From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Astrophysical Sciences

2011

Particle Acceleration and Nonthermal Emission in Relativistic Astrophysical Shocks

Sironi, Lorenzo

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01b2773v701

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Particle Acceleration and Nonthermal Emission in Relativistic Astrophysical Shocks

Lorenzo Sironi

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF

ASTROPHYSICAL SCIENCES

Advisor: Prof. Anatoly Spitkovsky

November, 2011

© Copyright 2011 by Lorenzo Sironi
All rights reserved.

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Prof. Anatoly Spitkovsky
(Principal Advisor)

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Prof. James M. Stone

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Prof. Andrei M. Beloborodov

Approved for the Princeton University Graduate School:

Dean of the Graduate School

Abstract

The common observational feature of Pulsar Wind Nebulae (PWNe), gamma-ray bursts (GRBs), and AGN jets is a broad nonthermal spectrum of synchrotron and inverse Compton radiation, extending from the radio up to the gamma-ray band. In these sources, it is usually assumed that the emitting electrons are accelerated to a power-law distribution by bouncing back and forth across a relativistic shock, via the so-called Fermi mechanism. Despite decades of research, the nonlinear coupling between Fermi-accelerated particles and magnetic turbulence — which is generated by the particles, and at the same time governs their acceleration — is still not understood from first principles. In turn, uncertainties in the acceleration efficiency, i.e., in the fraction of particles injected into the Fermi process, can cast doubts even on the nature of the radiation process (synchrotron versus inverse Compton) that contributes to a given frequency band. An assessment of the micro-physics of relativistic shocks, with particular focus on the efficiency of particle acceleration, is then of paramount importance to unveil the properties of relativistic astrophysical sources, and it is the subject of this dissertation.

In the first part of this thesis (Chapters 2 and 3), I explore by means of fully-kinetic first-principle particle-in-cell (PIC) simulations the properties of relativistic shocks that propagate in electron-positron (Chapter 2) and electron-proton (Chapter 3) plasmas carrying uniform magnetic fields. I find that nonthermal particle acceleration in highly relativistic flows (bulk Lorentz factor $\gamma_0 \gtrsim 5$) only occurs if the upstream magnetization is weak ($\sigma \lesssim 10^{-3}$), or if the pre-shock field is nearly aligned with the shock direction of propagation (quasi-parallel shocks). Relativistic shocks in PWNe, GRBs and AGN jets are usually thought to be appreciably magnetized ($\sigma \gtrsim 0.01$) and quasi-perpendicular, yet they need to be efficient particle accelerators, in order to explain the prominent nonthermal signatures of these sources. Motivated by this discrepancy, in Chapter 4 I relax the assumption of *uniform pre-shock fields*, and investigate the acceleration efficiency of perpendicular shocks that propagate in high-σ flows with *alternating magnetic fields*. This is the magnetic geometry expected at the termination shock of pulsar winds, but it could also be

relevant for Poynting-dominated jets in GRBs and AGNs. I show by means of PIC simulations that compression of the flow at the shock will force annihilation of nearby field lines, a process known as shock-driven reconnection. Magnetic reconnection can efficiently transfer the energy of alternating fields to the particles, generating flat power-law tails containing most of the particles (i.e., with acceleration efficiency of nearly 100%). In chapter 5, I directly relate the results of my PIC simulations to observations of nonthermal sources, by presenting a numerical technique that I have developed in order to extract *ab initio* photon spectra from PIC simulations of shocks. With this powerful technique, I have modeled the emission from GRB jets, ruling out a class of models that relied on the so-called "jitter radiation." This reinforces the idea that a detailed understanding of the micro-physics of particle acceleration in relativistic shocks is required in order to correctly interpret the emission signatures of astrophysical nonthermal sources.

Acknowledgments

The work presented in this thesis is the result of three years of research in the Department of Astrophysical Sciences at Princeton University. I am grateful to several people, whose advice and support contributed to the good development of this thesis. In particular, I wish to thank:

First of all, my PhD advisor, Prof. Anatoly Spitkovsky, for his constant guidance and help, for sharing the scientific excitement of new discoveries, and for motivating me to persevere in the difficult moments. To him I am indebted for all the knowledge I have been able to attain in the field of high-energy and plasma astrophysics, and for giving an inspiring example of commitment to scientific research.

To Prof. Jeremy Goodman, Prof. Bruce Draine, Prof. Scott Tremaine, and Dr. Aristotle Socrates, for encouragement and guidance while doing research in different fields of astrophysics.

To Prof. Jim Stone and Prof. Andrei Beloborodov, for being willing to serve as my thesis readers. To all the members of my thesis committee.

To Dr. Luís Gargaté and Dr. Dimitrios Giannios, for being great colleagues and friends, for many inspiring (or just funny) conversations over lunch and coffee, and for sharing many experiences both inside and outside the department.

To the Director of Graduate Studies Prof. Jill Knapp, for her unconditional support in many professional and personal circumstances.

To Princeton University, for five years of financial support of my PhD studies. To the PICSciE-OIT High Performance Computing Center and Visualization Laboratory, for access to almost unlimited computing time and resources.

Furthermore, I wish to thank my parents, for supporting me throughout my entire life, especially in the hardest times, and for encouraging me even in the difficulty of living 4,000 miles apart. To my beloved girlfriend Teresa, for her continuous support and love, for sharing my life in the last six years and all the years to come.

Contents

Abstract **iv**

Acknowledgments **vi**

1 Introduction **1**

1 The PIC Method 5

1.1 Specifics of TRISTAN-MP 5

2 Thesis Outline 6

2 Particle Acceleration in Relativistic Magnetized Collisionless Pair Shocks **13**

1 Introduction 14

2 Simulation Setup 16

3 Shock Structure 19

3.1 Oblique Shocks: General Overview 23

3.2 Subluminal Shocks: $0 \leq \theta \lesssim \theta_{\rm crit}$ 28

3.3 Superluminal Shocks: $\theta_{\rm crit} \lesssim \theta \leq 90^\circ$ 34

4 Particle Energy Spectra 36

4.1 Dependence on Magnetic Obliquity 37

4.2 Time Evolution 40

5 Particle Acceleration Mechanisms 45

5.1 $\theta = 0^\circ$ 46

5.2 $\theta = 15^\circ$ 49

5.3 $\theta = 30^\circ$ 51

6 Discussion and Summary 53

7 Appendix (A): 2.5D Simulations with In-Plane Upstream Background Magnetic Field 58

7.1 Shock Structure 58

7.2 Particle Energy Spectra 62

8 Appendix (B): The Sharp Transition in Acceleration Efficiency between Subluminal and Superluminal Shocks 63

Bibliography 67

3 Particle Acceleration in Relativistic Magnetized Collisionless Electron-Ion Shocks 70

1 Introduction 71

2 Simulation Setup 74

3 Shock Structure 77

3.1 Magnetized Relativistic Electron-Ion Shocks: General Overview 78

3.2 Subluminal Shocks: $0^\circ \leq \theta < \theta_{\rm crit}$ 80

3.3 Superluminal Shocks: $\theta_{\rm crit} < \theta \leq 90^\circ$ 87

4 Particle Acceleration 91

4.1 Subluminal Shocks: $0^\circ \leq \theta < \theta_{\rm crit}$ 92

4.2 Superluminal Shocks: $\theta_{\rm crit} < \theta \leq 90^\circ$ 100

5 Dependence on Upstream Conditions 106

5.1 Dependence on Magnetic Obliquity θ 107

5.2 Dependence on Magnetization σ . . . 111
5.3 Dependence on Bulk Lorentz Factor γ_0 . . . 117
6 Summary and Discussion . . . 121
7 Appendix (A): Dependence on the Mass Ratio . . . 125
8 Appendix (B): Dependence on the Field Orientation with respect to the Simulation Plane . . . 127
9 Appendix (C): Ion Counter-Streaming Instabilities . . . 129
Bibliography . . . 131

4 Acceleration of Particles at the Termination Shock of a Relativistic Striped Wind 135

1 Introduction . . . 136
2 Simulation Setup . . . 138
3 Shock Structure and Particle Acceleration . . . 143
3.1 Shock Structure . . . 144
3.2 Particle Spectrum and Acceleration . . . 148
4 Dependence on the Wind Properties . . . 154
4.1 Dependence on the Stripe Wavelength . . . 155
4.2 Dependence on the Wind Magnetization . . . 160
4.3 Dependence on the Stripe-Averaged Field . . . 164
4.4 Dependence on the Bulk Lorentz Factor . . . 174
5 Three-Dimensional Simulations . . . 174
6 Summary and Discussion . . . 178
7 Appendix (A): Dependence on the Transverse Size of the Simulation Box . . . 183
8 Appendix (B): Time Evolution of the Shock . . . 185
9 Appendix (C): Relativistic Striped Winds in the Presence of a Guide Field . . . 187

10 Appendix (D): The Termination Shock of Electron-Ion Striped Winds 191

Bibliography 194

5 Synthetic Spectra from PIC Simulations of Relativistic Collisionless Shocks 197

1 Introduction 198

2 Shock Structure and Particle Energy Spectrum 199

3 Synthetic Photon Spectrum 201

3.1 Numerical Technique 201

3.2 Results 202

4 Discussion 207

Bibliography 209

6 Conclusions and Future Prospects 211

Chapter 1

Introduction

What is the origin of nonthermal emission from relativistic astrophysical objects? Relativistic outflows are present in a variety of astrophysical sources, e.g., Pulsar Wind Nebulae (PWNe) powered by the relativistic wind of pulsars, gamma-ray bursts (GRBs) from the death of massive stars, and Active Galactic Nuclei (AGNs) at the center of galaxies. In such sources, signatures of nonthermal processes are revealed by power-law radiation spectra spanning a surprisingly wide range of wavelengths, from radio to X-rays and beyond, usually explained as synchrotron or inverse Compton emission from a nonthermal population of particles. It is often assumed that the emitting electrons are accelerated to a power-law distribution at shock fronts, compressions created where a relativistic outflow collides with the surrounding medium.

Despite decades of research, a fully self-consistent theory of particle acceleration in relativistic shocks is still lacking. The slope p of the power-law tail (such that $dN/d\gamma \propto \gamma^{-p}$, where γ is the particle Lorentz factor) and the acceleration efficiency, i.e., the fraction of particles and energy stored in the tail, are usually parameterized *ad hoc* to fit the observations. It is customary, especially in modeling the emission from GRBs and AGN jets, to employ the dimensionless parameters ζ_e and ϵ_e, defined respectively as the fraction of particles and flow energy stored in the accelerated electrons. The fraction of flow energy contained in magnetic fields is usually indicated as ϵ_B. The values of p, ζ_e, ϵ_e, and ϵ_B are constrained by fitting the broadband spectral signature of astrophysical sources, rather than by a detailed understanding of the micro-physics of relativistic shocks. Degeneracies among the different parameters can lead to serious ambiguities in the interpretation of the emission in a given frequency band, with potential uncertainties even about the nature of the relevant radiation process (synchrotron as opposed to inverse Compton, see Bošnjak et al. 2009, for GRBs). Clearly, by understanding from first principles the physics

of astrophysical shocks, with special focus on the process of particle acceleration, one could provide physically-grounded inputs for the phenomenological parameters introduced above, thus contributing to unveil the properties of astrophysical relativistic outflows.

The efficiency of particle acceleration in relativistic shocks may depend on the composition of the pre-shock flow (electron-positron, electron-proton or electron-positron-proton plasma), as well as its bulk Lorentz factor γ_0 and magnetization σ (defined as the ratio of Poynting to kinetic flux in the pre-shock medium). If the upstream flow is magnetized, an additional parameter is the obliquity angle that the pre-shock magnetic field makes with the shock direction of propagation.

Pulsar winds are thought to be dominated by electron-positron pairs. MHD models of PWNe (Del Zanna et al. 2004) require $\sigma \gtrsim 0.01$ behind the wind termination shock,[1] where the momentum flux of the relativistic pulsar wind ($\gamma_0 \gtrsim 10^2$) is balanced by the confining pressure of the nebula. Polarization measurements (Wilson 1972; Schmidt et al. 1979) indicate that the nebular magnetic field should be toroidal around the symmetry axis of the system, so that the termination shock is "perpendicular" (i.e., with the field orthogonal to the flow direction). Efficient acceleration of particles at the termination shock (with ζ_e and ϵ_e approaching 100%) is required to explain the broadband spectrum of PWNe, extending from the radio up to the X-ray band (and even GeV and TeV energies, as recently shown by the Fermi telescope and TeV experiments, Abdo et al. 2010). However, the flatness of the radio spectrum ($F_{\nu_r} \propto \nu^{-0.3}$ for the Crab Nebula, Bietenholz et al. 1997) is hard to reconcile with the steeper optical and X-ray slope ($F_{\nu_X} \propto \nu^{-1.1}$ for the Crab, Mori et al. 2004), unless the electron distribution is more complicated than a single power law (naively, the radio band would require $p \simeq 1.6$, whereas $p \gtrsim 2$ would be needed for the optical and X-ray emission). Even more fundamentally, how the pulsar wind termination shock can accelerate particles to the required energies is still an unsolved problem.

Synchrotron and inverse Compton emission from the core of AGN jets is usually attributed to high-energy electrons accelerated in mildly-relativistic internal shocks ($\gamma_0 \sim$ a few in the jet comoving frame), resulting from the collision of shells of matter intermittently ejected by the central engine. Such shocks should be magnetized ($\sigma \gtrsim 0.01$) and quasi-perpendicular, since polarization measurements of radio knots (Gabuzda et al. 2004; Pushkarev et al. 2005) indicate that the magnetic field is mostly transverse to the jet axis. Electron acceleration in internal shocks

[1] We remark that existing estimates of σ in pulsar winds refer to the post-shock flow. The magnetization ahead of the shock could be much larger, if a significant fraction of field energy is transferred to the particles at the shock.

is also invoked to explain the prompt emission of GRBs (e.g., Piran 2004). If the jet flow carries a substantial toroidal field, e.g., advected from the central engine, internal shocks in GRBs will be magnetized and quasi-perpendicular. For internal shocks in both AGNs and GRBs, it is inferred that the emitting electrons should contain a substantial fraction ($\epsilon_e \simeq 10\%$) of the flow energy, and their spectrum is usually modeled as a power law extending over a few decades in energy, with slope $2 \lesssim p \lesssim 3$. Similar values are required to explain the radiative signature of GRB afterglows, when the relativistic ejecta ($\gamma_0 \simeq 100$) driven by the central explosion are slowed down by the interaction with the circumburst medium. Here, the flow ahead of the GRB *external* shock (as opposed to the *internal* shocks discussed above) is the interstellar medium, a weakly magnetized ($\sigma \lesssim 10^{-7}$) electron-proton plasma. Typical values of $\epsilon_e \simeq 10\%$, $\epsilon_B \simeq 1\%$, and $2 \lesssim p \lesssim 3$ are required to model the spectrum and the light curve of GRB afterglows. It is still an open question whether shocks in GRBs and AGN jets can provide the level of magnetic turbulence (via ϵ_B) and the acceleration efficiency (via ζ_e and ϵ_e) required to explain the observations.

Relativistic shocks in AGNs and GRBs are also invoked as possible sources of Ultra-High Energy Cosmic Rays (UHECRs), hadronic particles detected on Earth (e.g., by the Pierre Auger Observatory) with energies in excess of 10^{19} eV. The Larmor radius of these particles is larger than the size of our Galaxy, suggesting an extragalactic origin. The requirements on the magnetic field, compactness, and energy budget of the putative sources are extremely stringent, and they pose serious challenges to current models of particle acceleration in shocks.

Charged particles are accelerated in shock-dominated systems via the so-called Fermi process, or Diffusive Shock Acceleration (e.g., Blandford & Ostriker 1978; Bell 1978; Drury 1983; Blandford & Eichler 1987). In the Fermi mechanism, particles are energized by bouncing back and forth across the shock, like a tennis ball trapped between two approaching players. The source of scattering on the two sides of the shock is provided by magnetic turbulence embedded in the converging flows, rather than by Coulomb collisions. In fact, due to the low densities of astrophysical plasmas, most astrophysical shocks are "collisionless," i.e., mediated by the nonlinear interaction between particles and collective electromagnetic fields, rather than by Coulomb processes. The theory of Fermi acceleration is well understood in the non-relativistic regime. In the test-particle approximation (in which the energy flux of accelerated particles is neglected in the overall energetics of the shock), the spectrum of shock-accelerated particles is described by a power law of index $p = (R+2)(R-1)$, where R is the compression ratio of the shock ($R = 4$ for strong non-relativistic shocks, so $p = 2$).

The formalism of Fermi acceleration is significantly more complicated in the relativistic limit. The non-relativistic derivation assumes that the particle distribution function is approximately isotropic in the local plasma rest frame, and the diffusion equation is expanded in powers of $v_{sh}/c \ll 1$, where v_{sh} is the shock velocity. The analysis in the relativistic regime is complicated by the fact that the bulk flow speed is comparable to the particle velocities, and the angular distribution function of accelerated particles becomes highly anisotropic near the shock. Most progress in studying particle acceleration in relativistic shocks has been made using semi-analytic kinetic theory methods (e.g., Kirk & Heavens 1989; Ballard & Heavens 1991; Kirk et al. 2000; Achterberg et al. 2001; Keshet & Waxman 2005) or Monte Carlo test-particle simulations (e.g., Bednarz & Ostrowski 1998; Niemiec & Ostrowski 2004; Ellison & Double 2004). Both approaches study particle acceleration with simplifying (and somewhat arbitrary) assumptions about the nature and strength of magnetic turbulence near the shock and the details of wave-particle interactions.

The major achievement of such studies is that in the limit $\gamma_0 \to \infty$ they predict spectral indices $p \simeq 2.2$ which are close to the ones usually inferred from observations. However, analytic methods and Monte-Carlo simulations cannot determine what fraction of the incoming particles will be injected into the acceleration process, i.e., they cannot provide physically-grounded estimates for ζ_e and ϵ_e. Moreover, most studies prescribe *a priori* the strength and spectrum of magnetic irregularities, thus giving no hint for the value of ϵ_B to be employed in modeling the observations. The highly nonlinear coupling between accelerated particles and magnetic turbulence — which is generated by the particles, and at the same time governs their acceleration — is extremely hard to incorporate in analytic studies or Monte-Carlo simulations, and needs to be addressed with *ab initio* plasma simulations.

Fully-kinetic particle-in-cell (PIC) codes provide a self-consistent description of particle acceleration and magnetic turbulence generation in collisionless plasmas. In this thesis, I study particle acceleration in relativistic astrophysical shocks by means of multi-dimensional PIC simulations. I describe the PIC method in §1, and outline the structure of my thesis in §2. My results will provide physically-grounded inputs for astrophysical models of high-energy emission, which will help, in concert with dedicated high-energy experiments (Chandra, Fermi, and TeV telescopes), to unveil the nature of astrophysical nonthermal sources.

1. The PIC Method

PIC codes model astrophysical plasmas from first principles, as a collection of charged macro-particles that are moved by integrating the Lorentz force (Birdsall & Langdon 1985; Hockney & Eastwood 1981). The currents associated with the macro-particles are deposited on a cartesian grid, on which Maxwell's equations are also discretized. Electromagnetic fields are then advanced via Maxwell's equations, with particle currents as the source term. Finally, the updated fields are extrapolated to the locations of individual particles and used for the computation of the Lorentz force, so the loop is closed self-consistently.

For accuracy and stability, PIC codes have to resolve the plasma oscillation frequency of electrons $\omega_{\rm pe} = \sqrt{4\pi e^2 n/m_e}$, and the corresponding plasma skin depth $c/\omega_{\rm pe}$. Here, n is the plasma density, and e and m_e are the electron charge and mass. Resolving the electron plasma scale is essential to describe correctly the kinetic processes at the basis of the Fermi acceleration mechanism. On the other hand, the need to follow the evolution of the shock up to the long time and length scales relevant for astrophysics requires massive computational resources. Only in the last few years, thanks to important advances in numerical algorithms and computer capabilities, plasma simulations have been able to tackle the problem of particle acceleration in astrophysical plasmas from first principles.

I employ the 3D PIC code TRISTAN-MP, which I contributed to develop at Princeton under the supervision of my PhD advisor Prof. Anatoly Spitkovsky (Spitkovsky 2005). It is a massively-parallel descendant of the publicly available code TRISTAN (Buneman 1993), that was optimized to handle ultra-relativistic flows.

