript *sh* indicates that the quantities are computed in the shock frame of reference and $B_{M,sh}$ denotes the value of the magnetic field at the pulse peak (Nakazawa and Ohsawa 1997). The larger the magnetic field strength, parameterised by the magnetisation σ_i , the stronger the resulting shock potential. Likewise, a stronger magnetic-field compression, $B_{M,sh}/B_{0,sh}$, leads to a larger potential. Using this formula, and adopting a magnetic-field compression of $B_{M,sh}/B_{0,sh} \approx 1.7$, the predicted potential amplitude for the first magnetosonic peak is $\tilde{\phi} \approx 0.76$, which is higher than the value obtained in the simulation.

If the magnetic field is oblique, then the electrostatic potential is shaped not only by the $\mathbf{E} \times \mathbf{B}$ drift, as in the strictly perpendicular case, but also by the electron motion parallel to the magnetic field, $\phi = \phi_d + \phi_{\parallel}$ (Miyahara et al. 2003). The component of the potential arising from the drift is similar to that in the strictly perpendicular case, but here the contribution comes from the z -component of the magnetic field:

$$e\phi_{sh,d} = \frac{B_{0,z,sh}^2}{\mu n_{0,sh}} \left(\frac{B_{M,z,sh}}{B_{0,z,sh}} - 1 \right). \quad (6.2)$$

The magnetic field at the shock peak is $B_{M,z,sh}/B_{0,z,sh} \approx 2.2$, and therefore one obtains $\tilde{\phi}_{sh,d} = 0.62$, which is slightly smaller than the value obtained using Equation 6.1. The parallel potential component can be obtained by assuming the guiding-centre theory (Eq. 4.22), in which case the final expression is

$$e\phi_{sh,\parallel} = \frac{B_{0,x,sh}^2}{\mu_0 n_{0,sh} B_{0,sh}} \int_{\infty}^{x_M} \frac{B_{sh}}{B_{z,sh}} \frac{dB_{z,sh}}{dx} dx, \quad (6.3)$$


where the integral goes from the far upstream $x \rightarrow \infty$ to the shock x_M . Note that $\phi_{sh,\parallel} > 0$. The more perpendicular the shock, the smaller the contribution from the parallel component. Since the shock is a large-amplitude wave pulse, the magnetic field at the shock is predominantly out-of-plane, with $B_{M,sh}/B_{0,z,sh} \approx 1.1 \sim 1$. In this case, the integral in Equation 6.3 becomes trivial, and the total cross-shock potential is given by

$$e\phi_{sh} = e\phi_{sh,d} + e\phi_{sh,\parallel} = \left(1 + \frac{B_{0,x,sh}^2}{B_{0,z,sh}B_{0,sh}}\right) \frac{B_{0,z,sh}^2}{\mu_0 n_{0,sh}} \left(\frac{B_{M,z,sh}}{B_{0,z,sh}} - 1\right), \quad (6.4)$$

which only depends on the magnetic field compression at the shock and the upstream magnetic-field components. Using the simulation values in Equation 6.4, the cross-shock potential becomes $\phi \approx 1.1$, which is, however, significantly larger than the value obtained in the simulation.

A caveat of the theoretical presented above analysis is that it neglects ion dynamics. Typically, ions in shocks are considered unmagnetised, $\sigma_i \ll 1$. However, this approximation is not valid in simulation $\theta 30$, for which the ion magnetisation is $\sigma_i \approx 1$. This is evident from Figure 6.13, which shows the ion phase-space for x - and y -components at the location of the first pulse. The magnetic field is sufficiently strong that the entire upstream ion beam is deflected and perform a cycloid-like motion along the x -axis (Nakazawa and Ohsawa 1997).

Figure 6.13: Ion phase-space for the $p_{x,i}$ and $p_{y,i}$ momentum components around the shock front at $x \approx 154.8\lambda_{si}$ (marked with the vertical dashed line) at time $t\Omega_i = 314.1$. The upstream ion beam enters from the right with $u_x/c = -2$ and $u_y/c = 0$, and is modulated as it encounters the shock. Since the number of ions with positive p_y exceeds those with negative p_y at the shock, the resulting net ion current is positive.



The total current is determined by the $\mathbf{E} \times \mathbf{B}$ drift. If ions are magnetised, their drift becomes comparable to that of electrons. Because the ion current opposes the electron current, the net current is reduced and the resulting electric potential decreases accordingly.

Although Nakazawa and Ohsawa 1997 derive the ion contribution for a magnetosonic wave in a plasma with a strictly perpendicular magnetic field – unlike the oblique configuration in simulation $\theta 30$ – their result still provides a useful estimate for comparison with the potential of the first magnetosonic pulse. Expressing the potential from Equation 30 in Nakazawa and Ohsawa 1997 in the shock frame of reference, one obtains:

$$e\phi_{sh} = \frac{B_{0,sh}^2}{\mu_0 n_{0,sh}} \left(\frac{B_{M,sh}}{B_{0,sh}} - 1 \right) - \frac{2A\gamma_{up,sh} m_i v_{up,sh}^2}{B_{M,sh}/B_{0,sh} + 1}, \quad (6.5)$$

where v_{up} is the upstream speed and A is a parameter related to the cycloid motion. If the motion is a perfect cycloid, A takes the value of $2/3$ and is larger for a realistic ion cycloid motion. Using the same simulation parameters, the potential evaluates to $\tilde{\phi} = 0.13$. This indicates that the ion current can reduce the drift-related component of the potential by more than 80%. One can use this estimate, derived for a strictly perpendicular shock, to correct the potential in the oblique shock geometry described by Equation 6.4. Combining Equations 6.4 and 6.5 then yields:

$$e\phi_{sh} = \left(1 + \frac{B_{0,x,sh}^2}{B_{0,z,sh} B_{0,sh}} \right) \frac{B_{0,z,sh}^2}{\mu_0 n_{0,sh}} \left(\frac{B_{M,z,sh}}{B_{0,z,sh}} - 1 \right) - \frac{2A\gamma_{up,sh} m_i v_{up,sh}^2}{B_{M,sh}/B_{0,sh} + 1}. \quad (6.6)$$

Assuming $A = 2/3$, this now gives $\tilde{\phi} \approx 0.47$, which is slightly larger than the average value of $\tilde{\phi}$. To achieve strict agreement with $\tilde{\phi} \approx 0.38$, it is necessary to assume a value of $A = 0.77$. While not strictly rigorous, Equation 6.6 provides a qualitative understanding of the expected potential level at strongly magnetised shocks and of the main contributions to this potential.

6.3.2 Late-stage evolution of the particle phase-space distributions

Figure 6.14 shows phase-space for ions and electrons at the end of simulation at time $t\Omega_i = 314.1$. Overplotted are the magnetic field fluctuations, $\delta B_z/B_{0z}$, scaled up for better visualisation. By comparing with Figure 6.11, one can see that particle acceleration continues at the shock, and particles reach even larger momenta. For instance, the maximum value of $p_{z,e}$ at $t\Omega_i = 113.08$ is $p_{z,e}/m_e c \approx 2.5 \times 10^3$ while at $t\Omega_i = 314.1$, the maximum value is $p_{z,e}/m_e c \approx 5 \times 10^3$. Accelerated electrons appear at the locations of the magnetosonic wave peaks, with the strongest acceleration occurring near the wave front and the maximum electron energies decreasing at subsequent peaks. The electron phase-space in the shock closely resembles the phase-space structure observed in simulations of nonrelativistic oblique magnetosonic waves by Bessho and Ohsawa 2002, showing a strong peak at the position of the magnetosonic wave. The main difference relative to simulation $\theta 30$ is that the mildly relativistic wave pulse contains several wave peaks, and electron acceleration occurs at each of them.

The other important difference with earlier simulations by Bessho & Ohsawa is the occurrence of ion acceleration in the shock wave. While the highest energy electrons are found near the positions of the magnetic field peaks, at which the electric potential also peaks, the




Figure 6.14: Phase-space for the p_x (top) and p_z (bottom) momentum components for ions (left) and electrons (right) for simulation $\theta 30$ at $t\Omega_i = 314.1$ (compare Figure 6.11). Overplotted is the y -averaged profile of the compressed magnetic field fluctuations, $\delta B_z/B_{0z}$, scaled for visualisation purposes.

energetic ions concentrate around the potential minima, i.e. the potential troughs. In addition, the ion phase-space shows a feature that is not present at an earlier shock stage shown in Figure 6.11. Ahead of the shock, at $x \approx 156\lambda_{si}$, diffuse peaks appear in both $p_{x,i}$ and $p_{z,i}$, reaching even higher momenta than the ions located at the shock front. This structure is a signature of *shock surfing acceleration* (SSA), and is discussed in detail in Section 6.3.6.

Another notable feature in the ion phase-space is an elongated "blob" in the $p_{x,i}$ and $p_{z,i}$ distributions upstream of the magnetosonic pulse, located at $x/\lambda_{si} \approx 55$ – 60 at $t\Omega_i = 113.08$ (Figure 6.11) and at $x/\lambda_{si} \approx 120$ – 125 at $t\Omega_i = 314$ (Figure 6.14). This structure forms very early within the first peak of the newly developing magnetosonic pulse. It originates from the ion interaction within the pulse, but is also affected by the initial plasma beam reflection at the simulation boundary. During the subsequent shock evolution, the blob drifts upstream and begins to form a distinct ion population in phase space. At the same time freshly incoming particles become trapped and accelerated within the first peak through the physical mechanism described in Section 4.3. Some of these accelerated ions are then reflected upstream, undergo SSA (see above), and merge with the existing blob in phase space.

By $t\Omega_i \approx 50$ (not shown), the blob has fully separated from the first peak, yet it continues to be replenished by ions accelerated at the leading pulse. As a result, the blob evolves in phase space, extending to higher momenta over time, as seen when comparing Figure 6.11 and Figure 6.14. Additional broadening arises from interactions with the very weak fields upstream of the shock. Because the shock is subluminal and the ambient magnetic-field angle lies very close to the threshold angle, the relative speed between the blob and the shock is very small. Consequently, these particles during the early evolution remain close enough to the shock to continue feeling both the electrostatic potential and the magnetic-field gradient. At later stages, however, the blob gradually moves farther upstream, becomes decoupled from the shock, and SSA-accelerated ions begin to form their own distinct upstream phase-space features, as seen at $t\Omega_i = 314$, Figure 6.14. Since this feature is, at least to some extent, influenced by the boundary conditions, it is excluded from the analysis below, along with the corresponding structure in the electron phase-space. Importantly, the presence of these populations does not affect the physical processes occurring within the magnetosonic wave or in the downstream region.

6.3.3 Electron trapping and acceleration in the magnetosonic shock

Electron acceleration in the magnetosonic wave, evidenced in phase-space distributions, proceeds through the acceleration mechanism described in Section 4.3. In contrast to the superluminal shock, electrons that are reflected off the wave potential minimum can temporarily move towards the potential peak and ahead of the peak, where they decelerate before moving downstream again. This motion repeats continuously, and the electrons become *trapped* as they cannot escape easily. While electrons are trapped within the wave pulse, the combined action of the motional electric field and $\mathbf{E} \times \mathbf{B}$ drift accelerates them to ultra-high energies.

Figure 6.15 shows transversely-averaged profiles of the key quantities that govern electron trapping and acceleration: the normalised electrostatic potential energy, $\tilde{\phi} = e\phi/(\gamma - 1)m_i c^2$, and normalised pseudo-potential, $\tilde{F}$, in panel (a), E_y component of the electric field, corresponding to the motional electric field E_0 (panel b), and three components of the $\mathbf{E} \times \mathbf{B}$ drift (panel c). All quantities are shown in the shock frame of reference. Vertical dashed lines show the positions of the first wave pulse peak at $x = x_M$ and the trough at $x = x_m$. Similar to the superluminal shock (compare Figure 6.8 in Section 6.2.4), the pseudo-potential, $\tilde{F}$, exhibits oscillations of larger amplitude than the corresponding oscillations of the electrostatic potential in the wave pulse structure. This difference in Figure 6.15a is most pronounced for the first pulse, with its maximum at $x = x_M$, and decreases for subsequent oscillations. However, as noted above, the average profiles fluctuate considerably due to wave nonstationarity and transverse modulations. The E_y electric field oscillates around its constant upstream value E_0 , as expected from the Rankine-Hugoniot conditions (Section 3.2.1). Upstream fluctuations are small, whereas in the shock and downstream regions the fluctuations are larger.

The x -component of the $\mathbf{E} \times \mathbf{B}$ drift remains approximately constant along the shock transition, although small increases are visible at locations of the peaks due to the com-




Figure 6.15: Vicinity of the shock in the shock frame of reference. Profiles of normalised electrostatic potential, $\tilde{\phi}$, and pseudo-potential, $\tilde{F}$ (a). The vertical dashed lines show the position of the first peak at $x_M \approx 154.7\lambda_{si}$, where the pseudo-potential has a maximum, $\tilde{F}(x_M)$, and the position of the trough at the end of the first pulse at $x_m \approx 153\lambda_{si}$, where $\tilde{F}$ is minimum, $\tilde{F}(x_m)$. E_y component of the electrostatic field (b). The horizontal dashed line shows the upstream value, E_{y0} . Components of the $\mathbf{E} \times \mathbf{B}$ drift, $\mathbf{v}_{E \times B} = (\mathbf{E} \times \mathbf{B})/|\mathbf{B}|^2$, in the speed of light units (c).

pressed magnetic field there,

$$v_{E \times B, x} = (E_y B_z - E_z B_y)/B^2 \approx E_{0y} B_z/B^2 > E_{0y} B_{0z}/B_0^2. \quad (6.7)$$

The y -component shows the strongest variation due to the strong E_x gradient, becoming negative in front of each pulse and positive behind it, most prominently for the first pulse. The z -component remains positive throughout. It decreases at the shock front and maintains this approximate value downstream.

Upstream electrons first interact with the forward pulse in the magnetosonic wave at $x/\lambda_{si} \approx 153$ – 156) and are accelerated by the electrostatic potential. Being strongly magnetised, electrons gyrate along magnetic field lines as the magnetic field suddenly increases. Most electrons pass the first wave peak at $x = x_M$, but a small fraction are reflected at the end of the pulse at $x = x_m$, back towards the potential peak. Although the x -component of the $\mathbf{E} \times \mathbf{B}$ drift – which is negative – opposes this motion, the potential is strong enough to prevent the electrons from being advected towards subsequent pulses. This condition is

satisfied when the pseudo-potential is negative, $F < 0$. Note that Figure 6.15 a shows a positive $\tilde{F}$, which is an average. The pseudo-potential can locally achieve negative values along the y -direction due to electromagnetic field perturbations. Moreover, trapped electrons modify F via charge fluctuations that alter the electric field, an effect discussed further in Section 6.3.3.

The energy gain obtained by an electron during the first peak crossing towards the end of the pulse and at subsequent reflection (the path labeled $A \rightarrow B \rightarrow C$ in Figure 4.3) is almost zero, but positive. Following Equation 4.37, the energy gain by the electric potential ϕ – which is always positive – is positive, as ϕ is larger at the end of the pulse than in the far upstream. The contribution from the parallel velocity is negligible due the magnetic field compression, which keeps $B_y \sim 0$, and the gain from the $\mathbf{E} \times \mathbf{B}$ drift is also negligible because the y -component changes sign, as one can see in Figure 6.15c.

Once the electrons are reflected, they go back towards the potential peak, following the path $C \rightarrow D$. The electric potential transfers potential energy to the reflected electrons as kinetic energy, accelerating the electrons. At the same time, the negative x -component of the $\mathbf{E} \times \mathbf{B}$ drift opposes this motion, and the electron takes a longer time to reach the peak. This combined effect of the drift and the potential is crucial for accelerating electrons to very high energies, as it prolongs the interaction time with the motional electric field while the electron moves from the trough to the peak (Fig. 6.15b). The force due to E_{y0} is positive, $F_{L,y} = -eE_{0y} > 0$, causing the electron to move in the positive y -direction, while this motion is facilitated by the positive y -component of the drift.

Once the reflected and accelerated electrons move ahead of the peak, $x/\lambda_{si} \gtrsim 154.8$, the y -component of their velocity now acts against the motional electric field, $v_{E \times B, y} < 0$, so the electron either loses energy or its energy gain is smaller, corresponding to the path $D \rightarrow E$. Since most of the electron's kinetic energy is in parallel velocity, the x -component of its velocity is small, and the potential once again reflects it back towards the peak, albeit this time moving from the pulse front towards the peak. As shown later in Section 6.3.5, electrons rarely move ahead of the peak and remain within the pulse width ($x/\lambda_{si} \approx 153.5$ –156). These repeated reflections trap the electrons within the pulse, and once trapped, they hardly escape, as discussed below.

Electrons are also trapped in subsequent pulses ($x/\lambda_{si} \approx 147.6, 149.7, 152.1$) as evidenced by electron phase-space (Fig. 6.14). Among the electrons trapped across all pulses, most accumulate in the first pulse due to its strongest electric potential difference, $\phi - \phi_0$. Electrons that escape the first pulse are slightly decelerated by its potential and the $\mathbf{E} \times \mathbf{B}$ drift, and may then be trapped in subsequent pulses, despite a weaker potential difference within them. Moreover, although the motional electric field and x -component of the drift remain constant, the y -component fluctuation at the pulses becomes weaker and affects the efficiency of the mechanism. This, combined with the small amplitude of the subsequent pulses, reduces the efficiency of the trapping mechanism and, consequently, the resulting electron acceleration.




Figure 6.16: Evolution of the pseudo-potential $\tilde{F}(x, t)$ over $t\Omega_i = 37.7\text{--}314.1$ at two distinct locations of the first pulse: the peak, $\tilde{F}(x_M, t)$ (blue), and the trough, $\tilde{F}(x_m, t)$ (orange). The positions x_M and x_m are defined in Fig. 6.15 at the latest time $t\Omega_i = 314.1$ for reference. The thick dashed line marks $F_\infty = 0$, and the dotted lines show the average values of $F(x_M)$ and $F(x_m)$.

Variability of F

As explained in Section 4.3, due to the high electron magnetisation, $\sigma_e \gg 1$, electrons closely follow magnetic field lines while the drift and motional electric fields act upon them. Therefore, the pseudo-potential F is the key quantity in understanding the mechanism. The overall structure of F resembles that of the electric potential, including the relative positions of minima and maxima, as shown in Figure 6.15a. Note that $F(x_M) > \phi(x_M)$, indicating that the trapping mechanism is more efficient in the direction of the magnetic field than when considered solely in the x -direction. Consequently, magnetised electrons feel the pseudo-potential rather than the electric potential.

When an electron interacts with the pulse, it may either extract energy from the potential or transfer energy to it. Due to the non-stationarity of the pulse, the pseudo-potential may occasionally fall below zero, $F(x_m) < 0$, allowing a small fraction of electrons to be reflected. These temporal variations in the pseudo-potential are shown in Figure 6.16, in which the blue line shows the value of $\tilde{F}$ at the peak, $\tilde{F}(x_M)$, while the orange line shows the value at the trough, $\tilde{F}(x_m)$. The dotted horizontal lines mark the average values, which remain positive at both x_M and x_m , while the instantaneous values oscillate significantly around these averages. Whenever $\tilde{F}(x_m) < 0$ (indicated by the dashed line at zero), electrons are reflected. Although $\tilde{F}(x_m)$ rarely appears negative in the figure, the plotted value represents an average over the transverse y -direction. As already discussed, the shock structure fluctuates along the transverse direction due to wave nonstationarity and turbulence, and thus both the potential and the pseudo-potential may fluctuate considerably. Therefore, although the averaged $\tilde{F}(x_m)$ remains positive, the local value can become negative, allowing reflections to occur more frequently than the averaged profile suggests.

Let us assume an upstream electron that arrives at the pulse and is reflected at the trough at $x = x_m$ and $t = t_r$. At this location, F is negative to allow the electron reflection. Assume also that spatial variations of the pseudo-potential, $\partial F / \partial x$, are significantly slower than the

time scale of the perturbations in the pseudo-potential, $\partial F/\partial t$. In this case, the second term on the right-hand side of Equation 4.31 can be neglected:

$$\frac{d}{dt} (\gamma m_e c^2 - eF) = \frac{d\varepsilon}{dt} = -e \frac{\partial F}{\partial t}, \quad (6.8)$$

where ε is the parallel energy of the electron. After the electron becomes reflected at $t = t_r$ and $F = F(x_m, t_r)$, F starts to increase, so that $-e \partial F/\partial t < 0$ and the parallel energy decreases, $d\varepsilon/dt < 0$. Since the electron loses parallel energy, it becomes trapped. If the electron returns to x_m again, even if $F(x_m) > 0$, the electron has a reduced parallel energy and can be reflected again. Once the electron becomes deeply trapped, its kinetic energy subsequently grows and it is accelerated to ultra-relativistic energies.

Once $F(x_m)$ becomes positive, electrons are no longer reflected and are simply advected downstream. These advected electrons lose part of their energy, which is transferred to the potential, and so $F(x_m)$ is restored to lower values and can drop below zero again. This cycle then repeats, as observed in Figure 6.16, with fast variations on a characteristic timescale of $\sim 2.5\Omega_i^{-1}$.

6.3.4 Electron spectra

Figure 6.17a shows temporal evolution of the electron spectrum at the shock region, taken in the range $\Delta x_{sh} = [-24, 2]\lambda_{si}$ with respect to the instantaneous shock position, x_{sh} (compare, e.g., Figures 6.10 and 6.14). Figure 6.17b presents the maximum electron kinetic energy, computed as $(\gamma_{\max} - 1)(dN_e/d\gamma_{\max})/N = 1 \times 10^{-7}$. The broad peak at $\gamma \approx 40$ corresponds mostly to electrons that are not trapped in the shock, are advected downstream, and partially thermalise. The spectrum exhibits a high-energy tail that evolves, as the shock develops, towards a power-law distribution $dN/dE \propto E^{-p}$ with spectral index $p \approx 2$. At $t\Omega_i = 314.1$, this tail extends from $\gamma \approx 150$ up to an exponential cutoff at $\gamma \approx 5500$. Upstream electrons are continuously trapped at the shock, increasing the number of high-energy electrons (raising the tail), while electrons already trapped, continue to gain energy (extending the tail). As long as electrons remain trapped, the exponential cutoff of the spectrum shifts to higher energies. Figure 6.17b shows that the maximum kinetic energy increases approximately linearly over time (black line), reaching $\max(\gamma - 1) \approx 5.5 \times 10^3$ at $t\Omega_i \approx 314.1$, although significant variations around this trend are present. Note that during the early stage of the simulation, $t\Omega_i \lesssim 100$, the shock is still forming, so the number of trapped electrons is small and the corresponding maximum energies are low. As trapping becomes more efficient, the maximum energy increases rapidly before transitioning to the approximately steady linear growth indicated by the black line.

Figure 6.18 shows the time evolution of the number of trapped electrons. The values are computed from electron spectra by counting all high-energy electrons, i.e. those with $\gamma - 1 > 150$. The number of trapped electrons grows continuously and nearly linearly until $t\Omega_i \approx 200$, after which the growth slows. By the end of the simulation, high-energy electrons constitute approximately 10% of the upstream population.

This behaviour resembles that observed in simulations by Bessho and Ohsawa 2002, where a constant increase in the number of trapped electrons was reported. Superimposed




Figure 6.17: Panel a): Evolution of the electron spectrum around the shock $\Delta x_{sh} = [-24, 2]\lambda_{si}$ in simulation $\theta 30$ (compare, e.g., Figures 6.10 and 6.14). The spectrum is shown at six different times $t\Omega_i = \{62.8, 113.1, 163.3, 213.6, 263.8, 314.1\}$. The black line shows a power-law dependence, $dN/dE \propto E^{-p}$ with index $p = 2$, for comparison with the shape of the high-energy spectral component. Panel b): Maximum normalised electron kinetic energy versus time. The vertical dashed red lines, including the plot limits, mark the times corresponding to the spectra in panel a). The black line shows a linear fit, $\max(\gamma - 1) \propto 6.7t\Omega_i$.

on this overall increase, the number of trapped electrons displays small fluctuations over time with a characteristic timescale of $\sim 2.5\Omega_i^{-1}$. These quick variations are directly correlated with the similar variations in the pseudo-potential F shown in Figure 6.16. However, unlike in simulation $\theta 30$, which has a constant influx of fresh upstream plasma and thus a continuously increasing system energy, the model used by Bessho and Ohsawa 2002 utilises a *bounded plasma model*. In their case, both the number of particles and the total energy remain fixed. Consequently, although both simulations initially show a linear increase in trapped electrons, the growth in simulations by Bessho and Ohsawa 2002 eventually stalls, whereas in simulation $\theta 30$ it merely slows. This slowing can be attributed to the finite free energy available: as the magnetosonic shock evolves, its amplitude gradually decreases while transferring energy to the trapped electrons, reducing its ability to trap additional particles. At late times in Bessho and Ohsawa 2002, the shock becomes too weak to trap electrons at all.

Although the accelerated electron spectra in Figure 6.17 show a power-law distribution with index $p \approx 2$, this slope is not consistent with the theoretically expected value of $p = 1$ (see Section 4.3.2). This is because the spectrum at the magnetosonic wave pulse is formed from the combined contribution of several potential peaks, each of which traps its own population of electrons. As an example, Figure 6.19a shows the electron spectra




Figure 6.18: Number of high-energy electrons with $\gamma - 1 \geq 150$ over time. Compare Figure 6.17.




Figure 6.19: Panel a): Electron spectra at $t\Omega_i = 314.1$ in the regions of three consecutive phase-space peaks shown in panel b). Panel b): Zoom-in of the electron phase-space for the $p_{z,e}$ momentum component around the first three pulses of the magnetosonic wave. Compare Figure 6.14. Only regions with high phase-space density are shown to clearly highlight the locations of the phase-space pulses.

at time $t\Omega_i = 314.1$ around three consecutive regions corresponding to the first three potential peaks, as shown in Figure 6.19b. The bulk of the spectrum for each peak is formed around $\gamma \approx 40$ and corresponds to the incoming upstream beam that has been heated by the cross-shock potential, with the peaks becoming broader the farther they are from the shock front. Spectra of electrons accelerated at each peak are power laws with index $p \approx 1$ and the exponential cutoff, as expected. The amplitudes of the subsequent potential peaks decrease the farther they are located downstream (see Figure 6.15) and therefore the maximum electron kinetic energy is lower, as expected from Equation 4.45. This results in a narrower energy range in the power-law portions with increasing distance from the shock front. The combined contribution of all peaks that constitute the magnetosonic shock produces the spectrum with index $p \approx 2$ observed in Figure 6.17.




Figure 6.20: Trajectories (top left and right and bottom left panels) and energy evolution (bottom right panel) of three electrons accelerated near the first peak of the magnetosonic shock, x_{sh} . The particle trajectories are overlaid on the map of the normalised longitudinal electric field E_x/E_0 , and are color-coded by particle kinetic energy according to the color bar in the bottom-right panel of the figure. The field remains approximately constant throughout the entire period, with only minor variations in the shock front position. Particle motions start at $t\Omega_i \approx 38$, marked with a cross, and finish at $t\Omega_i \approx 239$, marked with a filled circle. Note that particle B ends its trajectory outside the plotted region and thus the circle is not visible. Black arrowheads at the top and bottom of the panels indicate where particle cross the periodic boundaries and their direction of motion. The vertical black lines in the energy-evolution plot correspond to $t\Omega_i = 72, 129$ and 183 and are marked as empty triangles indexed with (1, 2, 3) in the other three panels.

6.3.5 Electron tracing

Figure 6.20 show trajectories and energy evolution of three representative electrons that have been trapped by the electric potential at the magnetosonic shock front.

Trajectories of all three electrons start in the downstream region after being trapped, moving slowly towards the shock front while being accelerated in the positive y -direction.