1.1. Specifics of TRISTAN-MP

TRISTAN-MP uses space- and time-centered finite difference schemes to calculate spatial and temporal derivatives, so that the code is second-order accurate in both space and time. Electric and magnetic fields are staggered on a cartesian Eulerian mesh of Yee type (Yee 1966), and two different spatial stencils (second-order [Yee (1966)] and fourth-order [Greenwood et al. (2004)] accurate) are implemented in the code to discretize the terms $\nabla \times \vec{E}$ and $\nabla \times \vec{B}$ in Maxwell's equations. The fourth-order stencil is better suited to handle ultra-relativistic flows, where particle

velocities may be larger than the speed of electromagnetic waves on the grid, leading to numerical Cerenkov radiation.

From the computational mesh, the fields are extrapolated to the positions of individual particles via a tri-linear interpolation function (linear in each spatial dimension), and then used for the computation of the Lorentz force. The code allows to choose between two different integrators for the particle momentum via the Lorentz force: the so-called Boris algorithm (Boris 1970), and the Vay pusher (Vay 2008). We find that the Vay pusher is more accurate than the standard Boris algorithm in dealing with relativistic drift velocities of electron-ion magnetized flows. Currents deposited as each computational particle moves through the grid are first filtered (with a three-point digital filter) to significantly reduce noise, and then used as the source term in Maxwell's equations.

The code is parallelized using domain decomposition with the Message Passing Interface (MPI), and particle sorting is employed to improve cache performance. The simulation box is evenly subdivided across the transverse dimensions (y and z for 3D, only y for 2D), and each processor evolves fields and particles along the entire longitudinal extent of the box (i.e., along x), but only in a small part of the domain in y and z. In application to the study of shocks, special performance improvements have been implemented in TRISTAN-MP, including a dynamically adjustable domain size, so the size of the computational box can grow in time along x as the system evolves. This allows to save memory and computing time, while following the evolution of all regions that are causally connected with the initial stage of the system. This method also helps to reduce particle heating by the numerical Cerenkov instability, which occurs when cold relativistic plasma propagates for large distances over the numerical grid.

TRISTAN-MP allows to save, at selected time intervals, a suite of field and particle quantities for post-processing and data analysis. Output files are written to disk via Hierarchical Data Format 5 (HDF5) libraries. At given times specified by the user, the full state of a simulation can also be stored on disk, which permits to restart the simulation from that point and extend it to longer times, as long as the computational resources allow.

2. Thesis Outline

In this thesis, I study via multi-dimensional PIC simulations the process of particle acceleration in relativistic astrophysical shocks, and its implications for the emission signatures of shock-dominated systems.

In the first part, I explore the properties of relativistic shocks that propagate in electron-positron (Chapter 2) and electron-proton (Chapter 3) plasmas carrying *uniform magnetic fields*. I assess how the internal structure of the shock and the efficiency of particle acceleration depend on the conditions of the pre-shock flow, such as bulk Lorentz factor, magnetic field strength and orientation. It turns out that, for highly relativistic flows (bulk Lorentz factor $\gamma_0 \gtrsim 5$), the main parameter that controls the shock physics is the magnetization σ. Weakly magnetized shocks ($\sigma \lesssim 10^{-3}$) are mediated by electromagnetic plasma instabilities (the so-called Weibel instability, Weibel 1959; Medvedev & Loeb 1999; Gruzinov & Waxman 1999), that generate magnetic fields stronger than the background field. Such shocks do self-consistently accelerate particles up to nonthermal energies, as observed in the strictly unmagnetized case (i.e., $\sigma = 0$) discussed by Spitkovsky (2008). Accelerated particles populate a power-law tail of slope $p \simeq 2.4$, containing a few percent of particles and a few tens of percent of flow energy.

For high magnetizations ($\sigma \gtrsim 10^{-3}$), the shock structure and acceleration properties depend critically on the inclination angle θ between the pre-shock field and the shock direction of propagation. If the magnetic obliquity is larger than a critical angle $\theta_{\rm crit} \simeq 34^\circ$ (measured in the post-shock frame), charged particles would need to move along the field faster than the speed of light in order to outrun the shock ("superluminal" configurations, Begelman & Kirk 1990). We find that "subluminal" shocks ($\theta < \theta_{\rm crit}$) are efficient particle accelerators. Shock-accelerated particles can stream ahead of the shock along the field lines, and their counter-streaming with respect to the incoming flow triggers plasma instabilities that generate strong magnetic turbulence. In turn, such magnetic waves govern the acceleration process, by providing the scattering required for the Fermi mechanism. As a result, the downstream particle spectrum shows a pronounced nonthermal tail of shock-accelerated particles, containing a few percent of particles and a few tens of percent of flow energy. The power-law slope of the nonthermal tail is $2 \lesssim p \lesssim 3$. In contrast, superluminal shocks ($\theta > \theta_{\rm crit}$) show negligible particle acceleration. Here, due to the lack of significant self-generated turbulence, charged particles are forced to slide along the background field lines, whose orientation

prohibits repeated crossings of the shock. This will inhibit the Fermi process, and in fact the downstream particle distribution in superluminal shocks is purely thermal.

Overall, my results in Chapters 2 and 3 imply that nonthermal particle acceleration in relativistic flows ($\gamma_0 \gtrsim 5$) only occurs if the upstream magnetization is weak ($\sigma \lesssim 10^{-3}$), or if the pre-shock field is nearly aligned with the shock direction of propagation (within an angle $\lesssim 34^\circ/\gamma_0$, as seen from the pre-shock frame). On the other hand, relativistic shocks in PWNe, GRBs and AGN jets are usually thought to be appreciably magnetized ($\sigma \gtrsim 0.01$) and superluminal, as discussed above. Based on the results of my PIC simulations, they should be poor particle accelerators, in sharp contrast with the clear nonthermal signatures of these sources.

In Chapter 4, motivated by this discrepancy, I relax the assumption of *uniform pre-shock fields*, and investigate the acceleration efficiency of perpendicular shocks that propagate in high-σ flows with *alternating magnetic fields*. In the equatorial region of pulsar winds, if the rotational and magnetic axes of the pulsar are misaligned, the flow consists of stripes of opposite field polarity, alternating with the pulsar rotation period. Compression of the flow at the wind termination shock will force annihilation of nearby field lines (Lyubarsky 2003), a process known as shock-driven magnetic reconnection. I show that shock-driven reconnection can efficiently transfer the energy of alternating fields to the particles, whose average Lorentz factor increases by a factor of σ across the shock. In the limit $\lambda/(r_L\sigma) \gg 1$ — where λ is the wavelength of the striped wind and r_L the relativistic Larmor radius in the wind — the particle spectrum downstream from the shock approaches a flat power-law tail with slope $p \simeq 1.5$, populated by particles accelerated by the reconnection electric field. The resulting synchrotron slope $\simeq -0.25$ is close to the radio spectral index of the Crab Nebula. Around the equatorial plane of the wind, where the stripes are symmetric, the highest energy particles accelerated by the reconnection process can escape ahead of the shock, and be injected into a Fermi-like acceleration mechanism. In the post-shock spectrum, they populate a power-law tail with slope $p \simeq 2.5$, that extends beyond the flat component produced by reconnection. Synchrotron emission from these particles may explain the optical and X-ray signature of the Crab Nebula.

As a further step in relating the details of particle acceleration in shocks to astrophysical observations, I have developed a numerical technique that allows to extract *ab initio* photon spectra from PIC simulations, as I describe in Chapter 5. The two basic ingredients for the radiation, namely accelerated particles and magnetic fields, are produced self-consistently as part of the

shock evolution. The photon spectrum can then be computed via Fourier transform of the electric far-field from a large number of particles, sampled directly from the simulation.

In Chapter 5, I apply this technique to relativistic unmagnetized shocks, which I have shown in Chapters 2 and 3 to be efficient particle accelerators. Such shocks are mediated by the Weibel instability, that generates small-scale magnetic fluctuations. It has been speculated that emission from these shocks may occur in so-called "jitter" regime, if the scale $\lambda_{\rm turb}$ of the turbulence is such that $K \equiv qB\lambda_{\rm turb}/mc^2 \ll 1$ (with standard synchrotron radiation occurring for $K \gg 1$). Jitter radiation has been proposed as a solution for the so-called "line of death" puzzle in GRB spectra, which below the peak frequency are sometimes harder than expected from synchrotron radiation. However, the synthetic spectra I extract from my shock simulations are entirely consistent with synchrotron radiation in the fields generated by the Weibel instability. I can recover the so-called "jitter" regime only if I artificially reduce the strength of the fields, such that the parameter K becomes much smaller than unity. So, if the GRB prompt emission results from relativistic unmagnetized shocks, it seems that resorting to the jitter regime is not a viable solution for the "line of death" puzzle.

At frequencies above the peak, my results show, somewhat unexpectedly, that the contribution of upstream particles to the total emission is not negligible, yet it is omitted in most models. It causes the radiation spectrum to be flatter than the corresponding downstream spectrum, thus partly masking the contribution of downstream thermal particles. This could potentially explain the absence of clear signatures of thermal emission in GRB spectra. Overall, the results presented in Chapter 5 show that, by extracting *ab initio* radiation spectra from fully-kinetic self-consistent PIC simulations of shocks, we could directly constrain the properties of relativistic astrophysical flows, thus shedding light on the origin of nonthermal emission in the Universe.

Each of the four chapters that follow is based on a paper that has been published or accepted for publication, as detailed here:

Chapter 2. *Particle Acceleration in Relativistic Magnetized Collisionless Pair Shocks: Dependence of Shock Acceleration on Magnetic Obliquity*
Lorenzo Sironi & Anatoly Spitkovsky
ApJ, 698, 1523 (2009)

Chapter 3. *Particle Acceleration in Relativistic Magnetized Collisionless Electron-Ion Shocks*
Lorenzo Sironi & Anatoly Spitkovsky
ApJ, 726, 75 (2011)

Chapter 4. *Acceleration of Particles at the Termination Shock of a Relativistic Striped Wind*
Lorenzo Sironi & Anatoly Spitkovsky
ApJ accepted, astro-ph/arXiv:1107.0977

Chapter 5. *Synthetic Spectra from PIC Simulations of Relativistic Collisionless Shocks*
Lorenzo Sironi & Anatoly Spitkovsky
ApJL, 707, L92 (2009)

Bibliography

Abdo, A. A., et al. 2010, ApJ, 708, 1254

Achterberg, A., Gallant, Y. A., Kirk, J. G., & Guthmann, A. W. 2001, MNRAS, 328, 393

Ballard, K. R. & Heavens, A. F. 1991, MNRAS, 251, 438

Bednarz, J. & Ostrowski, M. 1998, Physical Review Letters, 80, 3911

Begelman, M. C. & Kirk, J. G. 1990, ApJ, 353, 66

Bell, A. R. 1978, MNRAS, 182, 147

Bietenholz, M. F., Kassim, N., Frail, D. A., Perley, R. A., Erickson, W. C., & Hajian, A. R. 1997, ApJ, 490, 291

Birdsall, C. K., & Langdon, B. 1985, Plasma Physics via Computer Simulation

Blandford, R. & Eichler, D. 1987, Phys. Rep., 154, 1

Blandford, R. D. & Ostriker, J. P. 1978, ApJ, 221, L29

Boris, J. P. 1970, Relativistic plasma simulation-optimization of a hybrid code, In Fourth Conf. Num. Sim. Plasmas, page 3, Washington DC, Naval Res. Lab

Bošnjak, Ž., Daigne, F., & Dubus, G. 2009, A&A, 498, 677

Buneman, O. 1993, Computer Space Plasma Physics (Terra Scientific, Tokyo), 67

Del Zanna, L., Amato, E., & Bucciantini, N. 2004, A&A, 421, 1063

Drury, L. O. 1983, Reports on Progress in Physics, 46, 973

Ellison, D. C. & Double, G. P. 2004, Astroparticle Physics, 22, 323

Gabuzda, D. C., Murray, É., & Cronin, P. 2004, MNRAS, 351, L89

Greenwood, A. D., Cartwright, K. L., Luginsland, J. W., & Baca E. A. 2004, On the elimination of numerical Cerenkov radiation in PIC simulations, JCP 201, 665

Gruzinov, A. & Waxman, E. 1999, ApJ, 511, 852

Hockney, R. W., & Eastwood J. W. 1981, Computer Simulation using Particles

Keshet, U. & Waxman, E. 2005, Physical Review Letters, 94, 111102

Kirk, J. G., Guthmann, A. W., Gallant, Y. A., & Achterberg, A. 2000, ApJ, 542, 235

Kirk, J. G. & Heavens, A. F. 1989, MNRAS, 239, 995

Lyubarsky, Y. E. 2003, MNRAS, 345, 153

Medvedev, M. V. & Loeb, A. 1999, ApJ, 526, 697

Mori, K., Burrows, D. N., Hester, J. J., Pavlov, G. G., Shibata, S., & Tsunemi, H. 2004, ApJ, 609, 186

Niemiec, J. & Ostrowski, M. 2004, ApJ, 610, 851

Piran, T. 2004, Reviews of Modern Physics, 76, 1143

Pushkarev, A. B., Gabuzda, D. C., Vetukhnovskaya, Y. N., & Yakimov, V. E. 2005, MNRAS, 356, 859

Schmidt, G. D., Angel, J. R. P., & Beaver, E. A. 1979, ApJ, 227, 106

Spitkovsky, A. 2005, in AIP Conf. Ser., Vol. 801, Astrophysical Sources of High Energy Particles and Radiation, ed. T. Bulik, B. Rudak, & G. Madejski, 345

—. 2008, ApJ, 682, L5

Vay, J.-L. 2008, Simulation of beams or plasmas crossing at relativistic velocity, Phys. of Plasmas 15, 056701

Weibel, E. S. 1959, Physical Review Letters, 2, 83

Wilson, A. S. 1972, MNRAS, 157, 229

Yee, K. S. 1966, Numerical Solution of Initial Boundary Value Problems Involving Maxwell's Equations in Isotropic Media, IEEE Trans. on Ant. and Prop., 14, 3, 202

Chapter 2

Particle Acceleration in Relativistic Magnetized Collisionless Pair Shocks

Abstract

We investigate shock structure and particle acceleration in relativistic magnetized collisionless pair shocks by means of 2.5D and 3D particle-in-cell simulations. We explore a range of inclination angles between the pre-shock magnetic field and the shock normal. We find that only magnetic inclinations corresponding to "subluminal" shocks, where relativistic particles following the magnetic field can escape ahead of the shock, lead to particle acceleration. The downstream spectrum in such shocks consists of a relativistic Maxwellian and a high-energy power-law tail with exponential cutoff. For increasing magnetic inclination in the subluminal range, the high-energy tail accounts for an increasing fraction of particles (from $\sim 1\%$ to $\sim 2\%$) and energy (from $\sim 4\%$ to $\sim 12\%$). The spectral index of the power law increases with angle from -2.8 ± 0.1 to -2.3 ± 0.1. For nearly parallel shocks, particle energization mostly proceeds via the Diffusive Shock Acceleration process; the upstream scattering is provided by oblique waves which are generated by the high-energy particles that escape upstream. For larger subluminal inclinations, Shock-Drift Acceleration is the main acceleration mechanism, and the upstream oblique waves regulate injection into the acceleration process. For "superluminal" shocks, self-generated shock turbulence is not strong enough to overcome the kinematic constraints, and the downstream particle spectrum does

not show any significant suprathermal tail. As seen from the upstream frame, efficient acceleration in relativistic (Lorentz factor $\gamma_0 \gtrsim 5$) magnetized ($\sigma \gtrsim 0.03$) flows exists only for a very small range of magnetic inclination angles ($\lesssim 34^\circ/\gamma_0$), so relativistic astrophysical pair shocks have to be either nearly parallel or weakly magnetized to generate nonthermal particles. These findings place constraints on the models of AGN jets, Pulsar Wind Nebulae and Gamma Ray Bursts that invoke particle acceleration in relativistic magnetized shocks.

1. Introduction

Collisionless shocks have been implicated in most nonthermal phenomena in the Universe. Models of nonthermal emission from Pulsar Wind Nebulae (PWNe), jets from Active Galactic Nuclei (AGN), Gamma-Ray Bursts (GRBs) and supernova remnants (SNRs) require a population of high-energy particles with power-law spectra, and collisionless shocks are the primary candidates to produce scale-free distributions of energetic particles. However, despite decades of research, particle acceleration in astrophysical shocks is still not understood from first principles.

Shock acceleration of charged particles is thought to occur by means of two main mechanisms: Diffusive Shock Acceleration (DSA, or first-order Fermi acceleration) and Shock-Drift Acceleration (SDA).[1] In first-order Fermi acceleration, particles stochastically diffuse back and forth across the shock front and gain energy by scattering off magnetic turbulence embedded in the converging flows (e.g., Blandford & Ostriker 1978; Bell 1978; Drury 1983; Blandford & Eichler 1987). SDA is the process whereby particles gain energy from the background motional electric field $\mathbf{E} = -\boldsymbol{\beta} \times \mathbf{B}$ when gyrating across the shock front, while they drift parallel (or anti-parallel, according to their charge) to the electric field (e.g., Chen & Armstrong 1975; Webb et al. 1983; Begelman & Kirk 1990). SDA is a very fast process compared to DSA (Jokipii 1982); however, it is not very efficient without a means for the particles to remain close to the shock and encounter it more than once, thereby being repeatedly accelerated. It is generally accepted that when a sufficient amount of magnetic field turbulence exists on both sides of the shock, both mechanisms operate in tandem making the overall acceleration process very efficient.

[1] In electron-ion magnetized shocks, Shock-Surfing Acceleration is also important, a mechanism by which ions repelled by the shock electric potential are accelerated by the background motional electric field (e.g., Lee et al. 1996). In this chapter, we consider only electron-positron shocks, where the shock-surfing mechanism does not operate.

The efficiency of particle acceleration, namely the fraction of post-shock particles and energy stored in a suprathermal tail of the particle spectrum, may depend on the composition of the upstream flow (electron-positron, electron-ion or pairs-ion plasma), as well as its bulk Lorentz factor and magnetization. If the upstream medium is magnetized, an additional parameter is the obliquity angle that the upstream magnetic field makes with the shock direction of propagation. For strongly magnetized shocks, cross-field scattering is suppressed and the particle gyro-centers are constrained to slide along the field lines, which are advected downstream from the shock. It is then natural to define a critical magnetic obliquity above which particles sliding along the magnetic field should be moving faster than the speed of light in order to return upstream (Begelman & Kirk 1990). Since efficient acceleration requires that particles repeatedly cross the shock, such "superluminal" geometries should be very inefficient for particle acceleration (Kirk & Heavens 1989; Ballard & Heavens 1991), unless there is strong scattering across field lines or appreciable fluctuations of the magnetic field at the shock, that may create local "subluminal" configurations. If the upstream medium is not turbulent by itself, magnetic field fluctuations may be generated by the high-energy particles accelerated at the shock that escape upstream. In turn, such magnetic waves could scatter the particles and mediate their acceleration. The highly nonlinear problem of wave generation and particle acceleration should then be addressed self-consistently.

Most progress in studying particle acceleration in magnetized shocks has been made using semi-analytic kinetic theory methods (e.g., Kirk & Heavens 1989; Ballard & Heavens 1991; Kirk et al. 2000; Achterberg et al. 2001; Keshet & Waxman 2005) or Monte Carlo test-particle simulations (e.g., Bednarz & Ostrowski 1998; Niemiec & Ostrowski 2004; Ellison & Double 2004). Both approaches study particle acceleration with simplifying (and somewhat arbitrary) assumptions about the nature of magnetic turbulence near the shock and the details of wave-particle interactions. In particular, it is found that superluminal shocks might accelerate as efficiently as subluminal configurations, but the required level of magnetic turbulence must be unreasonably large (Ostrowski & Bednarz 2002). The open question is then if superluminal shocks can self-consistently produce such a high degree of turbulence.

Fully kinetic particle-in-cell (PIC) simulations provide a powerful tool for exploration of the fluid and kinetic structure of collisionless shocks on all scales, thus determining the nature of the shock-generated electromagnetic turbulence simultaneously with the spectrum of accelerated particles. PIC simulations of relativistic *unmagnetized* shocks have been presented by Spitkovsky (2005, 2008a), Chang et al. (2008) and Keshet et al. (2008); Spitkovsky (2008b) has shown

that unmagnetized shocks naturally produce accelerated particles as part of the shock evolution. One-dimensional PIC simulations of *magnetized* perpendicular pair shocks, i.e., with the magnetic field along the shock surface, have shown negligible particle acceleration (Langdon et al. 1988; Gallant et al. 1992), unless a sizable fraction of the upstream bulk kinetic energy is carried by heavy ions (Hoshino et al. 1992; Amato & Arons 2006).