Figure 6.21: Electron motion in $u_{\parallel}$ - $u_{\perp}$ momentum space with colour-coded time for the same three particles as in Figure 6.20, with the same markers (cross, circle and triangles). The four-velocity is computed with respect to the local magnetic field and in the shock frame of reference.

Once they cross the potential peak, they lose some energy and quickly return downstream. The longer the electrons remain downstream, the more energy they gain, while local perturbations in E_x can either hinder or enhance this energisation. This cycle repeats until at some point, the electrons begin losing energy. For particle B, this moment is marked by the vertical line (2) in the bottom-right panel. At this stage, the electrons are strongly scattered and the perturbation of their trajectories is sufficient to deflect them into the negative y -direction. Electrons still follow magnetic field lines, but now in the anti-parallel direction relative to their initial motion. Consequently, the motional electric field E_y decelerates them significantly, as they continue to oscillate within the potential.

After the initial energisation phase, particle A is scattered on the downstream side of the pulse and loses some energy while moving towards the upstream region, where it becomes detrapped and is unlikely to return to the shock. Particle B loses a comparable amount of energy after scattering. After $t\Omega_i \approx 183$, it begins surfing in the upstream region, steadily losing nearly all of its energy due to motion against E_y . Eventually, the electron crosses the shock front and later ends up far downstream (not shown).

Particle C gains energy at the beginning in a similar way to the other two particles. After scattering at $t\Omega_i \approx 129$ the particle loses a significant amount of energy, reaching a minimum of $\gamma - 1 \approx 300$ at $t\Omega_i \approx 183$ – vertical line (3), where it starts gaining energy again and moving in the positive y -direction. However, at $t\Omega_i \approx 200$ the particle appears to be stuck in the downstream, $(x - x_{sh}) \approx -1.5\lambda_{si}$, moving slowly in the positive y -direction. The electron becomes trapped by a perturbation in B_y while it is accelerated by E_y , slowly gaining energy. It is finally detrapped out of this transient perturbation and starts gaining energy by the usual mechanism, evidenced by the zig-zag pattern in $\gamma - 1$ at the latest time.

Figure 6.21 shows the motion of the same particles in $u_{\parallel}$ - $u_{\perp}$ momentum space. Initially,

all particles gain energy from the motional electric field (triangle 1). Since E_{y0} is along the y -direction, the energy gain occurs in momentum components both parallel and perpendicular to the magnetic field. Around $t\Omega_i \approx 129$ (triangle 2), all three particles undergo strong scattering, seen as arc-like trajectories in the $u_{\parallel}$ - $u_{\perp}$ plane at nearly constant total four-velocity. This scattering changes the direction of motion along the y -axis (see Fig. 6.20) and leads to a decrease in both parallel and perpendicular components of the four-velocity. At $t\Omega_i \approx 183$ (triangle 3), the energy of particle A decreases slowly, whereas particle B loses energy much more rapidly. Particle C has the lowest energy at this time, but its energy then increases quickly until it undergoes trapping again, similar to the initial phase of the tracing period.




Figure 6.22: Evolution of different quantities along the trajectory of particle A (see Figure 6.20) for two different intervals, $t\Omega_i = 63$ –101 (left) and $t\Omega_i = 188$ –239 (right). All quantities are given in the shock frame of reference. From top to bottom: electron kinetic energy, $\gamma_e - 1$, normalised longitudinal electrostatic field at particle position, E_x/E_0 , x -component of the electron velocity, v_x/c , and y -component of the four-velocity, u_y/c . Vertical dashed lines and letters on the top panel mark points along the electron trajectory corresponding to the mechanism described in Section 4.3.2, using the same notation as in Figure 4.3.

Figure 6.22 shows the evolution of $\gamma_e - 1$, E_x , v_x and u_y for particle A over two time intervals: $t\Omega_i = 63$ –101 (left) and $t\Omega_i = 188$ –239 (right). The first interval corresponds to the portion of the electron trajectory in Figure 6.20 where the electron gains energy, while in the second interval it loses energy. The letters indicated on the top panel with $\gamma_e - 1$, follow the electron's motion according to the mechanism described in Section 4.3.2, following the

same notation as in Figure 4.3. Sudden changes in E_x correspond directly to extrema in v_x/c . The y -component of the four-velocity shows periodic motion associated with the helical motion along the magnetic field, with gyrofrequency $\Omega_e = eB/(\gamma m_e)$ that is small because, although the magnetic field is strong at the magnetosonic wave, the electron is highly relativistic. The overall energy gain correlates with u_y : for positive u_y , the electron gains energy, while for negative u_y , it loses energy.

The electron whose trajectory is shown in Figure 6.22 is trapped, so part of its motion from A to B is not displayed. The electron is reflected at point C , corresponding to a change of sign in v_x/c , and starts moving towards D . Its energy and v_x gradually increase as it is drawn towards the potential peak. Upon reaching D , E_x rapidly changes sign and v_x decreases as the electron moves towards E , where it is reflected again. During $D \rightarrow E$, the electron continues to gain energy because the motional electric field overcomes the energy loss from $\Delta\phi = \phi_E - \phi_D < 0$, which is small because $x_D \approx x_E$. The electron then moves back towards the potential peak located at B_2 , on the downstream side; E_x changes sign again, decelerating the electron through the change in the potential, but since v_{dx} is negative, the electron remains longer downstream of the peak and the net energy gain temporarily stalls, $\gamma_{C2} - \gamma_{B2} \sim 0$. Finally, the electron is reflected at C_2 and returns towards the peak at D_2 , gaining energy similarly to the $C \rightarrow D$ segment.

At a later time (right panels), the electron begins to lose energy. It spends a longer time on the upstream-side, where E_x is positive and v_x decreases slower, e.g. during $t\Omega_i \approx 200$ –216. When the particle has a large negative v_y , the Lorentz force driving it towards the shock is reduced:

$$F_{L,x} = -e(E_x + v_y B_z - v_z B_y). \quad (6.9)$$


On the upstream side of the peak, the first and third terms on the RHS are positive, while the second term is negative for $v_y < 0$. However, due to E_y , v_y increases ($|v_y|$ decreases), causing the particle to cross the shock front towards downstream slightly faster. Note the interval $t\Omega_i \approx 230$ –239, when the electron starts gaining energy even though v_x remains positive and $E_x \sim 0$. Here, the electron moves far enough upstream that the effects of E_x and the compressed magnetic field B_z are small, while E_y continues to act on it, increasing its energy.

6.3.6 Ion trapping and acceleration

Figure 6.23 compare the ion and electron spectra at $t\Omega_i = 314.1$ within the region $x/\lambda_{si} = 130$ –158. The x -axis differs from that in Figure 6.17, as it represents the particle kinetic energy $E_k = (\gamma - 1)mc^2$. Both species exhibit two dominant populations: a low-energy component ($E_k \approx 10$), formed by particles that have crossed the shock and are subsequently advected downstream; and a trapped high-energy component ($E_k \gtrsim 50$) that forms a power-law distribution. In addition, the ion spectrum shows a distinct bump at very high energies ($E_{k,i} \approx 10^3$), which is absent in the electron spectrum. This feature is shown below to originate from ions undergoing SSA.

The low-energy ion population does not follow a Maxwellian distribution, as the downstream plasma is not yet thermalised. Instead, the upstream beam visible at $E_k \approx 31$ is mod-

Figure 6.23: Ion and electron energy spectra at $t\Omega_i = 314.11$ for the region $x/\lambda_{si} = 130\text{--}158$. This region excludes the upstream "blob" observed in ion phase-space (Fig. 6.14). Note that the x -axis shows the kinetic energy $(\gamma - 1)mc^2$, in contrast to Figure 6.17.



erately heated which produces a broad low-energy spectrum. The high-energy populations of both electrons and ions form approximate power-law distributions. Electrons develop a power-law tail with index $p \approx 2$ and a cutoff at $E_{k,e} \approx 250$, containing approximately 10% of the upstream electron population, (compare Figure 6.17). In contrast, the high-energy ion population exhibits a steeper spectrum with a power-law index $p \approx 2.5$ and a cutoff extending to higher energies, $E_{k,i} \sim 500$. This ion population is much less populated, comprising roughly two orders of magnitude fewer particles than the electrons, i.e., $\sim 0.1\%$ of the upstream ions..

Ions, like electrons, can be trapped by the electrostatic potential; however, ions are trapped between potential peaks, i.e., within potential troughs. This is evident in Figure 6.14, in which the first two populations of trapped electrons are located near the potential peaks at $x \approx 154.8\lambda_{si}$ and $x \approx 152\lambda_{si}$, whereas the first trapped and accelerated ions are found between these positions, at $x \approx 153\lambda_{si}$. The potential troughs are, however, shallower than the peaks, so the available potential energy for ions is smaller. Consequently, the maximum kinetic energy attained by ions is slightly lower than that reached by electrons. One can see that ions reach kinetic energies of approximately $E_{k,i} \approx 750$, just before the very-high-energy bump at $E_{k,i} \approx 10^3$, while the electron energy reached in a similar process is $E_{k,e} \approx 1200$.

The condition for ion trapping is analogous to that for electrons, but instead of comparing the upstream pseudo-potential $F_\infty = 0$ with its minimum $F(x_m)$, one compares the values at two consecutive peaks, $F(x_M)$ and $F(x_{M,2})$ (see Section 4.4). If $F(x_{M,2}) > F(x_M)$, the ion becomes trapped. In practice, however, ions may be perturbed while crossing the first shock pulse and undergo scattering, making the trapping condition less strict and harder to satisfy. Due to such scattering, ions may approach the second peak with a lower parallel velocity than needed for trapping and be advected downstream instead. Furthermore, because ions are less magnetised than electrons ($\sigma_i \approx (m_e/m_i)\sigma_e$), they do not follow magnetic field lines as

closely. In addition to the more stringent trapping condition, ions can readily escape either downstream or upstream. Since their gyroradius is comparable to the separation between consecutive peaks – of order $\sim 2.5\lambda_{si}$ – they may become detrapped if their perpendicular velocity increases sufficiently.

Figure 6.24 shows three representative ions that are initially trapped at the first potential trough but eventually escape downstream. The ions are accelerated by the motional electric field, moving in the negative $-y$ -direction while gaining energy. Note that all three ions are accelerated near the second peak ($x - x_{sh} \approx -2.5$). As in the case of electrons, the x -component of the drift velocity is negative, pushing ions towards the second peak. More importantly, the y -component of the drift velocity is negative near the second peak, $v_{d,y} \approx -E_x B_z$ – i.e., in the same direction as $F_L = eE_{0y}$ – while it is positive near the first peak ($x - x_{sh} \approx 0$). Consequently, ions remain closer to the second peak while gaining energy and moving in the negative y -direction, whereas they stay closer to the first peak while losing energy and moving in the positive y -direction. The same behavior is observed for electrons, although their trapping occurs around the peaks.




Figure 6.24: Trajectories and energy evolution of three downstream-escaping ions around the shock, $\Delta x_{sh}/\lambda_{si} = -5$ – -1.5 , overlaid on the map of the E_x/E_0 electric field (compare Figure 6.20).




Figure 6.25: Trajectories and energy evolution of three upstream-escaping ions around the shock, $\Delta x_{sh}/\lambda_{si} = -3-5$, see Figure 6.24. After the ions overcome the cross-shock potential, they undergo SSA at the shock front.

At $t\Omega_i \approx 120$, the ions scatter off magnetic field perturbations and begin to lose energy. Their parallel velocity decreases, and they start moving opposite to E_{0y} . Once the ions lose a significant amount of parallel velocity, their perpendicular velocity is still relatively high and they start to gyrate in the x - y plane (see Fig. 6.24 at $t\Omega_i \gtrsim 220$). The ions are advected downstream after they lose most of their energy.

A different case of ions that are trapped and accelerated, but eventually escape upstream rather than downstream, is illustrated in Figure 6.25. The three ions initially are accelerated by E_{0y} , moving in the negative y -direction. As in Figure 6.24, they subsequently scatter and continue moving in the positive y -direction while losing energy. The key difference is in their perpendicular velocity: although these upstream-escaping ions lose a substantial fraction of their parallel velocity, their perpendicular velocity remains large. At a certain point in their gyromotion, the perpendicular energy becomes sufficient for the ion to overcome the electric potential peak. This is visible in the bottom-right panel for all three particles, where their kinetic energies reach minima (e.g., at $t\Omega_i \approx 160$ for particle A). Once the ions cross the peak, they are pushed further upstream by E_x and begin to be accelerated again by E_{0y} , executing a cycloidal motion and undergoing SSA.




Figure 6.26: Evolution of different quantities along the trajectory of ion A (see Figure 6.25) for two different intervals, $t\Omega_i = 63\text{--}111$ (left) and $t\Omega_i = 160\text{--}239$ (right). The quantities are given in the shock frame of reference, see Figure 6.22.

The Lorentz force acting on the escaping ions is:

$$F_{L,x} = e\left(-\frac{d\phi}{dx} + v_y B_z - v_z B_y\right) \quad (6.10)$$

$$F_{L,y} = e(E_{0y} + v_z B_{0x} - v_x B_z), \quad (6.11)$$

where $B_y \approx 0$ upstream. The x -component of the force oscillates as the particle approaches and recedes from the shock. As the y -component of the velocity increases due to E_{0y} , the cycloidal trajectory extends over a larger distance in the x -direction, enabling the ion to escape upstream.

In the y -direction, $v_z \approx c$ because the escaping ions are highly relativistic and follow the magnetic field lines, which are approximately aligned with the z -direction. The y -component of the Lorentz force can therefore be approximated as

$$F_{L,y} \approx e(-v_{sh} B_{0z} + c B_{0x} - v_x B_z) = ec B_{0x} \left(1 - \frac{\tan \theta_{Bn}}{\tan \theta_{th}} - \beta_x \tan \theta\right), \quad (6.12)$$

where $E_{0y} = -v_{sh} B_{0z}$ and $v_{sh} \tan \theta_{th} = c$. In a subluminal shock, the combined first and second terms are positive because $\tan \theta_{Bn} / \tan \theta_{th} < 1$, while the third term depends on the particle motion along x . Thus, $F_{L,y}$ is positive on average. Since the wave activity in front of the shock is considerably lower than on the downstream side, ions undergo SSA unimpeded. Over time, energisation via SSA gradually weakens (compare the slopes in the

bottom-right panel of Figure 6.25 at $t\Omega_i \approx 180$ and $t\Omega_i \approx 250$). Ultimately, an ion either escapes definitively upstream or is carried back downstream.

The energisation of ion A from Figure 6.25 is shown in more detail in Figure 6.26 over two distinct acceleration intervals. The left panels correspond to the interval in which the ion is trapped between the first and second potential peaks. The evolution of v_x resembles that of a trapped electron, albeit with a clear difference in magnitude (compare Figure 6.22, third row). However, E_x rarely drops below zero, because the minimum of the electric potential does not necessarily coincide with the pseudo-potential minimum (compare $\tilde{\phi}$ and $\tilde{F}$ near $x = x_m$). Consequently, E_x does not have to become negative. Moreover, the negative x -component of the $\mathbf{E} \times \mathbf{B}$ drift pushes the ion towards the second peak, aiding the motion.

The right panels show the interval during which the ion undergoes SSA. In this case the ion energy grows continuously while E_x oscillates as the ion moves away or approaches the shock. The velocity v_x and four-velocity u_y exhibit a behavior similar to that in the left panels, but with larger oscillation amplitudes. Note, however, that the energy growth slows down as the four-velocity also decreases and approaches zero by the end of the tracing.

6.3.7 Estimates of the maximum particle energies

The acceleration mechanism that is observed at the shock formed in simulation $\theta 30$ was previously observed in simulations of non-relativistic magnetosonic shock waves by Bessho and Ohsawa 1999; Bessho and Ohsawa 2002; Bessho and Ohsawa 2000; Zindo et al. 2005. In their simulations, only electrons are trapped and accelerated by this process. This work shows that the same mechanism occurs in mildly relativistic magnetised shocks. Moreover, it is demonstrated that, unlike in the non-relativistic case, ions can also be trapped and accelerated in a similar manner in mildly relativistic shocks. Indeed, ion trapping is expected to be effective only in magnetised relativistic shocks: the ion magnetisation must be sufficiently large, $\sigma_i \gtrsim 1$, for ions to follow magnetic field lines and become trapped within potential troughs. In such a case, the Alfvén velocity is relativistic, and the shock must therefore also be relativistic.

The maximum energy of electrons forming the power-law tail in the spectra in Figure 6.17 can be estimated from the pulse width. For an ultra-relativistic electron, for which the perpendicular velocity approaches the speed of light ($v_\perp \approx c$), and a pulse width of $\sim 2.5 \lambda_{si}$, the maximum energy follows from

$$r_{g,e,\max} = \frac{\gamma_{e,\max} m_e c}{eB} \approx 2.5 \lambda_{si} \rightarrow \gamma_{e,\max} \approx 2.5 \gamma_{\text{up}} \frac{B}{B_0} \frac{\Omega_{e,0}}{\omega_{pi,0}} \approx 2.5 \gamma_{\text{up}} \frac{m_i}{m_e} \frac{B_M}{B_0}, \quad (6.13)$$

where the subscript 0 denotes upstream quantities. This result shows that the maximum attainable electron energy scales linearly with the upstream Lorentz factor, the ion-electron mass ratio and the shock-compressed magnetic field. Using the measured compression $B_M/B_0 \approx 1.7$ in the simulation frame, the estimate yields $\gamma_{e,\max} \approx 950$, in agreement with the energy at which the exponential cutoff begins in Figure 6.17. Electrons can reach higher energies due to shock nonstationarity and nonuniformity, which modify both the pulse width and the magnetic compression.

The linear dependence on the mass ratio is also confirmed in the simulation $\theta 30$ - $m400$ with a larger $m_i/m_e = 400$ (see Appendix A.4). In this run the electrons reach a maximum energies roughly four times higher than in simulation $\theta 30$, in line with Equation 6.13. Extrapolating to the physical mass ratio $m_i/m_e = 1836$ therefore suggests an expected maximum electron energies of $\gamma_{e,\max} \approx 2 \times 10^4$ in this process.

The same approach can be used to estimate the maximum ion energy. Equating the ion gyroradius to the average pulse width yields a maximum ion energy of $\gamma_{i,\max} \approx 9.5$. This value is smaller than the observed cutoff at $\gamma_i \approx 20$. However, the perpendicular ion velocity is expected to be smaller than c , since most of the energy is directed along the magnetic field lines, and the corresponding maximum energy is therefore larger. Ions are further accelerated via SSA at the shock front.

6.3.8 Discussion on oblique subluminal shocks

Appendix A.3 compares the results presented above with those of simulation $\theta 30$ - $\varphi_B 0$, which employs an in-plane background magnetic field configuration. Both simulations exhibit magnetosonic dispersive wave trains with comparable average amplitudes of electrostatic gradients and magnetic field oscillations. However, in simulation $\theta 30$ - $\varphi_B 0$ the wave train is more strongly dissipated, resulting in only two to three well-defined strong peaks, in contrast to simulation $\theta 30$, where the shock transition typically consists of four to five strong peaks.

Despite these differences, the particle-acceleration efficiency in both simulations is comparable. The maximum ion energies attained within the wave pulse are similar, and nearly the same number of ions undergo acceleration in both cases. Electrons, however, are accelerated more efficiently in the out-of-plane configuration, where they reach nearly twice the maximum momentum obtained in the in-plane run. This difference is primarily attributed to the faster dissipation of the magnetosonic wave train in simulation $\theta 30$ - $\varphi_B 0$. Nevertheless, the wave nonstationarity and transverse fluctuations are significant in both simulations. These effects occasionally lead to the local development of several strong oscillations along the transverse extent of the wave. Such transient enhancements of the electrostatic potential facilitate electron trapping and acceleration across multiple peaks, enabling efficient energisation despite on average faster damping of the wave train.

It is therefore concluded that particle acceleration is expected to occur with similarly high efficiency in fully 3D setups. A two-dimensional simulation with an increased mass ratio, $m_i/m_e = 400$, further demonstrates that the shock structure is largely independent of m_i/m_e , at least within the range explored. It is also confirmed that, in this mass-ratio regime, the ion and electron acceleration processes at oblique subluminal shocks scale approximately linearly with m_i/m_e (see Section 6.3.7).

Note that the maximum electron energies may also be limited by synchrotron losses, for which the cooling time depends on both the electron energy and the magnetic field strength. When synchrotron losses are significant, trapped electrons cannot easily escape the shock potential, allowing their acceleration to be sustained over longer timescales. Furthermore, under conditions of a heated upstream plasma, the trapping condition is more readily satisfied, enabling a larger fraction of electrons to be trapped and subsequently accelerated to high energies.

6.4 Quasi-parallel shock (simulation $\theta 10$)

FIGURE 6.27 shows the late-stage ($t\Omega_i = 314.1$) structure of a quasi-parallel mildly relativistic magnetised shock transition obtained for simulation $\theta 10$ with the mean magnetic field inclined at angle $\theta_{Bn} = 10^\circ$ to the shock normal (see Table 6.1). This is far from the threshold angle, $\theta_{th} \approx 36.1$. The field obliquity in the upstream rest frame is $\theta_{Bn,up} \approx 4.5^\circ$.

A magnetic field compression pulse – whose leading edge is referred here to as the shock front – with amplitude $\delta B_z/B_0 \approx 0.9$ is visible at $x \approx 105\lambda_{si}$, moving at a speed of $\beta_{sh} \approx 0.33$, corresponding to a Lorentz factor of $\gamma_{sh} \approx 1.06$. The corresponding upstream-frame values are $\beta_{sh,up} \approx 0.95$ and $\gamma_{sh,up} \approx 3.1$. Although the magnetic field is compressed to $r_B \approx 3$ to the left of the magnetic pulse (bottom panel), the plasma density (top panel) increases gradually rather than through a sharp jump, reaching compression ratio of $r \approx 4$ far downstream of the magnetic pulse. Note, however, that the shock in simulation $\theta 10$ is not yet fully formed (see below). Evidently, the shock structure in this case is unique and does not resemble the typical shock profiles usually discussed.

The Alfvén Mach number of the shock in the upstream frame, calculated from Equation 3.21 and using the magnetic field strength $B_{0,up}$, is $M_A = 1.035$. Thus, the shock is weak and subcritical. It is shown in Section 6.4.2 that the shock is also subcritical with respect to the whistler critical Mach number, M_w .




Figure 6.27: Maps of n_i/n_0 (top panel), E_x/E_0 (middle panel), and B_z/B_0 (bottom panel) for simulation $\theta 10$ at $t\Omega_i = 314.1$. Compare Figures 6.2 and 6.10.

Upstream of the magnetic pulse, a train of waves is visible (see B_z and E_x maps in Figure 6.27). It will be argued in Section 6.4.2, that these waves are dispersive whistlers emitted by the weak shock. Their group speed – different for different whistler modes – is larger than the shock speed, allowing them to propagate ahead of the shock. The highest-amplitude whistlers ($x/\lambda_{si} \approx 110$ – 150) move in phase with the pulse. The presence of





Figure 6.28: Phase-space for ions (left) and electrons (right) at $t\Omega_i = 314.1$ for simulation $\theta 10$ in the region shown in Figure 6.27). Compare Figure 6.3.

corresponding E_x coherent fluctuations results from an oblique propagation angle of the waves with respect to the background magnetic field, $\mathbf{B}_0$.

In the region of the upstream whistlers, undulating ion density filaments are observed whose amplitude increases toward the pulse. Similar filaments appear in the electron density (not shown), though more diffuse. These structures are a consequence of charge-separation effects by the whistler waves acting differently on ions and electrons due to their mass difference. Associated with the filaments are oblique waves observed in both the longitudinal electric field and B_z in Figure 6.27, with much shorter wavelengths than the upstream whistlers. Their orientations coincide with undulations of the ion density filaments. For example, the wavevectors of the waves at $x \approx 125\lambda_{si}$ and $x \approx 130\lambda_{si}$ have different directions, consistent with their location relative to minima or maxima in E_x . Both the filaments and the oblique short-scale waves are discussed in detail in Section 6.4.3.

Figure 6.29: Ion distribution for the x -component of the normalised momentum $p_{i,x}/m_i c$ for the shock downstream region $x/\lambda_{si} = 60\text{--}70$ (compare Figure 6.27). The vertical lines indicate the maximum values of the double-peaked distribution at $p_{i,x}/m_i c = -1.26$ and $p_{i,x}/m_i c = 0.82$.



6.4.1 Ion and electron phase-space

Figure 6.28 shows the ion and electron phase-space with the same layout as in Figure 6.3. The position of the pulse is discernable from the ion distributions at around $x \approx 105\lambda_{si}$. The $p_{x,i}$ ion phase-space shows a clear two-beam structure both upstream and downstream of the pulse. In front of the pulse, a beam with negative $p_{x,i}$ is visible, corresponding to the cold incoming plasma, which far from the shock flows with $p_{x,i}/m_i c = -2$ and becomes slightly perturbed by the upstream whistlers as it approaches the shock ramp. In the opposite direction, a hot beam with positive $p_{x,i}$ is present, corresponding to ions leaking from the downstream region.

The incoming ion beam decelerates at the shock front only to $p_{i,x}/m_i c \approx -1.3$ and continues to counterstream with the beam of ions reflected from the left boundary of the simulation box that move with $p_{i,x}/m_i c \approx 0.8$. This double-peaked structure in $p_{x,i}$ is illustrated in Figure 6.29 for the downstream region of the shock at $x/\lambda_{si} = 60\text{--}70$ (compare Figure 6.28). In the same region a finite drift with negative $p_{i,z}$ exists – a feature of low-Mach number shocks (see Section 6.2.2). In simulations of strong shocks with large M_A , the upstream plasma nearly fully thermalises downstream and the resulting distribution becomes close to isotropic (e.g., Bohdan et al. 2017). In contrast, for very low Mach number shocks the strong magnetic field forces ions and electrons to follow field lines, which inhibits thermalisation even over long simulation times. This behaviour was observed in simulations $\theta 30$ and $\theta 40$. In the present quasi-parallel case, where the field is nearly aligned with x , the resulting two-beam structure becomes even more pronounced.

Note that a similar double-beam velocity structure across the entire shock transition is known from hybrid-kinetic simulations of low Mach number (quasi-)parallel shocks under solar-wind conditions (see, e.g., Omidi et al. 1990, and references therein). In particular, Omidi et al. 1990 showed that a strictly parallel 1D shock consists of large-amplitude electromagnetic waves on the magnetosonic-whistler branch, which scatter incoming ions only

weakly. As a result, the shock does not decouple from the simulation box boundary, and the downstream ion distribution required to satisfy conservation laws is instead formed by the superposition of two distinct populations: incident and wall-reflected ions. This distribution is unstable to the ion-ion beam instability at the shock front and in near downstream, generating the magnetosonic-whistler waves that fill the upstream. However, in contrast to Omidi et al. 1990, where the two beams have similar temperatures and the velocity decrease across the shock is mainly due to the negative velocity of the reflected ions, in simulation $\theta 10$ (as well as in the 1D simulation reported in Appendix A.6) the incident ions decelerate and the reflected beam becomes significantly heated (compare Figure 6.29). The ions are further heated in the region populated by dispersive whistlers; however, nearly all ions heated in this region leak from the downstream. Thus, ion reflection at the wall continues to play an important role in the shock structure, as in Omidi et al. 1990.

A prominent feature in the ion phase-space is the wave-like structure that appears ahead of the shock front, in the region populated by upstream whistlers, $x/\lambda_{si} \approx 110$ –230. Modulations in $p_{x,i}$ are clearly visible in the incoming beam, and are also present in the hot reflected beam. The oscillations of the reflected beam appear approximately 180° out of phase with those of the incoming beam. The oscillations in $p_{y,i}$ and $p_{z,i}$ are centered at $p_{y,i} = p_{z,i} = 0$. This structure results from the large-amplitude waves driving a periodic quiver motion of the upstream plasma along the x -direction. Because the upstream magnetic field is strong, this quiver motion remains coherent, producing a well-defined wave patterns that become progressively diffused closer to the magnetic pulse, evidencing progressively stronger wave-particle interactions.