In this chapter, we investigate via first-principles two- and three-dimensional PIC simulations the acceleration properties of relativistic *magnetized* pair shocks. We fix the upstream magnetic field strength and explore how the efficiency and mechanism for particle acceleration vary with magnetic obliquity. We find that subluminal configurations produce a population of energetic particles, and the downstream particle spectrum develops a high-energy power-law tail. We confirm that particle acceleration in superluminal shocks is extremely inefficient, meaning that self-generated turbulence is not strong enough to allow for significant cross-field diffusion of the particles. For subluminal shocks, we find that the acceleration efficiency increases as the magnetic obliquity approaches the boundary between subluminal and superluminal configurations. Also, by studying the trajectory of high-energy particles extracted from our simulations, we show that at low magnetic obliquities, where the motional electric field is still small, the acceleration process is mostly controlled by DSA, whereas particles are mainly accelerated via SDA when the magnetic field inclination is close to the superluminality threshold.

This chapter is organized as follows: in §2 we discuss the setup of our simulations and the magnetic field geometry; in §3 the shock structure and internal physics is investigated for some representative magnetic obliquities. The main results of our work, concerning the efficiency of shock acceleration as a function of magnetic obliquity, are presented in §4. In §5 we comment on the relative importance of different acceleration mechanisms (namely, DSA and SDA) as a function of the magnetic inclination angle. We summarize our findings in §6 and comment on the application of our results to astrophysical scenarios.

2. Simulation Setup

We perform numerical simulations using the three-dimensional (3D) electromagnetic particle-in-cell (PIC) code TRISTAN-MP (Spitkovsky 2005), which is a parallel version of the publicly available code TRISTAN (Buneman 1993) that was optimized for studying collisionless shocks.

The shock is set up by reflecting a stream of pair plasma (the "upstream" plasma) off a conducting wall. The interaction between the injected and reflected beams triggers the formation of a shock, which moves away from the wall. This setup is equivalent to the head-on collision of two identical relativistic pair streams, which would form a forward and reverse shock and a contact discontinuity. In our simulations we follow one of these shocks, and replace the contact discontinuity with the conducting wall. The simulation is performed in the "wall" frame, where the "downstream" plasma behind the shock has zero bulk velocity along the direction of the upstream flow. We stress that the wall frame of our simulations does not in general coincide with the downstream fluid frame, since for oblique magnetic configurations the downstream plasma may have a velocity component transverse to the upstream flow, as we show in §3.

We perform simulations in both two- and three-dimensional computational domains, and we find that most of the shock and acceleration physics is captured extremely well by two-dimensional simulations (see §6 and Appendix 7 for a discussion). Therefore, to follow the shock evolution for longer times with fixed computational resources, we mainly utilize two-dimensional runs. All three components of particle velocities and electromagnetic fields are tracked, however.

We use a rectangular simulation box in the xy plane, with periodic boundary conditions in the y direction (Fig. 2.1). The box is 1024 cells wide (along y), corresponding to approximately $100\,c/\omega_{\rm p}$, where the plasma skin depth $c/\omega_{\rm p} = 10$ cells. We also performed limited experiments with larger boxes, up to 4096 cells wide, obtaining similar results. In the expression above, c is the speed of light and $\omega_{\rm p} \equiv \sqrt{4\pi e^2 n_{\rm u}/(\gamma_0 m)}$ is the relativistic plasma frequency of the upstream flow; $n_{\rm u}$ is the upstream number density measured in the wall frame, γ_0 the Lorentz factor of the

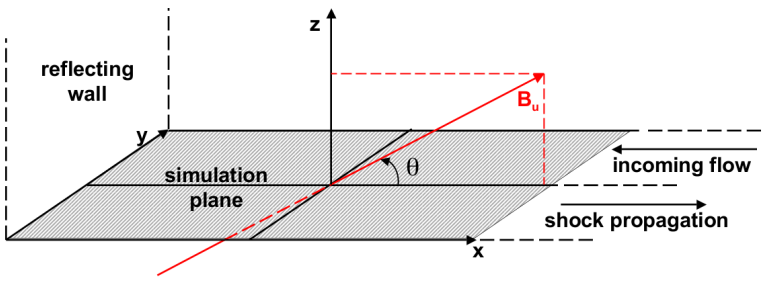


Fig. 2.1.— Simulation geometry and configuration of the upstream magnetic field.

injected plasma, e and m the positron charge and mass. We fix $c = 0.45$ cells/timestep, so that the temporal resolution of our simulations is $0.045\,\omega_{\rm p}^{-1}$ per timestep. At the final time $9000\,\omega_{\rm p}^{-1}$ of our simulation runs, the simulation domain is up to ~ 50000 cells ($5000\,c/\omega_{\rm p}$) long along x.

The incoming pair stream is injected along $-\hat{\mathbf{x}}$ with bulk Lorentz factor $\gamma_0 = 15$. The injected plasma is cold, with energy spectrum in the proper frame given by a three-dimensional Maxwellian $f(\gamma) \propto \gamma\sqrt{\gamma^2-1}\exp(-\gamma/\Delta\gamma)$ with thermal spread $\Delta\gamma = 10^{-4}$. Each computational cell is initialized with two electrons and two positrons. We also performed limited experiments with a larger number of particles per cell (up to 8 per species), obtaining essentially the same results. The magnetization of the upstream plasma is $\sigma = 0.1$, where $\sigma \equiv B_{\rm u}^2/(4\pi\gamma_0 m n_{\rm u} c^2)$ is the squared ratio of relativistic Larmor frequency $\Omega_{\rm L} \equiv eB_{\rm u}/(\gamma_0 mc^2)$ to plasma frequency $\omega_{\rm p}$ for the injected flow ($B_{\rm u}$ is the upstream magnetic field measured in the wall frame). The upstream magnetization is chosen such that perpendicular shocks are not mediated by the relativistic Weibel (1959) instability. As discussed by Spitkovsky (2005), the growth of this instability is suppressed in perpendicular shocks if $\sigma \gtrsim 10^{-3}$; as we show in §3.1, for parallel shocks the Weibel instability may still be important for magnetizations as large as $\sigma = 0.1$. The regime we explore may be relevant for PWNe and internal shocks in GRBs and AGN jets, as we comment in §6.

Keeping the magnetization fixed, we explore a range of magnetic inclinations by varying the angle θ between the shock direction of propagation $+\hat{\mathbf{x}}$ and the upstream magnetic field $\mathbf{B}_{\rm u}$, which is initialized to point towards the first quadrant of the xz plane, as sketched in Fig. 2.1 (i.e., $B_{\rm x,u}, B_{\rm z,u} \geq 0$ and $\theta \equiv \arctan[B_{\rm z,u}/B_{\rm x,u}]$). The magnetic obliquity angle θ is measured in the wall frame. We vary θ from $\theta = 0°$, which corresponds to a "parallel" shock, with magnetic field aligned with the shock normal, up to $\theta = 90°$, i.e., a "perpendicular" shock, with magnetic field along the shock front. For $\theta \neq 0°$, in the upstream medium we also initialize a motional electric field $\mathbf{E}_{\rm u} = -\boldsymbol{\beta}_0 \times \mathbf{B}_{\rm u}$ along $-\hat{\mathbf{y}}$. Here, $\boldsymbol{\beta}_0 = -\sqrt{1-1/\gamma_0^2}\,\hat{\mathbf{x}} = -\beta_{\mathbf{0}}\,\hat{\mathbf{x}}$ is the three-velocity of the injected plasma.

As a byproduct of the shock evolution, particles and electromagnetic waves may propagate upstream from the shock at the velocity of light. A fixed simulation box with a reflecting boundary on the left and an open boundary on the right would not be suitable to study the long-term evolution of the shock, since particles and fields would be removed from the simulation domain when they cross the open boundary, thus artificially suppressing any feedback they may have on the shock. Moreover, in the first stages of shock evolution, physical quantities in a large portion of

the simulation box have not been affected yet by the presence of the shock; following their time evolution would imply a waste of computational resources, in terms of both time and memory. Instead, we inject particles and fields into the simulation box from a "moving injector", which recedes from the wall (at $x_{\rm wall} = 0$) along $+\hat{\mathbf{x}}$ at the speed of light, so that its position at time t is roughly $x_{\rm inj} \sim c\,t$. This ensures that the simulation domain consists only of those regions which are in causal contact with the initial setup, thus saving on computational time. Also, our simulation box expands (via discrete jumps along $+\hat{\mathbf{x}}$) whenever the injector approaches its right boundary, thus allowing to save memory. In summary, a moving injector with an expanding domain permits to follow the shock evolution as far as the computational resources allow, preserving all the particles and waves generated by the shock.

However, since the shock velocity is smaller than the light speed, the distance between the injector and the shock increases linearly with time. In all PIC codes, a numerical heating instability arises when cold relativistic plasma propagates for large distances over the numerical grid. The instability is a mixture of grid-Cerenkov instability – due to the fact that electromagnetic waves on a discrete grid have subluminal phase velocity for large wavenumbers – and spurious plasma oscillations (M. Dieckmann, private communication). We reduce this numerical instability by using the fourth-order solver for Maxwell's equations (Greenwood et al. 2004), which improves the numerical phase speed of the shortest wavemodes. To further get rid of the instability, we let our moving injector jump backwards (i.e., towards the wall) every 10000 time steps, resetting the electromagnetic fields beyond the injector to their initial values $\mathbf{B}_{\rm u}$ and $\mathbf{E}_{\rm u}$ and discarding the particles to the right of the injector; in the following, we refer to this as a "jumping injector". The first jump happens after 40000 time steps and we tune the jump according to the shock speed so that the injected plasma never propagates farther than $1000\,c/\omega_{\rm p}$ from the injector before encountering the shock. Although this suppresses the growth of the numerical instability – which arises if the injected plasma moves farther than $\sim 1500\,c/\omega_{\rm p}$ from the injector – it is not a completely satisfactory solution. In fact, each jump may remove from the computational domain some of the particles and waves propagating upstream, which would not happen if the injector were always receding at the speed of light. Also, electromagnetic waves moving away from the shock may be reflected back. We have tested the effects of a jumping injector by letting the first jump happen after 30000, 40000 and 50000 timesteps; in §4.2 we show that a jumping injector does not significantly affect our results.

3. Shock Structure

In this section we describe the structure and internal physics of relativistic magnetized pair shocks. We fix the upstream bulk Lorentz factor ($\gamma_0 = 15$) and magnetization ($\sigma = 0.1$) and vary only the obliquity angle θ, i.e., the angle between the upstream magnetic field and the shock normal, as measured in the wall frame. According to the criterion anticipated in §1, a shock is "superluminal" if particles cannot escape ahead of the shock by sliding along the magnetic field, and they are inevitably advected into the downstream medium. The critical boundary between subluminal and superluminal configurations should satisfy $\cos\theta'_{\rm crit} = \beta'_{\rm sh}$ in the *upstream* fluid frame, if the accelerated particles are assumed to move along the magnetic field lines at a velocity comparable to the speed of light; here, $\beta'_{\rm sh}$ is the shock speed as measured in the upstream frame. In the *wall* frame of our simulations, the critical obliquity angle between the upstream magnetic field and the shock normal will be $\theta_{\rm crit} = \arctan\left[\gamma_0 \tan\theta'_{\rm crit}\right]$, due to the Lorentz transformation of the magnetic field. Since $\beta'_{\rm sh} = (\beta_{\rm sh} + \beta_0)/(1 + \beta_{\rm sh}\beta_0)$, where $\beta_{\rm sh}$ is the shock speed in the simulation frame, the expression for the critical angle $\theta_{\rm crit}$ reduces to

$$\theta_{\rm crit} = \arctan\left[\frac{1}{\gamma_{\rm sh}\left(\beta_0 + \beta_{\rm sh}\right)}\right] \quad , \tag{2-1}$$

where $\gamma_{\rm sh} = 1/\sqrt{1-\beta_{\rm sh}^2}$ is the shock Lorentz factor in the simulation frame. For fixed upstream Lorentz factor γ_0 and magnetization σ, eq. (2-1) is an implicit equation for $\theta_{\rm crit}$, since the shock speed $\beta_{\rm sh}$ depends itself on the magnetic obliquity. For $\gamma_0 = 15$ and $\sigma = 0.1$, we find $\theta_{\rm crit} \approx 34^\circ$ by solving the MHD jump conditions. As shown in Fig. 2.2, the critical angle $\theta_{\rm crit}$ in the simulation frame decreases with increasing γ_0 and σ, but it stays confined within a relatively narrow range (between $\sim 26^\circ$ and $\sim 42^\circ$) for relativistic ($\gamma_0 \gtrsim 2$) flows with moderate magnetization ($\sigma \lesssim 1.0$). The rationale for such a weak dependence of $\theta_{\rm crit}$ on γ_0 or σ is that for $\gamma_0 \gtrsim 2$ and $\sigma \lesssim 1.0$ the shock speed does not depart significantly from the value $\beta_{\rm sh} = 1/3$ appropriate in the limit of ultrarelativistic flows ($\gamma_0 \to \infty$) with negligible magnetization ($\sigma \to 0$), when eq. (2-1) would formally yield $\theta_{\rm crit} \approx 35^\circ$.[2]

Among the representative sample of magnetic obliquities that we present in this section, $\theta = 0^\circ$ (Figs. 2.3 and 2.4), 15° (Figs. 2.5 and 2.6) and 30° (Figs. 2.7 and 2.8) are subluminal shocks, the last one being close to the critical obliquity angle $\theta_{\rm crit} \approx 34^\circ$; instead, $\theta = 45^\circ$ (Figs. 2.9 and 2.10) is a representative superluminal shock. We anticipate (see §6 for further discussion) that the main

[2]Of course, if the magnetization is too small the particles are no longer constrained to follow the magnetic field lines and no significant distinction between subluminal and superluminal configurations should persist.

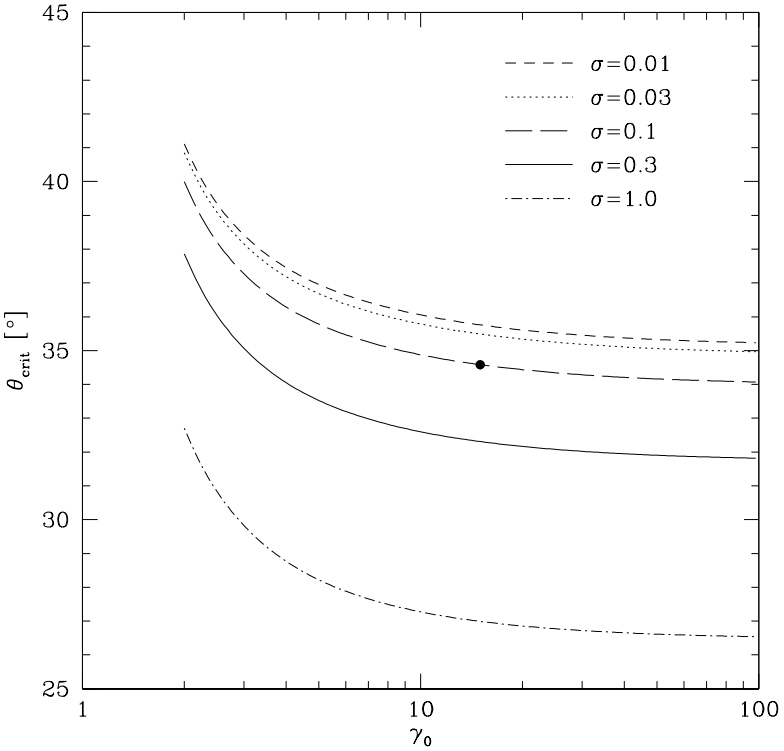


Fig. 2.2.— Critical obliquity angle $\theta_{\rm crit}$ (in the simulation frame) between subluminal and superluminal magnetic configurations as a function of the upstream Lorentz factor γ_0, for a representative sample of magnetizations of the upstream flow: $\sigma = 0.01$ (dashed line), $\sigma = 0.03$ (dotted line), $\sigma = 0.1$ (long-dashed line), $\sigma = 0.3$ (solid line) and $\sigma = 1.0$ (dot-dashed line). The solid circle is the value $\theta_{\rm crit} \approx 34.5°$ corresponding to the choice of $\gamma_0 = 15$ and $\sigma = 0.1$ adopted in this work.

results presented below for $\gamma_0 = 15$ and $\sigma = 0.1$ hold across a larger range of upstream Lorentz factors (we tried $\gamma_0 = 5$ and $\gamma_0 = 50$, for fixed $\sigma = 0.1$) and magnetizations (we tested $\sigma = 0.03$ and $\sigma = 0.3$, for fixed $\gamma_0 = 15$).

Figs. 2.3, 2.5, 2.7 and 2.9 present the shock internal structure as a function of the longitudinal coordinate x at time $\omega_{\rm p}t = 2250$. We include both two-dimensional plots in the xy simulation plane and transversely-averaged (along y) plots for the following quantities: particle number density (panel *a* for the two-dimensional plot and panel *b* for the corresponding y-average), magnetic energy fraction $\epsilon_{\rm B}$ (defined as the ratio of magnetic energy density to kinetic energy density of the incoming flow, panels *c* and *d*), transverse magnetic field $B_{\rm z}$ (panels *e* and *f*), transverse electric field $E_{\rm y}$ (panels *g* and *h*), longitudinal electric field $E_{\rm x}$ (panels *i* and *j*). We also plot the y-averaged electric pseudo-potential $-\int \langle E_{\rm x} \rangle \, {\rm d}x$ in panel *k*, normalized to the kinetic energy of injected particles. The positron phase space density projected onto the $x - \gamma\beta_{\rm x}$, $x - \gamma\beta_{\rm y}$ and $x - \gamma\beta_{\rm z}$ planes is shown in panels *l, m, n* respectively. Figs. 2.4, 2.6, 2.8, 2.10 show the positron $\gamma\beta_{\rm x} - \gamma\beta_{\rm z}$ momentum space and electron and positron energy spectra at three different locations through the flow, as marked by arrows at the bottom of Figs. 2.3, 2.5, 2.7, 2.9 respectively.

In all the simulation runs, a collisionless shock is formed which propagates away from the wall. The jump in density and electromagnetic fields at the shock, as well as the shock velocity, are in agreement with MHD calculations, as we show in §3.1. For oblique configurations ($0^\circ < \theta < 90^\circ$, i.e., with the exception of parallel and perpendicular shocks), the magnetic field refracts at the shock. In the frame of the simulation, the longitudinal speed of the incoming flow drops to zero after the shock; however, for oblique configurations, the downstream fluid preserves a small transverse velocity $\beta_{\rm z,d} < 0$, in agreement with MHD predictions. In this respect, the wall frame of our simulations does not strictly coincide with the downstream fluid frame, but for most purposes $\beta_{\rm z,d}$ is small enough that the two frames are almost undistinguishable. In the shock transition layer, the incoming cold beam isotropizes and thermalizes to a downstream temperature that relates to the upstream kinetic energy, reduced by the fraction of upstream particle energy converted to downstream magnetic energy. For low obliquities ($\theta \lesssim 45^\circ$), the downstream plasma is fully isotropic in three dimensions. For larger obliquities, the plane perpendicular to the downstream magnetic field, where particle orbits would lie in the absence of significant motion along the field, is almost degenerate with the direction of propagation of the incoming plasma. In this case, particle motions will mostly be confined to the plane orthogonal to the downstream field, which prevents efficient isotropization along the direction of the field. In other words, for low obliquities the effective

adiabatic index of the downstream plasma is appropriate for a three-dimensional relativistic gas ($\Gamma = 4/3$), whereas for $\theta \gtrsim 45^\circ$ the adiabatic index tends to that of a two-dimensional relativistic fluid ($\Gamma = 3/2$).

This section is structured as follows. In §3.1 we present a general description of the structure of oblique shocks. We discuss that Weibel-like filamentation instabilities are likely to mediate the shock formation for low magnetic obliquities, whereas magnetic reflection from the shock-compressed magnetic field triggers high-obliquity shocks. For $0^\circ < \theta < 90^\circ$, the longitudinal profile of particle density, magnetic energy and transverse magnetic field shows a fluid structure with two shocks, a strong "fast" shock (which we identify with the main shock) and a weak "slow" shock. The velocity and jump conditions of the fast shock are in agreement with MHD calculations. At the end of §3.1 we highlight some differences in the shock structure between our two-dimensional simulations and previous one-dimensional studies of perpendicular shocks.