The electron phase-space shows that electrons are also strongly heated in the shock transition, both downstream and in the whistler wave region. Upstream of the shock front, quiver-motion modulations are visible, though there is no clear distinction between the upstream and higher-energy electrons. For both ions and electrons, along the x -direction and in the whistler precursor region, the momentum distribution at a fixed position x periodically widens and contracts, following the whistler wave structure. This behavior is most evident in the p_y and p_z distributions, but also in p_x phase-spaces. As discussed further below, this heating results from wave-particle interactions. The quiver motion of ions and electrons produces a finite bulk velocity difference in the transverse components, which drives instabilities and associated waves that facilitate heating (see Sections 6.4.3 and 6.4.4).

6.4.2 Upstream whistler waves

Figure 6.30 shows profiles of the electromagnetic field oscillations at the shock front and in the upstream region. The amplitudes of the magnetic field perturbations are largest for the waves closest to the shock, at $x/\lambda_{si} \approx 100$ –150, where they reach $\delta B_z/B_0 \approx 0.58$, while the B_y amplitudes are lower by $\delta(B_z - B_y)/B_0 \approx 0.12$ in this region. Farther upstream the amplitudes decrease, with their envelope varying from $\delta B_z/B_0 \approx \delta B_y/B_0 \approx 0.2$ to ≈ 0.4 at $x/\lambda_{si} \approx 150$ –220 and declining farther upstream. The electric field perturbations have lower amplitudes than their magnetic field counterparts.

These electromagnetic perturbations exhibit an almost perfect phase difference by 90° between components, indicating right-hand circular polarization. However, a closer analysis,




Figure 6.30: Transversely-averaged profiles of the normalised electromagnetic field perturbations at the same time as 6.27. a) Perturbations of the magnetic field, $\delta B_y = B_y$ and $\delta B_z = B_z - B_{0z}$ and b) perturbations of the electric field $\delta B_y = E_y - E_{0y}$ and $\delta E_z = E_z$

together with a comparison to the simulation with an in-plane magnetic-field configuration (see Appendix A.5), reveals that these waves are elliptically polarised and propagate along the x -direction.

A consequence of the quasi-parallel propagation is the excitation of a longitudinal electric field wave, E_x . Since the whistler wave propagates along the x -direction, the x -component of the Poynting flux, $\mathbf{S} = \mathbf{E} \times \mathbf{B}$ has a finite value, but the transverse components are zero, $S_y = S_z = 0$. Since the z -component of the magnetic field is composed of the background magnetic field and the wave field $B_z = B_{0z} + \delta B_z$, an additional E_x wave field with the same frequency and wavenumber is required. In the particular case of an out-of-plane magnetic field – such as in simulation $\theta 10$ – the E_x wave is proportional to B_y and has an amplitude of $\delta E_x/E_0 \approx 0.6$ (see Figure 6.27), which is large and comparable to the transverse electric field of the whistler, $\delta E_y/E_0$ and $\delta E_z/E_0$.

Figure 6.31 shows the y -averaged profiles of the normalised out-of-plane magnetic field, B_z/B_{0z} , stacked at various times between $t\Omega_i = 37.7$ and 239 to show the evolution of the waves at the shock and upstream. In front of the magnetic pulse, whose motion is indicated by the black line with $v = 0.33c \approx v_{sh}$, a train of waves is observed. Most of these are whistlers propagating away from the shock. However, the long-wavelength, large-amplitude whistlers near the shock front have phase speeds nearly equal to the shock speed

and are thus phase-standing with the shock. Some waves in this region appear to have phase velocities directed toward the shock, particularly during the early stages of evolution. Higher-frequency whistlers advance faster, producing a smooth transition from high- k at the leading edge of the whistler wave train to low- k near the pulse at late times. Ahead of the large-amplitude whistlers, high-frequency, small-amplitude waves are observed (highlighted by an ellipse in Figure 6.31). These waves are whistlers generated during the initial reflection of ions at the left boundary. It is argued below that these dispersive whistlers are generated at the shock and in its immediate downstream region, and waves with group velocities greater than the shock speed can outrun the shock and propagate upstream.

Ahead of the whistler waves, small-scale, low-amplitude waves moving at the speed of light are observed (marked by the black line in Figure 6.31). These correspond to electromagnetic precursor waves excited by the SMI. Since the magnetic field is quasi-parallel, the SMI is very weak, resulting in wave amplitudes that are significantly smaller than those observed in simulations $\theta 30$ and $\theta 40$ of oblique shocks.

Critical whistler Mach number M_w in the relativistic regime

The very low Alfvénic Mach number of the shock in the upstream frame, $M_{A,up} = 1.035$ (see Table 6.1), suggests that the shock is subcritical to the *whistler critical Mach number*, $M_{w,up}$; this is verified below. The analysis in this section is done in the upstream frame of reference. The numerical value of $M_{w,up}$ can be calculated analytically in the non-relativistic regime (see Equation 3.51 in Section 3.2.5). However, as the Alfvén speed approaches relativistic values (see Equation 3.21), Equation 3.51 can no longer be applied directly. Nevertheless, $M_{w,up}$ can still be obtained numerically from the dispersion relation for a circularly polarised wave propagating parallel to the magnetic field.

In the upstream frame of reference, the upstream plasma is at rest and cold. A circularly polarised wave propagating parallel to the magnetic field has the dispersion relation:

$$\omega^2 - k^2 c^2 - \omega_{p,i}^2 \frac{\omega}{\omega + \Omega_i} - \omega_{p,e}^2 \frac{\omega}{\omega - \Omega_e} = 0. \quad (6.14)$$

This equation can be arranged to obtain a quartic equation in ω :

$$\begin{aligned} \omega^4 + \omega^3 (\Omega_i - \Omega_e) - \omega^2 (\Omega_i \Omega_e + k^2 c^2 + \omega_{p,i}^2 + \omega_{p,e}^2) + \\ \omega (-k^2 c^2 (\Omega_i - \Omega_e) + \omega_{p,i}^2 \Omega_e - \omega_{p,e}^2 \Omega_i) + k^2 c^2 \Omega_i \Omega_e = 0 \end{aligned} \quad (6.15)$$

which has four solutions for real wavenumbers, $\forall k \in \mathbb{R}$. These solutions are shown in Figure 6.32 for the parameters of simulation $\theta 10$. They correspond to two right-handed waves (whistler and R) and two left-handed waves (ion cyclotron and L), which are plotted in the same quadrant for convenience. The right cutoff frequency, ω_R , is the frequency above which the R wave can propagate and is obtained in the limit $\omega \gg \Omega_i$ as

$$\omega_R = \frac{\Omega_e}{2} \left(1 + \sqrt{1 + \frac{4}{\sigma_e}} \right). \quad (6.16)$$




Figure 6.31: Profiles of the normalised magnetic field B_z/B_{0z} , averaged over the y -direction and stacked over the time interval $t\Omega_i = 37.7\text{--}239$. On the left, the black line indicates the position of the shock, moving at $v_{sh} \approx 0.33, c$. On the right, the black line marks the initial SMI-generated precursor waves propagating at the speed of light ($v = c$), while the circled region highlights the initial whistlers generated during the first reflection of ions off the left boundary.

Note that $\omega_R \rightarrow \Omega_e$ in the highly magnetised limit $\sigma_e \rightarrow \infty$. In simulation $\theta 10$, the electron magnetisation in the upstream frame is $\sigma_{e,up} = \gamma^2 (1 - \beta^2 \sin^2 \theta_{Bn} \sigma_{e,sim}) \approx 500$, and thus ω_R lies very close to the electron cyclotron frequency. In contrast, the cutoff frequency for the L wave, ω_L , cannot be computed in the same manner because, in the highly magnetised regime, it lies much closer to the ion cyclotron frequency. In fact, $\omega_L \rightarrow \Omega_i$ for $\sigma_e \rightarrow \infty$ which can also be observed in the figure.

As explained above, it is not possible to analitically compute the whistler critical Mach number in the highly magnetised regime due to the overlap between the L wave and the whistler branch, which spans roughly two orders of magnitude in kc/Ω_i . More precisely, this overlap occurs for $\Omega_i \lesssim kc \lesssim \Omega_e$. Therefore, the whistler critical Mach number must be computed directly from its definition (see Section 3.2.5). Recall that the whistler critical Mach number is given by the ratio of the maximum phase speed, $v_{ph} = \omega/k$, to the relativistic




Figure 6.32: Dispersion relation of a circularly polarised wave propagating parallel to the magnetic field, calculated using the upstream plasma parameters from simulation $\theta 10$ in the proper frame of reference. The four solutions correspond to the L wave (blue), ion-cyclotron wave (orange), whistler wave (green), and R wave (red). Left-handed waves (L and IC) are plotted in the same quadrant as the right-handed waves for convenience. A zoom-in of two selected regions of the left panel is shown on the right. The horizontal dashed lines indicate characteristic frequencies of the dispersion relation: in descending order, these are ω_R , Ω_e , ω_L , and Ω_i . The black line marks the maximum phase speed of the whistler branch, used to determine the whistler critical Mach number.

Alfvén speed (Krasnoselskikh et al. 2002; Burgess and Scholer 2015). In the upstream frame of reference, the Alfvén speed and its relativistic counterpart are

$$v_{A,rel,up} = \frac{v_{A,up}}{\sqrt{1 + v_{A,up}^2}} = 0.911c, \quad v_{A,up} = \frac{B_{0,up}}{\sqrt{\mu_0 n_{i,up} m_i}} = 2.21c, \quad (6.17)$$

, where $B_{0,up} \approx B_0$ since the magnetic field obliquity is close to 0, $\theta_{Bn} = 10^\circ$ and $n_{i,up} = n/\gamma_{up}$. For $M_{A,up} > M_{w,up}$, whistlers cannot phase-stand with the shock and are advected downstream, whereas for $M_{A,up} < M_{w,up}$, they can outrun the shock and form a wave train ahead of it. In Figure 6.32, the maximum phase speed is $v_{ph,max} = 0.996c$ (black line), occurring at frequency and wavenumber of $\omega_{max}/\Omega_{i,up} \approx 49.6$, $k_{max}c/\Omega_{i,up} \approx 49.7$. Given $v_{ph,max}$ and the relativistic Alfvén speed, $v_{A,rel,up}$, the whistler critical Mach number in the

upstream rest frame is

$$M_{w,up} = \frac{v_{ph,max}}{v_{A,rel,up}} \cos \theta_{up} \approx 1.09, \quad (6.18)$$

where a correction term $\cos \theta_{up}$ is included since the simulation is quasi-parallel, $\theta_{up} \approx 4.51^\circ$. This correction term is valid only in the low-obliquity limit. The Alfvénic Mach number in the upstream rest frame is $M_{A,up} = v_{sh,up}/v_{A,rel,up} \approx 1.035$, which is lower than the whistler critical Mach number. Thus, the shock in simulation $\theta 10$ is subcritical with respect to the whistler critical Mach number $M_{A,up} < M_{w,up}$. Note that the only way to obtain a subcritical shock in a highly magnetised plasma is for the shock to be mildly relativistic. If the shock were highly relativistic, $v_{sh} \approx c$, $M_{A,up}$ would increase and could exceed $M_{w,up}$. Therefore, the shock speed must be of the order of the Alfvén speed, i.e., $M_{A,up} \sim 1$. Note, that although the Alfvén Mach number is used here, the whistler critical Mach number is formally defined with respect to the fast magnetosonic speed v_f rather than the Alfvén speed. However, in a cold plasma, the sound speed is zero ($c_s = 0$), so $v_f = v_{A,rel}$ and using $M_{A,up}$ is appropriate (Kang et al. 2019).

All velocities involved in the calculations of the Mach numbers in the upstream rest frame are close to the speed of light – $v_{sh,up} \approx 0.945c$, $v_{ph,max} = 0.996c$ and $v_{A,rel,up} \approx 0.911c$ – and so the Alfvén and critical whistler Mach numbers have comparable values ($M_{A,up} \approx 1.035$ and $M_{w,up} \approx 1.09$). This might suggest that the observed precursor whistlers should be confined to high frequencies and wavenumbers, near those corresponding to the maximum phase speed, $\omega_{max}/\Omega_{i,up} \approx 49.6$ and $k_{max}c/\Omega_{i,up} \approx 49.7$ (see above). However, this is not the case: the waves span a broad range of ω and k , extending down to $\omega \approx kc \approx 1$. Evaluating the frequency and wavenumber at which the phase speed equals the shock speed, $v_{ph} = v_{sh,up} < v_{ph,max}$, yields $\omega_{sh}/\Omega_{i,up} \approx 0.94$, $k_{sh}c/\Omega_{i,up} \approx 0.99$, which are much smaller than the values associated with the maximum phase speed.

Whistler waves whose phase speed matches the shock speed in simulation $\theta 10$ can phase-stand with the shock. Their frequency and wavenumber correspond to the location in the dispersion relation where $v_{ph} = v_{sh,up}$. At these values, k_{sh} and ω_{sh} , the whistler group speed, $v_g = \frac{\partial \omega}{\partial k}$, exceeds the shock speed, allowing the waves to propagate upstream. There also exists a range of k and ω for which $v_g > v_{sh,up}$ while $v_{ph} < v_{sh,up}$. In this regime, whistlers propagate upstream ahead of the shock, but their wave crests move toward the shock (see, e.g., Omidi et al. 1990). The lower bound of this range is defined by $v_{sh,up} = v_g$, which occurs at $kc/\Omega_i \approx 0.43$ and $\omega/\Omega_i \approx 0.4$, while the upper bound is set by $v_{sh,up} = v_{ph}$, corresponding to $k_{sh}c/\Omega_{i,up} \approx 0.99$ and $\omega_{sh}/\Omega_{i,up} \approx 0.94$.

It is possible to consider a different definition of the critical whistler Mach number. An alternative definition of the critical whistler Mach number can be considered. Instead of using the maximum phase speed (Eq. 3.51), may define the threshold at which waves cannot propagate ahead of the shock. It is given by the maximum group speed, $v_{g,max}$. Following a procedure analogous to that used to compute the maximum phase speed, the maximum group speed is found to be $v_{g,max} = 0.998c$, occurring at $k_{g,max}c/\Omega_{i,up} = 17.23$ and $\omega_{g,max}/\Omega_{i,up} = 17.11$ in the dispersion relation. The subscript g distinguishes these values from those associated with $v_{ph,max}$. Since the shock speed is below this limit, whistlers are again able to propagate upstream. Note that $k_{g,max} < k_{max}$ and $\omega_{g,max} < \omega_{max}$. Therefore,

there is a range of k and ω for which the whistlers propagate toward the shock while their wave crests move upstream (see, e.g., Krauss-Varban and Omidi 1991). In this case, a short whistler precursor forms ahead of the shock.

Origin of precursor whistlers

As discussed in Section 6.4.1, the downstream plasma is not thermalised and instead exhibits a persistent two-beam structure. A two-beam instability is therefore expected. Given that the excited waves are whistlers with large amplitudes, the most likely mechanism is the ionion right-hand resonant instability (Omidi et al. 1990; Gary et al. 1984). In this scenario, the instability develops at the pulse and in the downstream region ($x/\lambda_{si} \approx 70$ –100 in Figure 6.27). Whistlers are excited and continuously contribute to the nonlinear steepening of the shock, but those with group speeds exceeding the shock speed do not participate in the steepening; instead, they escape upstream and form the whistler precursor. When the whistler group speed exceeds the shock speed, the precursor grows over time as additional waves emerge ahead of the shock. This behaviour is clearly visible in Figure 6.31, where the whistler wave train grows progressively longer.

The excited whistlers that propagate slower than the shock remain downstream. As the shock evolves, these slow whistlers contribute to the increase of the shock width, which grows over time (Omidi et al. 1990). This effect is visible in Figure 6.31. At early times of the simulation ($t\Omega_i \approx 100$, $x/\lambda_{si} \approx 30$) the shock has a relatively low amplitude and appears as a single, narrow pulse. As the system evolves, the pulse gradually steepens and broadens. At later times ($t\Omega_i \approx 230$), the pulse becomes significantly wider, trailing behind the precursor whistlers and producing a noticeably flatter B_z/B_0 profile.

These slower whistlers in the shock are, however, less apparent in the figure compared to the upstream precursor whistlers. Since they propagate slower than the shock, they accumulate immediately behind it rather than forming an extended wave train. Moreover, their amplitudes are smaller than those of the upstream whistlers, and they merge with the compressed downstream magnetic field, resulting in a smoother overall pulse instead of distinct wave peaks, as observed, e.g., in Omidi et al. 1990.

Precursor whistlers have been identified in quasi-parallel, low-Mach-number shocks and have even been observed in the Earth's bow shock. In particular, Schwartz et al. (1992) study a quasi-parallel bow shock with obliquity $\theta_{Bn} \approx 10$ –15 and Alfvénic Mach number of $M_A = 7.5$. Such a shock is subcritical with respect to the whistler critical Mach number. Note that for a non-relativistic shock – such as the bow shock – $M_w \propto \sqrt{m_i/m_e}$, so the critical value is significantly higher than in a simulation employing a reduced ion-electron mass ratio. In addition to the dispersive whistlers observed upstream of this shock, short large-amplitude magnetic structures (SLAMS) are also detected. SLAMS are commonly observed together with dispersive whistlers (Wilson et al. 2013; Giacalone et al. 1993) and have been reproduced in PIC simulations as well (Tsubouchi and Lembège 2004). These structures are characterised by an ion overdensity associated with magnetic field amplification occurring ahead of the whistlers. They arise from a nonlinear instability between reflected ions and upstream solar wind (upstream) ions mediated by the electrostatic potential of oblique whistlers (Scholer et al. 2003). Such structures, however, do not appear in Figure 6.31. The

tip of the whistler train simply corresponds to the fastest excited whistlers, and no noticeable ion-ion instability develops at that location.




Figure 6.33: Power spectrum of $\tilde{B} = (B_y - iB_z)/\sqrt{2}$ in the first quadrant, corresponding to right-hand polarised waves propagating rightward (R^+). The spectrum is computed in the region $x/\lambda_{si} = 80\text{--}160$ with a spatial resolution of two grid cells, $\delta x = 2\Delta = 0.005\lambda_{si}$, over the interval $t\Omega_i = 38\text{--}239$ and using a time resolution of $\delta t\Omega_i \approx 1.26$ (compare Figure 6.31). The overplotted curves correspond to the full circularly polarised wave dispersion relation (grey, Equation C.1), the dispersion relation in the cold plasma low-frequency limit (white, Equation 6.19), the shock speed (yellow), and the ion cyclotron resonance of the rightward diffuse reflected ion beam including thermal effects (red). The Fourier spectrum was smoothed using a Hanning window to mitigate finite-domain artefacts at $\omega = k = 0$.

Fourier analysis of upstream precursor whistlers

It is convenient to analyse the precursor whistlers in Fourier space by examining their frequency ω -wavenumber k power spectra. The Fourier analysis follows the convention described by Terasawa et al. 1986, which is reviewed in Appendix B.

Figure 6.33 shows the power spectrum of the quantity $\tilde{B} = (B_y - iB_z)/\sqrt{2}$. The spectrum was obtained using a fixed spatial window in the near-upstream region of size $\Delta x = 80\lambda_{si}$, with a spatial resolution corresponding to twice the grid cell size, $\delta x = 2\Delta = 0.005\lambda_{si}$. The temporal interval used for the analysis spans $t\Omega_i = 38\text{--}239$ with a time resolution of

$\delta t, \Omega_i \approx 1.26$. For visualisation purposes, the spectrum was smoothed with a Hanning window to suppress artificial power at $\omega = 0$ and $k = 0$ caused by the non-periodicity of the data in both space and time. The power is concentrated in the low-frequency, low-wavenumber regime, close to the ion cyclotron frequency. The electron cyclotron frequency, Ω_e , lies roughly two orders of magnitude above the range displayed in the figure.

The grey curve in Figure 6.33 shows the dispersion relation for a circularly polarised wave, obtained from Equation C.1. Additional details on this dispersion relation are provided in Appendix C. The parameters entering Equation C.1 are listed in Table 6.2. The upstream ion and hot ion-beam parameters are extracted directly from the ion phase space, where the two populations are clearly distinguishable and can be fitted independently with Maxwellians (see $p_{x,i}/mc$ in Figure 6.28). In contrast, the electron phase space does not separate cleanly into distinct components and cannot be reliably fitted. Instead, the parameters of the upstream electron population are inferred from the charge neutrality and current neutrality conditions, Equations C.4 and C.5 (see Appendix C for details).

Component	$n_{rest}(\text{ppc})$	$n_{up}(\text{ppc})$	β_{up}	γ_{up}	β_{th}
Upstream ions	4.47	4.47	0	1	0
Upstream electrons	8.14	9.48	0.51	1.17	0
Hot ion beam	1.18	5.01	0.97	4.23	0.77

Table 6.2: Parameters of the three plasma components upstream of the strongly magnetised quasi-parallel shock (simulation $\theta 10$), consisting of upstream ions, hot reflected ions and heated upstream electrons. Here, n_{rest} denotes the proper rest-frame density in units of particles per cell (ppc) and $\beta_{th} = v_{th}/c$ the thermal speed, while quantities with the subscript up are computed in the upstream frame of reference.

Additionally, Figure 6.33 shows a white line representing the dispersion relation for a cold plasma uniform collisionless plasma in a uniform magnetic field in the low-frequency limit, $\omega \ll \Omega_e$:

$$n^2 = \frac{\omega_{pi}^2}{\Omega_i(\omega + \Omega_i)}, \quad (6.19)$$

where the initial upstream values of Ω_i and ω_{pi} are used, and $\mathbf{n} = \mathbf{k}c/\omega$ is the index of refraction (Kulsrud 2005). In the non-relativistic limit, Equation 6.19 reduces to the well-known whistler dispersion relation

$$\omega^2 = k^2 v_A^2 \left(1 + \frac{\omega}{\Omega_i} \right). \quad (6.20)$$

One can see in Figure 6.33, that the difference between the grey and white lines is negligible at low frequencies and becomes slightly more pronounced at higher frequencies.

Both dispersion relations (grey and white lines) are initially computed in the upstream frame of reference for convenience and subsequently Lorentz-transformed to the simulation frame:

$$\begin{aligned} \omega' &= \gamma(\omega - \beta kc) \\ k'c &= \gamma(kc - \beta\omega). \end{aligned} \quad (6.21)$$

The yellow line in Figure 6.33 corresponds to the shock propagation, $\omega = v_{sh}k$. The red line represents the ion cyclotron resonance condition of the hot ion beam:

$$\omega = \beta_b kc - \frac{\Omega_i}{\gamma_T}, \quad (6.22)$$

where β_b is the speed of the beam in the simulation frame and $\gamma_T(\beta_b, v_{th}) > \gamma(\beta_b)$ is the Lorentz factor accounting for both the bulk velocity, β_b , and the thermal velocity, v_{th} , of the hot ion beam. This effectively incorporates the increase of energy due to the thermal motion of the ions.

As shown in Figure 6.33, both the grey and white lines closely match the observed signal, demonstrating that these waves are whistlers, as expected. The close agreement between the grey and white lines also indicates that the contribution of electrons and the hot ion beam is not required to generate the waves, which are maintained by the background upstream plasma. This implies that whistlers are generated at the shock and in the downstream region, consistent with the wave evolution analysis presented above. This finding aligns with hybrid kinetic simulations by Omidi et al. 1990, which showed that the reflected ion beam does not contribute significantly to shock dissipation or wave generation. Nevertheless, in the present case, weak effects of the reflected hot ion beam interacting with the waves are observed, as evidenced by the low power at $kc/\Omega_i \approx 1.6$ – 2.4 and $\omega/\Omega_i \approx 0.6$ – 1.2 , closely matching the ion cyclotron resonance (red line).




Figure 6.34: Power spectrum of $\tilde{B} = B_y - iB_z$ computed at the shock and close downstream, at $(x - x_{sh})/\lambda_{si} = 30$ – 60 (compare Figure 6.33). The yellow line shows the shock speed.

The largest power of the spectrum occurs approximately where the shock (yellow line) and the whistler branch (white line) intersect, at $k_x c / \Omega_i \approx 1$ and $\omega / \Omega_i \approx 0.45$. At this point,

the shock and whistler waves have similar phase speeds, so the growth rate of the whistlers is expected to be maximal at this wavenumber. Waves with higher phase speeds exhibit lower power. Whistlers with lower phase speeds are also excited, and their group speeds remain larger than the shock speed, allowing them to escape upstream. However, there exists a wavenumber at which the whistler group speed matches the shock speed, approximately at $k_x c / \Omega_i \approx 0.3$ and $\omega / \Omega_i \approx 0.1$. Whistlers with lower group speeds emitted by the shock cannot outrun it and are thus absent upstream, as indicated by the sudden drop in power around this wavenumber and frequency.

Note that in Figure 6.33, at the highest frequencies and wavenumbers ($\omega / \Omega_i \approx 1.7\text{--}2.4$) and $k_x / \Omega_i \approx 2.4\text{--}3.0$), a weak signal appears slightly above the expected dispersion relation indicated by the grey line. These waves are whistlers excited during the initial reflection of the cold upstream plasma beam at the wall. As a result, these boundary-affected whistlers follow the dispersion relation without the cosine correction for the quasi-parallel magnetic field (not shown). They are only visible at high frequencies, where they are not dominated by the larger-amplitude precursor whistlers.

Figure 6.34 shows the wave power spectrum over the same time interval and resolution as Figure 6.33, but computed for a downstream region including the shock, using a window width of $\Delta x = 30\lambda_{si}$. Here, the whistler branch is again visible, along with low-frequency waves whose group speeds are smaller than the shock speed. Since the two-beam structure persists downstream (see Section 6.4.1), whistlers are excited not only in the region of the compressed magnetic field but also downstream of this region. Nevertheless, the highest power remains located where the shock intersects the whistler branch.

6.4.3 Nonlinear dynamics of large-amplitude precursor whistlers

As discussed in Section 6.4.1 and evident in the phase-space plot, Figure 6.28, the presence of large-amplitude whistlers in the shock precursor is associated with ion and electron energisation. This section examines the nonlinear dynamics of these large-amplitude precursor whistlers as they interact with the upstream plasma, including wave dissipation through the generation of instabilities and the excitation of small-scale wave modes.

Quiver motion and transverse bulk drift

As noted in Section 6.4.1, as the upstream plasma moves toward the shock, both ions and electrons interact with the large-amplitude whistlers and perform a quiver motion, characterised by oscillations in all momentum components. Figure 6.35 shows a zoom-in of Figure 6.28, displaying the p_z and p_y components of the ion and electron momenta in the whistler precursor region. The profiles of the magnetic-field fluctuations $-\delta B_y / B_0$ and $-\delta B_z / B_0$ are overplotted in the top and bottom panels, respectively (not to scale). One can see that the quiver oscillations coincide remarkably well with the magnetic field profiles, which is consistent with the drift velocities obtained from the ion and electron equations of motion coupled to Maxwell's equations in a large-amplitude wave field, $\mathbf{v}_\perp \propto \delta \mathbf{B}_\perp / B_0$ (see below and also, e.g., Nariyuki et al. 2008; Sakai et al. 2005). For the hot ion and electron




Figure 6.35: Zoom-in of the p_y (top) and p_z (bottom) components of the ion (left) and electron (right) phase space (see Figure 6.28) for the region $x/\lambda_{si} \approx 90$ –250, encompassing the shock ramp and the whistler precursor. The solid black lines show the corresponding magnetic-field components, $-B_y/B_0$ (top panels) and $-\delta B_z/B_0$ (bottom panels). For clarity, the magnetic-field profiles are scaled by a factor of 1.5 for ions and 15 for electrons.

components, a corresponding phase-space modulation associated with the quiver motion is also evident.

For a transverse electromagnetic wave with magnetic perturbation $\delta \mathbf{B}_\perp$ and associated electric field $\delta \mathbf{E}_\perp$, the perpendicular quiver velocity of species j is $\delta v_{qj,\perp} \sim (q_j/m_j)\delta E_\perp/\omega$, where ω is the wave frequency. For whistlers with $\omega \ll \Omega_e$, this can be written as $\delta v_{qj,\perp} \sim v_{ph}\delta B_\perp/B_0$, which explains why the quiver oscillations observed in phase space closely track $\delta \mathbf{B}_\perp$. Because of the mass dependence, electrons exhibit much larger quiver velocities and amplitudes than ions. This produces a relative ion-electron drift and an associated transverse current with respect to the background magnetic field $\mathbf{B}_0$. The ion quiver amplitude observed in Figure 6.35 spans $\gamma_i v_i \approx 0.2$ –1.2, while the electron amplitude is approximately an order of magnitude larger, resulting in an ultrarelativistic relative drift. The resulting transverse current is therefore strong and provides a source of the *modified two-stream instability* (MTSI), as discussed below.

An alternative interpretation of the ion and electron modulation in circularly polarised waves is provided by the Walén relation (Hollweg et al. 1993). To sustain a whistler wave, the plasma requires a finite perpendicular bulk velocity of both ions and electrons, which is given by

$$v_{j,\perp} = -v_{ph} \frac{\Omega_j}{\omega \pm \Omega_j} \frac{\delta B_\perp}{B_0}, \quad (6.23)$$

where index $j = i, e$ denotes ions or electrons, respectively, and the sign depends on the




Figure 6.36: Maps of normalised charge density (a), high-pass filtered E_x electric field (b), E_y electric field (c), and B_y magnetic field (d) in the near-upstream region of the shock, $x/\lambda_{si} = 125$ – 155 (compare Fig. 6.27). The high-pass filtering applied to the E_x map is performed using a Gaussian filter, such that $E_{x,hp}(x, y) = E_x(x, y) - (E_x * g(\sigma = 20))(x, y)$, where the subscript hp denotes *high-pass*, $g(\sigma)$ is a Gaussian with standard deviation σ , and $*$ denotes the convolution operator.

particle charge. In the MHD regime ($\omega \ll \Omega_i \ll \Omega_e$), ions and electrons have the same bulk perpendicular speed, corresponding to the $\mathbf{E} \times \mathbf{B}$ drift. In the kinetic regime with $\omega \sim \Omega_i$, however, the Walen relation yields a relative drift, creating conditions for the excitation of MTSI.


Oblique electric field waves

Figure 6.36 shows the distributions of charge density, the electric-field components E_x and E_y , and the B_y component of the magnetic field in a part of the near-upstream region of the shock, $x/\lambda_{si} = 125$ – 155 , where whistler waves dominate (compare Fig. 6.27). The longitudinal electric field E_x has been high-pass filtered, which effectively removes the predominantly low-frequency whistler contribution.

The structure of the E_x electric field in Figure 6.36b is dominated by short-scale oblique waves. A similar pattern is observed in E_y in Figure 6.36c, where these waves are superim-

posed on the background of long-wavelength whistlers. The oblique waves in E_x alternate their wavevector direction between approximately 70° and 110° with respect to the x -axis. Their orientation is closely linked to the periodic modulation of the magnetic-field component B_y . At B_y maxima (e.g., around $x/\lambda_{si} \approx 130$), the wavevector is inclined at an angle of approximately 75° , whereas at B_y minima (e.g., at $x/\lambda_{si} \approx 135$) the orientation is about 105° . The waves disappear where $B_y \approx 0$. This provides evidence that the observed waves are driven by the transverse current associated with the ion-electron drift. The observed structure is consistent with modulations of the relative ion-electron drift velocity, which is proportional to the transverse magnetic field fluctuations $\delta \mathbf{B}_\perp$, as discussed above.

Figure 6.37: Profile of the ion phase-space component $p_{i,y}$ at $t\Omega_i \approx 314.1$ across the upstream region near the shock front, $x/\lambda_{si} \approx 112$ – 119 , where the ion beam undergoes a single quiver oscillation in the electromagnetic field of the whistler wave (see Fig. 6.35). The orange and green curves show Maxwellian fits to the negative and positive $p_{i,y}$ peaks, respectively.



The relative drift between ions and electrons across the background magnetic field, $\mathbf{B}_0$, can drive the MTSI if the drift velocity exceeds the ion thermal velocity, even in the presence of finite-temperature electrons (see, e.g., Matsukiyo and Scholer 2003, and references therein). The ion $p_{y,i}$ distribution across a narrow upstream region, in which the ion beam undergoes a single quiver oscillation, is shown in Figure 6.37. This distribution demonstrates that the ion thermal velocity is indeed a factor of 2–3 smaller than the quiver amplitude. Note that the beam profiles in Figure 6.37 are extracted from the surrounding hot ion population.

Figures 6.38a–b show the Fourier power spectra in $k_x - k_y$ wave number space of E_x and E_y computed in the whistler precursor region, $x/\lambda_{si} = 105$ – 220 . The MTSI-driven short-scale oblique waves are visible in $P(E_x)$ at wavenumbers $|k_x|c/\omega_{pi} \approx 2$ – 4 and $|k_y|c/\omega_{pi} \approx 7$ – 18 , nearly symmetric with respect to positive and negative wavenumbers. This corresponds to a wavevector modulus $kc/\omega_{pi} \approx 7.3$ – 18.5 , corresponding to wavelengths $\lambda \approx 0.35$ – $0.85\lambda_{si}$, with dominant modes at $|k_x|c/\omega_{pi} \approx 3$ and $|k_y|c/\omega_{pi} \approx 10$ – 15 , i.e., $\lambda \approx 0.4$ – $0.6\lambda_{si}$ (compare Figure 6.36). The wavevectors are inclined at angles $\theta_{Bk} \approx 73^\circ - 79^\circ$ for with $k_x > 0$, and $\theta_{Bk} \approx 101^\circ - 107^\circ$ for waves with $k_x < 0$, relative to the x -axis. A similar wavevector range is seen in the power spectrum of E_y . The wave frequencies are of the order $\omega \approx 1$ – $2\omega_{pi} \approx 1$ – $2\Omega_i$. These properties are therefore consistent with MTSI-driven lower-hybrid modes, with $\Omega_i < \omega_{LH} < \Omega_e$ (see Nariyuki et al. 2008; Sakai et al. 2005).

Dissipation of large-amplitude electromagnetic waves propagating parallel to the background magnetic field through the excitation of MTSI-driven modes and subsequent para-




Figure 6.38: Fourier power spectra in $k_x - k_y$ space for E_x (a), E_y (b) and the charge density n_c (c) computed at $t\Omega_i = 314.1$ in the whistler precursor region $x/\lambda_{si} = 105\text{--}220$ (compare Fig. 6.35). The k_x axis is plotted using a symmetric logarithmic scale with a linear threshold at $k_x c/\Omega_i = 5$.

metric decay instabilities has been studied for low-frequency Alfvén waves by Sakai et al. 2005 and Nariyuki et al. 2008. In their 2D PIC simulations, they observed electron heating parallel to the magnetic field due to quasi-electrostatic lower-hybrid waves generated by the MTSI, while ions were heated by parallel-propagating ion-acoustic waves. The wave properties observed in simulation $\theta 10$ are in good agreement with these works. However, the waves generated here do not appear to contribute significantly to the electron and ion heating observed in the whistler precursor region. This is consistent with the fact that, in the present simulation, the MTSI-driven short-scale oblique waves are very weak and constitute only a small perturbation, with $\delta E_x/E_0 \approx 0.02$, $\delta E_y/E_0 \approx 0.05$, superposed on the large-amplitude whistlers. Despite their low amplitude, these waves are not damped by electron Landau damping because the precursor whistlers are continuously emitted from the shock front, maintaining the quiver-induced cross-field current throughout the entire precursor region.

Note that, by performing simulations with different polarisations of the parent Alfvén waves and varying background plasma conditions, Sakai et al. 2005 and Nariyuki et al. 2008 demonstrated that the nonlinear dynamics of large-amplitude waves is highly sensitive to wave properties, and that the conditions required to excite the MTSI can be easily violated by modest changes in parameters. Further studies over a broad parameter range are therefore required to identify the typical processes that govern the nonlinear wave dynamics in the precursor region dominated by large-amplitude dispersive whistlers in mildly relativistic, quasi-parallel, strongly magnetised shocks.

Density filaments

Top panel of Figure 6.27 shows undulating filaments in ion density in front of the magnetic pulse, for $x/\lambda_{si} \gtrsim 105$. They appear at the tip of the upstream whistlers ($x/\lambda_{si} \approx 230$), where their amplitude is low, and continue along the x -direction with increasing undulation and amplitude. Their semi-coherent structure is destroyed at the shock ramp, $x/\lambda_{si} \approx 105$.

The filament wavelength matches that of the background whistlers, and the undulation becomes more pronounced as the whistler amplitude increases. A similar pattern, coinciding with the filaments, is visible in the B_z magnetic field (bottom panel in Figure 6.27), but its amplitude is very weak and only discernible near the shock front, at $x/\lambda_{si} \approx 105\text{--}150$. A more detailed view of the filamentary structure is provided by the charge-density map in Figure 6.36a. The filamentary structures are also present in the electron density (not shown), co-spatial with the ion density filaments. However the electron density perturbations are weaker and the filaments are more diffuse than the sharp ion filaments. This behaviour is expected, since electrons are strongly magnetised and have a high temperature, allowing them to rapidly neutralise charge and current perturbations.

Comparison of the Fourier power spectrum of the charge density in Figure 6.38c with the spectra of the electric-field components in Figures 6.38ab shows that the charge-density spectrum closely resembles that of E_y . This indicates that the filaments are driven by the same mechanism that generates the electric field perturbations, namely the MTSI. Therefore, the density filaments are interpreted as a nonlinear outcome of the MTSI. The relative ion-electron drift induced by quiver motion produces charge imbalances in the direction of the drift, leading to charge-separation electric fields and, consequently, filamentary density perturbations.

In a 3D simulation, the filaments would look like helices since the quiver motion occurs both in the y - and z -direction. Nevertheless, these whistlers do not significantly increase the temperature of the upstream ions, evidenced by the thin wave structure observed in Figure 6.35

Although filaments are observed along the x -direction, the actual waves excited by the instability have a wavenumber directed in the y -direction. To be precise, the instability has the largest growth rate for a perpendicular propagation of the magnetic field, and so the filaments form along the magnetic field. This is better observed in a simulation with an in-plane configuration of the magnetic field (see Appendix A.5 for a detailed analysis). There, since the magnetic field lies in the simulation plane, the filaments appear oblique in the direction of the background magnetic field and so the waves are excited perpendicularly.

Small-scale waves by the parametric decay instability

In addition to the oblique MTSI-driven waves and density filaments, the field and charge density distributions shown in Figures 6.27 and 6.36 exhibit even shorter-scale oscillations in the longitudinal E_x electric field throughout the shock precursor. These waves propagate with the plasma beam and are also imprinted in the transverse electromagnetic field components and the density. They appear in the Fourier power spectra of Figure 6.38 with dominant wavevectors $k \approx k_x \approx 20\Omega_i/c$, corresponding to a wavelength $\lambda \approx 0.3\lambda_{si}$. As argued below, these waves are identified as ion acoustic waves generated by the large-amplitude whistlers via the *parametric decay instability* (PDI).

In a PDI the parent wave – typically a whistler or Alfvén wave – decays into two daughter waves: a density wave (an ion acoustic wave) and a backward-propagating Alfvén-like wave (Goldstein 1978). Figure 6.39 shows a sketch of the PDI in the $\omega - k$ space. The black lines show the parent wave dispersion relation which, in the plasma rest frame, is symmetric with

respect to the ω -axis. The highest power of the parent wave is located at k_0 . An Alfvén-like wave with wavenumber k_1 and an ion acoustic wave with wavenumber k_2 are excited such that

$$k_0 = k_1 + k_2 \quad (6.24)$$

In the cold plasma limit ($\beta_p = 0$), the sound wave dispersion relation is simply $\omega = 0$. Therefore, the daughter waves have wavenumber $k_1 = -k_0$ and $k_2 = 2k_0$. The values k_1 and k_2 can be seen in Figure 6.39. In the case of a warm plasma (see red lines in Fig. 6.39), the sound speed is finite and the wavenumber of the daughter waves decreases according to Equation 6.24, noting that k_2 is located along $\omega = kc_s$.




Figure 6.39: Sketch of the parametric instability. The highest power of the parent wave is at k_0 . In the limit of a cold plasma ($\beta_p = 0$), the sound speed is zero and so the sound wave dispersion relation is along the k -axis. In result of the PDI two daughter waves with wavenumbers $-k_0$ and $2k_0$ are generated. Solution for a warm plasma is shown in red. In this case, the sound wave moves at a speed c_s and has a dispersion relation $\omega = kc_s$. The wavenumber of the resultant daughter waves are indicated with two red dashed circles.

It is convenient to analyse the short-scale waves observed in the precursor in Fourier space. Figure 6.40 shows the power spectrum of $\tilde{B} = \frac{1}{\sqrt{2}}(B_y - iB_z)$ and charge density $n_c = n_i - n_e$ for high frequency and wavenumber waves. The data has been first filtered using a high-pass filter to remove the lower-frequency lower-wavenumber region and so the whistler and oblique MTSI-driven waves are not visible. In the left panel, the power spectrum of $\tilde{B}$ shows a circled region with highest power at $k_1 c / \Omega_i \approx -15.8$, $\omega_1 / \Omega_i \approx 14.9$, while in the right panel there are two circles located at $k_2 c / \Omega_i \approx \pm 22.3$, $\omega_2 / \Omega_i \approx \mp 19.8$. The same signal in $\tilde{B}$ appears for the transverse electric field (not shown) and so this wave is an electromagnetic wave that propagates towards the shock ($\partial\omega/\partial k < 0$ and $\omega/k < 0$). Note the signal in the right panel is real and thus the power spectrum is symmetric $P_{n_c}(k, \omega) = P_{n_c}(-k, -\omega)$.

To study the instability, it is necessary to move to the upstream frame of reference. Using Equations 6.21, wavenumbers and frequencies of the density and daughter waves identified above can be calculated in the upstream rest frame as:

$$\begin{aligned} k_{1,\text{up}} c / \Omega_i &\approx -5.64, & \omega_{1,\text{up}} / \Omega_i &\approx +1.61 \\ k_{2,\text{up}} c / \Omega_i &\approx \pm 10.16, & \omega_{2,\text{up}} / \Omega_i &\approx \pm 0.22 \end{aligned} \quad (6.25)$$

Following the sketch in Figure 6.39, and assuming that the wave at $k \approx 2k_0$ corresponds to the sound wave $k_{2,\text{up}}$, it is evident that the plasma is not strictly cold since $\omega_{2,\text{up}}$ is finite.




Figure 6.40: Power spectrum of $\tilde{B}$ (left) and the charge density n_c (right) showing the high- ω , high- k space for the region at the end of the precursor whistler ($\Delta x = 30\lambda_{si}$) and a time interval of $t\Omega_i = 239\text{--}245$ with space and time resolution of $\delta x = 2\Delta = 0.005\lambda_{si}$ and $\delta t/\Omega_i = 0.05$, respectively. Note that charge density is real and therefore the power spectrum in the right panel is symmetric. A high-pass filter based on a Gaussian kernel $g(\sigma = 4)$ was applied to the data prior to Fourier transforming to the $\omega - k$ space. The red line denotes the upstream beam dispersion relation $\omega = v_{up}k$, while the red circles mark the locations of maximum power at $k_1c/\Omega_i \approx -15.8$, $\omega_1/\Omega_i \approx 14.9$ for the magnetic field and $k_2c/\Omega_i \approx \pm 22.3$, $\omega_2/\Omega_i \approx \mp 19.8$ for charge density.

Nevertheless, the sound speed is very small ($c_s = \omega_{2,up}/k_{2,up} \approx 0.02c$), indicating a low plasma temperature. Therefore, the cold plasma approximation is reasonable.

To fit the momentum conservation formula, Equation 6.24, the first daughter wave can be $k_{1,up}$. The resulting wavenumber of the parent wave can thus be derived as $k_{0,up} = -k_1 \approx 5.64\Omega_i/c$ or $k_{0,up} = k_2/2 \approx 5.08\Omega_i/c$. The discrepancy in these estimates of ~ 0.56 can be attributed to the finite plasma temperature and a small but nonzero wavenumber component in the y -direction, k_y . Therefore, an average value is chosen, $k_{0,up}c/\Omega_i \approx 5.36$. Using Equation 6.14, the corresponding frequency of the parent wave is $\omega_{0,up}/\Omega_i \approx 5.19$. In the simulation frame, this corresponds to $k_0c/\Omega_i \approx 1.61$ and $\omega_0/\Omega_i \approx 0.88$.

The parent wave is therefore not the whistler with the highest power ($kc/\Omega_i \approx 1$ and $\omega/\Omega_i \approx 0.45$) in the wave spectrum in Figure 6.33, but instead corresponds to a higher-wavenumber and higher-frequency wave with slightly lower power. This discrepancy arises because the higher- k , higher- ω waves propagate faster than the high-power waves, which remain approximately in phase with the shock. While whistler waves are continuously emitted from the shock, the faster waves travel farther into the precursor region and PDI has more time to develop with these high- k parent whistlers.

Since one of the resulting waves is an ion acoustic wave, a longitudinal electric field wave appears as observed in Figure 6.36b. These small-scale waves interact with plasma ions and

electrons and accelerate them along the x -direction towards the shock. This is discussed in the next section.

6.4.4 Particle heating

The shock upstream region is populated by several waves ions and electrons can interact with. The identified potential sources of heating are the precursor whistlers, the oblique electrostatic MTSI-driven waves and the x -aligned small-scale waves. While ion distributions show a clear distinction between the heating provided to the upstream ion beam and the reflected ion beam, electrons form a unified distribution in the upstream region and so it is difficult to distinguish which electrons (either upstream beam or reflected) are affected.

Ion trapping by elliptically polarised whistlers

The oblique propagation of whistler waves with respect to the background magnetic field generates a longitudinal electric field component, E_x , which gives rise to an electric potential in which ions can be trapped within the potential troughs. Figure 6.41 shows the phase-space distribution of the x -component of the ion momentum, $p_{i,x}/m_i c$, with the longitudinal electric field E_x and the corresponding normalised electric potential $\tilde{\phi}$ overplotted. The potential is computed using only the E_x electric field, which is justified since modulations of E_y along the y -direction are significantly weaker than the whistlers amplitude.

By removing the low-density regions of phase-space, clear accumulations of ions are visible at locations where the electric potential reaches minima. Because the whistler waves propagate upstream, ions trapped by the potential must move approximately at the same phase speed as the waves in order to remain trapped. This is confirmed by comparing the dashed blue line in Figure 6.41, which represents the phase speed of the whistler waves closest to the shock (see Section 6.4.2), with the momentum of the accumulated ions. Near the shock, both ions and whistlers propagate approximately in phase. Farther upstream, where the whistler amplitude decreases, fewer ions are trapped. Moreover, these whistlers propagate at a faster phase speed and thus the trapped ions must reach higher momenta to satisfy the trapping condition, as visible, for example, around $x \approx 165\lambda_{si}$.




Figure 6.41: Phase-space distribution of the x -component of the ion momentum, $p_{i,x}/m_i c$, with low-density regions removed to highlight ion accumulations. The transversely-averaged longitudinal electric field E_x/E_0 (black curve) and the normalised electric potential $\tilde{\phi}$ (red curve) are overplotted. The dashed blue line indicates the phase speed of the whistler waves closest to the shock. Compare Figure 6.28.




Figure 6.42: Momentum-space $p_{i,y} - p_{i,z}$ averaged along the x -direction over a distance of $\delta x = \lambda_{si}$ for four selected locations along the x -direction: $x/\lambda_{si} = \{122, 126, 130, 134\}$. Low-density momentum-space has been removed to highlight the reflected beam. The arrows show the direction of the electromagnetic field components of the whistler wave.

The width of the ion momentum distribution at these locations correlates with the amplitude of the electric potential. The energy associated with the potential is comparable to the ion kinetic energy,

$$e\Delta\phi \sim (\Delta\gamma_i - 1)m_i c^2 \quad \Rightarrow \quad \Delta\tilde{\phi} \sim \frac{\Delta\gamma_i - 1}{\gamma_{i,\text{up}} - 1}. \quad (6.26)$$

For a potential amplitude of $\Delta\tilde{\phi} \approx 0.28$ at $x \approx 122\lambda_{si}$, this relation yields $\Delta\gamma_i \approx 1.35$, corresponding to a momentum spread of $\Delta p_{i,x}/m_i c \approx 0.91$, in agreement with the observed width of the phase-space ion accumulations. Since the ions are mildly relativistic, the distribution is slightly skewed towards larger $p_{i,x}$, elongating these ion structures in phase-space. In addition, the electric potential is stronger closer to the shock, leading to asymmetric poten-




Figure 6.43: Same as Figure 6.41 but for the x -component of the electron momentum, $p_{x,e}/m_e c$.

tial troughs and, consequently, asymmetric momentum distributions, with broader spreads closer to the shock (e.g. $x \approx 120\lambda_{si}$) and narrower ones farther upstream (e.g. $x \approx 124\lambda_{si}$).

This ion trapping mechanism is specific to quasi-parallel shocks in which the upstream oblique whistler waves are not purely electromagnetic and have electrostatic E_x field component. If the background magnetic field were strictly parallel to the shock normal, the whistler waves would lack a longitudinal electric field component, preventing particle trapping.

Following similar analysis, electrons should accumulate at maxima of the electric potential. However, the distribution of the $p_{e,x}$ electron momentum (Figure 6.43), shows that electrons preferentially populate regions near potential troughs, opposite of what is expected. At these locations, electrons are unstable with respect to the electric potential and would tend to move away from the troughs. Nevertheless, the strong ion accumulation within the troughs generates a significant charge-separation electric field as electrons attempt to escape the trough, pulling electrons back towards these regions. This charge-neutralisation effect allows electrons to remain near the potential troughs.


Another potential mechanism for ion trapping in circularly polarised waves is phase bunching (Hoshino and Terasawa 1985). Figure 6.42 shows the ion momentum distribution in the $p_{i,y} - p_{i,z}$ plane at several locations within a single wavelength of the whistler wave. Phase bunching requires force-free motion, $\mathbf{v} \cdot \mathbf{E} = 0$, which occurs when the transverse particle momentum is perpendicular to the electric field perturbation and thus aligned with the magnetic field fluctuation. For instance, at $x \approx 134, \lambda_{si}$ (bottom-right panel of Figure 6.42), the electric field perturbation is oriented along $\hat{z}$, such that the force-free direction corresponds to $p_{i,y}$. However, only part of the force-free direction is stable: momentum antiparallel to the magnetic field perturbation leads to stable bunching, whereas parallel configurations are unstable.

The upstream ion beam (indicated by a high-density in momentum-space) remains almost perfectly phase-locked to the magnetic fluctuations, showing a clear anti-phase correlation with δB and thus stable momentum-space bunching. In contrast, the reflected ion population (lower density phase-space) displays no preferred orientation and remains approxi-

mately isotropic in transverse momentum space. This indicates that, while phase bunching operates for the upstream beam – and corresponds to quiver motion in the wave field– it does not for the reflected ions and thus trapping for this ion population is due to the electric potential shown above. Such electrostatically trapped particles are distributed more or less uniformly in the $p_{i,y} - p_{i,z}$ space which results in phase-space modulations visible in these components in Figures 6.28 and 6.35.

Electrons act in a similar way. The upstream electron beam remains phase-locked with respect to the electromagnetic field components, while the diffused electrons remain approximately isotropic and are not trapped by the perpendicular components of the whistler wave. The modulations in $p_{e,y}$ and $p_{e,z}$ components visible in Figures 6.28 and 6.35 result from distribution of electrostatically trapped populations, as in the case of the ions.

Figure 6.44: Space-time diagram, $x - t$, showing the motion along the x -direction ($x/\lambda_{si} = 0-120$) for a time interval of $t\Omega_i = 37.7-314.1$ of 50 ions. The dashed line shows the position of the shock front over time. The background shows the transversely-averaged profile of the z -component of the magnetic field, B_z/B_0 .



Shock-front interactions and scattering off the precursor whistlers

Figure 6.44 shows trajectories of 50 ions along the x -direction as a function of time. The trajectories are color-coded with the Lorentz factor and are overplotted on the background of the transversely-averaged profile of the B_z magnetic field component. At the latest time shown, thirty of these ions have Lorentz factors $\gamma > 8$ and are located upstream of the shock front (marked with dashed black line in the figure), while the remaining twenty ions have $\gamma > 5$ and are found downstream.

Initially, all fifty ions are reflected at the conducting wall. As discussed in Section 6.4.1, thermalisation is inefficient in very low Mach number shocks. As a result, dissipation remains weak and the shock does not detach from the wall, even at late times. Because the ions are strongly magnetised and the magnetic field is almost aligned in the x -direction, no ions are

reflected at the shock front itself. Instead, they are advected downstream and subsequently reflected at the wall, as illustrated in the figure.

Following the initial reflection, most ions interact with the shock and experience a modest energy increase. A subset of these ions then propagates upstream, where they interact with the precursor whistler waves and become trapped (see previous Section). Once they are trapped, particle-wave interactions may slowly energise the ions, as indicated through the colour-coded trajectories. Other ions remain confined within the magnetic pulse and are accelerated downstream of the shock front, while some are advected downstream without undergoing significant energisation. Notably, the ions reaching the highest energies are preferentially located upstream of the shock, where sustained interaction with the whistler waves is possible.




Figure 6.45: Evolution of the Lorentz factor γ (a) and the work done by the parallel (b) and perpendicular (c) electric field over time normalised to the ion rest energy $W_{\parallel}/m_i c^2$, $W_{\perp}/m_i c^2$ for 5 distinct ions for the time interval $t\Omega_i = 37.7\text{--}314.1$. The work is computed with respect to the background magnetic field, $\mathbf{B}_0$.

Figure 6.45 shows the temporal evolution of the Lorentz factor and the work done by the electric field, parallel and perpendicular to the background magnetic field $\mathbf{B}_0$, for five representative ions interacting with upstream whistlers. All ions exhibit an initial increase in their Lorentz factor at $t\Omega_i \lesssim 90$ (panel a), corresponding to the interaction with the shock front (see Figure 6.44). Afterwards, their trajectories diverge, reflecting distinct interactions with the whistler waves. Four ions remain trapped in the electric potential, exhibiting strong Lorentz factor oscillations around a constant value of the Lorentz factor. The fifth ion (blue curve) is initially trapped at $t\Omega_i \approx 240$ it scatters off the wave, escaping the potential trough and moving ahead of the wave.

Panels (b) and (c) show the work done by the electric field. The initial shock interaction

produces a clear signature in the parallel work, likely due to the electric potential associated with the magnetic pulse. Subsequently, the parallel work oscillates over long timescales, reflecting the motion of ions back and forth within a potential trough, while remaining approximately constant on average. In contrast, the perpendicular work exhibits larger-amplitude, higher-frequency oscillations and dominates the net energy exchange, particularly for ions that remain coupled to the whistlers. Note that upon the scattering of the fifth ion off the wave, both parallel and perpendicular work contributions increase slightly. Overall, although the perpendicular electric field dominates the energisation, the net ion energy gain remains moderate.