In §3.2 and §3.3 we describe in detail some representative examples of subluminal ($0 \leq \theta \lesssim \theta_{\rm crit}$) and superluminal ($\theta_{\rm crit} \lesssim \theta \leq 90^\circ$) shocks. In subluminal configurations, some shock-accelerated particles escape ahead of the shock, and they populate a diffuse beam of high-energy particles propagating upstream. These "returning" particles trigger the generation of upstream waves, which propagate towards the shock with a wavevector oblique to the background magnetic field. In superluminal shocks, no returning high-energy particles or upstream waves are observed. The importance of these oblique modes for particle acceleration will be discussed in §5, where we present a detailed investigation of the acceleration mechanisms for subluminal shocks.

3.1. Oblique Shocks: General Overview

Unmagnetized relativistic shocks are mediated by Weibel-like filamentation instabilities (Weibel 1959; Medvedev & Loeb 1999; Gruzinov & Waxman 1999). Incoming particles are scattered by the Weibel-generated fields and the incoming flow is isotropized, decelerated and compressed (e.g., Spitkovsky 2005). The same mechanism is likely to mediate moderately magnetized ($\sigma \lesssim 1$) parallel shocks. In fact, we find that the width of the transition layer in $\sigma = 0.1$ parallel shocks, as evaluated from the density profile (Fig. 2.3 *b* at $x_{\rm sh} \approx 730\, c/\omega_{\rm p}$), is of the order of a few tens of

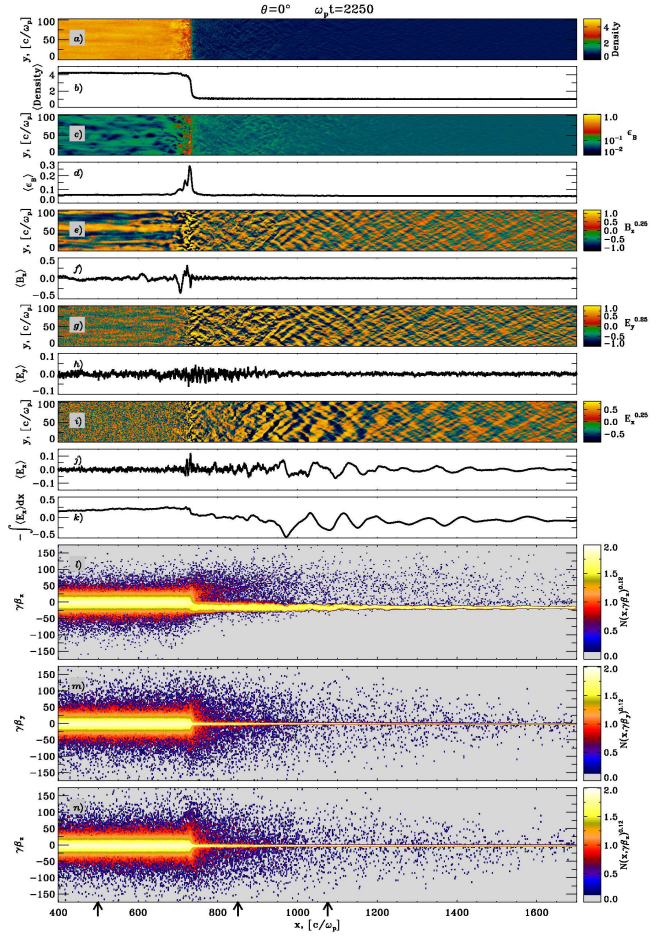


Fig. 2.3.— Shock structure and positron phase space at time $\omega_{\rm p}t = 2250$ for $\theta = 0°$: *a*) Particle number density in the simulation plane, normalized to the upstream value; *b*) Transversely-averaged number density; *c*) Magnetic energy fraction $\epsilon_{\rm B}$, i.e., the ratio of magnetic energy density to the kinetic energy density of the incoming flow; *d*) Transversely-averaged $\langle\epsilon_{\rm B}\rangle$; *e*)-*f*) $B_{\rm z}$ and transversely-averaged $\langle B_{\rm z}\rangle$, in units of the upstream background magnetic field; *g*)-*h*) $E_{\rm y}$ and transversely-averaged $\langle E_{\rm y}\rangle$, in units of the upstream background magnetic field; *i*)-*j*) $E_{\rm x}$ and transversely-averaged $\langle E_{\rm x}\rangle$, in units of the upstream background magnetic field; *k*) Electric pseudo-potential $-\int\langle E_{\rm x}\rangle\,{\rm d}x$, normalized to the kinetic energy of incoming particles; *l*)-*n*) Longitudinal phase space plots $x - \gamma\beta_i$ ($i = x, y, z$) of positrons, with phase space density plotted as a two-dimensional histogram. For the two-dimensional plots of $B_{\rm z}$, $E_{\rm y}$ and $E_{\rm x}$, q^α stands for $(q/|q|)|q|^\alpha$.

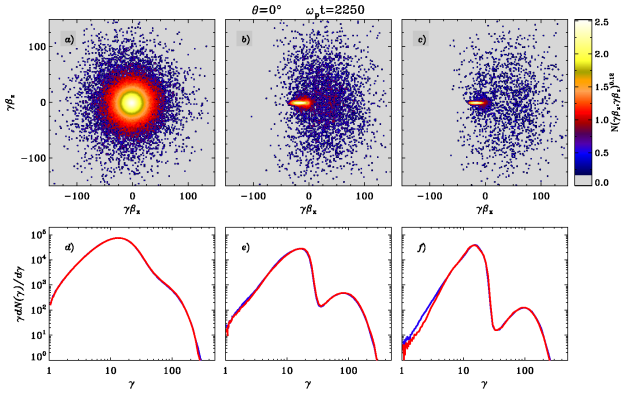


Fig. 2.4.— Positron momentum space and electron (blue) and positron (red) spectra at time $\omega_{\rm p}t = 2250$ for $\theta = 0°$, at three slices through the flow as marked by arrows at the bottom of Fig. 2.3. *a*)-*c*) Momentum space $\gamma\beta_{\rm x} - \gamma\beta_{\rm z}$ in the three slabs, with momentum space density plotted as a two-dimensional histogram; *d*)-*f*) Electron (blue) and positron (red) spectra $\gamma\,{\rm d}N/{\rm d}\gamma$ in the three slices, normalized to the total number of particles within the leftmost slab (apart from a constant multiplication factor).

skin depths, as observed in unmagnetized shocks.[3] Also, the magnetic energy density peak at the shock (reaching $\sim 30\%$ of the upstream kinetic energy density, see Fig. 2.3 *d*) is a characteristic signature of magnetic field amplification by a Weibel-like process. The field then decays to the plateau predicted by MHD calculations on a length scale $\sim 50\,c/\omega_{\rm p}$.

Oblique shocks ($\theta \neq 0°$) may be triggered by another process, namely magnetic reflection of the incoming flow from the shock-compressed magnetic field (Alsop & Arons 1988), as reported for perpendicular shocks by previous one-dimensional (e.g., Langdon et al. 1988; Gallant et al. 1992) and three-dimensional (Spitkovsky 2005) PIC simulations. The synchrotron maser instability, excited by coherent gyration of the incoming particles in the shock magnetic field, mediates the thermalization of the upstream plasma (Hoshino & Arons 1991). Indeed, the peak in the transversely-averaged density profile observed for $\theta = 45°$ at the shock location (see Fig. 2.9 *b* at $x_{\rm sh} \approx 920\,c/\omega_{\rm p}$) is due to Larmor gyration of the incoming particles in the compressed shock fields, which decelerates the upstream flow and therefore increases its density. The density rise is as sharp as a few Larmor radii in the compressed field (i.e., a few plasma skin depths), much thinner than for parallel shocks (compare Fig. 2.9 *b* with Fig. 2.3 *b*). In just a few Larmor scales downstream

[3]We point out that the transition region in strongly magnetized ($\sigma \gtrsim 1$) parallel shocks is much wider ($\sim 100\,c/\omega_{\rm p}$), since cross-field motion is suppressed and the Weibel instability is quenched.

from the shock, the density profile flattens to a plateau in reasonable agreement with MHD jump conditions. Shocks with larger obliquities ($\theta \gtrsim 45^\circ$) are also triggered by magnetic reflection.

Both Weibel instabilities and magnetic reflection play a role in mediating shocks with intermediate obliquities ($0^\circ < \theta \lesssim 45^\circ$). Filamentation instabilities dominate for lower obliquities, such as $\theta = 15^\circ$, whose fluid structure resembles $\theta = 0^\circ$ in the magnetic overshoot at the shock (Fig. 2.5*d* at $x_{\rm sh} \approx 720\, c/\omega_{\rm p}$). Instead, such Weibel-induced magnetic energy excess is not seen for $\theta = 30^\circ$ (Fig. 2.7*d* at $x_{\rm sh} \approx 770\, c/\omega_{\rm p}$) or larger obliquities; rather, investigation of early times ($\omega_{\rm p} t = 225$) for $\theta = 30^\circ$ shows that this shock started out with the prominent density peak characteristic of highly-oblique shocks, which suggests that this shock is mainly triggered by magnetic reflection. For $\theta = 30^\circ$, the sharp density peak observed at early times was later smoothed by the increased pressure in shock-accelerated particles.

For all oblique magnetic configurations ($0^\circ < \theta < 90^\circ$), a fluid structure with two shocks (a strong "fast" shock and a weak "slow" shock) is seen in our PIC simulations. The presence of two shocks is predicted for the parameters ($\gamma_0 = 15$ and $\sigma = 0.1$) that we adopt (Landau et al. 1984) and has been confirmed by MHD simulations of the collision of a magnetized shell against a wall (S. Komissarov, private communication). The particle number density rises across both shocks; however, while the fast shock involves an amplification of the magnetic field, across the slow shock the magnetic energy decreases. The slow shock propagates into the downstream medium of the fast shock. Its transition layer in our PIC simulations is not as narrow as MHD would predict, due to insufficient dissipation along the magnetic field. Rather, the slow-shock transition region stretches self-similarly with time; the jump in fluid quantities between the wall and the fast-shock downstream stays roughly constant in time. We observe that the slow shock is slower and weaker as θ approaches either 0° or 90°, and relatively stronger and faster for intermediate obliquities ($\theta \approx 30^\circ$). For oblique magnetic configurations, we shall refer to the fast shock simply as "the shock" and to its downstream region as "the downstream medium".

For the main (i.e., fast) shock, the velocity and jump conditions extracted from our PIC simulations agree within a few percent with the solution of relativistic MHD equations (Appl & Camenzind 1988) supplemented by the *three-dimensional* Synge (1957) equation of state, provided that MHD results are Lorentz-transformed from the standard shock frame into the simulation

frame.[4] In Table 2.1 we compare the results from PIC simulations and MHD calculations for the shock velocity $\beta_{\rm sh}$ and the jump across the shock in particle number density ($n_{\rm d}/n_{\rm u}$), magnetic energy density ($B_{\rm d}^2/B_{\rm u}^2$), transverse magnetic field ($B_{\rm z,d}/B_{\rm z,u}$) and transverse electric field ($E_{\rm y,d}/E_{\rm y,u}$). We remark that, for oblique shocks ($0^\circ < \theta < 90^\circ$), the background motional electric field (last column in Table 2.1) does not drop to zero after the shock since, as discussed above, the downstream flow preserves a small transverse velocity $\beta_{\rm z,d} < 0$ which gives a motional field $E_{\rm y,d} = -\beta_{\rm z,d} B_{\rm x,d} > 0$.

Table 2.1: Comparison between PIC simulations (upper row) and MHD calculations (lower row)

θ	$\beta_{\rm sh}$	$n_{\rm d}/n_{\rm u}$	$B_{\rm d}^2/B_{\rm u}^2$	$B_{\rm z,d}/B_{\rm z,u}$	$E_{\rm y,d}/E_{\rm y,u}$
0°	0.32	4.1	1.1		
	0.31	4.2	1.0		
15°	0.32	4.0	2.4	4.3	-0.08
	0.32	4.1	2.3	4.5	-0.16
30°	0.34	3.8	5.2	4.3	-0.10
	0.35	3.9	5.1	4.2	-0.11
45°	0.41	3.2	6.5	3.6	-0.06
	0.37	3.7	7.8	3.8	-0.07

As shown in Table 2.1, for $\theta \lesssim 30^\circ$ our kinetic calculations are in remarkable agreement with *three-dimensional* MHD results, suggesting that for low magnetic obliquities the downstream plasma in our *two-dimensional* simulation box has a *three-dimensional* effective adiabatic index, as anticipated above. For $\theta \gtrsim 45^\circ$, the adiabatic index of the downstream fluid in our simulations tends to that of a two-dimensional gas, which explains the discrepancy between PIC simulations and three-dimensional MHD calculations for $\theta = 45^\circ$ in Table 2.1.[5]

Comparison with Previous 1D Simulations

A common feature of all oblique and perpendicular shocks ($0^\circ < \theta \leq 90^\circ$) that we investigate is the presence of an electromagnetic precursor (seen in $B_{\rm z}$ and $E_{\rm y}$ for the magnetic configuration

[4]Shock speed and jump conditions from our PIC simulations are computed at relatively early times ($\omega_{\rm p} t = 2250$), since for later stages of the shock evolution the pressure by accelerated particles may significantly affect the shock structure.

[5]We point out that the downstream adiabatic index in high-obliquity shocks tends to that of a two-dimensional fluid also in *three-dimensional* PIC simulations, since particle isotropization along the field is still harder than orthogonal to the field. However, in our two-dimensional runs, the simulation plane for high obliquities is almost degenerate with the plane orthogonal to the downstream field, where particle gyro-orbits lie, and this makes such isotropization even harder.

shown in Fig. 2.1) propagating at the speed of light ahead of the shock, as reported in previous one-dimensional (e.g., Gallant et al. 1992; Hoshino et al. 1992) and short three-dimensional (Spitkovsky 2005) PIC simulations of perpendicular shocks. This is a transverse electromagnetic wave which is generated by coherent bunching of particles on Larmor orbits near the shock front. In the one-dimensional simulations by Gallant et al. (1992) and Hoshino et al. (1992), the electromagnetic precursor is generated continuously by the shock, and it is speculated that its radiation pressure may affect the acceleration of particles (Hoshino 2008). In our long multi-dimensional simulations, we observe that the generation of this electromagnetic precursor is confined to the initial stages of shock evolution, and it ceases as soon as the coherence of particle reflections between different locations along the shock surface is lost. The transient precursor generated at early times survives as a dispersive wave packet propagating upstream, farther from the shock than Figs. 2.5, 2.7, 2.9 show. We find that its power increases with magnetic obliquity and its Poynting flux may be comparable or in excess of $c\,B_{\rm z,u}^2/4\pi$. Although the electromagnetic precursor seems to have little influence on the structure of pair plasma shocks, its effect may be more pronounced in electron-ion shocks, because of the electrostatic charge separation induced in the upstream plasma by the radiation push on the electrons (Lyubarsky 2006).

3.2. Subluminal Shocks: $0 \leq \theta \lesssim \theta_{\rm crit}$

The internal structure at time $\omega_{\rm p}t = 2250$ of shocks with magnetic obliquities $\theta = 0^\circ$ (Figs. 2.3, 2.4), 15° (Figs. 2.5, 2.6) and 30° (Figs. 2.7, 2.8) shows some features common to all subluminal magnetic configurations. Subluminal shocks are characterized by the presence of a diffuse stream of high-energy shock-accelerated particles that propagate ahead of the shock and trigger the growth of oblique waves in the upstream medium. In turn, as we explain in §5, such waves are of fundamental importance for particle acceleration.

The longitudinal phase space plots of positrons in panels *l*-*n* of Figs. 2.3, 2.5, 2.7 show the injected plasma as a cold dense beam propagating with $\gamma\beta_{\rm x} \approx -15$ and broadening while it approaches the shock, which is located at $x_{\rm sh} \approx 730\,c/\omega_{\rm p}$ for $\theta = 0^\circ$, at $x_{\rm sh} \approx 720\,c/\omega_{\rm p}$ for $\theta = 15^\circ$ and at $x_{\rm sh} \approx 770\,c/\omega_{\rm p}$ for $\theta = 30^\circ$. In the shock transition layer, the incoming plasma is isotropized and thermalized. A population of high-energy particles moving ahead of the shock appears in the upstream region as a hot diffuse beam with $\gamma\beta_{\rm x} > 0$ (panel *l* in Figs. 2.3, 2.5, 2.7); as we show in §5, these particles have been accelerated at the shock. Additional insight is provided by the positron

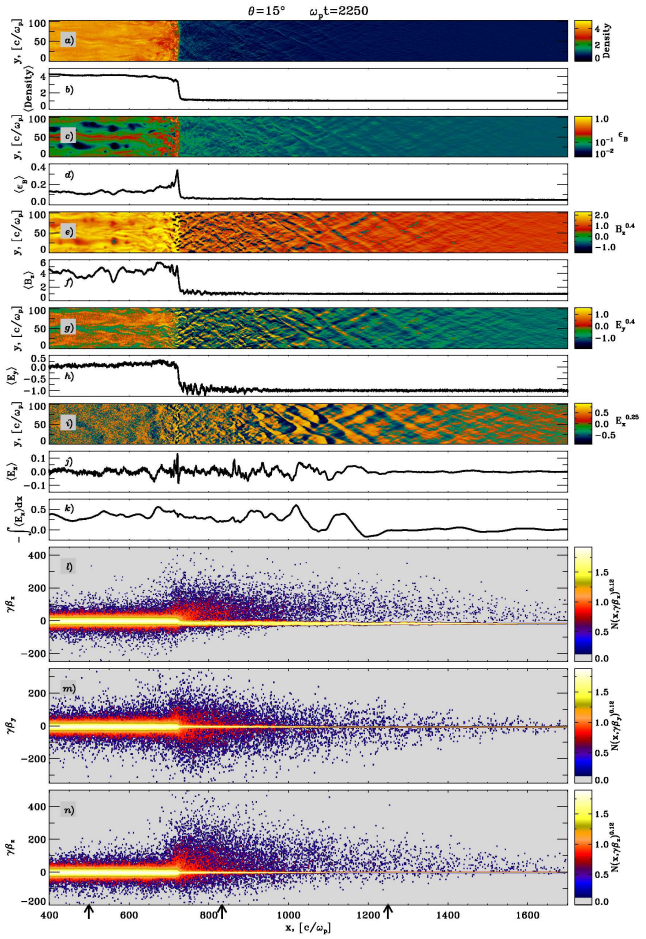


Fig. 2.5.— Shock structure and positron phase space at time $\omega_{\rm p}t = 2250$ for $\theta = 15°$. The fluid quantities are normalized as in Fig. 2.3, but here $B_{\rm z}$ and $E_{\rm y}$ are in units of the transverse component $B_{\rm z,u}$ of the upstream background magnetic field. The excess of magnetic energy at the shock (panel d at $x_{\rm sh} \approx 720\, c/\omega_{\rm p}$) does not show up in the transversely-averaged $\langle B_{\rm z}\rangle$ (panel f) due to cancellation between magnetic fluctuations of opposite sign.

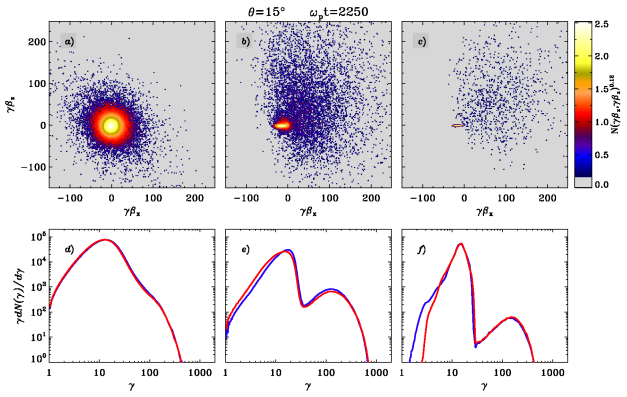


Fig. 2.6.— Positron momentum space and electron (blue) and positron (red) spectra at time $\omega_{\rm p}t = 2250$ for $\theta = 15°$, at three slices through the flow as marked by arrows at the bottom of Fig. 2.5. See the caption of Fig. 2.4 for details.