Electron heating by electron-whistler interaction and the MTSI

Figure 6.46 shows the space-time diagram as in Figure 6.44 but for a sample of 30 electrons. At the latest time shown, the electrons have Lorentz factors $\gamma > 150$. All electrons start their trajectory after their reflection off the shock during its formation. Electrons move ahead and later scatter off the precursor whistlers. From here onwards, electrons either: are advected downstream, where they can gain energy at the location of the magnetic pulse; are advected towards the shock but do not go past it and are accelerated at the shock front, where a continuous stream of whistler waves is generated; are scattered but remain within the precursor whistlers. A small portion of electrons do not scatter significantly and move upstream, gaining less energy than those electrons that stay within the near upstream.

Figure 6.46: Same as Figure 6.44, but for a sample of 30 electrons.




Figure 6.47 shows the trajectories of five representative electrons selected from Fig-

ure 6.46. The evolution of the Lorentz factor significantly resembles the interaction of ions with whistlers (see previous Section 6.4.4). The blue curve corresponds to an electron that propagates ahead of the precursor whistler waves, while the orange curve represents an electron that is scattered by the whistlers upstream and subsequently advected downstream of the shock. The remaining trajectories correspond to electrons located close to the shock front or within the whistler precursor. A modest initial increase in the Lorentz factor is observed at early times ($t\Omega_i \lesssim 70$), which is attributed to the first interaction with the forming shock, similar to ions. At later times, the electron energy either remains approximately constant or shows a slow increase.

The electron propagating ahead of the precursor (blue curve) exhibits a notably different energy evolution. Regardless of the initial increase, the energy remains constant along its trajectory, only showing small oscillations. This is supported with the work done by the electric field. The parallel work shows oscillations around zero associated with the electron crossing the longitudinal E_x wave. The perpendicular work, although positive, shows small oscillations and remains constant over time after the initial interaction.




Figure 6.47: Same as Figure 6.45 but for 5 representative electrons.

The downstream-scattered electron (orange curve) gains energy as it crosses the shock at $t\Omega_i \approx 140$. Unlike the other electrons, its parallel work drops sharply at the shock and then slowly recovers, likely due to interaction with the electric potential. In contrast, the perpendicular work increases and then gradually decreases, consistent with strong scattering by waves in the magnetic pulse. Overall, the net energy gain is negligible, as the total work remains approximately zero.

The remaining electrons exhibit a behaviour similar to that of ions, characterised by a gradual, albeit slow, increase in the Lorentz factor, as illustrated by the purple, green and red curves. Note, however, that the contribution of the parallel work over time for all

electrons is negative as opposed to ions, which showed positive parallel work for some ions. This discrepancy in the parallel work shows how ion heating is more efficient than electron heating and it will become more evident in spectra (Section 6.4.4).

The MTSI-driven oblique waves may also contribute to electron heating (see Section 6.4.3). However, the structures observed in the transverse components of the electron momentum cannot be directly attributed to the MTSI. This instability predominantly heats electrons along the direction parallel to the background magnetic field (Sakai et al. 2005; Matsukiyo and Scholer 2003). In simulation $\theta 10$, the magnetic field is nearly aligned with the x -axis; consequently, the MTSI primarily affects the longitudinal momentum component $p_{e,x}$, rather than the transverse components $p_{e,y}$ or $p_{e,z}$.

Electrons interact with the MTSI-driven oblique waves mainly through Landau damping, which leads to heating of the upstream electron population while simultaneously damping the associated electric field fluctuations. However, the amplitude of these oblique waves is significantly smaller than that of the precursor whistlers. As a result, their overall contribution to electron heating is expected to be modest or even negligible. In fact, no characteristic feature of interaction with MTSI-driven waves is observed in phase-space nor spectra.

Effect of the small-scale waves

The impact of the previously analysed small-scale waves (see Section 6.4.3) on ions and electrons is illustrated in Figure 6.48. The top panel shows the x -component of the ion momentum $p_{i,x}$, focusing on the upstream beam with $u_{0,x}/c = -2$, while the middle panel shows the same component of the electron momentum $p_{e,x}$ over the range $p_{e,x}/m_e c \in [-60, 30]$. The bottom panel presents the corresponding longitudinal electric field E_x/E_0 . The reflected ion beam is not shown, as it does not exhibit any discernible interaction with the small-scale waves.

Small-scale waves superimposed on the whistler waves are clearly visible in the electric field and are directly associated with the spikes observed in both ion and electron momentum. The spacing of these spikes matches the wavelength of the small-scale waves. Overall, however, the impact on the upstream ion beam remains modest and results in only mild ion heating.

For electrons, the effect is considerably more pronounced, with high-density spikes belonging to the upstream electron beam extending to more negative momentum values. In addition to these spikes, elongated, low-density diffused structures appear at negative electron momenta ($p_{e,x} < 0$) at the locations of the small-scale waves, for example around $x/\lambda_{si} \approx 181$. These low-density structures belong instead to the reflected and diffused electron population that, unlike the reflected ion beam, is influenced by the small-scale waves. These electrons scatter off the small-scale waves, resulting in weak acceleration and a steep power-law feature, contributing to the overall electron heating (see next Section).

Energy spectra

Figure 6.49 shows the ion and electron energy spectra measured at the location of a single ion accumulation. Maxwellian fits are overlaid on both spectra, yielding normalised tem-




Figure 6.48: Phase-space for the x -component of ion (top) and electron (middle) momenta near the end of the precursor whistler region, $x/\lambda_{si} = 180$ – 192 , focusing on the upstream beam ($u_{0,x}/c = -2$). The bottom panel shows the normalised x -component of the electric field, E_x/E_0 . The spikes in phase-space coincide with the small-scale waves observed in E_x .




Figure 6.49: Ion and electron spectra at $t\Omega_i = 314.1$ with corresponding Maxwellian fits (black curves) at the location of a single ion accumulation $x/\lambda_{si} = 119$ – 124 . The dashed lines correspond to the location of the upstream beam.

peratures of $\Theta_i \approx 0.37$ and $\Theta_e \approx 2.33$, indicating that ions reach a significantly higher temperature than electrons, with $T_i \approx 16T_e$.

The ion spectrum exhibits a narrow peak at $\gamma_i - 1 \approx 1.24$, corresponding to the moderately heated upstream ion beam. Apart from this upstream peak, the Maxwellian fit reproduces the low-to-mid energy range well. At higher energies, a small non-thermal component is visible. As discussed in previous sections, the dominant heating mechanism is particle-wave interaction with whistler waves. The whistlers observed upstream have wavelengths comparable to the ion cyclotron frequency (see the power spectrum in Figure 6.33), allowing ions to approach near-resonant conditions. Although the ion cyclotron resonance condition does not actually intersect the whistler branch (red and white lines in Figure 6.33), the proximity of the two branches leads to ion heating rather than efficient acceleration, shaping the spectrum into a Maxwellian as seen in Figure 6.49. By contrast, the electron cyclotron resonance lies far from the whistler branch, leading to substantially weaker electron heating,

consistent with the pronounced temperature difference between ions and electrons.

A small fraction of the most energetic ions in the hot ion beam may nevertheless resonate with the whistler waves. For highly relativistic ions, the cyclotron resonance condition steepens and the effective ion cyclotron frequency decreases, allowing the resonance condition to intersect the whistler branch and ions become accelerated, producing the non-thermal component in the shape of an exponential cutoff observed at the highest energies. However, the number of ions satisfying this condition is expected to be small, and their overall contribution to the spectrum is small.

The electron spectrum displays qualitatively different features. In addition to a thermalised component only spanning a limited energy range ($10 \lesssim \gamma_e - 1 \lesssim 30$), the spectrum exhibits a comparatively flat low-energy component. This feature arises from the superposition of the upstream electron beam, whose peak would be located at $\gamma_e - 1 \approx 1.24$, and a more diffuse reflected electron population, which together produce a largely featureless low-energy distribution.

At mid-to-high energies, a distinct non-thermal population emerges, composed of two power-law-like segments followed by an exponential cutoff. The lower-energy power law ($20 \lesssim \gamma_e - 1 \lesssim 50$) is attributed to electron interactions with small-scale waves (see Section 6.4.4). Although the momentum gain from this process is modest, it produces clear momentum spikes in phase-space that shape this portion of the spectrum into a distinct power law. The higher-energy power law and cutoff ($\gamma_e \gtrsim 60$) are instead associated with electron interactions with whistler waves. While this non-thermal component is more evident than in the ion spectrum, the maximum electron energy ($\gamma_e - 1 \approx 200$) exceeds the Maxwellian peak ($\gamma_e - 1 \approx 10$) by only about one order of magnitude, and the resulting power law is steep. Consequently, the overall electron acceleration is not significant.

Figure 6.50: Ion and electron energy spectra at $t\Omega_i = 314.1$ for the region encompassing the entire precursor whistler upstream of the shock and the magnetic pulse, $x/\lambda_{si} = 70, 237$. Along the x -axis the kinetic energy of the particles $E_k = (\gamma - 1)mc^2$ is shown, normalised to the upstream ion kinetic energy, $E_{k,i,up} = (\gamma_{up} - 1)m_i c^2$.




Figure 6.50 shows the ion and electron energy spectra for the region encompassing the precursor whistlers and the downstream magnetic pulse. The overall spectral shape resem-

bles that in Figure 6.49, although some of the previously visible features are less pronounced.

The global heating at the shock is significantly lower for electrons than for ions. The respective maxima of each spectrum – excluding the upstream ion peak – are located at $E_k/E_{k,i,\text{up}} \approx 0.05$ for electrons and $E_k/E_{k,i,\text{up}} \approx 1$ for ions. This discrepancy indicates that ion heating is, in general, more efficient than electron heating, consistent with spectra in Figure 6.49.

6.4.5 Discussion on quasi-parallel shocks

Simulation $\theta 10$ exhibits a subluminal shock that does not produce the strong particle acceleration seen in simulation $\theta 30$, but instead results primarily in efficient particle heating. The magnetic pulse at the shock front is significantly weaker than in shocks with higher magnetic obliquity, leading to a correspondingly weak electric potential. Consequently, shocks with intermediate obliquities, $10^\circ < \theta_{Bn} < 30^\circ$, may show the combined effect of both heating in the precursor whistler region and acceleration and trapping at the shock front, contributing to the overall energisation of particles.

A key parameter governing the qualitative behaviour of the shock is the whistler critical Mach number, M_w . As the magnetic field obliquity increases, the whistler branch progressively merges with the magnetosonic branch. This increased obliquity reduces the phase speed of the waves projected along the shock normal, eventually preventing whistlers from propagating upstream. As a result, the shock structure transitions into a train of decaying magnetosonic waves, as observed in simulations $\theta 40$ and $\theta 30$.

Appendix A.5 presents a comparison between simulation $\theta 10$ and simulation $\theta 10\text{-}\varphi_{B0}$, which shares identical parameters but features an in-plane background magnetic field, $B_0 = (B_{0x}, B_{0y}, 0)$. In both cases, a train of whistler waves develops upstream of the shock. Since the derivation of the whistler critical Mach number does not depend on the inclination of the magnetic field and the shock velocity is approximately the same as in simulation $\theta 10$, the shock in simulation $\theta 10\text{-}\varphi_{B0}$ is likewise subcritical with respect to M_w .

In addition to the precursor whistlers, shocks in both simulations excite the MTSI-driven oblique waves and small-scale waves associated with the PDI. While the PDI-driven waves propagate predominantly along the whistler direction, the orientation of the MTSI-driven waves differs substantially between the two simulations. Because the MTSI has its maximum growth rate when the wavevector is perpendicular to the background magnetic field, the direction of propagation changes when the magnetic field lies in the simulation plane. As a result, the density filaments – produced as a by-product of the MTSI – appear oblique within the simulation plane and are aligned with the background magnetic field.

Particle heating in both simulations is essentially identical, as expected since waves of similar nature and amplitude are excited in each case. This suggests that the shock structure and heating efficiency are largely insensitive to the magnetic field inclination, φ_B , in 2D simulations.

Appendix A.6 compares simulation $\theta 10$ with its one-dimensional counterpart, simulation $\theta 10\text{-}1D$. Also in this case precursor whistler waves are observed, PDI-driven waves are also present, as this instability requires only a single spatial dimension to develop. The MTSI is not triggered in simulation $\theta 10\text{-}1D$, as it requires at least two spatial dimensions to develop.

Nevertheless, since the contribution of MTSI-driven oblique waves to particle heating is weak in the 2D simulation, their absence in simulation $\theta 10$ -1D is expected to have a minor impact on the overall phase-space structure.

Phase space in the one-dimensional simulation exhibits a similar temperature throughout the shock structure. It should be noted, however, that the shock has not yet reached a steady state. Therefore, the plasma remains anisotropic and no Maxwellian component is observed in particle spectra. Overall, this comparison demonstrates that dimensionality plays a crucial role in capturing the particle-wave interaction that drives the particle heating observed upstream of the shock.

CHAPTER 7

Summary and discussion

IN this thesis, particle-in-cell simulations are used to study the structure of mildly relativistic collisionless shocks and associated particle acceleration processes in highly magnetised plasmas with oblique magnetic field configurations. As discussed in Chapter 2, such shocks may occur in astrophysical systems that host relativistic plasma outflows, including jets in AGN and microquasars. These objects are prominent sources of intense non-thermal radiation and are widely considered as potential acceleration sites for high-energy cosmic rays.

Particle acceleration at relativistic shocks depends on the conditions of the environment, including plasma magnetisation, the shock obliquity, and the Lorentz factor (see Sections 3.3 and 4.1.1). While DSA produces extended-energy-range power-law spectra in non-relativistic shocks, its efficiency regime is restricted to weakly magnetised, subluminal conditions and drops as plasma magnetisation increases. At large magnetisations, WFA can accelerate ions and electrons to high energies in ultrarelativistic shocks. However, this mechanism is inefficient in the mildly relativistic regime (Section 4.2). The goal of this thesis is thus to investigate whether mildly relativistic strongly magnetised shocks with a subluminal configuration of the magnetic field can support alternative mechanisms of efficient particle heating and acceleration.

Simulation studies performed in this thesis use the PIC simulation code THATMPI, that utilises a two-dimensional, three-velocity-component (2D3V) configuration (Section 5.1). For the purpose of this studies, the code was optimised and overhauled while including new features such as MPI subarrays for faster communication between computational processes, parallelised I/O, the implementation of the Higuera-Cary pusher (Higuera and Cary 2017), relativistic particle injection (Zenitani and Nakano 2022) and improved particle tracing (Section 5.5).

In addition to the main simulation results on mildly relativistic magnetised shocks, this thesis presents for the first time a neural-network-based framework for analysing particle tracing data from kinetic plasma simulations (Section 5.6). The dataset utilised to perform this analysis is taken from PIC simulations of high-Mach-number shocks in which the Buneman instability is triggered and electrons undergo SSA. The CNN and MLP neural network architectures are employed to predict the maximum kinetic energy attained by electrons using supervised and unsupervised learning strategies.

Using particle momentum as input data, supervised learning methods – both classification and regression – achieve very good performance despite the highly non-linear nature of the dataset. Anomaly detection is also explored using both semi-supervised and fully unsupervised approaches, which yield acceptable results but perform worse than supervised methods due to the absence of labels. When only a single component of the electric field is used as input data, supervised learning still achieves reasonable performance despite a noisy dataset, whereas anomaly detection degrades substantially. CNNs are identified as the most suitable neural network architecture for handling noisy data in future applications. These

results have been published in Torralba Paz et al. 2025. They provide guidance for future applications on other particle acceleration processes and highlight the potential of neural networks as efficient analysis tools for kinetic-based simulations.

The PIC simulations presented in this thesis investigate physical shock parameters that have never been explored before – mildly relativistic regime with Lorentz factor of $\gamma_{sh,up} \sim 3$ and strong plasma magnetisation, $\sigma = 1$, corresponding to an ion magnetisation of $\sigma_i \approx 1$, and highly magnetised electrons, $\sigma_e \approx 100$. Three magnetic field configurations are examined: a superluminal oblique shock with $\theta_{Bn} = 40^\circ$ in simulation $\theta 40$, a subluminal shock with obliquity just below the threshold angle, $\theta_{Bn} = 30^\circ$ in simulation $\theta 30$, and a quasi-parallel shock ($\theta_{Bn} = 10^\circ$, run $\theta 10$). The results can be summarised as follows.

1. The oblique shock transition in simulations $\theta 40$ (Section 6.2) and $\theta 30$ (Section 6.3) feature a magnetosonic wave composed of a series of sharp, large-amplitude wave pulses with decaying amplitude, which develop pronounced electrostatic potential peaks at their crests. The shock structure in both simulations is reminiscent of that of weak sub-critical shocks observed in the non-relativistic regime in Earth’s Magnetosheath and in the interplanetary space (Section 6.2.3). It is set by the ion and electron gyration in the wave’s magnetic field with the $\mathbf{E} \times \mathbf{B}$ drift. The wave oscillations are damped by some anomalous dissipative and wave dispersion processes which, in both simulations, generate entropy-mode waves providing additional means of dissipation.
2. In upstream regions of oblique shocks, X-mode and O-mode precursor waves generated by the SMI are observed. However, the wave amplitude is small and the resulting wakefield acceleration leads only to a mild electron heating in the region upstream of the magnetosonic pulse, consistent with previous studies of mildly relativistic perpendicular shocks (Ligorini et al. 2021b).
3. Due to the large amplitudes of the magnetosonic waves, the propagation of oblique magnetised shocks is intrinsically non-stationary. 2D simulations presented in this thesis further demonstrate that this non-stationarity leads to the development of pronounced transverse non-uniformities. These structures evolve on spatial and temporal scales of the ion gyromotion and reach amplitudes comparable to the local mean fields.
4. The wave non-stationarities and transverse non-uniformities create local conditions in the shock allowing for particle reflection and trapping inside the magnetosonic pulse and their associated heating and acceleration. This process corresponds to the electron acceleration mechanism previously demonstrated by Bessho & Ohsawa in a series of studies of non-relativistic and trans-relativistic shocks (with maximum investigated shock velocities of $v_{sh} \approx 0.75c$, $\gamma_{sh} \approx 1.5$) in plasmas with moderately magnetised electrons, $\sigma_e = 3 - 9$ (Section 4.3). In this mechanism, magnetised upstream electrons that cross the shock while streaming along magnetic field lines can be reflected at a minimum of the pseudo-potential, $F' < 0$ – i.e., the electric potential along the magnetic field – near the downstream edge of a magnetosonic pulse, and are subsequently driven back toward the potential maximum. During this return motion, the x -component of the $\mathbf{E} \times \mathbf{B}$ drift opposes their motion, reducing their parallel velocity

and enabling prolonged acceleration by the motional electric field E_{0y} , aided by the y -component of the $\mathbf{E} \times \mathbf{B}$ drift. It is demonstrated in this thesis that a similar process operates in mildly relativistic, strongly magnetised shocks and is capable of energising both electrons and ions. Ion reflection and trapping require the presence of potential troughs within a train of large-amplitude magnetosonic pulses, as observed in simulations $\theta 40$ and $\theta 30$. In addition, ions must be magnetised, i.e., $\sigma_i \gtrsim 1$. Under these conditions, electrons are correspondingly very strongly magnetised, $\sigma_e \gg 1$. Owing to this strong electron magnetisation and the substantially larger shock Lorentz factor considered in the present simulations, the Bessho & Ohsawa mechanism operates with significantly higher efficiency than in previously studied non-relativistic shock regimes.

5. According to theoretical predictions, the efficiency of the Bessho & Ohsawa process depends sensitively on the shock obliquity with respect to the threshold angle between subluminal and superluminal magnetic field configurations. In superluminal shocks, the strong $\mathbf{E} \times \mathbf{B}$ drift advects electrons and ions downstream, preventing multiple reflections and permitting only brief, transient trapping. Although shock non-stationarity and transverse non-uniformities lead to strong temporal and spatial variability in the trapping and acceleration efficiency – occasionally allowing short-lived episodes of enhanced energisation – the overall particle energisation remains inefficient in superluminal shocks (Section 6.2.4). In contrast, subluminal shocks allow electrons and ions to be repeatedly reflected within the wave potential, leading to deep trapping and substantial energy gains. Large-amplitude transverse wave fluctuations further facilitate particle trapping and acceleration across multiple wave pulses, enabling efficient energisation. This highlights the crucial role of two-dimensional effects in accurately capturing particle acceleration processes at mildly relativistic, strongly magnetised shocks.
6. In the case of the subluminal shock, simulation $\theta 30$, the number of trapped electrons, as well as the maximum energy attained, increase over time. The electron spectrum develops a power-law distribution at high energies. Individual magnetosonic pulses generate hard spectra with $p \approx 1$ – expected theoretically, while the full shock transition yields an overall spectral index of $p \approx 2$ (Section 6.3.4). However, while electrons continue to reach increasing maximum energies, the power-law spectral part – in which the bulk of high-energy electrons reside – extends only over roughly one order of magnitude in kinetic energy and is followed by the spectral cutoff. This limitation is expected, since the cutoff energy is constrained by the shock width. The cutoff energy is found to scale linearly with the upstream Lorentz factor, wave amplitude, and the ion-to-electron mass ratio. The latter scaling is confirmed by an additional simulation performed with a mass ratio of $m_i/m_e = 400$. For the physical mass ratio $m_i/m_e = 1836$, the expected maximum electron energy is therefore $\gamma_{e,\max} \approx 2 \times 10^4$. Electrons may reach even higher energies due to shock nonstationarity and transverse nonuniformities, which locally modify both the magnetosonic pulse width and the magnetic-field compression, but the number of such particles is relatively low. Previous studies in the non-relativistic regime (Bessho and Ohsawa 1999) did not report maximum electron energies or electron energy spectra. As a result, a direct quantita-

tive comparison with the non-relativistic case is not possible.

7. Ion acceleration operates analogous to that of electrons, but with efficiency reduced roughly two orders of magnitude. This is because the wave troughs are comparable in scale to the ion gyroradius and ions escape more readily. Some of them can overcome the shock potential and propagate upstream, where they undergo further energisation through the SSA process.
8. Although the simulations presented here are two-dimensional, the trapping and acceleration mechanism is expected to persist with similar efficiency in three dimensions. This expectation is supported by a comparison between simulation $\theta 30$ and a run with identical physical parameters but an in-plane magnetic field (simulation $\theta 30\text{-}\phi_B 0$; Appendix A.3), which reveals magnetosonic dispersive wave trains with comparable average amplitudes and similar cutoff energies in the spectra.
9. The shock structure in the simulation of a quasi-parallel shock, $\theta 10$, differs substantially from that of oblique shocks. In a strongly magnetised plasma, the shock is subcritical with respect to the whistler critical Mach number, allowing continuous emission of whistler waves from the shock front. In effect, the extended precursor to the shock is formed with large-amplitude dispersive whistlers, reminiscent of analogous structures identified in quasi-parallel, low-Mach-number shocks and observed in the Earth's bow shock. The nonlinear dynamics of precursor whistlers interacting with the upstream plasma lead to their dissipation via the excitation of the MTSI, which generates oblique electric-field fluctuations and produces density filaments aligned with the background magnetic field, as well as through the PDI, which excites small-scale ionacoustic wave modes. Particle-whistler interactions dominate the energy dissipation and lead primarily to efficient particle heating. PDI- and MTSI-driven waves play a secondary role in the heating processes. Whistler phase velocities in similar shock configurations reported in the non-relativistic regime in observations of the Earth's bow shock (Krauss-Varban and Omidi 1991) and hybrid simulations (Omidi et al. 1990) are lower than the shock speed and the resulting particle heating is weak. In contrast, the faster whistlers in simulation $\theta 10$ enable sustained waveparticle interactions, leading to strong heating upstream.

The results presented in this thesis demonstrate that an efficient particle acceleration mechanism can operate in strongly magnetised, mildly relativistic oblique shocks and can energise both ions and electrons to high energies. In astrophysical sources that can host such shocks, such as jets of AGN and microquasars, these processes provide a viable alternative to DSA and WFA, which are expected to be inefficient in this parameter regime. However, the results also show that the mechanism identified here operate only within a restricted region of the parameter space – the plasma magnetisation, shock speed, and magnetic field obliquity.

Efficient energetic particle production – and generation of intense electromagnetic emission – is expected at shocks with the magnetic field obliquity close to the threshold angle.

Although in such configurations the identified trapping and acceleration mechanism cannot provide extended power-law distributions in either the electron nor ion spectra required to account for the wide-band emission spectra of astrophysical sources, the total energy gain remains substantial. This mechanism may therefore act as an efficient injection process, in which ions and electrons can subsequently undergo further acceleration via other mechanisms, reaching even higher energies. Nonstationary shocks of the type discussed here, particularly when accompanied by significant transverse nonuniformities, can provide additional enhanced diffusion perpendicular to the background magnetic field, thereby further increasing particle energisation at the shock. However, owing to the restriction of the present simulations to a 2D geometry, cross-field particle diffusion is artificially suppressed, since an ignorable coordinate exists in the description of the electromagnetic fields in this case (Jokipii et al. 1993). The acceleration processes must therefore be studied in fully 3D simulations.

Note, that because the magnetic field amplitude is large, synchrotron losses by highly energetic trapped electrons may become important and limit the maximum attainable energy. Therefore, the main process that constrains the maximum electron energy is given by the shortest timescale of the processes: the synchrotron cooling time, τ_{sync} , in which case electrons remain trapped, or the escape time, τ_{esc} , in which case the strong modulations at the shock front allow electrons to escape.

A natural question arising from the results presented in this thesis concerns the dependence of the acceleration efficiency and particle heating on the total plasma magnetisation. A higher magnetisation is expected to yield a more efficient acceleration. However there is an upper limit to the magnetisation beyond which a shock cannot form, since the shock speed would fall below the Alfvén speed, equivalent to $M_A \lesssim 1$. A lower limit on the magnetisation is imposed through the requirement that, electrons need to be magnetised, $\sigma_e \gtrsim 1$, to undergo trapping and acceleration. A systematic parameter scan in magnetisation is therefore required to quantify the dependence of particle acceleration efficiency on the plasma magnetisation.

For quasi-parallel shocks with whistler-dominated upstream, low plasma magnetisation raises the Alfvénic Mach number, making the shock supercritical to M_w at some magnetisation level. This sets a lower limit $M_A = M_w$; for $M_A > M_w$ upstream precursor whistlers would not be observed and particle heating is suppressed. In this case, other mechanisms may play a role. More importantly, changes in magnetisation alter the phase speed of the standing whistlers, potentially affecting secondary instabilities and particle heating.

Shock speed is another critical factor. As shown by Bessho and Ohsawa 1999 and detailed in Section 4.3, faster shocks enhance acceleration efficiency at angles θ_{Bn} close to the threshold angle separating super- and sub-luminal shocks. Slower shocks increase the threshold angle, but reduce efficiency. Conversely, for higher shock speeds the threshold angle is reduced, and may fall in the regime of quasi-parallel configuration. Future studies should quantify how shock speed and magnetic field obliquity affect the shock structure and efficiency of particle acceleration processes.

APPENDIX A

Comparison of the main simulations for different parameters

This appendix presents a comparison of the main simulation runs, $\theta 40$, $\theta 30$, and $\theta 10$, discussed in Chapter 6, with the results of additional simulations performed using different parameter settings. These include simulations with the mean magnetic field lying in the simulation plane, $\varphi_B = 0^\circ$ (runs marked $^*\varphi_B 0$), a simulation with larger ion-to-electron mass ratio, $m_i/m_e = 400$ (marked with *m400), and simulations performed in one dimension (runs *1D). Table A.1 lists parameters of all simulations included in the analysis (compare Table 6.1).