$\gamma\beta_{\rm x} - \gamma\beta_{\rm z}$ momentum space and electron (blue) and positron (red) spectra in Figs. 2.4, 2.6, 2.8, extracted from three different locations across the flow as marked by arrows at the bottom of Figs. 2.3, 2.5, 2.7 respectively. Upstream slabs (panels *c* and *f* in Figs. 2.4, 2.6, 2.8) show two components: the cold incoming plasma appears in particle spectra (panel *f* in Figs. 2.4, 2.6, 2.8) as a relatively narrow bump peaking at low energies, around $\gamma \approx 15$, which broadens as the flow approaches the shock; the returning diffuse beam of shock-accelerated particles populates the high-energy bump, which is more and more depleted (especially at intermediate energies) as we move farther upstream. Far downstream slabs (panels *a* and *d* in Figs. 2.4, 2.6, 2.8) show roughly isotropic particle distributions with thermal spectra supplemented by high-energy tails. For $\theta = 30°$, the downstream population of energetic particles significantly varies along x, with regions closer to the shock showing more particles of higher energies than slices farther downstream (Fig. 2.7*l*-*n*, and compare panels *a*,*d* with panels *b*,*e* in Fig. 2.8). Since fluid elements which are farther downstream passed through the shock at earlier times, this suggests that the shock acceleration efficiency increased over time, as we discuss in §4.2.

For oblique magnetic configurations (Fig. 2.5 for $\theta = 15°$ and Fig. 2.7 for $\theta = 30°$), the phase space plots $x - \gamma\beta_{\rm y}$ (panel *m*) and $x - \gamma\beta_{\rm z}$ (panel *n*) are not symmetric around $\gamma\beta_i = 0$ ($i = y$, z), when focusing on the high-energy particles located in the vicinity of the shock (see also panels *b*, *c* in Figs. 2.6, 2.8). The high-energy positrons with $\gamma\beta_{\rm z} > 0$ in the shock and upstream regions are shock-accelerated particles whose gyro-center is heading upstream along the oblique magnetic

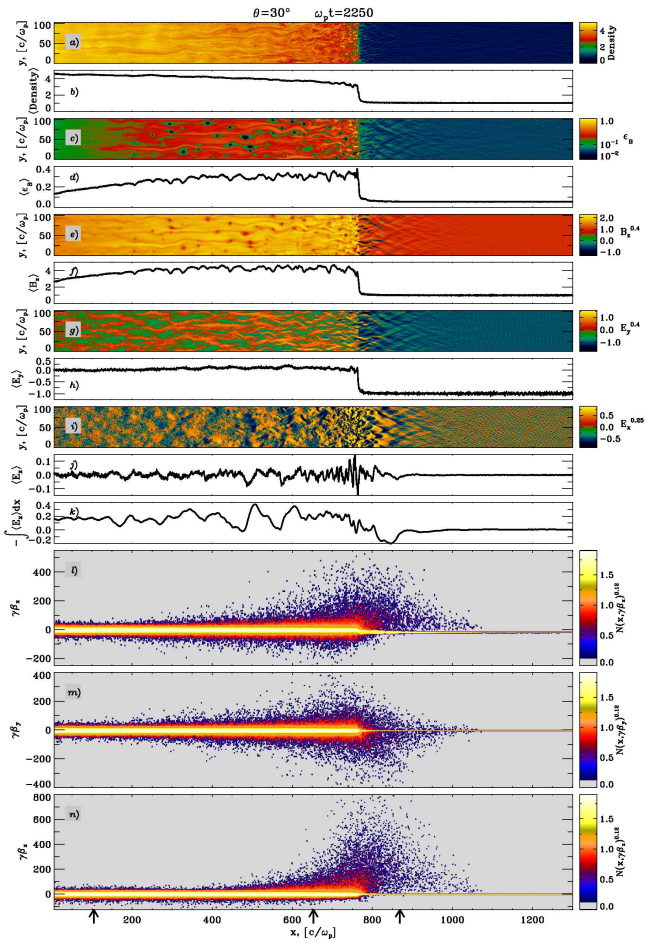


Fig. 2.7.— Shock structure and positron phase space at time $\omega_p t = 2250$ for $\theta = 30°$. The fluid quantities are normalized as in Fig. 2.5.

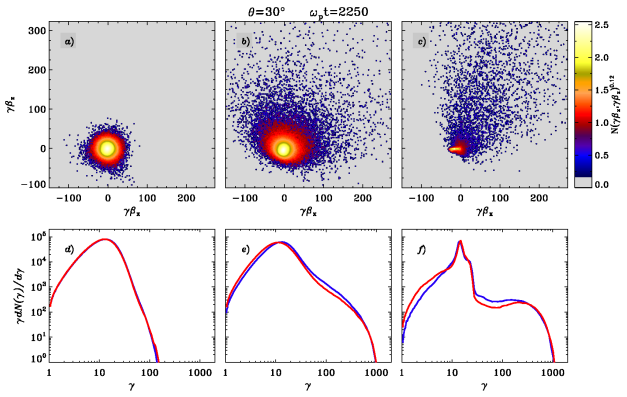


Fig. 2.8.— Positron momentum space and electron (blue) and positron (red) spectra at time $\omega_{\rm p}t = 2250$ for $\theta = 30°$, at three slices through the flow as marked by arrows at the bottom of Fig. 2.7. See the caption of Fig. 2.4 for details.

field (we remind that $B_{\rm x}, B_{\rm z} > 0$).[6] With increasing magnetic obliquity from $\theta = 15°$ to $30°$, the asymmetry around $\gamma\beta_{\rm z} = 0$ becomes more evident and the returning high-energy positrons show increasingly larger $\gamma\beta_{\rm z}$ than $\gamma\beta_{\rm x}$ (compare panels l and n between Fig. 2.5 and Fig. 2.7). The electron phase space $x - \gamma\beta_{\rm z}$ (not plotted) shows the same asymmetry.

The excess of high-energy positrons with $\gamma\beta_{\rm y} < 0$ observed in the vicinity of the shock (panel m of Figs. 2.5 and 2.7) is a natural consequence of the ∇B drift experienced by particles accelerated at the shock. For positively-charged particles, the drift is directed along $-\hat{\mathbf{y}}$, which explains the sign of the asymmetry in Figs. 2.5 m, 2.7 m. Since the magnetic field jump at the shock is larger for increasing obliquities (Table 2.1), and the drift velocity is then higher, we expect the asymmetry to be more significant for $\theta = 30°$ than for $15°$, as observed. Moreover, since electrons and positrons move in opposite directions under ∇B drift, the electron $x - \gamma\beta_{\rm y}$ phase space (not shown) exhibits a reversed asymmetry, with an excess of particles having $\gamma\beta_{\rm y} > 0$.

The returning beam of high-energy particles triggers in the upstream medium the generation of waves whose wavevector is oblique to the background magnetic field. Such waves have an electromagnetic component ($B_{\rm z}$ in panel e and $E_{\rm y}$ in panel g of Figs. 2.3, 2.5, 2.7) and a smaller

[6]For $\theta = 30°$, the x-velocity of a particle sliding at the speed of light along the upstream magnetic field would be $\approx 0.5\,c$ in the simulation frame. This is compatible with the longitudinal location $x \approx 1100\,c/\omega_{\rm p}$ at time $\omega_{\rm p}t = 2250$ for the returning particles which are farthest upstream (Fig. 2.7l-n). The same is true for smaller obliquities, but the leading edge of the population of returning particles is outside the x-range plotted in Figs. 2.3 and 2.5.

electrostatic component (E_x in panel i of Figs. 2.3, 2.5, 2.7). The electrostatic component also appears as upstream oscillations in the transversely-averaged $\langle E_x \rangle$ (panel j of Figs. 2.3, 2.5, 2.7) and in the electric pseudo-potential $-\int \langle E_x \rangle \, dx$ (panel k of Figs. 2.3, 2.5, 2.7). The upstream waves grow in amplitude while propagating towards the shock, where the regular oblique pattern observed far upstream (with characteristic wavelength $\sim 20\, c/\omega_p$ at time $\omega_p t = 2250$) changes to a more complicated structure (e.g., compare $x \gtrsim 1000\, c/\omega_p$ with $750\, c/\omega_p \lesssim x \lesssim 900\, c/\omega_p$ for $\theta = 15^\circ$ in Fig. 2.5*e,g,i*). The ratio between the magnetic energy of the oblique waves (mostly contributed by their B_z) and of the background upstream field is of the order of a few tens of percent, roughly constant with magnetic obliquity.

By comparing panels *e,g,i* with panels *l-n* of Fig. 2.7 for $\theta = 30^\circ$, the growth of the upstream oblique modes appears to be spatially coincident (at $x \approx 1000\, c/\omega_p$) with the leading edge of the population of returning particles. The same holds true for smaller obliquities, but for $\theta = 0^\circ$ and 15° the returning particles that are farthest upstream are outside the x-range plotted in Figs. 2.3 and 2.5. This suggests that interplay between the unperturbed incoming plasma and the returning high-energy beam may be responsible for triggering the oblique waves.

The growth of the upstream waves is neither an unphysical consequence of our jumping injector nor an effect of the numerical heating associated with the propagation of a cold beam through the grid. In fact, if we replace the reflecting wall at $x_{wall} = 0$ with an open boundary, everything else being equal, we observe that no upstream waves are produced, nor is there a shock. In order to test the role played by the returning high-energy particles in generating the upstream waves, we have performed a suite of simulations in which we artificially suppress particle acceleration. In these runs, we introduced cooling, so that particles with $\gamma > \gamma_{cool}$ lose a random fraction of their nonthermal energy (a similar technique was used in Keshet et al. 2008). By varying γ_{cool} (we tried $\gamma_{cool} = 50$, 80, 140 for $\theta = 0^\circ$), we conclude that particle cooling suppresses the growth of the upstream waves, and by decreasing γ_{cool} (i.e., by forcing more particles to cool) both the wave amplitude and the distance from the shock where they are triggered decrease.

In summary, the oblique waves seen in the upstream medium are generated by the returning beam of high-energy particles; they may result from a mixed-mode streaming instability which couples the electrostatic two-stream instability with the electromagnetic filamentation (Weibel) instability, as envisioned by Bret & Dieckmann (2008). A complete analysis of the wave dispersion relation and generation mechanism is deferred to a future work.

Finally, we point out that the electrostatic component of the upstream waves induces oscillations in the $\gamma\beta_{\rm x}$-velocity of the incoming plasma, as seen in the positron phase space of Fig. 2.3*l*. The electron $x-\gamma\beta_{\rm x}$ phase space (not shown) presents oscillations of opposite sign. The difference at low energies between electron and positron spectra in the upstream medium (panel *f* in Figs. 2.4, 2.6, 2.8) can be correlated with the sign of the electric pseudo-potential at the same longitudinal location (panel *k* in Figs. 2.3, 2.5, 2.7 respectively), since a net electric potential accelerates incoming charges of one sign and decelerate charges of the opposite sign. We further comment on the effects on particle spectra of the electrostatic component of upstream oblique waves in §4.2.

3.3. Superluminal Shocks: $\theta_{\rm crit} \lesssim \theta \leq 90^\circ$

Figs. 2.9 and 2.10 show the shock structure, phase space plots and particle spectra for a representative superluminal shock ($\theta = 45^\circ$) at time $\omega_{\rm p}t = 2250$. No significant particle acceleration occurs in superluminal shocks, and the upstream plasma is not appreciably perturbed by the presence of the shock.

In the phase space plot of Fig. 2.9 *l*, the incoming plasma is seen as a cold beam with $\gamma\beta_{\rm x} \approx -15$, which is slightly heated while approaching the shock and then sharply thermalized in the shock transition layer at $x_{\rm sh} \approx 920\, c/\omega_{\rm p}$. As discussed by Hoshino & Arons (1991) and anticipated in §3.1, the incoming plasma thermalizes due to the absorption of cyclotron waves produced by coherent particle gyration in the compressed shock fields. Thermalization proceeds with time in the downstream medium: panels *l* and *m* of Fig. 2.9 show that the plasma farther downstream, which had more time to thermalize after passing through the shock, has a broader $x-\gamma\beta_{\rm x}$ and $x-\gamma\beta_{\rm y}$ distribution than the fluid closer to the shock.

Downstream particles preferentially move in a plane perpendicular to the downstream oblique field; indeed, the oblique boxy shape of their downstream $\gamma\beta_{\rm x}-\gamma\beta_{\rm z}$ momentum space (Figs. 2.10 *a*, *b* for two different downstream slabs) is suggestive of particle gyro-motion around the oblique field. As discussed at the beginning of §3, for magnetic obliquities $\theta \gtrsim 45^\circ$ it is harder for downstream particles to isotropize along the magnetic field rather than in the plane orthogonal to it. Since for high obliquities such plane approaches the xy simulation plane, this should result in the downstream thermal width of the $x-\gamma\beta_{\rm z}$ phase space being smaller than the $x-\gamma\beta_{\rm x}$ or $x-\gamma\beta_{\rm y}$ distributions, as observed (compare panel *n* with panels *l* and *m* in Fig. 2.9).

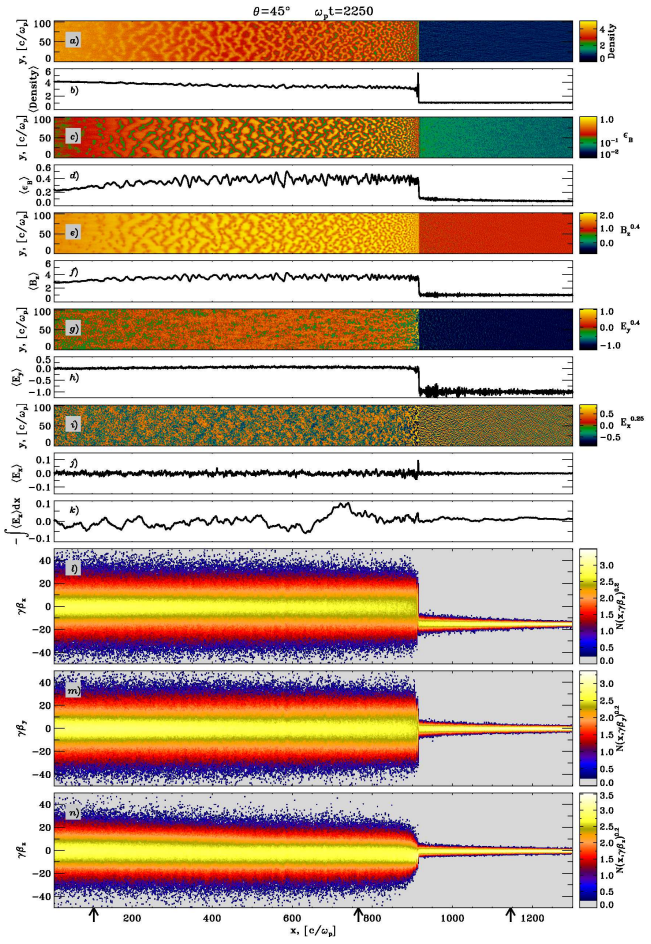


Fig. 2.9.— Shock structure and positron phase space at time $\omega_p t = 2250$ for $\theta = 45°$. The fluid quantities are normalized as in Fig. 2.5.

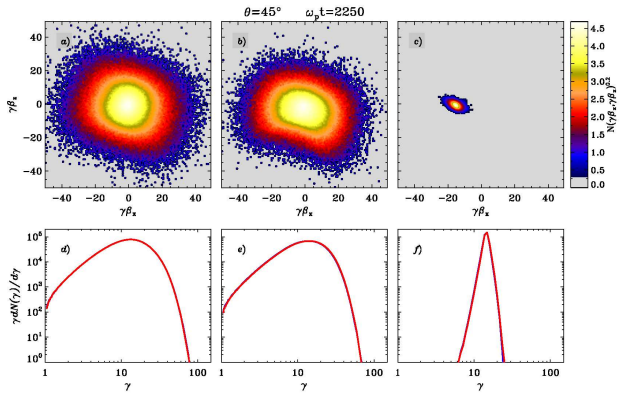


Fig. 2.10.— Positron momentum space and electron (blue) and positron (red) spectra at time $\omega_{\rm p}t = 2250$ for $\theta = 45°$, at three slices through the flow as marked by arrows at the bottom of Fig. 2.9. See the caption of Fig. 2.4 for details.

The net fluid velocity $\beta_{\rm z,d} < 0$ behind the shock in Fig. 2.9 *n* is responsible for the residual motional electric field $E_{\rm y,d} > 0$ observed in the simulation frame downstream from the shock (hard to see in Fig. 2.9 *h*, but see Table 2.1), a feature common to all oblique magnetic configurations, as discussed in §3. This residual fluid velocity vanishes farther downstream, in regions already overtaken by the slow shock. The slow shock appears as a smooth increase in density (Fig. 2.9 *b*) and decrease in magnetic energy (Fig. 2.9 *d*) and $B_{\rm z}$ (Fig. 2.9 *f*) behind the main shock towards the wall. Similar properties for the slow shock are observed for $\theta = 30°$ in Fig. 2.7, whereas the slow shock for $\theta = 15°$ is outside the x-range shown in Fig. 2.5.

Most notably, no evidence for particle acceleration is present for $\theta = 45°$. No high-energy particles are observed to propagate upstream from the shock (Fig. 2.9 *l*-*n*), and the particle energy spectrum in upstream slices is just the thermal distribution of the injected plasma Lorentz-boosted along $-\hat{\mathbf{x}}$ (Fig. 2.10 *f*). The range of downstream 4-velocities is much smaller than for subluminal shocks (compare panels *l*-*n* in Fig. 2.9 with the same panels in Figs. 2.3, 2.5, 2.7), and the downstream energy spectrum (Fig. 2.10 *d*, *e*) resembles a purely thermal distribution without any suprathermal tail. The lack of returning particles explains why no oblique waves are seen in the upstream medium of superluminal shocks (compare panels *e*, *g*, *i* in Fig. 2.9 with the same panels in Figs. 2.3, 2.5, 2.7).

We remark that the main features discussed here for $\theta = 45°$, and most importantly the absence of shock-accelerated particles, are common to all superluminal configurations up to the

perpendicular case $\theta = 90^\circ$ investigated via PIC simulations by, e.g., Gallant et al. (1992) and Spitkovsky (2005).

4. Particle Energy Spectra

We now investigate particle acceleration in subluminal and superluminal shocks by studying the downstream energy spectra of our representative sample of magnetic obliquities: $\theta = 0^\circ$, 15°, 30° (subluminal shocks) and $\theta = 45^\circ$ (superluminal shock). We remind that the boundary between subluminal and superluminal configurations is $\theta_{\rm crit} \approx 34^\circ$.

In §4.1 we show that subluminal shocks ($0 \leq \theta \lesssim \theta_{\rm crit}$) produce a population of high-energy particles as part of the shock evolution. Such shock-accelerated particles populate a suprathermal power-law tail in the energy spectrum behind the shock. As the magnetic obliquity increases within the subluminal range, the high-energy tail becomes flatter and it accounts for an increasing fraction of downstream particle number and energy. In other words, the acceleration efficiency increases with magnetic obliquity for $\theta \lesssim \theta_{\rm crit}$. Instead, the downstream particle spectrum for superluminal shocks, $\theta_{\rm crit} \lesssim \theta \leq 90^\circ$ (including perpendicular shocks), does not show any significant suprathermal tail, meaning that particle acceleration is extremely inefficient for superluminal configurations.

In §4.2 we follow the time evolution of downstream particle spectra and show that the slope and normalization of the high-energy tail in high-obliquity – yet still subluminal – shocks varies with time much more than for low obliquities. This suggests that a different acceleration mechanism is operating in shocks close to the superluminality threshold as compared to quasi-parallel shocks, as we discuss in §5.

In the following, all the particle spectra are computed in a slab $200\,c/\omega_{\rm p}$-wide which is centered $400\,c/\omega_{\rm p}$ downstream from the shock, regardless of time or magnetic obliquity. Strictly speaking, our particle spectra are not computed in the downstream fluid frame, due to the net fluid velocity $\beta_{\rm z,d}$ observed in the simulation frame behind oblique shocks; however, $|\beta_{\rm z,d}| \lesssim 0.1$, so that our wall-frame particle spectra are an excellent proxy for their downstream-frame counterparts. In order to estimate the acceleration efficiency of subluminal shocks, we fit the downstream particle spectrum with a three-dimensional Maxwellian (as appropriate for $\theta \lesssim 45^\circ$ shocks, see §3) with spread $\Delta\gamma_{\rm MB}$ plus a power law starting at $\gamma_{\rm min}$ with exponential cutoff at $\gamma_{\rm cut}$ whose width is $\Delta\gamma_{\rm cut}$. The fitting function is $f(\gamma) = N_{\rm MB}\gamma\sqrt{\gamma^2-1}\exp(-\gamma/\Delta\gamma_{\rm MB}) + N_{\rm PL}\gamma^{-p}\min[1, \exp(-(\gamma-\gamma_{\rm cut})/\Delta\gamma_{\rm cut})]$, with

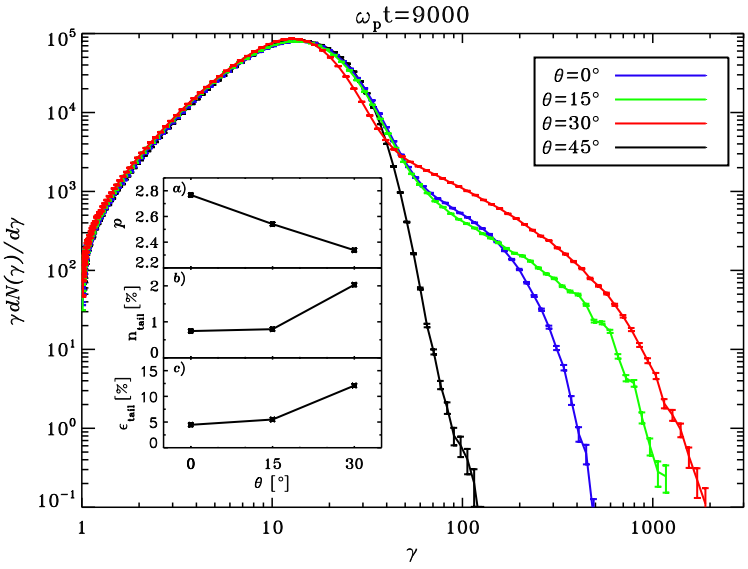


Fig. 2.11.— Downstream particle energy spectra at time $\omega_p t = 9000$ for different magnetic obliquities: $\theta = 0°$ (blue), 15° (green), 30° (red) and 45° (black). All spectra are computed in a slab $200\,c/\omega_p$-wide centered $400\,c/\omega_p$ downstream from the shock. The spectrum is normalized to the total number of particles within the considered slab, apart from a constant multiplication factor; it is supplemented with Poissonian error bars. Subpanels: *a*) Power-law slope of the high-energy tail as a function of the obliquity angle θ. The tail is defined to start at the intersection between the fitting Maxwellian and the fitting power law. *b*) Fraction of particles (in percent) in the high-energy tail. *c*) Fraction of energy (in percent) stored in the high-energy tail.

$N_{\rm PL} = 0$ for $\gamma < \gamma_{\rm min}$. The high-energy tail is defined to start at the intersection of the fitting Maxwellian and the fitting power law.

4.1. Dependence on Magnetic Obliquity

Fig. 2.11 shows the main result of our work, namely the dependence of the acceleration efficiency on magnetic obliquity: downstream spectra for $\theta = 0°$ (blue), 15° (green), 30° (red) and 45° (black) are shown at the final time $\omega_p t = 9000$ of our simulation runs. Subpanel *a* shows the slope p of the fitting power law as a function of magnetic obliquity, while subpanels *b* and *c* present how the fraction of particles and energy stored in the suprathermal tail depends on the magnetic inclination.

For subluminal shocks, the downstream energy spectrum consists of a low-energy thermal bump and a high-energy suprathermal tail. The low-energy part, which accounts for most of the energy, is populated by the particles that are directly transmitted downstream, where they are isotropized and thermalized. A fraction of the incoming particles may be scattered or reflected back upstream when they first encounter the shock; they remain in the vicinity of the shock for longer times and may be significantly accelerated (see the high-energy population at the shock in panels *l*-*n* of Figs. 2.3, 2.5, 2.7). If they finally escape upstream, they may contribute to the beam of returning particles which trigger the generation of the upstream oblique waves; if transmitted downstream, they populate the high-energy suprathermal tail of downstream particle spectra.

From Fig. 2.11 it is apparent that the suprathermal tail in subluminal shocks (blue, green and red lines for $\theta = 0°$, $15°$ and $30°$ respectively) is much more pronounced than in superluminal shocks (black line for $\theta = 45°$), whose particle spectrum does not significantly deviate from a thermal distribution. The suprathermal tail becomes flatter and more populated as the magnetic obliquity increases within the subluminal range, meaning that shocks close to the superluminality threshold $\theta_{\rm crit} \approx 34°$ accelerate more efficiently than quasi-parallel shocks. As θ increases from $0°$ to $30°$, the slope of the power-law tail increases from -2.8 ± 0.1 to -2.3 ± 0.1 (subpanel *a*), the fraction of particles in the tail grows from $\sim 1\%$ to $\sim 2\%$ (subpanel *b*) and the energy fraction they account for increases from $\sim 4\%$ to $\sim 12\%$ (subpanel *c*); at the end of our simulations the spectrum extends up to $\gamma_{\rm max} \approx 500$ for $\theta = 0°$ and reaches $\gamma_{\rm max} \approx 2000$ for $\theta = 30°$. The negligible suprathermal tail observed for $\theta = 45°$ is much steeper than for subluminal shocks; as a conservative upper limit, the fraction of particles and energy it accounts for is smaller by two orders of magnitude than in subluminal configurations.

The dependence of the acceleration efficiency on magnetic obliquity should affect the temperature $\Delta\gamma_{\rm MB}$ of the low-energy Maxwellian; we expect that, the larger the fraction of energy stored in the suprathermal tail, the cooler the low-energy bump, if the mean downstream particle Lorentz factor $\bar{\gamma}$ were constant with magnetic obliquity. However, with increasing θ, a larger fraction of the upstream kinetic energy is converted to downstream magnetic energy (see Table 2.1), so that the average downstream particle energy monotonically decreases with θ. In summary, we expect that, within the subluminal range, $\Delta\gamma_{\rm MB}$ should systematically decrease with θ, both because $\bar{\gamma}$ decreases and because a larger fraction of downstream particle energy is stored in the suprathermal tail; across the boundary between subluminal and superluminal configurations, $\Delta\gamma_{\rm MB}$ should sharply increase, if the drop in acceleration efficiency prevails over the decrease of $\bar{\gamma}$ with

θ. Indeed, the thermal spread of the fitting Maxwellian decreases from ~ 4.6 for $\theta = 0^\circ$ to ~ 4.2 for $\theta = 30^\circ$, and then jumps back to ~ 4.5 for $\theta = 45^\circ$; in Fig. 2.11, such trend can be seen in the position of the peak of the low-energy bump.

In summary, the high-energy tail becomes harder and more populated as the magnetic obliquity increases within the subluminal range, in agreement with semi-analytical (e.g., Kirk & Heavens 1989; Ballard & Heavens 1991) and Monte-Carlo (e.g., Bednarz & Ostrowski 1998; Niemiec & Ostrowski 2004; Ellison & Double 2004) calculations. For superluminal shocks, we find that self-generated turbulence is not sufficiently strong to trap a significant fraction of particles in the vicinity of the shock for efficient acceleration. As we show in §5, the upstream oblique waves discussed in §3.2 regulate injection into the acceleration process for $\theta \lesssim \theta_{\rm crit}$. Since no waves are present in the upstream medium of superluminal shocks, due to the absence of the stream of returning particles that trigger their growth, a negligible fraction of the incoming particles enters the acceleration process, and the acceleration efficiency is correspondingly small. The suprathermal tail barely seen for $\theta = 45^\circ$ in Fig. 2.11 is probably populated by particles accelerated during a single shock crossing from upstream to downstream, and then advected downstream. We further explore the transition between subluminal and superluminal shocks with regards to downstream particle spectra in Appendix 8, where we show that the drop in acceleration efficiency across $\theta_{\rm crit} \approx 34^\circ$ is rather abrupt.

Finally, we remind that in Fig. 2.11 we are comparing downstream spectra computed in a slab whose distance behind the shock is the same for all magnetic obliquities. This may be misleading, since the acceleration process in subluminal shocks is not in steady state yet (see §4.2), and downstream particle spectra at a given time may vary among different longitudinal locations (e.g., compare panels *d* and *e* in Fig. 2.8). The particle energy spectrum in a slab at fixed distance downstream from the shock is affected by the efficiency both in accelerating particles and in transmitting downstream the freshly-accelerated particles. Since the latter decreases with increasing obliquity if cross-field diffusion is of minor importance, shocks of larger obliquities may appear, as evaluated from a slab at fixed distance behind the shock, to accelerate less efficiently just because particle diffusion downstream from such shocks is slower. By comparing particle spectra in a slab at constant post-shock distance (Fig. 2.11), we find that subluminal shocks with higher obliquities accelerate more efficiently; this will hold true *a fortiori* when taking into account the obliquity dependence of the downstream diffusion probability.

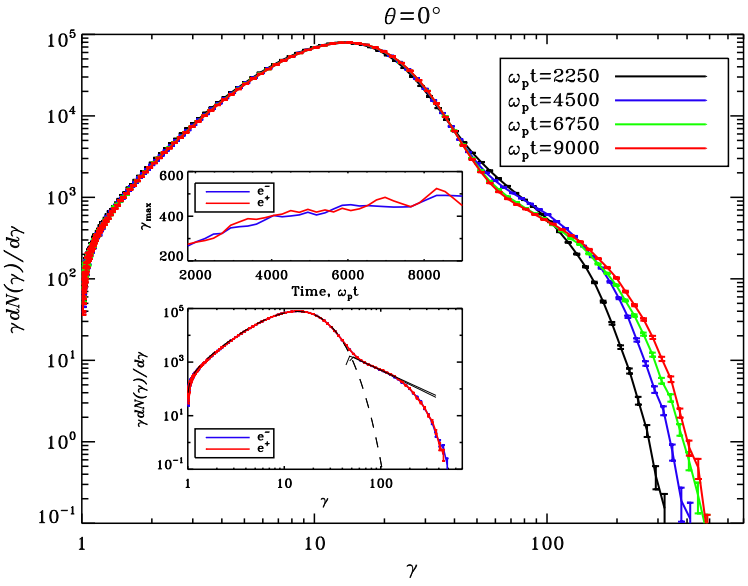


Fig. 2.12.— Time evolution of the particle energy spectrum for $\theta = 0°$: $\omega_p t = 2250$ (black), $\omega_p t = 4500$ (blue), $\omega_p t = 6750$ (green), $\omega_p t = 9000$ (red). All spectra are computed in a slab $200\, c/\omega_p$-wide centered $400\, c/\omega_p$ downstream from the shock. Upper subpanel: electron (blue) and positron (red) maximum Lorentz factor within the slab where the spectrum is computed, as a function of time. Lower subpanel: electron (blue) and positron (red) spectra at time $\omega_p t = 9000$ with Poissonian error bars, supplemented by the best-fitting Maxwellians (black dashed lines) and power laws (black solid lines).

4.2. Time Evolution

Figs. 2.12, 2.13 and 2.14 follow the time evolution of downstream particle spectra for our representative sample of subluminal shocks ($\theta = 0°$, $15°$ and $30°$ respectively), whereas Fig. 2.15 refers to the superluminal shock $\theta = 45°$. In the upper subpanel of each figure, we plot the time evolution of the electron (blue) and positron (red) maximum Lorentz factor in the downstream slab where the spectrum is computed. The lower subpanel shows the electron (blue) and positron (red) spectra at the final time of our simulations ($\omega_p t = 9000$), together with the best-fitting low-energy Maxwellians (black dashed lines) and high-energy power laws (black solid lines).

For both subluminal and superluminal magnetic configurations, the low-energy thermal bump does not significantly evolve with time. Instead, the high-energy tail grows with time as more and more particles are shock-accelerated, but its evolution is much more significant for subluminal shocks than for superluminal configurations. For $\theta = 45°$ (Fig. 2.15), the particle spectrum for

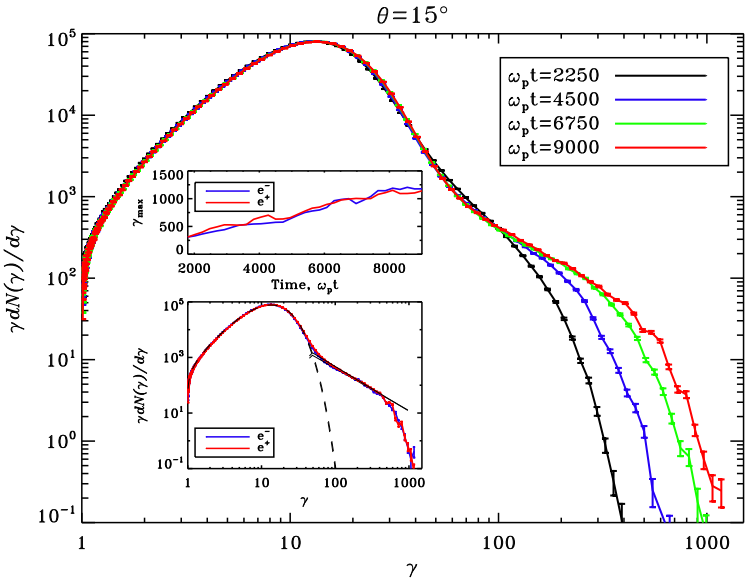


Fig. 2.13.— Time evolution of the particle energy spectrum for $\theta = 15°$. See the caption of Fig. 2.12 for details.

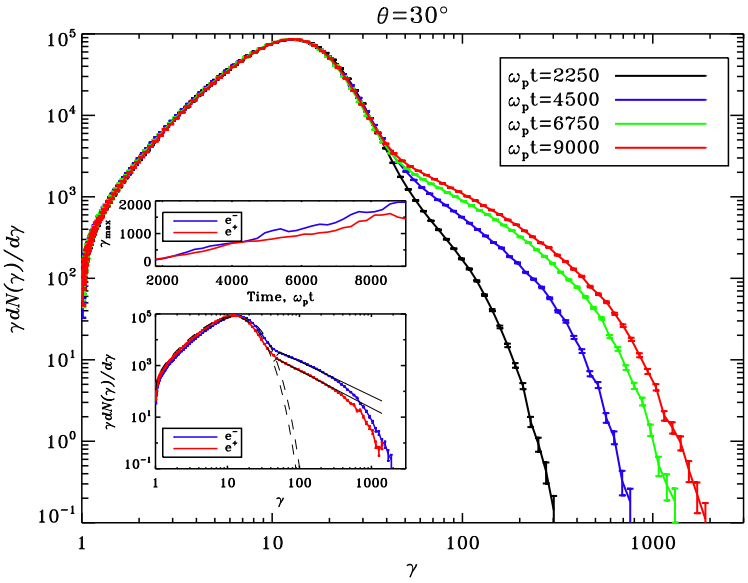


Fig. 2.14.— Time evolution of the particle energy spectrum for $\theta = 30°$. See the caption of Fig. 2.12 for details.

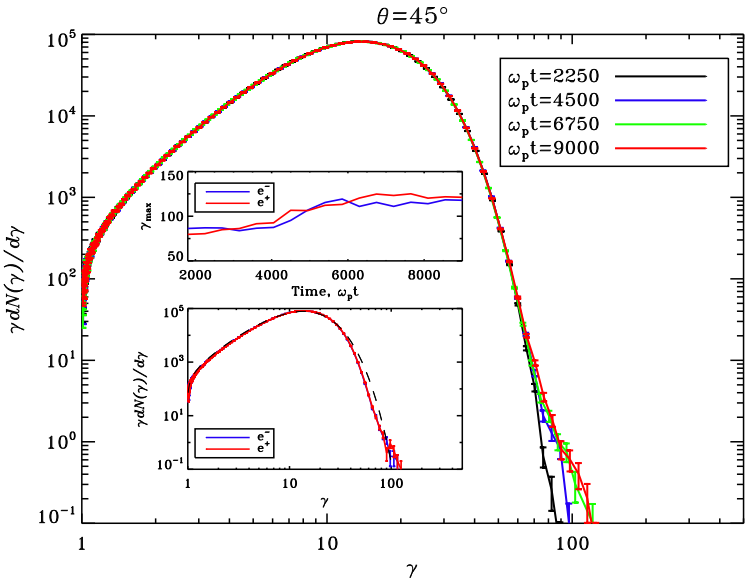


Fig. 2.15.— Time evolution of the particle energy spectrum for $\theta = 45°$. See the caption of Fig. 2.12 for details. Here, we do not attempt to fit the suprathermal tail.

early times (black line) is remarkably thermal, and the minor suprathermal tail visible at later times stops evolving after $\omega_{\rm p}t = 6650$ (compare green and red lines in Fig. 2.15). The upper subpanel in Fig. 2.15 confirms that the maximum particle Lorentz factor saturates at $\gamma_{\rm max} \approx 120$ after $\omega_{\rm p}t \approx 5000$. We are then reasonably confident that for $\theta = 45°$ our simulation captures the full time evolution of the downstream particle spectrum, so that the claimed inefficiency of particle acceleration in superluminal shocks should still hold true with a longer simulation timespan.

For subluminal shocks, the high-energy tail stretches and flattens with time. Its time evolution is moderate for low obliquities ($\theta = 0°$ in Fig. 2.12 and $\theta = 15°$ in Fig. 2.13) and more dramatic for magnetic inclinations close to the superluminality threshold ($\theta = 30°$ in Fig. 2.14). For $\theta = 0°$ the maximum particle Lorentz factor increases from $\gamma_{\rm max} \approx 300$ at $\omega_{\rm p}t = 2250$ to $\gamma_{\rm max} \approx 500$ at $\omega_{\rm p}t = 9000$, whereas for $\theta = 30°$ it grows much faster, from $\gamma_{\rm max} \approx 300$ to $\gamma_{\rm max} \approx 2000$. The fraction of particles and energy stored in the tail also increases during our simulations. Therefore, for subluminal obliquities (and especially for angles $\theta \lesssim \theta_{\rm crit}$, which show the most significant time evolution), the shock acceleration process is not in steady state yet, and the values we quote for the fractional contribution by the suprathermal tail to the downstream particle and energy census should be taken as lower limits. On the other hand, for late times (compare the green and red lines

in Figs. 2.12, 2.13, 2.14), the shape of the spectrum does not change significantly, apart from the linear increase with time of the upper energy cutoff. Since the power-law spectral index is always $p > 2$, this means that the values we quote for the acceleration efficiency at $\omega_{\rm p}t = 9000$ will not significantly increase at later times.

Concerning the time evolution of the maximum particle Lorentz factor, the upper subpanels of Figs. 2.13 and 2.14 (for $\theta = 15^\circ$ and 30° respectively) show a roughly linear increase with time, whereas for $\theta = 0^\circ$ it saturates at $\gamma_{\rm max} \approx 500$ after $\omega_{\rm p}t \approx 6000$ (upper subpanel in Fig. 2.12). By running simulations of parallel shocks in which the first backward jump of the injector happens at different times (we tried 30000, 40000 and 50000 timesteps), we observe that, the sooner the injector starts jumping, the quicker the acceleration process saturates. Each jump unphysically removes the returning particles which are farthest upstream. This has two consequences: first, we artificially prevent these particles from returning to the shock and possibly contributing to the downstream spectrum; also, since the stream of returning particles triggers the growth of the upstream oblique waves, which play a major role in the acceleration process (see §5), we are unphysically affecting the whole process of shock-acceleration. In summary, the spectral time-evolution for $\theta = 0^\circ$ saturates due to our jumping injector, and the acceleration efficiency and maximum Lorentz factor we report for $\theta = 0^\circ$ should be taken as lower limits. Instead, for higher obliquities ($\theta = 15^\circ$ and 30°), no unphysical saturation is observed in the spectrum, since the returning particles remain closer to the shock, so that fewer particles are removed by our jumping injector. The jumping injector does not seem to cause other unphysical effects on particle acceleration.

Finally, we point out that the difference between the downstream spectra of positrons and electrons for $\theta = 30^\circ$ (lower subpanel of Fig. 2.14, at $\omega_{\rm p}t = 9000$), which is reminescent of an analogous difference observed at earlier times (Fig. 2.8*e*, at $\omega_{\rm p}t = 2250$), is an artifact of the small transverse size of our simulations. For magnetic obliquities near $\theta = 30^\circ$, the stream of returning particles, which triggers the growth of the upstream oblique waves, does not propagate very far from the shock. As a consequence, the upstream pattern of electrostatic waves does not move far enough from the shock to let the shock electric potential fluctuate in time (Fig. 2.7 *k* for $\omega_{\rm p}t = 2250$). Particles of one charge will then enter the shock with energies systematically higher than particles of the opposite charge, and the spectral difference between electrons and positrons will persist over time. For smaller obliquities, the stream of returning particles recedes from the shock at higher speeds and the time-average of the shock electric potential vanishes, so that no systematic difference is observed between the downstream spectra of electrons and positrons for

$\theta = 0^\circ$ and $\theta = 15^\circ$ (lower subpanels in Figs. 2.12 and 2.13). For $\theta = 30^\circ$, we have tested that the downstream spectral difference disappears when the simulation box becomes larger (more than 4096 transverse cells) than the characteristic transverse coherence length of the upstream electrostatic oscillations. Therefore, no spectral asymmetry between electrons and positrons is expected in any realistic astrophysical scenario. Moreover, the sum of electron and positron spectra does not vary among runs with different box width, proving that the shock electric potential does not play any appreciable role in the overall acceleration process.[7]

5. Particle Acceleration Mechanisms

We explore the mechanisms responsible for particle acceleration in subluminal shocks by tracing the trajectories of representative high-energy particles extracted from the simulations. We find that particle acceleration is mostly mediated by Diffusive Shock Acceleration (DSA) for quasi-parallel shocks ($\theta \lesssim 10^\circ$), but Shock-Drift Acceleration (SDA) is the main acceleration mechanism for larger, yet still subluminal, magnetic obliquities.