Sim. #	θ_{Bn}	φ_B	m_i/m_e	$L_y (\lambda_{si})$	$T_{\max} (\Omega^{-1})$
$\theta 40$	40	90	100	7.2	113.1
$\theta 30$	30	90	100	4.8	314.1
$\theta 10$	10	90	100	4.8	314.1
$\theta 40\text{-}\varphi_B 0$	40	0	100	12.8	113.1
$\theta 30\text{-}\varphi_B 0$	30	0	100	12.8	188.5
$\theta 10\text{-}\varphi_B 0$	10	0	100	12.8	188.5
$\theta 30\text{-}m400$	30	0	400	7.2	68.84
$\theta 40\text{-}1D$	40	—	100	—	113.1
$\theta 10\text{-}1D$	10	—	100	—	163.3

Table A.1: Lists of simulation runs performed for the analysis of magnetised mildly relativistic shocks. The columns show, from left to right, the magnetic field obliquity angle, θ_{Bn} , the magnetic field inclination angle with respect to the 2D simulation plane, φ_B , the mass ratio m_i/m_e , and the simulation time, $T_{\max}$. The table is divided into groups marked by horizontal lines: the first group are the simulations presented in Chapter 6; the second group contains simulations with the mean magnetic field oriented in the simulation plane, $\varphi_B = 0^\circ$; the third group includes simulations with mass ratios $m_i/m_e = 10$ or 400 ; and the fourth group lists the 1D simulations performed for two obliquity angles. All other numerical parameters are the same. Compare Table 6.1.

A.1 Simulation $\theta 40\text{-}\varphi_B 0$

Figures A.1 and A.2 compare the results obtained for the oblique superluminal magnetised shock with $\theta_{Bn} = 40^\circ$ for two orientations of the background magnetic field with respect to the simulation plane: out-of-plane with $\varphi_B = 90^\circ$ (run $\theta 40$) and in-plane with $\varphi_B = 0^\circ$ (run $\theta 40\text{-}\varphi_B 0$). Figure A.1 compares the maps of the electron density, n_e , the normalised profiles of the E_x electric field, and the profiles of the shock-compressed components of the magnetic field ($\delta B_y/B_0$ for simulation $\theta 40\text{-}\varphi_B 0$ and $\delta B_z/B_0$ for simulation $\theta 40$) at time $t\Omega_i = 113.08$. One can see that the shock structure is very similar in both simulations. For both magnetic field configurations, a large-amplitude magnetosonic wave is formed that contains several




Figure A.1: Maps of electron density fluctuations, δn_e , and transversely averaged profiles of the E_x electric field, and shock-compressed magnetic field components, δB_y and δB_z , at time $t\omega_{ci} = 113.08$ for simulations $\theta 40$ and $\theta 40\text{-}\varphi_{B0}$, which differ in the configuration of the background magnetic field with respect to the 2D simulation plane. The map for simulation $\theta 40\varphi_{B0}$ is shown over a reduced transverse extent (smaller than the full width $L_y = 12.8$) to allow for a clearer comparison with the corresponding map from simulation $\theta 40$. All quantities are normalised to their respective upstream values.

pulses with decreasing amplitude, reminiscent of a weak shock with dispersive wave train that is gradually damped downstream due to dissipation. Both shocks feature the same number of wave oscillations with comparable average amplitudes, and in both cases the wave structure is strongly perturbed in the y -direction. As in simulation $\theta 40$, the shock transition in run $\theta 40\text{-}\varphi_{B0}$ also features SMI-driven upstream precursor waves, as well as entropy-mode waves downstream of the magnetosonic wave, at $x/\lambda_{si} \approx 40\text{--}55$. The main difference is that the entropy waves in the in-plane configuration propagate obliquely rather than along the x -direction. Note that the shock positions differ by approximately $\sim \lambda_{si}$, which reflects that the shocks fully developed at slightly different times.

Figure A.2 compares phase-space distributions for momentum components along the direction of the background magnetic field: p_x and p_y for the in-plane configuration (left panels) and p_x and p_z for the out-of-plane configuration (right panels), both at $t\Omega_i = 113.08$.




Figure A.2: Phase-space for ions (rows 1 and 2 from top) and electrons (rows 3 and 4) for selected momentum components at $t\Omega_i = 113.08$. The left panels show results for simulation $\theta40\text{-}\varphi_B0$, while the right panels for run $\theta40$.

The two simulations exhibit nearly identical particle distributions across the entire shock transition, although the in-plane case shows a slightly hotter electron population.

Overall, these results demonstrate that the shock properties are insensitive to the choice

of simulation-plane orientation, and that run $\theta 40$ reliably captures the physics of a superluminal magnetised shock.




Figure A.3: Profiles for $\delta n_i/n_0$, E_x/E_0 and $\delta B_z/B_0$ comparing simulations $\theta 40$ -1D and $\theta 40$ at $t\Omega_i = 113.08$. Profiles for simulation $\theta 40$ -1D are single-line and have been smoothed using a Butterworth filter due to the high level of noise; while profiles for simulation $\theta 40$ are transversely averaged.

A.2 Simulation $\theta 40$ -1D

This section aims to assess the impact of simulation dimensionality and determine whether 2D simulations are necessary. Figure A.3 compares the profiles of the normalised ion density profiles n_i/n_0 , the longitudinal electric field, E_x/E_0 , and the compressed magnetic field component, $\delta B_z/B_0$, obtained for one-dimensional simulation $\theta 40$ -1D and two-dimensional run $\theta 40$ at $t\Omega_i = 113.08$. The profile for simulation $\theta 40$ -1D is taken as a single-line profile while for simulation $\theta 40$ the profile is transversely averaged. One can see that shocks in both simulations display nearly identical structures in terms of wave amplitudes, compression ratio, and shock propagation speed. However, clear differences can be observed in phase-space distributions downstream of the magnetosonic wave pulse. As one can see in the left panels of Figure A.4, although both ions and electrons become trapped and accelerated in the magnetosonic wave potential structures, they reach slightly lower momenta. Moreover the downstream temperature is lower in the 1D simulation, as evidenced by the smaller phase-space widths of the distributions. This demonstrates the importance of 2D effects for accurately capturing particle heating.




Figure A.4: Phase-space for ions (rows 1 and 2 from top) and electrons (rows 3 and 4) for selected momentum components at $t\Omega_i = 113.08$. The left panels show results for run $\theta 40$ -1D, while the right panels for simulation $\theta 40$.




Figure A.5: Maps of the normalised ion density, n_i/n_0 (top panel), the longitudinal electric field, E_x/E_0 (middle panel), and in-plane component of the magnetic field, B_y/B_0 , for simulation $\theta 30\text{-}\varphi_B 0$ at $t\Omega_i = 188.5$. Compare Figure 6.10.

A.3 Simulation $\theta 30\text{-}\varphi_B 0$

Figure A.5 shows maps of the normalised ion density, n_i/n_0 , the longitudinal electric field, E_x/E_0 , and the magnetic field component B_y at the end of simulation $\theta 30\text{-}\varphi_B 0$, in which the background magnetic field at oblique angle $\theta = 30^\circ$ lies in the $x - y$ simulation plane, $\varphi_B = 0^\circ$. This should be compared with Figure 6.10 in Section 6.3, obtained for simulation $\theta 30$ with an out-of-plane magnetic-field configuration, $\varphi_B = 90^\circ$. Figure A.6 also presents the comparison of the y -averaged profiles of the electric field E_x , the normalised electrostatic potential ϕ , and the compressed magnetic-field component B_y (B_z) for both simulations at the same time, $t\Omega_i = 188.5$. In both simulations the shock appears as a magnetosonic wave pulse, and the shock velocities are comparable. However, while the wave pulse in simulation $\theta 30$ typically consists of 4–5 strong peaks – clearly visible in both the field maps and the averaged profiles, along with another 4–5 weaker peaks – the in-plane simulation $\theta 30\text{-}\varphi_B 0$ shows only 2–3 strong peaks, with an additional 2–3 weaker peaks visible only in the averaged profiles. The amplitudes of the first two peaks at $t\Omega_i = 188.5$ are comparable, with $\delta B_y/B_0$ ($\delta B_z/B_0$) ≈ 1.3 (Figure A.6, bottom). In simulation $\theta 30\text{-}\varphi_B 0$, the amplitude of the first peak varies only moderately, within the range $\delta B_y/B_0 \approx 1.25\text{--}1.35$, and the variation relative to the second peak is similarly small, not exceeding $0.15B_0$. The average amplitudes of the electrostatic gradients in the wave (Figure A.6, top) are smaller in simulation $\theta 30\text{-}\varphi_B 0$, with $\delta E_x/E_0 \approx 0.8$ (compared to $\delta E_x/E_0 \approx 1.0$ in run $\theta 30$), and vary also moderately within $\delta E_x/E_0 \approx 0.7\text{--}0.85$. The average electrostatic potential is comparable in both simulations (Figure A.6, middle), but the potential troughs in run $\theta 30\text{-}\varphi_B 0$ are shallower and clearly defined only between the first 3–4 peaks. The average potential varies moderately in the




Figure A.6: Profiles of E_x/E_0 , $\tilde{\phi}$, B_y/B_0 , and B_z/B_0 comparing simulations $\theta 30\text{-}\varphi_B 0$ and $\theta 30$ for the in-plane and out-of-plane configurations of the mean magnetic field, respectively, at time $t\Omega_i = 188.5$. The bottom panel shows the B_y profile for simulation $\theta 30\text{-}\varphi_B 0$ and the B_z profile for simulation $\theta 30$.

range $\tilde{\phi} \approx 0.35\text{--}0.4$. However, local variations along the transverse shock structure can be substantial, reaching $\Delta\tilde{\phi} \approx 0.3$ in the first peaks – an 80% fluctuation – significantly larger than in the out-of-plane simulation $\theta 30$.

As in simulation $\theta 30$ and also in run $\theta 40$, the variability in the field quantities and in the electrostatic potential is driven primarily by wave nonstationarity and by fluctuations in the transverse wave structure. Similar to the $\varphi_B = 90^\circ$ simulations, the transverse modulations evolve on spatial and temporal scales of the ion gyromotion and are chaotic, showing behaviour distinct from the ion-temperature-anisotropy-driven Alfvén-ion-cyclotron (AIC) instability that typically appears in in-plane magnetic-field configurations and produces shock-front rippling Matsukiyo and Matsumoto 2015; Kobzar et al. 2021 (see also run $\theta 40\text{-}\varphi_B 0$, Section A.1). Despite the presence of strong fluctuations in the subluminal shock for the in-plane configuration, its overall structure remains similar to the out-of-plane case. The substantial difference in the general magnetosonic wave morphology (e.g., the number of strong peaks) between the in-plane and out-of-plane simulations at $\theta_{Bn} = 30^\circ$ must therefore result from stronger dissipation of the dispersive wave train in simulation $\theta 30\text{-}\varphi_B 0$. Although a detailed investigation of the underlying cause is beyond the scope of this work, it is noted that the particle-acceleration efficiency in both simulations is comparable, as discussed next.

Figure A.2 compares phase-space distributions at $t\Omega_i = 188.5$ for momentum compo-




Figure A.7: Phase-space for ions (rows 1 and 2 from top) and electrons (rows 3 and 4) for selected momentum components at $t\Omega_i = 188.5$. The left panels show results for run $\theta30\text{-}\varphi_B0$, while the right panels for simulation $\theta30$.

nents along B_0 -direction: p_x and p_y for the in-plane configuration (left panels) and p_x and p_z for the out-of-plane configuration (right panels). Figure A.8 shows the corresponding ion (left) and electron (right) spectra in the shock region, $x/\lambda_{si} \approx 75\text{--}92$, for both simulations. The maximum ion energies reached in the wave pulse are comparable in both simulations, which is also reflected in the ion spectra. Similar overall shapes of these spec-




Figure A.8: Comparison of ion (left) and electron (right) normalised spectra at the shock for simulations $\theta30\text{-}\varphi_{B0}$ and $\theta30$ at $t\Omega_i = 188.5$.

tra indicate that a nearly equal number of ions undergo acceleration. The highest energies are produced ahead of the shock front via SSA, while most particles gain energy through trapping and acceleration within the wave train. However, unlike simulation $\theta30$, where energetic ions populate several consecutive potential troughs – with both phase-space density and maximum momentum gradually decreasing with distance from the wave front – in the in-plane simulation $\theta30\text{-}\varphi_{B0}$ they are concentrated primarily in the first trough, with a smaller contribution from the second. Ions accelerated in these locations subsequently diffuse downstream, where the weaker potential wells cannot trap them, and are scattered by turbulence. The nearly flat portion of the energetic ion spectrum in run $\theta30\text{-}\varphi_{B0}$ is consistent with a dominant contribution from a single potential trough, analogous to the case of electrons in simulation $\theta30$ discussed in Section 6.3.4.

Both the electron phase-space distributions and the spectra show that electrons are accelerated more efficiently in the out-of-plane configuration, where they reach nearly twice the maximum momentum obtained in the in-plane run. As with ions, electron acceleration in simulation $\theta30\text{-}\varphi_{B0}$ occurs predominantly within the first potential peak of the wave, producing a nearly flat accelerated spectrum. However, an analysis of the time evolution of the wave structure and the electron phase-space (not shown) reveals that the wave occasionally develops several strong oscillations locally along its transverse extent. These transient enhancements in the potential facilitate electron trapping and acceleration in multiple peaks. Once these structures are destroyed, the electrons are released and subsequently scatter, forming a broad thermal-like phase-space distribution downstream of the first two peaks and a corresponding broadened spectral component. Similar episodes of multi-peak acceleration also occur for ions, though with much lower efficiency. In contrast, in simulation $\theta30$ more electrons remain trapped for extended periods in the wave potential. Consequently, fewer are advected downstream, which results in a much narrower thermal-like component in both phase space and spectrum.




Figure A.9: Profiles of E_x/E_0 , $\tilde{\phi}$ and B_z/B_0 obtained in simulations $\theta 30$ and $\theta 30\text{-}m400$ with different ion-to-electron mass ratios, $m_i/m_e = 100$ and $m_i/m_e = 400$, respectively, at time $t\Omega_i \approx 69$.

A.4 Simulation $\theta 30\text{-}m400$

Figure A.9 compares profiles of E_x/E_0 , $\tilde{\phi}$, and B_z/B_0 obtained in simulations $\theta 30$ and $\theta 30\text{-}m400$ at similar times, $t\Omega_i = 68.84$ and $t\Omega_i = 69.10$, respectively. In both simulations, the total magnetisation is $\sigma = 1.0$, but the ion-to-electron mass ratios differ, $m_i/m_e = 100$ and $m_i/m_e = 400$. Consequently, the ion magnetisation remains unchanged ($\sigma_i \approx 1$), while the electron magnetisation differs, with $\sigma_e \approx 100$ and $\sigma_e \approx 400$, respectively. To maintain the same total magnetisation, the numerical values of the background magnetic field strength have been adjusted accordingly.

Both simulations exhibit very similar overall profiles and shock speeds. While the leading shock pulse has a comparable amplitude, the subsequent pulses appear weaker for the higher mass ratio, resulting in a smoother electrostatic potential. Note, however, that the comparison is made at a relatively early stage of the magnetosonic pulse development, when the wave structure has not yet fully evolved. The observed differences are also compatible with variability in the field components and the electric potential caused by nonstationarities and fluctuations in the transverse wave structure, which are present in run $\theta 30\text{-}m400$ with properties similar to those observed in runs $\theta 30$ and $\theta 30\text{-}\varphi_{B0}$.

Figure A.10 compares phase-space for the momentum components p_x and p_z aligned with the upstream magnetic field. Ion phase-space distributions are similar in both simulations, and in particular the maximum attained momenta are comparable. This is also




Figure A.10: Phase-space for ions (rows 1 for $p_{x,i}$ and 2 for $p_{z,i}$ from top) and electrons (rows 3 for $p_{x,e}$ and 4 for $p_{z,e}$) at $t\Omega_i \approx 69$. The left panels show results for run $\theta 30$ - $m400$, while the right panels for simulation $\theta 30$.

reflected in the nearly identical cutoff energies in the ion spectra shown in Figure A.11 (left). The widths of phase-space distributions downstream of the leading wave pulses at $x/\lambda_{si} \lesssim 25$ are comparable in both simulations. Consequently, the thermal-like components of the ion spectra are also similar.




Figure A.11: Comparison of ion (left) and electron (right) normalised spectra at the shock for simulations $\theta30\text{-}m400$ and $\theta30$ at $t\Omega_i \approx 69$. Particle kinetic energies are normalised to the upstream ion kinetic energy.

The only significant difference in the ion phase-space distributions is the number of ions trapped within the potential troughs, which is larger in simulation $\theta30\text{-}m400$, e.g., near $x/\lambda_{si} \approx 32$ and $x/\lambda_{si} \approx 29$. In run $\theta30$, only the first and third peaks are significantly populated with accelerated ions. Consequently, while the higher-mass-ratio case exhibits a nearly flat energetic spectral tail in the normalised ion energy range $\tilde{E}_{k,i} = E_{k,i}/(\gamma_{up}m_i c^2) \approx 1\text{--}10$, in simulation $\theta30$ a limited-energy power-law portion is formed at $\tilde{E}_{k,i} \approx 1\text{--}2$, with a spectral hump at higher energies, $\tilde{E}_{k,i} \approx 5\text{--}10$, corresponding to ions accelerated in the first trough. However, as noted above, these differences are transient, caused by early-phase wave disturbances, and over time the energetic ion spectra in both runs are expected to converge to power-laws, as shown for run $\theta30$ in Figure 6.23 in Section 6.3.6.

Electron phase-space distributions shown in the third and fourth row of Figure A.10 differ more substantially. The higher electron magnetisation in the $m_i/m_e = 400$ case allows electrons crossing the shock to gain more energy, within the wave pulses, resulting in broader momentum distributions. The difference in the maximum attained electron momenta between the two runs is approximately a factor of four, which is proportional to the ratio of the ion masses in both simulations. This scaling with the ion mass is evident in the scaled electron spectra presented in Figure A.11 (right). One can see that, already at this early evolution stage, the energetic electron spectra obtained for both mass ratios show a clear power-law dependence with a spectral index of $p \approx 2$, consistent with the late stage discussed in Section 6.3.4.

In conclusion, the results of simulation $\theta30\text{-}m400$ confirm that the ion and electron acceleration processes at oblique subluminal shocks discussed in Section 6.3 are essentially independent of the mass ratio and scale approximately linearly with m_i/m_e .




Figure A.12: Maps of different quantities obtained in simulation $\theta 10\text{-}\varphi_B 0$ at $t\Omega_i = 188.5$: charge density (a), high-pass filtered E_x electric field (b), E_y electric field (c), and B_y magnetic field (d) at the shock and in the near-upstream (compare Figs. 6.36 and 6.27). A high-pass filter is applied to E_x to eliminate the low-frequency whistler waves.

A.5 Simulation $\theta 10\text{-}\varphi_B 0$

Figure A.12 shows maps of the charge density, n_c/n_0 , the E_x/E_0 and E_y/E_0 electric fields, and B_y magnetic field at the end of simulation $\theta 10\text{-}\varphi_B 0$, with the background magnetic field at an angle $\theta = 10^\circ$ lying in the $x - y$ simulation plane, $\varphi_B = 0^\circ$. This should be compared with Figures 6.36 and 6.27 in Section 6.4, obtained for simulation $\theta 10$ with $\varphi_B = 90^\circ$. Note that a high-pass filter is applied to the E_x electric field and thus the whistler waves are not visible in this map.

The shock front is located at $x \approx 59\lambda_{si}$, where the magnetic field is compressed (Figure A.12d), and is preceded by a train of precursor waves similar to those observed in simulation $\theta 10$. The amplitude of the precursor waves varies with the distance from the shock front, ranging from $\delta B_y/B_0 \approx 0.6$ and $\delta B_z/B_0 \approx 0.35$ near the shock front ($x \approx 74\lambda_{si}$) down to $\delta B_y/B_0 \approx 0.4$ and $\delta B_z/B_0 \approx 0.3$ at the end of the precursor region ($x \approx 132\lambda_{si}$). These values are comparable to the precursor-wave amplitudes in simulation $\theta 10$, although in this case the largest amplitudes are associated with B_y rather than B_z , consistent with the in-plane magnetic-field configuration.

The shock Mach number in the upstream frame of reference is $M_{A,\text{up}} \approx 1.042$. Since the whistler critical Mach number, M_w , is the same as in simulation $\theta 10$ ($M_{w,\text{up}} \approx 1.09$), the shock in simulation $\theta 10\text{-}\varphi_B 0$ is also subcritical with respect to M_w , and the precursor waves are therefore dispersive whistlers. Because these whistlers are emitted by the shock –




Figure A.13: Phase-space for ions (rows 1 and 2 from top) and electrons (rows 3 and 4) for selected momentum components at $t\Omega_i = 188.5$. The left panels show results for run $\theta 10\text{-}\varphi_B 0$, while the right panels for simulation $\theta 10$.

which propagates strictly along the x -direction – their propagation direction is aligned with the shock normal and does not depend on the orientation of the background magnetic field. Consequently, these dispersive whistlers are elliptically polarised, although their polarisation is close to circular.

Also in the in-plane configuration oblique waves driven by the MTSI are observed in E_x (Figure A.12b) and E_y (Figure A.12c), with properties similar to those of the waves shown in Figure 6.36b for simulation $\theta 10$. Their orientation is correlated with the amplitude of the




Figure A.14: Comparison of ion (left) and electron (right) normalised spectra for the region $x/\lambda_{si} = 40\text{--}130$ for simulations $\theta 10\text{-}\varphi_{B0}$ and $\theta 10$ at $t\Omega_i \approx 188.5$.

in-plane magnetic field component B_y . In simulation $\theta 10\text{-}\varphi_{B0}$, when B_y is maximum, the waves are oblique, with $k_x c/\Omega_i \approx -4$ and $k_y c/\Omega_i \approx 7.5$. When B_y is at a minimum, the waves become nearly aligned with the y -direction, with $k_x c/\Omega_i \approx 1$ and $k_y c/\Omega_i \approx 7.5$. Nevertheless, these waves remain oblique with respect to the background magnetic field, $\mathbf{B}_0 = (B_{0x}, B_{0y}, 0)$, with $\theta_{Bk} \approx 73^\circ$ and $\theta_{Bk} \approx 107^\circ$. Thus, their properties are qualitatively similar in simulations $\theta 10\text{-}\varphi_{B0}$ and $\theta 10$, confirming that these oblique waves are excited by the MTSI.

As a consequence of the different orientation of the background magnetic field, the density filaments in Figure A.12 are now oblique with respect to the y -axis. In fact, they are aligned with the background magnetic field.

In addition to the precursor whistlers and the oblique MTSI-driven waves, the small-scale waves with $\lambda \approx 0.3\lambda_{si}$, excited by the PDI, are observed as in simulation $\theta 10$ (see Section 6.4.3). Two daughter electromagnetic waves generated through decay of the parent whistler wave are visible in Figures A.12c and A.12d, and an ion acoustic wave in Figure A.12a. The longitudinal electric field, E_x , is also perturbed as required by Gauss law.

Figure A.13 compares phase-space along the direction of the background magnetic field: p_x and p_y for the in-plane case (left panels) and p_x and p_z for the out-of-plane case (right panels). Overall, the phase-space structure is similar in both cases, showing the same localised heating due to interaction with the whistler waves, as well as particle interactions with the small-scale waves. The only visible distinction is slightly stronger heating in the in-plane configuration, both downstream and upstream. However, spectra shown in Figure A.14 are remarkably similar between the two simulations, indicating that the processes responsible for plasma heating are largely independent of the magnetic field inclination, φ_{B0} .

A.6 Simulation $\theta 10$ -1D

Figure A.15 compares single-line profiles of the ion density and the magnetic field fluctuations at $t\Omega_i = 163.33$ obtained from the 1D simulation $\theta 10$ -1D and the 2D run $\theta 10$. Precursor waves in front of the magnetic pulse at $x/\lambda_{si} \approx 50$ are present in both simulations, with comparable amplitudes ($\delta B_z/B_0 \approx 0.4$). The shock speed in simulation $\theta 10$ -1D is similar to that in simulation $\theta 10$ ($v_{sh} \approx 0.33c$), and therefore the upstream Alfvénic Mach number is approximately the same in both cases ($M_{A,up} \approx 1.035$). The shock in the 1D simulation is thus subcritical with respect to the Whistler critical Mach number $M_w = 1.09$, and the precursor waves are identified as dispersive whistlers.

In both simulations, small-scale waves are superimposed on the whistler wave structure, with similar wavelengths in the two runs, $k_x c/\Omega_i \approx -17$ in simulation $\theta 10$ -1D and $k_x c/\Omega_i \approx -20$ in simulation $\theta 10$. In 2D simulations, these waves were identified as arising from the PDI. It is therefore expected that the same instability excites these waves in simulation $\theta 10$ -1D, since the PDI requires only a single spatial dimension to operate.

Although both dispersive whistlers and small-scale waves are generated in both simulations, no oblique waves associated with the MTSI are present in simulation $\theta 10$ -1D, as this instability requires at least two spatial dimensions. The effect of the MTSI in the 2D simulation is visible in the density profile in Figure A.15 (top panel), where the orange curve shows pronounced peaks close to the shock front associated with the undulations due to density filaments. Peaks of such high amplitudes are absent in simulation $\theta 10$ -1D (blue line).




Figure A.15: Profiles for simulations $\theta 10$ -1D and $\theta 10$ at $t\Omega_i = 163.3$. The profile of the two-dimensional simulations is a single line at $y/\lambda_{si} = 2.9$. Top panel: normalised ion density n_i/n_0 for both simulations. Middle and Bottom panels: normalised transverse magnetic field B_y/B_0 and $\delta B_z/B_0$ for simulation $\theta 10$ -1D and $\theta 10$, respectively.




Figure A.16: Phase-space for ions (rows 1 and 2 from top) and electrons (rows 3 and 4) for selected momentum components at $t\Omega_i = 163.3$. The left panels show results for run $\theta 10$ -1D, while the right panels for simulation $\theta 10$.

A noticeable feature that appears only in simulation $\theta 10$ -1D is an overdensity in the ion density profile at $x/\lambda_{si} \approx 85 - 100$ (Figure A.15; top panel), which is co-spatial with a wave packet of enhanced amplitude visible in the magnetic field of (middle panel) over the same spatial range. This wave packet exerts a ponderomotive force that pushes ions forward, leading to their accumulation within the packet, as reflected in the ion density profile. The packet forms in early in the shock formation process.




Figure A.17: Comparison of ion (left) and electron (right) normalised spectra for the region $x/\lambda_{si} = 40\text{--}100$ for simulations $\theta 10\text{-}1\text{D}$ and $\theta 10$ at $t\Omega_i \approx 163.3$.

Figure A.16 compares the x - and z -components of ion and electron momenta for simulation $\theta 10\text{-}1\text{D}$ (left panels) and $\theta 10$ (right panels). The interaction of reflected ions and electrons with the whistlers is visible in both simulations, as well as the particle trapping by the E_x wave potential. Although particle heating may appear weaker in simulation $\theta 10\text{-}1\text{D}$ due to the narrower momentum distribution, this is a consequence of the smaller number of particles in the one-dimensional simulation that can reach higher energies. In fact, the energy spectra presented below show similar distributions in both simulations.

A notable feature in run $\theta 10\text{-}1\text{D}$ is the enhanced energisation observed through the x -component of the reflected ion beam at $x \approx 90\lambda_{si}$, which is caused by the high-amplitude wave packet.

The interaction of the small-scale waves with both ions and electrons appears qualitatively similar in the two simulations, and is visible as narrow spikes in the upstream ion and electron beams. This effect is, however, more visible in 1D simulation due to the colder upstream beam in this case.

Particle spectra, compared for both simulations in Figure A.17, provide further insight into the degree of heating. While the overall shapes of the spectra are broadly comparable between the two simulations, distinguishing between thermal and non-thermal components is challenging, as the particle distributions are not isotropised and the spectra deviate significantly from a Maxwellian. One remarkable difference can nevertheless be observed in the ion spectrum, which in simulation $\theta 10\text{-}1\text{D}$, shows a plateau at low energies ($0.2 < \gamma - 1 < 0.9$). This feature corresponds to the upstream ion beam as it decelerates while crossing the wave packet at $x \approx 90\lambda_{si}$. The same behaviour can be identified in phase space at the same location, where the upstream ion momentum $p_{x,i}$ increases temporarily before returning to approximately its upstream value, $p_{x,i}/m_i c \approx -2$.