The trajectories shown in Figs. 2.16, 2.17 and 2.18 (for $\theta = 0^\circ, 15^\circ$ and 30° respectively) are characteristic of positrons that end up with the highest energies. In each figure, we show the time evolution of the particle Lorentz factor (panel *a*) and 4-velocities (panel *d*) in the simulation frame. In panel *a*, red stands for negative y-velocity of the selected positron, yellow for positive y-velocity. Particle positions in the wall frame as a function of time are shown in panel *c*. To clarify the interaction between particles and electromagnetic fields, in panel *b* we plot the particle x-position relative to the shock (with the same color coding as in panel *a*), superimposed over two-dimensional strips of E_y stacked in time. For every piece of the particle trajectory that lasts $22.5\,\omega_\mathrm{p}^{-1}$, we plot a $5\,c/\omega_\mathrm{p}$-wide E_y strip which is centered along y at the particle transverse location after $11.25\,\omega_\mathrm{p}^{-1}$ since the start of that portion of the trajectory. Each strip extends along the whole x dimension. The electric field is still measured in the wall frame but shifted along x so that the shock appears stationary. We choose to plot E_y because it is the dominant electric component of the upstream oblique modes, and therefore it should be the primary agent of particle energization even for a parallel shock, where no uniform background $E_\mathrm{y,u}$ is present.

[7] As we comment in §6, this is no longer true for electron-ion shocks, where ion acceleration via the so called Shock-Surfing Acceleration mechanism depends on the shock electric potential.

In order to estimate the relative contribution of DSA and SDA to the acceleration process in oblique shocks ($\theta \neq 0°$), we make use of the fact that particles are shock-drift accelerated while they drift parallel (or anti-parallel, according to their charge) to the background motional electric field. Since in the wall frame of our simulations the downstream motional electric field $E_{\rm y,d}$ is much smaller than the upstream $E_{\rm y,u}$ (see last column in Table 2.1), we may assume that in the downstream medium the particle energy stays approximately constant and we can relate the energy gain $\Delta\gamma_{\rm SDA}$ due to SDA to the upstream drift $\Delta y_{\rm u}$ experienced by the particle:

$$\Delta\gamma_{\rm SDA} \sim \frac{q}{mc^2} E_{\rm y,u}\, \Delta y_{\rm u} \tag{2-2}$$

where q is the particle charge and $E_{\rm y,u} = -\beta_0 B_{\rm z,u} < 0$ the upstream background electric field. Since the right hand side of eq. (2-2) can be computed from particle orbits, as well as the overall energy gain $\Delta\gamma$, the SDA fractional contribution $\Delta\gamma_{\rm SDA}/\Delta\gamma$ will be a diagnostic of the relative importance of DSA and SDA for the acceleration process. In the extreme limit of scatter-free SDA we expect $\Delta\gamma_{\rm SDA}/\Delta\gamma \sim 1$.

5.1. $\theta = 0°$

The orbit of a representative high-energy positron from the simulation of a parallel shock is plotted in Fig. 2.16. For $\omega_{\rm p}t \lesssim 5800$, the particle is approaching the shock from upstream with velocity $\boldsymbol{\beta}(t) \approx -\beta_0\,\hat{\mathbf{x}}$ (see $x(t)$, the blue line in panel *c*). Since $y(t)$ stays approximately constant before the particle encounters the shock, the $E_{\rm y}$ strips stacked in panel *b* are all centered at the same y-location for $\omega_{\rm p}t \lesssim 5800$. Therefore, panel *b* shows for $\omega_{\rm p}t \lesssim 5800$ how the $E_{\rm y}$ component of the upstream oblique modes (Fig. 2.3 *g* for $\omega_{\rm p}t = 2250$) evolves in time within a fixed longitudinal slice. It is thus evident that such waves are moving towards the shock and their phase velocity in the wall frame is smaller than the speed $\beta_0 \approx 1$ of the incoming particles. The positron energy varies with time even before it encounters the shock (Fig. 2.16 *a*), in response to oscillations in $\gamma\beta_{\rm x}(t)$ (blue line in Fig. 2.16 *d*) induced by the electrostatic component of the upstream waves.

From $\omega_{\rm p}t \approx 5800$ to $\omega_{\rm p}t \approx 6600$, the selected positron scatters between the shock and the upstream medium, and its energy is rapidly boosted up to a Lorentz factor $\gamma \approx 250$ (Fig. 2.16 *a*). As the particle energy grows, its gyro-radius also proportionally increases, as seen in panel *c* for $y(t)$ and $z(t)$. A moderate overall energy gain occurs upstream from the shock between $\omega_{\rm p}t \approx 6600$ and $\omega_{\rm p}t \approx 7500$, and finally at $\omega_{\rm p}t \approx 7500$ the positron is transmitted downstream, where its energy

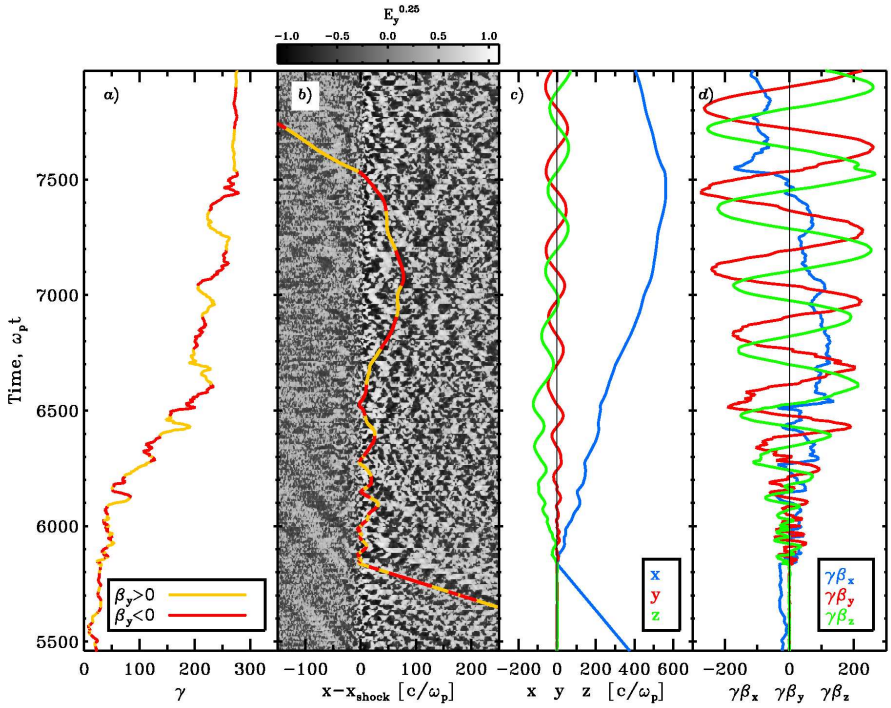


Fig. 2.16.— Trajectory of a representative high-energy positron for $\theta = 0°$, in the wall frame of the simulation: *a*) Particle Lorentz factor as a function of time; colors show when the particle y-velocity is positive (yellow) or negative (red). *b*) Longitudinal position relative to the shock as a function of time, with the same color coding as in panel *a*. The particle orbit is superimposed over two-dimensional strips of E_y stacked in time. For every piece of the particle trajectory that lasts $22.5\,\omega_p^{-1}$, we plot a $5\,c/\omega_p$-wide E_y strip which is centered along y at the particle transverse location after $11.25\,\omega_p^{-1}$ since the start of that portion of the trajectory. Each strip extends along the whole x dimension. E_y^α stands for $(E_y/|E_y|)\,|E_y|^\alpha$. E_y is measured in the simulation frame, but it is shifted in space so that the shock appears stationary. *c*) Particle orbit as a function of time ($x(t)$ blue, $y(t)$ red and $z(t)$ green), relative to the particle position at its first encounter with the shock. *d*) Particle dimensionless 4-velocities $\gamma\beta_x(t)$ (blue), $\gamma\beta_y(t)$ (red) and $\gamma\beta_z(t)$ (green) as functions of time.

remains constant. Most of the accelerated particles end up downstream, but the small fraction that escapes upstream populates the returning stream which triggers the formation of the upstream oblique waves.

As seen in Fig. 2.16 *b*, the acceleration process involves stochastic bouncings between the shock and the upstream medium; magnetic turbulence in the shock transition layer prevents the particle from being transmitted downstream, whereas the upstream scattering is provided by the oblique waves discussed in §3.2, whose amplitude is largest close to the shock. In particular, Fig. 2.16 *b* suggests that the positron gains energy whenever its y-velocity $\beta_{\rm y}(t)$ (red if negative, yellow if positive) has the same sign as the electromagnetic field $E_{\rm y}$ (black if negative, white if positive), and decelerates if they have opposite signs. The converse will be true for electrons (not shown). Overall boosting of the selected positron results from favorable $\beta_{\rm y}(t)E_{\rm y} > 0$ scatterings being more frequent than encounters in which $\beta_{\rm y}(t)E_{\rm y} < 0$.

The stochastic nature of such wave-particle interactions, which are ultimately responsible for particle acceleration, is a defining feature of DSA. So, particle energization in parallel shocks proceeds via DSA, at the very least because the absence of a uniform background motional electric field rules out SDA. However, in the traditional picture of first-order Fermi acceleration, accelerated particles scatter off magnetic turbulence *embedded* in the *upstream* and *downstream* media. Here, instead, high-energy particles bounce between the *upstream* medium and the *shock front*; moreover, the oblique modes responsible for the upstream scattering are not convected with the upstream plasma, as discussed above. Nevertheless, the stochastic character of the acceleration process points towards a diffusive Fermi-like acceleration mechanism. Indeed, as theoretically predicted for diffusive acceleration in relativistic magnetized shocks (e.g., Achterberg et al. 2001), we find that each acceleration cycle between the shock and the upstream medium involves an energy gain $\Delta\gamma \sim \gamma$, at least for the limited number of shock crossings ($\lesssim 10$) observed in our simulation.

These results for a parallel shock should also apply to quasi-parallel magnetic configurations, in which the obliquity is small enough that the upstream motional electric field does not significantly affect the acceleration process. However, in §5.2 we show that SDA provides most of the energization for obliquity angles as low as $\theta = 15^\circ$, so that DSA is the dominant acceleration mechanism only for $\theta \lesssim 10^\circ$.

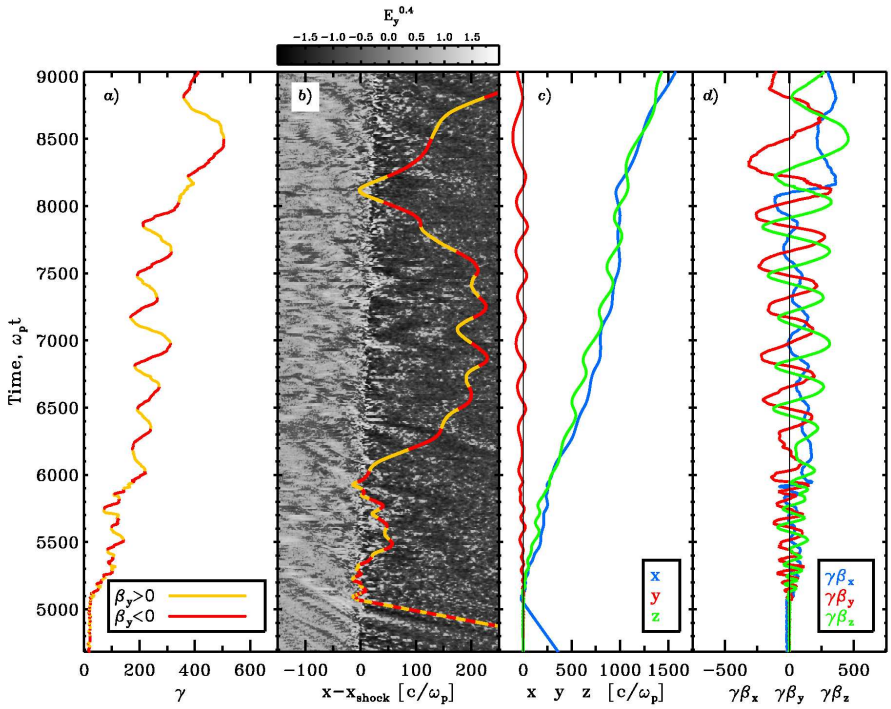


Fig. 2.17.— Trajectory of a representative high-energy positron for $\theta = 15°$. See the caption of Fig. 2.16 for details.

5.2. $\theta = 15°$

Fig. 2.17 presents the acceleration process for a representative positron in a shock with $\theta = 15°$. The particle is reflected back upstream on its first interaction with the shock ($\omega_{\rm p}t \approx 5100$) and then gains energy at each shock encounter ($\omega_{\rm p}t \approx 5200$, $\omega_{\rm p}t \approx 5900$ and $\omega_{\rm p}t \approx 8100$). The selected positron eventually escapes upstream at time $\omega_{\rm p}t \approx 8200$ with Lorentz factor $\gamma \approx 400$, contributing to the beam of returning high-energy particles. However, due to the limited timespan of our simulation, we cannot exclude the possibility that the positron will finally be deflected back and advected downstream.

The time evolution of the positron Lorentz factor plotted in panel *a* suggests that energy gain is always associated with $\beta_{\rm y}(t) < 0$ (red), i.e., when the positron y-velocity has the same sign as the uniform background electric field $E_{\rm y,u} < 0$. Whenever the positron gyro-orbit is fully contained in the upstream region (e.g., from $\omega_{\rm p}t \approx 5300$ to $\omega_{\rm p}t \approx 5800$, or from $\omega_{\rm p}t \approx 6000$ to $\omega_{\rm p}t \approx 8000$, as

well as after the selected positron escapes upstream), the energy gain within the "red" half of the gyro-period is balanced by an approximately equal energy loss during the "yellow" half, since the background motional electric field $E_{\rm y,u}$ stays approximately constant throughout the gyro-orbit. However, when the positron trajectory intersects the shock front, in the "upstream" portion of its clockwise Larmor gyration around the oblique field the positron y-velocity has always the same sign as the background motional electric field, thus resulting in systematic acceleration. Instead, the "downstream" portion of the orbit does not result in deceleration because in the simulation frame the motional electric field downstream is much smaller than in the upstream.[8] For electrons (not shown), whose Larmor gyration around the magnetic field is counter-clockwise, the y-velocity in the upstream medium will have opposite sign than $E_{\rm y,u}$, resulting again in net energization.

In summary, particles are systematically accelerated by the upstream motional electric field whenever their gyro-orbit intersects the shock surface, as expected for the Shock-Drift Acceleration process. SDA is inevitably associated with a ∇B drift along (or opposite, according to the particle charge) the background electric field, due to the magnetic field jump at the shock. This is seen, for the positron in Fig. 2.17, in the slow drift along $-\hat{\mathbf{y}}$ of $y(t)$ in panel *c*. Such a drift will be more significant for higher obliquities, since the magnetic field jump at the shock is larger.

A comparison between $x(t)$ and $z(t)$ in Fig. 2.17 *c* suggests that our representative positron is heading upstream along the oblique magnetic field. As compared to a parallel shock, it will take longer to escape ahead of the shock, the more so for higher obliquities. Many gyro-orbits could then intersect the shock during a single shock-crossing of the positron gyro-center, thus resulting in multiple energizations (as happens at $\omega_{\rm p}t \approx 5200$ and $\omega_{\rm p}t \approx 5900$). The particles that escape upstream are not systematically shock-drift accelerated any longer, until they return to the shock ($\omega_{\rm p}t \approx 8000$ in Fig. 2.17), due to deflection by the upstream background fields or scattering off the upstream oblique waves. In that respect, even though the upstream motional field $E_{\rm y,u}$ is the primary agent for particle energization, DSA cooperates with SDA to trap the accelerated particles at the shock, thus increasing the overall acceleration efficiency. Making use of eq. (2-2) with a sample of ~ 1000 high-energy particles, we are able to confirm that the ratio $\Delta\gamma_{\rm SDA}/\Delta\gamma$ is independent of the total energy gain $\Delta\gamma$ (as predicted by Jokipii 1982), and we estimate the

[8] We emphasize that a similar line of reasoning can be made in any other frame. In the particular frame of the simulation, the downstream motional electric field is $E_{\rm y,d} \ll E_{\rm y,u}$, which simplifies the argument. However, the shock is moving in this frame.

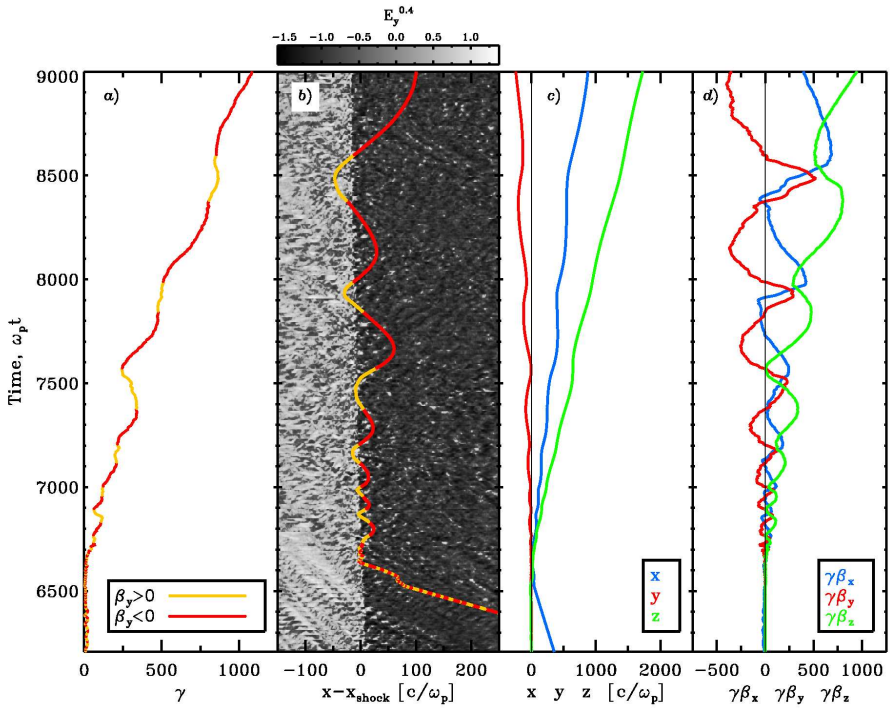


Fig. 2.18.— Trajectory of a representative high-energy positron for $\theta = 30°$. See the caption of Fig. 2.16 for details.

fractional contribution by SDA to be $\sim 60\%$. Therefore, for obliquity angles as small as $\theta = 15°$, SDA is already the main acceleration mechanism.

5.3. $\theta = 30°$

The acceleration process for a representative positron in a shock with $\theta = 30°$ is shown in Fig. 2.18. As the obliquity angle approaches the critical boundary between subluminal and superluminal configurations ($\theta_{\rm crit} \approx 34°$), particles reflected upstream can hardly escape ahead of the shock. They are effectively trapped at the shock while heading upstream along the oblique magnetic field, since for $\theta \lesssim \theta_{\rm crit}$ their gyro-center x-velocity is comparable to the shock speed. Before escaping upstream, their orbit may intersect the shock many times (many more than for $\theta = 15°$), thus resulting in repeated energizations via SDA.

As discussed in §5.2, for a particle gyrating across the shock any excursion into the upstream region results in efficient SDA by the upstream motional electric field $E_{\rm y,u}$ (compare panels *a* and *b* in Fig. 2.18), whereas in the downstream medium its energy stays roughly constant since $E_{\rm y,d} \ll E_{\rm y,u}$ in the simulation frame. The corresponding time evolution of the Lorentz factor $\gamma(t)$ resembles a stairway (see Fig. 2.18 *a*), with steps that grow linearly with γ since the time spent upstream is proportional to the particle energy. When a particle completes a full gyro-orbit in the upstream medium (between $\omega_{\rm p}t \approx 7200$ and $\omega_{\rm p}t \approx 7600$ for the positron in Fig. 2.18), its energy gain is balanced by an equal energy loss. Due to scattering off the incoming oblique waves or regular deflection by the upstream background fields, the particle may be re-captured by the shock, where its acceleration via SDA resumes. The representative positron in Fig. 2.18, as most of the high-energy particles extracted from our $\theta = 30^\circ$ simulation, is still being shock-accelerated at the final simulation time $\omega_{\rm p}t = 9000$; its final Lorentz factor $\gamma \approx 1000$, which is already much larger than in shocks with smaller obliquities, is therefore just a lower limit.