APPENDIX B


Wave polarisation convention

The numerical analysis of the circularly polarised waves presented in section 6.4.2 uses the convention used in Terasawa et al. 1986 and is summarised below:

For a magnetic wave propagating in the positive x -direction, $\mathbf{B} = B_y \hat{j} + B_z \hat{k}$, the rotation of the vector at a fixed position is *right-hand polarised* if the vector's rotation is positive, $d\theta/dt > 0$ with $\tan \theta = B_z/B_y$, and has a *right-hand helicity* since the wave has positive wavenumber, $k > 0$. This behaviour can be visualised in Fourier space for all real ω and k , as illustrated in Figure B.1. The polarisation of the wave is determined by the sign of the frequency ω : if $\omega > 0$, the wave is right-hand polarised, while if $\omega < 0$, it is left-hand polarised. The helicity, on the other hand, is determined by the sign of the wavenumber k : if $k < 0$, the wave has right-hand helicity, whereas if $k > 0$, it has left-hand helicity. The labels L/R indicate the handedness of the wave, and the signs $+/-$ indicate the direction of propagation along k .

Using this convention, it is convenient to define the quantity $\tilde{B} = B_y - iB_z$ (here, the normalisation factor $1/\sqrt{2}$ has been omitted for simplicity). A right-hand circularly polarised wave produces a spectrum with a signal in the top-right quadrant, consistent with Figure B.1. This is particularly useful for visualisation purposes. This convention is closely related to the Jones vector formalism (Jones 1941).

Figure B.1: $\omega - k$ diagram where k is the wavenumber and ω the frequency. The wave is right-hand polarised if $\omega > 0$ and left-hand polarised for $\omega < 0$. The wave has right-hand helicity if $k < 0$ and has left-hand helicity if $k > 0$. The letter L/R indicates the handedness while the sign $+/-$ indicates the direction of propagation along k . This diagram is based on the diagram from Terasawa et al. 1986.



APPENDIX C

Dispersion relation for a finite amplitude circularly polarised wave propagating along the magnetic field. Charge and current neutrality

The dispersion relation for a finite-amplitude circularly polarised wave propagating along the magnetic field in a relativistic cold plasma is

$$g(\omega, k) = \omega^2 - k^2 c^2 - \sum_j \omega_{p,j}^2 \frac{\omega - kv_j}{\omega - kv_j + \Omega_j/\gamma_j} = 0, \quad (\text{C.1})$$

where the sum runs over all plasma species' j , v_j is the species velocity along the field, $\gamma_j(v_j)$ is the corresponding Lorentz factor, and Ω_j is the non-relativistic cyclotron frequency (Kennel and Pellat 1976). This equation is a generalisation of Eq. 6.14 used to obtain the whistler critical Mach number. For parallel propagation, the cyclotron frequency is Lorentz invariant. In the quasi-parallel case, $\theta_{Bn} \gtrsim 0$, the frequency is modified by a factor $\cos \theta_{Bn}$. Since density perturbations are neglected in the derivation of this dispersion relation, finite and isotropic temperature effects can also be incorporated (Amano 2016).

The derivation follows that in Kennel and Pellat 1976, with the modification that the species mass includes the specific enthalpy, $\bar{m}_j = h_j m_j$, where

$$h_j = \frac{w_j}{\rho_j c^2} = 1 + \frac{\Gamma}{\Gamma - 1} \frac{k_B T_j}{m_j c^2} \quad (\text{C.2})$$

is the specific enthalpy of species j . A finite temperature increase can be understood as an increase in internal energy and thus an increase in the inertia of the fluid, which makes it less responsive to perturbations. Effectively, this corresponds to a more massive fluid. Accordingly, the plasma and cyclotron frequencies are modified as $\bar{\omega}_{p,j}^2 = \omega_{p,j}^2/h_j$, $\bar{\Omega}_j = \Omega_j/h_j$.

Because density perturbations are neglected, the longitudinal electric field caused by such perturbations is zero, $E_x = 0$. In the absence of such longitudinal component, the enthalpy is a constant, simplifying the derivation greatly.

It should be noted that the derivation of equation C.1 assumes charge and current neutrality; otherwise, a zero-order electrostatic field would need to be included. This neutrality condition can be recovered by performing a first-order Taylor expansion of the dispersion relation around $\omega = 0$ and $k = 0$:

$$g(\omega, k) \approx g(0, 0) + \omega \left. \frac{\partial g}{\partial \omega} \right|_{(0,0)} + k \left. \frac{\partial g}{\partial k} \right|_{(0,0)} = -\omega \sum_j \frac{\omega_{p,s}^2}{\Omega_j/\gamma_j} + k \sum_j v_j \frac{\omega_{p,s}^2}{\Omega_j/\gamma_j} = 0. \quad (\text{C.3})$$

APPENDIX C. DISPERSION RELATION FOR A FINITE AMPLITUDE CIRCULARLY POLARISED WAVE PROPAGATING ALONG THE MAGNETIC FIELD. CHARGE AND CURRENT NEUTRALITY

Discarding the trivial solution $\omega = k = 0$, the neutrality conditions are obtained:

$$\text{Charge neutrality : } \sum_j \gamma_j \frac{\omega_{p,s}^2}{\Omega_j} = 0 \quad (\text{C.4})$$

$$\text{Current neutrality : } \sum_j \gamma_j \frac{\omega_{p,s}^2}{\Omega_j} v_j = 0, \quad (\text{C.5})$$

where the plasma and cyclotron frequencies may be swapped by their finite-temperature counterparts, $\bar{\omega}_{p,j}$ and $\bar{\Omega}_j$. This derivation can be interpreted physically: a wave with $\omega = 0$ and $k = 0$ corresponds to a constant field in space and time, representing the unperturbed background. If the plasma is initially in equilibrium (charge and current neutral) and no perturbations are present, the plasma remains in equilibrium, consistent with $\omega = k = 0$.

In simulation $\theta 10$, the upstream ion and hot ion populations are well-defined and their parameters can be computed directly. The electron distribution, however, is more complex, making direct computation of its parameters non-trivial. Nevertheless, equations C.4 and C.5 can be used to derive the electron velocity, plasma frequency, and cyclotron frequency.

Bibliography

- Abbott, B. P. et al. (Oct. 2017). “Multi-messenger Observations of a Binary Neutron Star Merger”. In: *ApJ* 848.2, L12, p. L12.
- Abdo, A. A. et al. (Oct. 2013). “The Second Fermi Large Area Telescope Catalog of Gamma-Ray Pulsars”. In: *ApJS* 208.2, 17, p. 17.
- Abeysekara, A. U. et al. (July 2017). “Observation of the Crab Nebula with the HAWC Gamma-Ray Observatory”. In: *ApJ* 843.1, 39, p. 39.
- Abeysekara, A. U. et al. (Oct. 2018). “Very-high-energy particle acceleration powered by the jets of the microquasar SS 433”. In: *Nature* 562.7725, pp. 82–85.
- Abeysekara, A. U. et al. (July 2023). “The High-Altitude Water Cherenkov (HAWC) observatory in México: The primary detector”. In: *Nuclear Instruments and Methods in Physics Research A* 1052, 168253, p. 168253.
- Adams, Fred C., Katherine Freese, Gregory Laughlin, Nathan Schwadron, and Gregory Tarlé (Dec. 1997). “Constraints on the Intergalactic Transport of Cosmic Rays”. In: *The Astrophysical Journal* 491.1, p. 6.
- Adams, Scott M., C. S. Kochanek, John F. Beacom, Mark R. Vagins, and K. Z. Stanek (Dec. 2013). “Observing the Next Galactic Supernova”. In: *ApJ* 778.2, 164, p. 164.
- Alfaro, R. et al. (Oct. 2024). “Ultra-high-energy gamma-ray bubble around microquasar V4641 Sgr”. In: *Nature* 634.8034, pp. 557–560.
- Alfaro, R. et al. (2025). “A measurement of the all-particle energy spectrum of cosmic rays from 1013 to 1015eV using HAWC”. In: *Astroparticle Physics* 167, p. 103077.
- Alispach, C. et al. (Feb. 2025). “The SST-1M imaging atmospheric Cherenkov telescope for gamma-ray astrophysics”. In: *J. Cosmology Astropart. Phys.* 2025.2, 047, p. 047.
- Amano, Takanobu (Nov. 2016). “A Second-order Divergence-constrained Multidimensional Numerical Scheme for Relativistic Two-fluid Electrodynamics”. In: *ApJ* 831.1, 100, p. 100.
- Amano, Takanobu and Masahiro Hoshino (Oct. 2009). “Nonlinear evolution of Buneman instability and its implication for electron acceleration in high Mach number collisionless perpendicular shocks”. In: *Physics of Plasmas* 16.10, 102901, p. 102901.
- Amano, Takanobu et al. (Dec. 2022). “Nonthermal Electron Acceleration at Collisionless Quasi-Perpendicular Shocks”. In: *Reviews of Modern Plasma Physics* 6, p. 29.
- Ardaneh, Kazem, Dongsheng Cai, Ken-Ichi Nishikawa, and Bertrand Lembége (Sept. 2015). “COLLISIONLESS WEIBEL SHOCKS AND ELECTRON ACCELERATION IN GAMMA-RAY BURSTS”. In: *The Astrophysical Journal* 811.1, p. 57.
- Axford, W. I. (Feb. 1994). “The Origins of High-Energy Cosmic Rays”. In: *ApJS* 90, p. 937.
- Axford, W. I., E. Leer, and G. Skadron (Jan. 1977). “The Acceleration of Cosmic Rays by Shock Waves”. In: *International Cosmic Ray Conference*. Vol. 11. International Cosmic Ray Conference, p. 132.
- Baade, W. and F. Zwicky (July 1934). “Remarks on Super-Novae and Cosmic Rays”. In: *Physical Review* 46.1, pp. 76–77.
- Baixeras, C. et al. (Feb. 2004). “Commissioning and first tests of the MAGIC telescope”. In: *Nuclear Instruments and Methods in Physics Research A* 518.1-2, pp. 188–192.
- Ball, Lewis and D. B. Melrose (2001). “Shock Drift Acceleration of Electrons”. In: *Publications of the Astronomical Society of Australia* 18.4, pp. 361–373.

- Balogh, André and Rudolf A. Treumann (2013). *Physics of Collisionless Shocks: Space Plasma Shock Waves*. Vol. 12.
- Bardeen, James M. and Jacobus A. Petterson (Jan. 1975). “The Lense-Thirring Effect and Accretion Disks around Kerr Black Holes”. In: *ApJ* 195, p. L65.
- Batzofin, Rowan, Pierre Cristofari, Kathrin Egberts, Constantin Steppa, and Dominique M. - A. Meyer (July 2024). “The population of Galactic supernova remnants in the TeV range”. In: *A&A* 687, A279, A279.
- Bell, A. R. (Jan. 1978a). “The acceleration of cosmic rays in shock fronts - I.” In: *MNRAS* 182, pp. 147–156.
- (Feb. 1978b). “The acceleration of cosmic rays in shock fronts - II.” In: *MNRAS* 182, pp. 443–455.
- (Sept. 2004). “Turbulent amplification of magnetic field and diffusive shock acceleration of cosmic rays”. In: *MNRAS* 353.2, pp. 550–558.
- Bessho, Naoki and Yukiharu Ohsawa (Aug. 1999). “Electron Acceleration to Ultrarelativistic Energies in a Collisionless Oblique Shock Wave”. In: *Physics of Plasmas* 6.8, pp. 3076–3085.
- (Oct. 2000). “Energies of Ultrarelativistic Electrons Produced by an Oblique Shock Wave”. In: *Physics of Plasmas* 7.10, pp. 4004–4012.
- (Mar. 2002). “Electron Acceleration and Trapping by an Oblique Shock Wave”. In: *Physics of Plasmas* 9.3, pp. 979–986.
- Bethe, H. A. (Oct. 1990). “Supernova mechanisms”. In: *Reviews of Modern Physics* 62.4, pp. 801–866.
- Bhattacharjee, Pijushpani and Günter Sigl (2000). “Origin and propagation of extremely high-energy cosmic rays”. In: *Physics Reports* 327.3, pp. 109–247.
- Birdsall, C. K. and A. B. Langdon (1991). *Plasma Physics via Computer Simulation*.
- Blandford, R. D. and A. Königl (Aug. 1979). “Relativistic jets as compact radio sources.” In: *ApJ* 232, pp. 34–48.
- Blandford, R. D. and J. P. Ostriker (Apr. 1978). “Particle acceleration by astrophysical shocks.” In: *ApJ* 221, pp. L29–L32.
- Blandford, Roger, David Meier, and Anthony Readhead (Aug. 2019). “Relativistic Jets from Active Galactic Nuclei”. In: *ARA&A* 57, pp. 467–509.
- Blasi, Pasquale (Nov. 2013). “The origin of galactic cosmic rays”. In: *A&A Rev.* 21, 70, p. 70.
- Blasi, Pasquale and Elena Amato (Jan. 2012). “Diffusive propagation of cosmic rays from supernova remnants in the Galaxy. I: spectrum and chemical composition”. In: *J. Cosmology Astropart. Phys.* 2012.1, 010, p. 010.
- Bohdan, Artem, Jacek Niemiec, Oleh Kobzar, and Martin Pohl (Sept. 2017). “Electron Pre-acceleration at Nonrelativistic High-Mach-number Perpendicular Shocks”. In: *ApJ* 847.1, 71, p. 71.
- Bohdan, Artem, Jacek Niemiec, Martin Pohl, Yosuke Matsumoto, Takanobu Amano, and Masahiro Hoshino (June 2019a). “Kinetic Simulations of Nonrelativistic Perpendicular Shocks of Young Supernova Remnants. I. Electron Shock-surfing Acceleration”. In: *ApJ* 878.1, 5, p. 5.

-
- (Nov. 2019b). “Kinetic Simulations of Nonrelativistic Perpendicular Shocks of Young Supernova Remnants. II. Influence of Shock-surfing Acceleration on Downstream Electron Spectra”. In: *The Astrophysical Journal* 885.1, p. 10.
- Boris, Jay P (1970). “Relativistic plasma simulation-optimization of a hybrid code”. In: *Proc. Fourth Conf. Num. Sim. Plasmas*, pp. 3–67.
- Böttcher, Markus and Matthew G. Baring (Dec. 2019). “Multi-wavelength Variability Signatures of Relativistic Shocks in Blazar Jets”. In: *ApJ* 887.2, 133, p. 133.
- Box, G. E. P. and Mervin E. Muller (1958). “A Note on the Generation of Random Normal Deviates”. In: *The Annals of Mathematical Statistics* 29.2, pp. 610–611.
- Bromberg, Omer, Ehud Nakar, Tsvi Piran, and Re’em Sari (Oct. 2011a). “The Propagation of Relativistic Jets in External Media”. In: *ApJ* 740.2, 100, p. 100.
- (Oct. 2011b). “The Propagation of Relativistic Jets in External Media”. In: *ApJ* 740.2, 100, p. 100.
- Buneman, O. (Aug. 1959). “Dissipation of Currents in Ionized Media”. In: *Physical Review* 115.3, pp. 503–517.
- Buneman, Oscar (1993). “TRISTAN : The 3-D, E-M Particle Code”. In: *Computer Space Plasma Physics : Simulation Techniques and Software*.
- Burgess, David and Manfred Scholer (July 2015). *Collisionless Shocks in Space Plasmas: Structure and Accelerated Particles*. 1st ed. Cambridge University Press.
- Canfield, E., W. M. Howard, and E. P. Liang (Dec. 1987). “Inverse Comptonization by one-dimensional relativistic electrons”. In: *ApJ* 323, pp. 565–574.
- Cao, Zhen et al. (July 2025). “Observation of the γ -ray emission from W43 with LHAASO”. In: *Science China Physics, Mechanics, and Astronomy* 68.7, 279502, p. 279502.
- Cary, John R. and Alain J. Brizard (Apr. 2009). “Hamiltonian Theory of Guiding-Center Motion”. In: *Reviews of Modern Physics* 81, pp. 693–738.
- Cesarsky, C. J. (Jan. 1980). “Cosmic-ray confinement in the galaxy”. In: *ARA&A* 18, pp. 289–319.
- Chandrasekhar, S. (Mar. 1943). “Dynamical Friction. I. General Considerations: the Coefficient of Dynamical Friction.” In: *ApJ* 97, p. 255.
- Coll, Bartolomé, Joan Josep Ferrando, and Juan Antonio Sáez (Apr. 2019). “Relativistic hydrodynamic approach to the classical ideal gas”. In: *Phys. Rev. D* 99 (8), p. 084035.
- Collaboration, DAMPE et al. (2019). “Measurement of the cosmic ray proton spectrum from 40 GeV to 100 TeV with the DAMPE satellite”. In: *Science Advances* 5.9, eaax3793.
- Courant, R., K. Friedrichs, and H. Lewy (Dec. 1928). “Über die partiellen Differenzengleichungen der mathematischen Physik”. In: *Mathematische Annalen* 100.1, pp. 32–74.
- Cowley, A. P., P. C. Schmidtke, J. B. Hutchings, and David Crampton (Apr. 2003). “A Spectroscopic and Photometric Study of the Eclipsing Low-mass X-Ray Binary 2A 1822-371 (V691 Coronae Australis)”. In: *AJ* 125.4, pp. 2163–2172.
- Crumley, P., D. Caprioli, S. Markoff, and A. Spitkovsky (June 2019). “Kinetic Simulations of Mildly Relativistic Shocks – I. Particle Acceleration in High Mach Number Shocks”. In: *Monthly Notices of the Royal Astronomical Society* 485.4, pp. 5105–5119.
- Curtis, H. D. (Jan. 1918). “Descriptions of 762 Nebulae and Clusters Photographed with the Crossley Reflector”. In: *Publications of Lick Observatory* 13, pp. 9–42.
-

- Dawson, John (Apr. 1962). "One-Dimensional Plasma Model". In: *Physics of Fluids* 5.4, pp. 445–459.
- Dawson, John M. (Apr. 1983). "Particle Simulation of Plasmas". In: *Reviews of Modern Physics* 55.2, pp. 403–447.
- de la Cita, V. M., S. del Palacio, V. Bosch-Ramon, X. Paredes-Fortuny, G. E. Romero, and D. Khangulyan (July 2017). "Gamma rays from clumpy wind-jet interactions in high-mass microquasars". In: *A&A* 604, A39, A39.
- DePasquale, Joseph, Kimberly Arcand, and Peter Edmonds (Sept. 2015). "High Energy Vision: Processing X-rays". In: *Studies in Media and Communication* 3.
- Drury, L. Oc. (Aug. 1983). "REVIEW ARTICLE: An introduction to the theory of diffusive shock acceleration of energetic particles in tenuous plasmas". In: *Reports on Progress in Physics* 46.8, pp. 973–1027.
- Edmiston, J. P. and C. F. Kennel (Dec. 1984). "A Parametric Survey of the First Critical Mach Number for a Fast MHD Shock". In: *Journal of Plasma Physics* 32, pp. 429–441.
- Engelmann, J. J. et al. (July 1990). "Charge composition and energy spectra of cosmic-ray nuclei for elements from Be to Ni - Results from HEAO-3-C2." In: *A&A* 233, pp. 96–111.
- Escobar, G. J., L. J. Pellizza, and G. E. Romero (June 2021). "Cosmic-ray production from neutron escape in microquasar jets". In: *A&A* 650, A136, A136.
- Fabrika, S. (Jan. 2004). "The jets and supercritical accretion disk in SS433". In: *Astrophys. Space Phys. Res.* 12, pp. 1–152.
- Fanaroff, B. L. and J. M. Riley (May 1974). "The morphology of extragalactic radio sources of high and low luminosity". In: *MNRAS* 167, 31P–36P.
- Faussurier, G  rald (July 2024). "The sound speed of the relativistic free electron gas". In: *Physics of Plasmas* 31.7, 072708, p. 072708.
- Fehske, H., R. Schneider, and A. Wei  sse (2008). *Computational Many-Particle Physics*. Vol. 739.
- Fermi, E. (Jan. 1954). "Galactic Magnetic Fields and the Origin of Cosmic Radiation." In: *ApJ* 119, p. 1.
- Fermi, Enrico (Apr. 1949). "On the Origin of the Cosmic Radiation". In: *Physical Review* 75.8, pp. 1169–1174.
- Friedman, Alex (1990). "A second-order implicit particle mover with adjustable damping". In: *Journal of Computational Physics* 90.2, pp. 292–312.
- Gaensler, Bryan M. and Patrick O. Slane (Sept. 2006). "The Evolution and Structure of Pulsar Wind Nebulae". In: *ARA&A* 44.1, pp. 17–47.
- Gaisser, Thomas K. (1990). *Cosmic rays and particle physics*.
- Gary, S. Peter, Charles W. Smith, Martin A. Lee, Melvyn L. Goldstein, and David W. Forslund (July 1984). "Electromagnetic Ion Beam Instabilities". In: *The Physics of Fluids* 27.7, pp. 1852–1862.
- Ghisellini, G., F. Tavecchio, and M. Chiaberge (Mar. 2005). "Structured jets in TeV BL Lac objects and radiogalaxies. Implications for the observed properties". In: *A&A* 432.2, pp. 401–410.
- Giacalone, Joe, Steven J. Schwartz, and David Burgess (Jan. 1993). "Observations of suprathermal ions in association with SLAMS". In: *Geophys. Res. Lett.* 20.2, pp. 149–152.

- Ginzburg, V. L. and S. I. Syrovatsky (Jan. 1961). “Origin of Cosmic Rays”. In: *Progress of Theoretical Physics Supplement* 20, pp. 1–83.
- Goldstein, M. L. (Jan. 1978). “An instability of finite amplitude circularly polarized Alfvén waves.” In: *ApJ* 219, pp. 700–704.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. MIT Press.
- Grasso, D. and H. R. Rubinstein (July 2001). “Magnetic fields in the early Universe”. In: *Phys. Rep.* 348.3, pp. 163–266.
- Grebogi, Celso and Robert G. Littlejohn (Aug. 1984). “Relativistic Ponderomotive Hamiltonian”. In: *Physics of Fluids* 27, pp. 1996–2004.
- Greenstadt, E. W. and M. M. Mellott (May 1987). “Plasma wave evidence for reflected ions in front of subcritical shocks: ISEE 1 and 2 observations”. In: *J. Geophys. Res.* 92.A5, pp. 4730–4734.
- Greenwood, Andrew D., Keith L. Cartwright, John W. Luginsland, and Ernest A. Baca (2004). “On the elimination of numerical Cerenkov radiation in PIC simulations”. In: *Journal of Computational Physics* 201.2, pp. 665–684.
- Greisen, Kenneth (Apr. 1966). “End to the Cosmic-Ray Spectrum?” In: *Phys. Rev. Lett.* 16.17, pp. 748–750.
- Grieder, Peter K.F. (2001). “Chapter 5 - Primary Cosmic Radiation”. In: *Cosmic Rays at Earth*. Ed. by Peter K.F. Grieder. Amsterdam: Elsevier, pp. 669–891.
- H.E.S.S. Collaboration* et al. (Jan. 2024). “Acceleration and Transport of Relativistic Electrons in the Jets of the Microquasar SS 433”. In: *Science* 383.6681, pp. 402–406.
- Halzen, Francis and Spencer R. Klein (Aug. 2010). “Invited Review Article: IceCube: An instrument for neutrino astronomy”. In: *Review of Scientific Instruments* 81.8, pp. 081101–081101.
- Heber (Jan. 1968). “Yu. L. Klimontovich, The Statistical Theory of Non-Equilibrium Processes in a Plasma. (International Series of Monographs in Natural Philosophy, Volume 9). XV + 283 S. Oxford/London/Edinburgh/New York/Toronto/Sydney/Paris/Braunschweig 1967. Preis geb. 70s. net”. In: *Zeitschrift Angewandte Mathematik und Mechanik* 48.4, pp. 291–292.
- Hess, Victor (2018). *On the Observations of the Penetrating Radiation during Seven Balloon Flights*.
- HESS Collaboration et al. (Mar. 2016). “Acceleration of petaelectronvolt protons in the Galactic Centre”. In: *Nature* 531.7595, pp. 476–479.
- Higuera, A. V. and J. R. Cary (May 2017). “Structure-preserving second-order integration of relativistic charged particle trajectories in electromagnetic fields”. In: *Physics of Plasmas* 24.5, 052104, p. 052104.
- Hinton, J. A. and HESS Collaboration (Apr. 2004). “The status of the HESS project”. In: *New A Rev.* 48.5-6, pp. 331–337.
- Hjorth, Jens et al. (June 2003). “A very energetic supernova associated with the γ -ray burst of 29 March 2003”. In: *Nature* 423.6942, pp. 847–850.
- Hockney, R. W. and J. W. Eastwood (1988). *Computer simulation using particles*.
- Holder, Jamie (Jan. 2011). “VERITAS: Status and Highlights”. In: *International Cosmic Ray Conference*. Vol. 11. International Cosmic Ray Conference, p. 137.

- Hollweg, Joseph V., R. Esser, and V. Jayanti (Mar. 1993). “Modulational and decay instabilities of Alfvén waves: Effects of streaming He^{++} ”. In: *J. Geophys. Res.* 98.A3, pp. 3491–3500.
- Hörandel, Jörg R. (June 2004). “Models of the knee in the energy spectrum of cosmic rays”. In: *Astroparticle Physics* 21.3, pp. 241–265.
- Hoshino, Masahiro (Jan. 2008). “Wakefield Acceleration by Radiation Pressure in Relativistic Shock Waves”. In: *ApJ* 672.2, pp. 940–956.
- Hoshino, Masahiro and Jonathan Arons (Mar. 1991). “Preferential positron heating and acceleration by synchrotron maser instabilities in relativistic positron-electron-proton plasmas”. In: *Physics of Fluids B* 3.3, pp. 818–833.
- Hoshino, Masahiro and Toshio Terasawa (Jan. 1985). “Numerical Study of the Upstream Wave Excitation Mechanism: 1. Nonlinear Phase Bunching of Beam Ions”. In: *Journal of Geophysical Research: Space Physics* 90.A1, pp. 57–64.
- IceCube Collaboration et al. (July 2018). “Multimessenger Observations of a Flaring Blazar Coincident with High-Energy Neutrino IceCube-170922A”. In: *Science* 361, eaat1378.
- Iwamoto, Masanori, Takanobu Amano, Masahiro Hoshino, and Yosuke Matsumoto (May 2017a). “Persistence of Precursor Waves in Two-dimensional Relativistic Shocks”. In: *The Astrophysical Journal* 840.1, p. 52.
- (May 2017b). “Persistence of Precursor Waves in Two-dimensional Relativistic Shocks”. In: *ApJ* 840.1, 52, p. 52.
- (May 2018). “Precursor Wave Emission Enhanced by Weibel Instability in Relativistic Shocks”. In: *ApJ* 858.2, 93, p. 93.
- Iwamoto, Masanori, Takanobu Amano, Yosuke Matsumoto, Shuichi Matsukiyo, and Masahiro Hoshino (Jan. 2022). “Particle Acceleration by Pickup Process Upstream of Relativistic Shocks”. In: *ApJ* 924.2, 108, p. 108.
- Iwamoto, Masanori, Yosuke Matsumoto, Takanobu Amano, Shuichi Matsukiyo, and Masahiro Hoshino (Jan. 2024). “Linearly Polarized Coherent Emission from Relativistic Magnetized Ion-Electron Shocks”. In: *Phys. Rev. Lett.* 132.3, 035201, p. 035201.
- Iwamoto, Masanori et al. (Oct. 2019). “Precursor Wave Amplification by Ion-Electron Coupling through Wakefield in Relativistic Shocks”. In: *ApJ* 883.2, L35, p. L35.
- Jokipii, J. R., J. Kóta, and J. Giacalone (1993). “Perpendicular transport in 1- and 2-dimensional shock simulations”. In: *Geophysical Research Letters* 20.17, pp. 1759–1761.
- Jones, R. Clark (July 1941). “New calculus for the treatment of optical systems. I. Description and discussion of the calculus”. In: *Journal of the Optical Society of America (1917-1983)* 31.7, p. 488.
- Jüttner, Ferencz (1911). “Das Maxwellsche Gesetz der Geschwindigkeitsverteilung in der Relativtheorie”. In: *Annalen der Physik* 339.5, pp. 856–882.
- Kaaret, Philip et al. (Oct. 2024a). “Magnetic Field Geometry of the Gamma-Ray Binary PSR B125963 Revealed via X-Ray Polarization”. In: *ApJ* 974.1, L1, p. L1.
- Kaaret, Philip et al. (Jan. 2024b). “X-Ray Polarization of the Eastern Lobe of SS 433”. In: *ApJ* 961.1, L12, p. L12.

- Kajita, Takaaki (May 1999). “Atmospheric neutrino results from Super-Kamiokande and Kamiokande - Evidence for ν_μ oscillations”. In: *Nuclear Physics B Proceedings Supplements* 77.1-3, pp. 123–132.
- Kampert, Karl-Heinz and Michael Unger (May 2012). “Measurements of the cosmic ray composition with air shower experiments”. In: *Astroparticle Physics* 35.10, pp. 660–678.
- Kang, Hyesung, Dongsu Ryu, and Ji-Hoon Ha (May 2019). “Electron Preacceleration in Weak Quasi-perpendicular Shocks in High-beta Intracluster Medium”. In: *The Astrophysical Journal* 876, p. 79.
- Kataoka, Jun and ukasz Stawarz (Apr. 2005). “X-Ray Emission Properties of Large-Scale Jets, Hot Spots, and Lobes in Active Galactic Nuclei”. In: *ApJ* 622.2, pp. 797–810.
- Kennel, C. F. and R. Pellat (June 1976). “Relativistic nonlinear plasma waves in a magnetic field”. In: *Journal of Plasma Physics* 15.3, pp. 335–355.
- Khatoon, Rukaiya, Markus Böttcher, and Raj Prince (Oct. 2024). “Modeling Multiband SEDs and Light Curves of BL Lacertae Using a Time-dependent Shock-in-jet Model”. In: *ApJ* 974.2, 233, p. 233.
- Kirk, J. G. and P. Duffy (Aug. 1999a). “TOPICAL REVIEW: Particle Acceleration and Relativistic Shocks”. In: *Journal of Physics G Nuclear Physics* 25, R163–R194.
- (Aug. 1999b). “TOPICAL REVIEW: Particle Acceleration and Relativistic Shocks”. In: *Journal of Physics G Nuclear Physics* 25, R163–R194.
- Kobzar, Oleh et al. (Sept. 2021). “Electron Acceleration at Rippled Low-mach-number Shocks in High-beta Collisionless Cosmic Plasmas”. In: *The Astrophysical Journal* 919.2, p. 97.
- Komissarov, S. S. (Feb. 1999). “A Godunov-type scheme for relativistic magnetohydrodynamics”. In: *MNRAS* 303.2, pp. 343–366.
- Kounine, Andrei (Aug. 2012). “The Alpha Magnetic Spectrometer on the International Space Station”. In: *International Journal of Modern Physics E* 21.8, 1230005, p. 1230005.
- Krasnoselskikh, V. V., B. Lembège, P. Savoini, and V. V. Lobzin (Apr. 2002). “Nonstationarity of strong collisionless quasiperpendicular shocks: Theory and full particle numerical simulations”. In: *Physics of Plasmas* 9.4, pp. 1192–1209.
- Krauss-Varban, D. and N. Omid (Oct. 1991). “Structure of Medium Mach Number Quasi-Parallel Shocks: Upstream and Downstream Waves”. In: *Journal of Geophysical Research* 96, pp. 17715–17731.
- Krauss-Varban, D. and C. S. Wu (Nov. 1989). “Fast Fermi and gradient drift acceleration of electrons at nearly perpendicular collisionless shocks”. In: *J. Geophys. Res.* 94.A11, pp. 15367–15372.
- Kruer, William L., ed. (Jan. 1988). *The physics of laser plasma interactions*. Vol. 73.
- Krymskii, G. F. (June 1977). “A regular mechanism for the acceleration of charged particles on the front of a shock wave”. In: *Soviet Physics Doklady* 22, p. 327.
- Kulsrud, Russell M. (Jan. 2005). *Plasma Physics for Astrophysics*. Princeton University Press.
- Kulsrud, Russell M. and Ellen G. Zweibel (Apr. 2008). “On the origin of cosmic magnetic fields”. In: *Reports on Progress in Physics* 71.4, 046901, p. 046901.
- Laing, R. A. and A. H. Bridle (June 2013). “The spectra of jet bases in FR I radio galaxies: implications for particle acceleration”. In: *MNRAS* 432.2, pp. 1114–1132.
- Landau, Lev Davidovich and E. M. Lifshitz (1959). *Fluid mechanics*.

- Langdon, A. Bruce, Jonathan Arons, and Claire E. Max (Aug. 1988). “Structure of relativistic magnetosonic shocks in electron-positron plasmas”. In: *Phys. Rev. Lett.* 61.7, pp. 779–782.
- Lembege, B. and J. M. Dawson (May 1989). “Relativistic particle dynamics in a steepening magnetosonic wave”. In: *Physics of Fluids B* 1.5, pp. 1001–1010.
- Lemoine, Martin and Guy Pelletier (Oct. 2011). “Dispersion and Thermal Effects on Electromagnetic Instabilities in the Precursor of Relativistic Shocks: Instabilities at Relativistic Shock Waves”. In: *Monthly Notices of the Royal Astronomical Society* 417.2, pp. 1148–1161.
- Lemoine, Martin, Guy Pelletier, Laurent Gremillet, and Illya Plotnikov (June 2014a). “A fast current-driven instability in relativistic collisionless shocks”. In: *EPL (Europhysics Letters)* 106.5, 55001, p. 55001.
- (May 2014b). “Current-driven filamentation upstream of magnetized relativistic collisionless shocks”. In: *MNRAS* 440.2, pp. 1365–1378.
- Leroy, M. M. and A. Mangeney (Aug. 1984). “A theory of energization of solar wind electrons by the Earth’s bow shock.” In: *Annales Geophysicae* 2, pp. 449–456.
- LHAASO Collaboration (Oct. 2024a). “Ultrahigh-Energy Gamma-ray Emission Associated with Black Hole-Jet Systems”. In: *arXiv e-prints*, arXiv:2410.08988, arXiv:2410.08988.
- (Oct. 2024b). “Ultrahigh-Energy Gamma-ray Emission Associated with Black Hole-Jet Systems”. In: *arXiv e-prints*, arXiv:2410.08988, arXiv:2410.08988.
- Ligorini, Arianna et al. (Mar. 2021a). “Mildly relativistic magnetized shocks in electron-ion plasmas - I. Electromagnetic shock structure”. In: *MNRAS* 501.4, pp. 4837–4849.
- Ligorini, Arianna et al. (Apr. 2021b). “Mildly relativistic magnetized shocks in electron-ion plasmas - II. Particle acceleration and heating”. In: *MNRAS* 502.4, pp. 5065–5074.
- Lindner, Robert R. et al. (May 2014). “The Radio Relics and Halo of El Gordo, a Massive $z = 0.870$ Cluster Merger”. In: *ApJ* 786.1, 49, p. 49.
- Liu, Wei, Yi-Qing Guo, and Qiang Yuan (Oct. 2019). “Indication of nearby source signatures of cosmic rays from energy spectra and anisotropies”. In: *J. Cosmology Astropart. Phys.* 2019.10, 010, p. 010.
- Lyubarsky, Yuri (Dec. 2006). “Electron-Ion Coupling Upstream of Relativistic Collisionless Shocks”. In: *ApJ* 652.2, pp. 1297–1305.
- Mallick, Ritam (Dec. 20, 2011). “Solving the Relativistic Rankine-Hugoniot Condition in the Presence of a Magnetic Field in the Astrophysical Scenario of a Neutron Star”. In: *Physical Review C* 84.6, p. 065805.
- Marsaglia, G. and T. A. Bray (1964). “A Convenient Method for Generating Normal Variables”. In: *SIAM Review* 6.3, pp. 260–264.
- Massi, M. (Jan. 2007). “Neutron stars in Microquasars”. In: *WE-Heraeus Seminar on Neutron Stars and Pulsars 40 years after the Discovery*. Ed. by W. Becker and H. H. Huang, p. 185.
- Matsukiyo, S. and Y. Matsumoto (Sept. 2015). “Electron Acceleration at a High Beta and Low Mach Number Rippled Shock”. In: *Journal of Physics Conference Series*. Vol. 642. Journal of Physics Conference Series. IOP, 012017, p. 012017.

- Matsukiyo, Shuichi and Manfred Scholer (Dec. 2003). “Modified Two-stream Instability in the Foot of High Mach Number Quasi-perpendicular Shocks”. In: *Journal of Geophysical Research: Space Physics* 108.A12.
- Matsumoto, Yosuke, Takanobu Amano, and Masahiro Hoshino (Aug. 2012a). “ELECTRON ACCELERATIONS AT HIGH MACH NUMBER SHOCKS: TWO-DIMENSIONAL PARTICLE-IN-CELL SIMULATIONS IN VARIOUS PARAMETER REGIMES”. In: *The Astrophysical Journal* 755.2, p. 109.
- (Aug. 2012b). “Electron Accelerations at High Mach Number Shocks: Two-dimensional Particle-in-cell Simulations in Various Parameter Regimes”. In: *ApJ* 755.2, 109, p. 109.
- Matthews, J. H., A. R. Bell, K. M. Blundell, and A. T. Araudo (Feb. 2019). “Ultrahigh energy cosmic rays from shocks in the lobes of powerful radio galaxies”. In: *MNRAS* 482.4, pp. 4303–4321.
- McKinney, Jonathan C., Alexander Tchekhovskoy, and Roger D. Blandford (July 2012). “General relativistic magnetohydrodynamic simulations of magnetically choked accretion flows around black holes”. In: *MNRAS* 423.4, pp. 3083–3117.
- Medvedev, Mikhail V. and Abraham Loeb (Dec. 1999). “Generation of Magnetic Fields in the Relativistic Shock of Gamma-Ray Burst Sources”. In: *ApJ* 526.2, pp. 697–706.
- Melrose, D. B., R. G. Hewitt, and G. A. Dulk (Feb. 1984). “Electron-cyclotron maser emission: Relative growth and damping rates for different modes and harmonics”. In: *J. Geophys. Res.* 89.A2, pp. 897–904.
- Miller-Jones, J. C. A., D. G. McCormick, R. P. Fender, R. E. Spencer, T. W. B. Muxlow, and G. G. Pooley (Nov. 2005). “Multiple relativistic outbursts of GRS1915+105: radio emission and internal shocks”. In: *MNRAS* 363.3, pp. 867–881.
- Miyahara, Seiji, Takahiro Kawashima, and Yukiharu Ohsawa (Jan. 2003). “Field Strengths in Oblique Shock Waves”. In: *Physics of Plasmas* 10, pp. 98–102.
- Mollerach, S. and E. Roulet (2018). “Progress in high-energy cosmic ray physics”. In: *Progress in Particle and Nuclear Physics* 98, pp. 85–118.
- Morlino, G. and D. Caprioli (Feb. 2012). “Strong evidence for hadron acceleration in Tycho’s supernova remnant”. In: *A&A* 538, A81, A81.
- Motch, C. (Oct. 1998). “The transient jet of the galactic supersoft X-ray source RX J0925.7-4758”. In: *A&A* 338, pp. L13–L16.
- Mueller, Dietrich, Simon P. Swordy, Peter Meyer, Jacques L’Heureux, and John M. Grunsfeld (June 1991). “Energy Spectra and Composition of Primary Cosmic Rays”. In: *ApJ* 374, p. 356.
- Nakazawa, Shinji and Yukiharu Ohsawa (July 1997). “Electromagnetic Fields in Nonlinear Magnetosonic Waves with Relativistic Propagation Speeds”. In: *Journal of the Physical Society of Japan* 66.7, p. 2044.
- Nariyuki, Y., S. Matsukiyo, and T. Hada (Aug. 2008). “Parametric instabilities of large-amplitude parallel propagating Alfvén waves: 2D PIC simulation”. In: *New Journal of Physics* 10.8, 083004, p. 083004.
- Navas, S. et al. (Aug. 2024). “Review of particle physics*”. In: *Phys. Rev. D* 110.3, 030001, p. 030001.

- Neronov, A., F. Oikonomou, and D. Semikoz (May 2025). “Multimessenger signature of cosmic rays from the microquasar V4641 Sgr propagating along a Galactic magnetic field line”. In: *Phys. Rev. D* 111.10, 103025, p. 103025.
- Netzer, Hagai (Aug. 2015). “Revisiting the Unified Model of Active Galactic Nuclei”. In: *ARA&A* 53, pp. 365–408.
- Niemiec, Jacek, Micha Ostrowski, and Martin Pohl (Oct. 2006). “Cosmic-Ray Acceleration at Ultrarelativistic Shock Waves: Effects of Downstream Short-Wave Turbulence”. In: *ApJ* 650.2, pp. 1020–1027.
- Niemiec, Jacek, Martin Pohl, Thomas Stroman, and Ken-Ichi Nishikawa (Sept. 2008). “Production of Magnetic Turbulence by Cosmic Rays Drifting Upstream of Supernova Remnant Shocks”. In: *ApJ* 684.2, pp. 1174–1189.
- Nishikawa, K.-I. et al. (2009). “WEIBEL INSTABILITY AND ASSOCIATED STRONG FIELDS IN A FULLY THREE-DIMENSIONAL SIMULATION OF A RELATIVISTIC SHOCK”. In: *The Astrophysical Journal Letters* 698.1, p. L10.
- O’Dea, Christopher P and D. J. Saikia (Dec. 2021). “Compact steep-spectrum and peaked-spectrum radio sources”. In: *A&A Rev.* 29.1, 3, p. 3.
- Ohsawa, Yukiharu (Aug. 1986). “Solitary Wave of a Relativistic Magnetosonic Wave Propagating Perpendicularly to a Magnetic Field”. In: *Physics of Fluids* 29, pp. 2474–2478.
- (Mar. 2014). “Ultrarelativistic Particle Acceleration in Collisionless Shock Waves”. In: *Physics Reports* 536.4, pp. 147–254.
- Omidi, N., K. B. Quest, and D. Winske (Dec. 1990). “Low Mach Number Parallel and Quasi-Parallel Shocks”. In: *Journal of Geophysical Research* 95, pp. 20717–20730.
- Paczyski, Bohdan (Feb. 1998). “Are Gamma-Ray Bursts in Star-Forming Regions?” In: *ApJ* 494.1, pp. L45–L48.
- Paiano, Simona, Renato Falomo, Aldo Treves, and Riccardo Scarpa (Feb. 2018). “The Redshift of the BL Lac Object TXS 0506+056”. In: *ApJ* 854.2, L32, p. L32.
- Parizot, E., A. Marcowith, J. Ballet, and Y. A. Gallant (July 2006). “Observational constraints on energetic particle diffusion in young supernovae remnants: amplified magnetic field and maximum energy”. In: *A&A* 453.2, pp. 387–395.
- Pereira-Santaella, M. et al. (Sept. 2022). “Low-power jet-interstellar medium interaction in NGC 7319 revealed by JWST/MIRI MRS”. In: *A&A* 665, L11, p. L11.
- Plotnikov, Illya, Anna Grassi, and Mickael Grech (July 2018). “Perpendicular relativistic shocks in magnetized pair plasma”. In: *MNRAS* 477.4, pp. 5238–5260.
- Pohl, M., M. Hoshino, and J. Niemiec (2020). “PIC simulation methods for cosmic radiation and plasma instabilities”. In: *Progress in Particle and Nuclear Physics* 111, p. 103751.
- Pringle, J. E. (Jan. 1981). “Accretion discs in astrophysics”. In: *ARA&A* 19, pp. 137–162.
- Ptuskin, V. S. et al. (Feb. 1993). “Diffusion and drift of very high energy cosmic rays in galactic magnetic fields”. In: *A&A* 268.2, pp. 726–735.
- Rassel, Mohira, Patrick Kilian, Vito Aberham, Felix Spanier, Nicole Lloyd-Ronning, and Chris L. Fryer (Aug. 2023). “Particle-in-cell Simulations of Mildly Relativistic Outflows in Kilonova Emissions”. In: *ApJ* 952.2, 165, p. 165.
- Reipurth, Bo, John Bally, and David Devine (Dec. 1997). “Giant Herbig-Haro Flows”. In: *AJ* 114, p. 2708.

- Ressler, Sean M. et al. (Aug. 2014). “Magnetic Field Amplification in the Thin X-Ray Rims of SN 1006”. In: *ApJ* 790.2, 85, p. 85.
- Reville, B., J. G. Kirk, and P. Duffy (Dec. 1, 2006). “A Current-Driven Instability in Parallel, Relativistic Shocks”. In: *Plasma Physics and Controlled Fusion* 48, pp. 1741–1747.
- Romansky, V. I., A. M. Bykov, and S. M. Osipov (Nov. 2024). “Cosmic ray acceleration and non-thermal emission from fast luminous optical transient sources”. In: *Advances in Space Research* 74.9, pp. 4290–4301.
- Saavedra, Enzo A., Federico A. Fogantini, Gastón J. Escobar, Gustavo E. Romero, Jorge A. Combi, and Estefania Marcel (Dec. 2023). “NuSTAR and XMM-Newton observations of the binary 4FGL J1405.1-6119. A γ -ray-emitting microquasar?” In: *A&A* 680, A88, A88.
- Sagdeev, R. Z. (Jan. 1966). “Cooperative Phenomena and Shock Waves in Collisionless Plasmas”. In: *Reviews of Plasma Physics* 4, p. 23.
- Sakai, J I, W Yamamura, S Saito, H Washimi, D Tsiklauri, and G Vekstein (Nov. 2005). “Particle Simulation of Plasma Heating by a Large-Amplitude Shear Alfvén Wave through Its Transverse Modulation in Collisionless Plasmas”. In: *New Journal of Physics* 7, pp. 233–233.
- Scholer, Manfred, Harald Kucharek, and Iku Shinohara (July 2003). “Short large-amplitude magnetic structures and whistler wave precursors in a full-particle quasi-parallel shock simulation”. In: *Journal of Geophysical Research (Space Physics)* 108.A7, 1273, p. 1273.
- Schwartz, Steven J., David Burgess, William P. Wilkinson, Ramona L. Kessel, Malcolm Dunlop, and Herman Luehr (Apr. 1, 1992). “Observations of Short Large-Amplitude Magnetic Structures at a Quasi-Parallel Shock”. In: *Journal of Geophysical Research* 97, pp. 4209–4227.
- Sironi, L., U. Keshet, and M. Lemoine (Oct. 2015). “Relativistic Shocks: Particle Acceleration and Magnetization”. In: *Space Sci. Rev.* 191.1-4, pp. 519–544.
- Sironi, Lorenzo and Anatoly Spitkovsky (Jan. 2011). “Particle Acceleration in Relativistic Magnetized Collisionless Electron-Ion Shocks”. In: *ApJ* 726.2, 75, p. 75.
- Spada, M., D. Lazzati, G. Ghisellini, and A. Celotti (Jan. 2001a). “Internal shocks in the jets of blazars”. In: *Mem. Soc. Astron. Italiana* 72, pp. 157–159.
- Spada, Maddalena, Gabriele Ghisellini, Davide Lazzati, and Annalisa Celotti (Aug. 2001b). “Internal shocks in the jets of radio-loud quasars”. In: *MNRAS* 325.4, pp. 1559–1570.
- Spencer, R. E. (Nov. 1979). “A radio jet in SS433”. In: *Nature* 282.5738, pp. 483–484.
- Spitkovsky, Anatoly (Jan. 2008). “On the Structure of Relativistic Collisionless Shocks in Electron-Ion Plasmas”. In: *ApJ* 673.1, p. L39.
- Stockem, Anne, Frederico Fi?za, Ricardo A Fonseca, and Luis O Silva (2012). “The impact of kinetic effects on the properties of relativistic electron-positron shocks”. In: *Plasma Physics and Controlled Fusion* 54.12.
- Stone, E. C. et al. (July 1998). “The Advanced Composition Explorer”. In: *Space Sci. Rev.* 86, pp. 1–22.
- Stringer, T. E. (Jan. 1963). “Low-Frequency Waves in an Unbounded Plasma”. In: *Journal of Nuclear Energy* 5, pp. 89–107.

- Strong, Andrew W., Igor V. Moskalenko, and Vladimir S. Ptuskin (Nov. 2007). “Cosmic-Ray Propagation and Interactions in the Galaxy”. In: *Annual Review of Nuclear and Particle Science* 57.1, pp. 285–327.
- Swierc, Paxson, Damiano Caprioli, Luca Orusa, and Miha Cernetic (2025). *Deep Learning Analysis of Ions Accelerated at Shocks*.
- Takamoto, Makoto, Yosuke Matsumoto, and Tsunehiko N. Kato (June 2018). “Magnetic Field Saturation of the Ion Weibel Instability in Interpenetrating Relativistic Plasmas”. In: *ApJ* 860.1, L1, p. L1.
- Tchekhovskoy, Alexander, Ramesh Narayan, and Jonathan C. McKinney (Nov. 2011). “Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole”. In: *MNRAS* 418.1, pp. L79–L83.
- Terasawa, Toshio, Masahiro Hoshino, Jun-Ichi Sakai, and Tohru Hada (Apr. 1986). “Decay Instability of Finite-amplitude Circularly Polarized Alfvén Waves: A Numerical Simulation of Stimulated Brillouin Scattering”. In: *Journal of Geophysical Research: Space Physics* 91.A4, pp. 4171–4187.
- Thimmappa, R., . Stawarz, J. Neilsen, M. Ostrowski, and B. Reville (Dec. 2022). “X-Ray Spectral Analysis of the Jet Termination Shock in Pictor A on Subarcsecond Scales with Chandra”. In: *The Astrophysical Journal* 941.2, p. 204.
- Thoudam, S. et al. (2016). “Cosmic-ray energy spectrum and composition up to the ankle: the case for a second Galactic component”. In: *A&A* 595, A33.
- Thum, Clemens et al. (Jan. 2018). “POLAMI: Polarimetric Monitoring of Active Galactic Nuclei at Millimetre Wavelengths - II. Widespread circular polarization”. In: *MNRAS* 473.2, pp. 2506–2520.
- Torralba Paz, G., A. Bohdan, and J. Niemiec (Feb. 1, 2025). “Neural Networks for the Analysis of Traced Particles in Kinetic Plasma Simulations”. In: *AIP Advances* 15.2, p. 025116.
- Tran, Aaron and Lorenzo Sironi (Apr. 2024). “Electron Heating in 2D Particle-in-cell Simulations of Quasi-perpendicular Low-beta Shocks”. In: *ApJ* 965.1, 37, p. 37.
- Treumann, R. A. (Nov. 2009). “Fundamentals of Collisionless Shocks for Astrophysical Application, 1. Non-relativistic Shocks”. In: *The Astronomy and Astrophysics Review* 17.4, pp. 409–535.
- Tsubouchi, K. and B. Lembège (Feb. 2004). “Full particle simulations of short large-amplitude magnetic structures (SLAMS) in quasi-parallel shocks”. In: *Journal of Geophysical Research (Space Physics)* 109.A2, A02114, A02114.
- Ulrich, Marie-Helene, Laura Maraschi, and C. Megan Urry (Jan. 1997). “Variability of Active Galactic Nuclei”. In: *ARA&A* 35, pp. 445–502.
- Umeda, T, Y Omura, T Tominaga, and H Matsumoto (2003). “A new charge conservation method in electromagnetic particle-in-cell simulations”. In: *Computer Physics Communications* 156.1, pp. 73–85.
- Urry, C. Megan and Paolo Padovani (Sept. 1995). “Unified Schemes for Radio-Loud Active Galactic Nuclei”. In: *PASP* 107, p. 803.
- van Weeren, Reinout J., Huub J. A. Röttgering, Marcus Brüggen, and Matthias Hoeft (Oct. 2010). “Particle Acceleration on Megaparsec Scales in a Merging Galaxy Cluster”. In: *Science* 330.6002, p. 347.

- Vay, J. -L. (May 2008). "Simulation of beams or plasmas crossing at relativistic velocity". In: *Physics of Plasmas* 15.5, 056701, p. 056701.
- Vietri, Mario (Nov. 1995). "The Acceleration of Ultra-High-Energy Cosmic Rays in Gamma-Ray Bursts". In: *ApJ* 453, p. 883.
- Wieland, Volkmar, Martin Pohl, Jacek Niemiec, Iman Rafighi, and Ken-Ichi Nishikawa (Mar. 2016). "Nonrelativistic Perpendicular Shocks Modeling Young Supernova Remnants: Non-stationary Dynamics and Particle Acceleration at Forward and Reverse Shocks". In: *ApJ* 820.1, 62, p. 62.
- Wild, Wolfgang and Federico Ferrini (2022). "Cherenkov Telescope Array Observatory (CTAO): the worlds first and largest ground-based gamma-ray observatory". In: *Ground-based and Airborne Telescopes IX*. Ed. by Heather K. Marshall, Jason Spyromilio, and Tomonori Usuda. Vol. 12182. International Society for Optics and Photonics. SPIE, 121820P.
- Wilson III, L. B. et al. (Sept. 2017). "Revisiting the structure of low-Mach number, low-beta, quasi-perpendicular shocks". In: *Journal of Geophysical Research (Space Physics)* 122.9, pp. 9115–9133.
- Wilson, L. B. et al. (Mar. 2013). "Shocklets, SLAMS, and field-aligned ion beams in the terrestrial foreshock". In: *Journal of Geophysical Research (Space Physics)* 118.3, pp. 957–966.
- Wu, C. S. (Oct. 1984). "A fast Fermi process: Energetic electrons accelerated by a nearly perpendicular bow shock". In: *J. Geophys. Res.* 89.A10, pp. 8857–8862.
- Yee, Kane (May 1966). "Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media". In: *IEEE Transactions on Antennas and Propagation* 14.3, pp. 302–307.
- Yue, Chuan et al. (Dec. 2019). "Implications on the origin of cosmic rays in light of 10 TV spectral softenings". In: *Frontiers of Physics* 15.2, 24601, p. 24601.
- Zatsepin, G. T. and V. A. Kuz'min (Aug. 1966). "Upper Limit of the Spectrum of Cosmic Rays". In: *Soviet Journal of Experimental and Theoretical Physics Letters* 4, p. 78.
- Zatsepin, V. I. and N. V. Sokolskaya (Jan. 2007). "Energy spectra of the main groups of galactic cosmic rays in the model of three classes of sources". In: *Astronomy Letters* 33.1, pp. 25–33.
- Zech, Andreas and Martin Lemoine (Oct. 2021). "Electron-Proton Co-Acceleration on Relativistic Shocks in Extreme-TeV Blazars". In: *Astronomy and Astrophysics* 654, A96.
- Zenitani, Seiji and Shin'ya Nakano (Nov. 2022). "Loading a relativistic Kappa distribution in particle simulations". en. In: *Physics of Plasmas* 29.11, p. 113904.
- Zindo, Atsushi, Yukiharu Ohsawa, Naoki Bessho, and Richard Sydora (May 2005). "Deep Trapping of Electrons by Oblique Shock Waves". In: *Physics of Plasmas* 12, p. 052321.