Using eq. (2-2) with a sample of ~ 1000 high-energy particles, we find that on average SDA contributes a fraction $\sim 90\%$ to the overall energy gain, so that diffusive scattering plays a minor role in the acceleration process for $\theta = 30^\circ$. Thus, with increasing magnetic obliquity in the subluminal range, the acceleration process switches from DSA, which is the only mechanism acting for parallel shocks, to SDA, which for $\theta \lesssim \theta_{\rm crit}$ operates in an essentially scatter-free regime.

For $\theta = 30^\circ$ the upstream oblique waves are important to regulate injection into the acceleration process: all the particles that end up gaining the highest energies have been deflected from their upstream straight path before encountering the shock (see Fig. 2.18 *b* at $\omega_{\rm p}t \approx 6500$). We have tested that incoming particles enter the acceleration process only if their first shock encounter results in a reflection back upstream, whereas transmitted particles populate the low-energy bump of downstream spectra. In agreement with Ballard & Heavens (1991), we have verified that such reflected particles had encountered the shock with pitch angle cosine $\mu_{\rm HT}$ in the de Hoffmann-Teller frame (de Hoffmann & Teller 1950) which satisfies

$$|\mu_{\rm HT}| < \sqrt{1 - \frac{B_{\rm u,HT}}{B_{\rm d,HT}}} \tag{2-3}$$

where $B_{\rm d,HT}/B_{\rm u,HT}$ is the magnetic field jump in the de Hoffmann-Teller frame. In other words, only the incoming particles that are deflected by the upstream waves, such that their pitch angle satisfies eq. (2-3) when encountering the shock, will be reflected back upstream and participate in the SDA process.

In summary, while for quasi-parallel shocks ($\theta \lesssim 10^\circ$) the upstream oblique waves generated by the beam of returning particles provide the upstream scattering for DSA, for magnetic obliquities close to $\theta_{\rm crit}$ the same waves mediate the injection process into SDA.

6. Discussion and Summary

We have explored by means of 2.5D PIC simulations the internal structure and acceleration properties of relativistic collisionless shocks propagating in a magnetized pair plasma. A strong magnetic field may significantly suppress cross-field motion and therefore quench the growth of the Weibel instability (Weibel 1959; Medvedev & Loeb 1999; Gruzinov & Waxman 1999), which is believed to mediate unmagnetized or weakly magnetized shocks (e.g., Spitkovsky 2005, 2008a,b; Chang et al. 2008; Keshet et al. 2008). Instead, strongly magnetized perpendicular shocks are triggered by magnetic reflection of the incoming particles from the shock-compressed fields (e.g., Alsop & Arons 1988). Spitkovsky (2005) finds that the transition between weakly magnetized (Weibel-mediated) and strongly magnetized (reflection-mediated) perpendicular shocks happens around an upstream magnetization $\sigma \sim 10^{-3}$. In this chapter, we focus on strongly magnetized upstream flows ($\sigma = 0.1$) and we study the dependence of the shock properties on the inclination angle θ between the upstream magnetic field and the shock normal, as measured in the "wall" frame of our simulations.

In magnetized shocks, particle gyro-centers are constrained to move along the magnetic field. It is then natural to define a critical obliquity angle $\theta_{\rm crit}$ such that in "superluminal" shocks ($\theta > \theta_{\rm crit}$) particles cannot escape upstream along the magnetic field lines. If the upstream plasma has Lorentz factor $\gamma_0 = 15$ and magnetization $\sigma = 0.1$, MHD calculations yield $\theta_{\rm crit} \approx 34^\circ$ in the simulation frame, weakly dependent on both γ_0 and σ (see Fig. 2.2). Since for $\theta > \theta_{\rm crit}$ particles are quickly advected downstream from the shock, the acceleration efficiency should be strongly suppressed for superluminal configurations. In fact, we find that magnetized relativistic pair shocks produce a significant population of suprathermal particles only for "subluminal" magnetic geometries ($\theta < \theta_{\rm crit}$). This holds true for a wide range of upstream Lorentz factors (we also tried $\gamma_0 = 5$ and $\gamma_0 = 50$, for fixed $\sigma = 0.1$) and magnetizations (we also tested $\sigma = 0.03$ and $\sigma = 0.3$, for fixed $\gamma_0 = 15$). We conclude that only a narrow range of obliquity angles allows for efficient particle acceleration in relativistic ($\gamma_0 \gtrsim 5$) magnetized ($\sigma \gtrsim 0.03$) pair shocks: as measured from

the upstream frame, the magnetic field should lie within a cone of half-opening angle $\approx 34^\circ/\gamma_0$ around the shock normal to produce a significant population of nonthermal particles.[9]

For quasi-parallel magnetic configurations ($\theta \lesssim 10^\circ$), some of the incoming particles are scattered back upstream and then gain energy by bouncing between the shock and the upstream medium. The upstream scattering is provided by oblique waves which propagate towards the shock and are self-consistently triggered by the shock-accelerated particles that escape upstream; magnetic turbulence in the shock transition layer reflects the high-energy particles back upstream. The stochastic nature of the acceleration process suggests that DSA is the main energization mechanism in quasi-parallel magnetized pair shocks, at the very least because Shock-Surfing Acceleration does not operate in electron-positron shocks and the absence of a background motional electric field rules out SDA in parallel shocks. For $\theta = 0^\circ$, we find that at the final time of our simulation ($\omega_{\rm p}t = 9000$) the suprathermal tail of the downstream energy spectrum accounts for $\sim 1\%$ of the downstream particle number and $\sim 4\%$ of the energy, with a power-law spectral index -2.8 ± 0.1; however, these may be lower limits, since further evolution of the spectrum in our simulations of quasi-parallel shocks may be artificially suppressed by our choice of boundary conditions.

Whereas for quasi-parallel magnetic geometries ($\theta \lesssim 10^\circ$) the shock-generated turbulent fields are typically stronger than the uniform motional electric field, and so SDA plays a minor role in the acceleration process, its mean fractional contribution to the particle overall energy gain increases to $\sim 60\%$ for $\theta = 15^\circ$ and reaches $\sim 90\%$ for $\theta = 30^\circ$. Therefore, the acceleration process for $\theta \lesssim \theta_{\rm crit}$ is almost scatter-free: particles heading upstream along the oblique magnetic field hardly escape from the shock, and they can experience multiple shock-drift energizations as their orbit repeatedly crosses the shock surface. However, we cannot exclude that a fraction of the shock-drift accelerated particles may eventually gain enough energy to become effectively "unmagnetized", i.e., their gyrocenter would no longer be constrained to follow the background field. Such particles may be further accelerated by a diffusive Fermi-like mechanism, and the SDA process would then serve as a pre-acceleration stage before the onset of DSA. The upstream oblique waves, which provide the upstream scattering in quasi-parallel shocks, may also mediate this further step of diffusive

[9] We remark that this conclusion applies only to electron-positron shocks. For electron-ion magnetized shocks, the electric potential across the shock complicates the particle dynamics and we cannot exclude that shocks which are nominally superluminal may be moderate particle accelerators. Relativistic shocks propagating in magnetized electron-ion plasmas will be discussed in the next chapter.

acceleration for larger, but still subluminal, obliquities. We also find that in $\theta \lesssim \theta_{\rm crit}$ shocks these waves regulate the process of particle injection into SDA.

As predicted by Jokipii (1982), we find that SDA in $\theta \lesssim \theta_{\rm crit}$ shocks is a much faster process than DSA in quasi-parallel shocks, since each SDA cycle takes approximately one gyro-period, which is much shorter than any diffusive scattering timescale. As the acceleration mechanism switches from DSA to SDA with increasing obliquity, the acceleration efficiency increases as well. For $\theta = 30^\circ$, the number and energy fractions contributed by the suprathermal tail of the downstream particle spectrum are respectively $\sim 2\%$ and $\sim 12\%$ and its power-law spectral index is -2.3 ± 0.1. Since acceleration in $\theta \approx \theta_{\rm crit}$ shocks is close to a steady state around the final time of our simulations ($\omega_{\rm p} t = 9000$), the efficiencies quoted above should be near the saturation values.

For superluminal shocks ($\theta \gtrsim \theta_{\rm crit}$), we find that if the upstream background field is not highly turbulent on scales smaller than the transverse size of our simulations ($\sim 100\, c/\omega_{\rm p}$), the self-generated turbulence is not strong enough to trap a significant fraction of particles at the shock for efficient acceleration; rather, most of the particles are quickly advected downstream with the magnetic field lines. Monte-Carlo simulations of relativistic magnetized shocks found evidence for particle acceleration in perpendicular shocks (e.g., Bednarz & Ostrowski 1998; Niemiec & Ostrowski 2004; Ellison & Double 2004), but they required exceptionally high levels of turbulence (Ostrowski & Bednarz 2002). We show via PIC simulations with laminar upstream fields that superluminal shocks do not self-consistently generate such large degrees of turbulence, and they are inefficient particle accelerators. In Appendix 8 we show that the drop in acceleration efficiency at $\theta_{\rm crit} \approx 34^\circ$ between subluminal and superluminal shocks is rather abrupt.

We have tested that our results do not significantly depend on the number of particles per computational cell (we tried 2 and 8 particles per species per cell) or the transverse size of our two-dimensional simulation box (we tried 512, 1024, 2048 and 4096 transverse cells). We have verified that the geometry of the magnetic field with respect to the simulation plane (in-plane or out-of-plane magnetic field) does not change our main conclusions (see Appendix 7); moreover, our 2.5D results have been confirmed by running a limited suite of simulations in three-dimensional domains with 64, 128 or 256 cells along each transverse dimension.

A direct comparison of 2.5D simulations with 3D runs is possible, since for moderate obliquities ($\theta \lesssim 45^\circ$) the gyro-orbits of downstream particles mostly lie outside the two-dimensional simulation plane, and the flow thermalizes behind the shock to a *three-dimensional* distribution. For the

obliquity angles investigated with 3D simulations ($\theta = 30^\circ$ and $\theta = 35^\circ$), we see a remarkable agreement with the results of our 2.5D runs (for $\theta = 30^\circ$, see Appendix 7). Jones et al. (1998) found that numerical simulations in two spatial dimensions could artificially inhibit particle diffusion across field lines, which is required for efficient particle acceleration in superluminal shocks. Instead, we confirm with 3D simulations that superluminal relativistic shocks propagating in pair plasmas with laminar fields do not produce a significant population of nonthermal particles, at least for the domain sizes we tried in 3D.

In summary, our results underscore that particle acceleration in relativistic ($\gamma_0 \gtrsim 5$) magnetized ($\sigma \gtrsim 0.03$) pair flows is efficient only for magnetic obliquities $\theta \lesssim \theta_{\rm crit} \approx 34^\circ$. As seen from the upstream fluid frame, only nearly-parallel shocks (with obliquity $\lesssim 34^\circ/\gamma_0$) produce a significant suprathermal tail of accelerated particles. These findings place constraints on the models of Pulsar Wind Nebulae (PWNe), Active Galactic Nuclei (AGN) jets and Gamma-Ray Bursts (GRBs) that invoke particle acceleration in relativistic magnetized shocks.

Recent relativistic MHD simulations of PWNe (Komissarov & Lyubarsky 2003, 2004; Del Zanna et al. 2004) have shown that polar jets with the observed velocities can develop provided that the magnetization of the relativistic ($\gamma_0 \gtrsim 10^6$) electron-positron pulsar wind is large enough ($\sigma \gtrsim 0.01$) to collimate the jet by hoop stresses downstream of the wind termination shock. Since the asymptotic wind magnetic field is mostly toroidal around the pulsar axis, the termination shock will be quasi-perpendicular. If the wind magnetization is $\sigma \gtrsim 0.03$, our results predict that the termination shock should not produce a significant population of nonthermal particles, which is required instead by the observed synchrotron and inverse Compton emission. However, for oblique rotators, the pulsar wind in the equatorial plane should consist of stripes of alternating field, and the field may dissipate via magnetic reconnection of the opposite polarity stripes before the termination shock (Coroniti 1990; Lyubarsky & Kirk 2001); the effective magnetization of the shock would then be lower and particle acceleration would proceed like in an essentially unmagnetized shock. Alternatively, magnetic reconnection of the striped wind field when compressed at the termination shock may be responsible for particle energization (Pétri & Lyubarsky 2007; Lyubarsky & Liverts 2008).

Synchrotron and inverse Compton emission from the core of blazar jets is usually attributed to high-energy electrons and positrons accelerated in mildly-relativistic internal shocks ($\gamma_0 \sim 2$ in the jet comoving frame). Such shocks should be quasi-perpendicular, since polarization measurements

of radio knots (Gabuzda et al. 2004; Pushkarev et al. 2005) indicate that the magnetic field is mostly transverse to the jet axis. If the jet flow is significantly magnetized ($\sigma \gtrsim 0.03$), the presence of a substantial population of high-energy particles accelerated at such perpendicular magnetized shocks would be hard to explain, unless particle acceleration proceeds not in the central "spine" of the jet but rather at its edges; here, interaction of the jet with the surrounding medium may create subluminal configurations where particle acceleration would be more efficient. Otherwise, the jet flow should be weakly magnetized ($\sigma \lesssim 0.03$). Alternatively, particle acceleration may occur not in shocks but in reconnection layers of a Poynting-dominated jet, although Sikora et al. (2005) have shown that the jet is likely to be matter-dominated at the pc-scale emission region.

Similarly, in the standard "fireball" scenario of GRBs, most models of the prompt emission require a population of energetic particles accelerated at mildly-relativistic internal shocks. If the ejecta are strongly magnetized, possibly by the field of the progenitor star, and the magnetic field is mostly transverse to the outflow, we would predict inefficient particle acceleration and thus negligible nonthermal emission, contrary to observations. This could be circumvented if the relativistic outflow primarily takes the form of a large-scale Poynting flux, and particle acceleration occurs not in shocks but by dissipation of magnetic energy via current-driven instabilities (Lyutikov & Blandford 2003).

In summary, our results suggest that models of nonthermal emission in PWNe, AGN jets and GRBs which require particle acceleration in relativistic pair shocks should resort to either nearly-parallel or weakly magnetized shocks. Alternatively, other acceleration mechanisms, e.g., by dissipation of the field via magnetic reconnection, should be invoked, or the upstream flow should contain a significant ionic component (Hoshino et al. 1992; Amato & Arons 2006) or small-scale turbulence (Sironi & Goodman 2007).

7. Appendix (A): 2.5D Simulations with In-Plane Upstream Background Magnetic Field

In the 2.5D simulations presented in this chapter, the fact that the upstream background magnetic field lies in a plane which is perpendicular to the simulation plane ("out-of-plane" magnetic configurations) may be suspected in reducing turbulence along the field. Here we test this by analyzing the results of 2.5D simulations with an "in-plane" upstream blackground magnetic field, for each of the subluminal oblique configurations presented in §3, i.e., $\theta = 15^\circ$ and $\theta = 30^\circ$.

In the in-plane case, the upstream background magnetic field will have components $B_{\rm x,u} \geq 0$ and $B_{\rm y,u} \geq 0$ (so that $\theta \equiv \arctan[B_{\rm y,u}/B_{\rm x,u}]$), and the associated upstream motional electric field will be $\mathbf{E}_{\rm u} = \beta_0 B_{\rm y,u}\,\hat{\mathbf{z}}$. As regards to the electromagnetic fields, the direction $+\hat{\mathbf{z}}$ in out-of-plane configurations should then correspond to in-plane $+\hat{\mathbf{y}}$, and out-of-plane $+\hat{\mathbf{y}}$ to in-plane $-\hat{\mathbf{z}}$. Apart from this overall rotation of the coordinate system, we expect out-of-plane and in-plane geometries with the same magnetic obliquity to show comparable results in terms of shock structure and particle spectra. Any difference should be an artifact of the reduced dimensionality of our simulation box and would be a caveat against the generalization of our 2.5D results to three-dimensional problems.

7.1. Shock Structure

Figs. 2.19 and 2.20 show the shock structure and positron phase space at time $\omega_{\rm p}t = 2250$ for in-plane magnetic configurations with obliquity $\theta = 15^\circ$ and $\theta = 30^\circ$ respectively. Their out-of-plane counterparts are plotted in Figs. 2.5 and 2.7 respectively. We find that the shock internal structure does not significantly change by varying the orientation of the field with respect to the simulation plane.

Out-of-plane and in-plane configurations with the same obliquity show a comparable shock velocity and similar transversely-averaged density (panel *b*) and magnetic energy (panel *d*) profiles, as relates to the width of the shock transition region ($\sim 10\,c/\omega_{\rm p}$), the jump across the shock (in agreement with three-dimensional MHD calculations) and the presence of a fast and a slow shock (compare Fig. 2.7 and Fig. 2.20 for $\theta = 30^\circ$). For $\theta = 15^\circ$, the y-averaged magnetic energy profile peaks in the shock layer at $\sim 30\%$ of the upstream kinetic energy density, both for out-of-plane (Fig. 2.5*d*) and in-plane (Fig. 2.19*d*) configurations, and in both cases it decays downstream on a lenghtscale $\sim 50\,c/\omega_{\rm p}$. For $\theta = 30^\circ$, the overall fluid structure does not significantly change between in-plane and out-of-plane configurations, but a few minor differences are present: for the in-plane case, the magnetic energy in the shock layer slightly exceeds the downstream plateau (Fig. 2.20 *d*), which was not observed in its out-of-plane counterpart (Fig. 2.7 *d*); also, the speed of the slow shock, tracked from the point where the magnetic energy starts to decrease downstream from the main shock, seems to be higher for in-plane than out-of-plane configurations. For both $\theta = 15^\circ$ and $\theta = 30^\circ$, the motional electric field (panel *h* in Figs. 2.5 and 2.7, not shown in Figs. 2.19 and 2.20) does not drop to zero behind the main shock due to a residual downstream bulk velocity

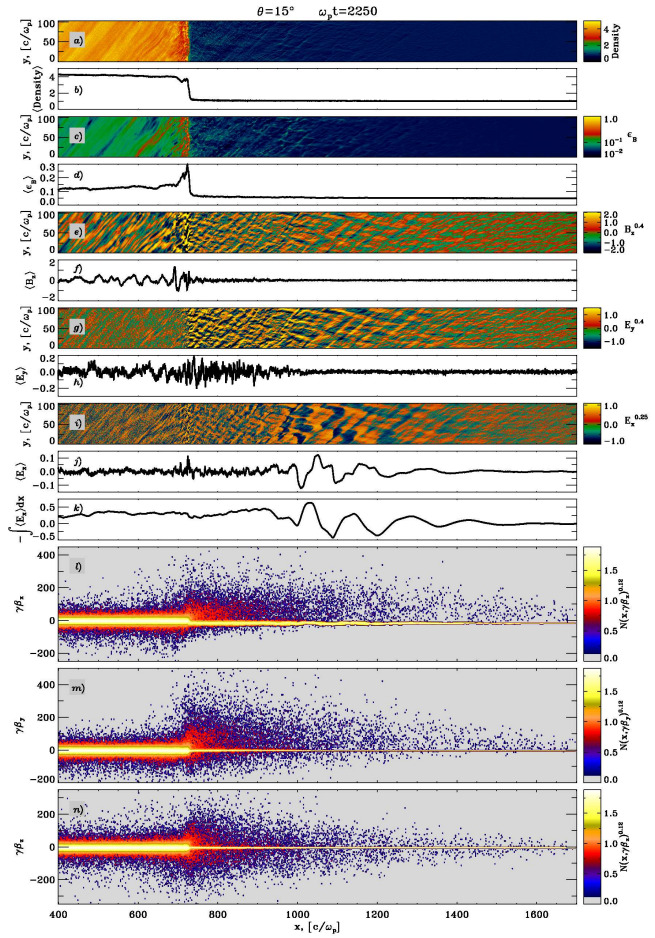


Fig. 2.19.— Shock structure and positron phase space at time $\omega_{\rm p}t = 2250$ for $\theta = 15°$, with upstream background magnetic field lying in the simulation plane. The fluid quantities are normalized as in Fig. 2.3, but here $B_{\rm z}$ and $E_{\rm y}$ are in units of the transverse component $B_{\rm y,u}$ of the upstream background magnetic field.

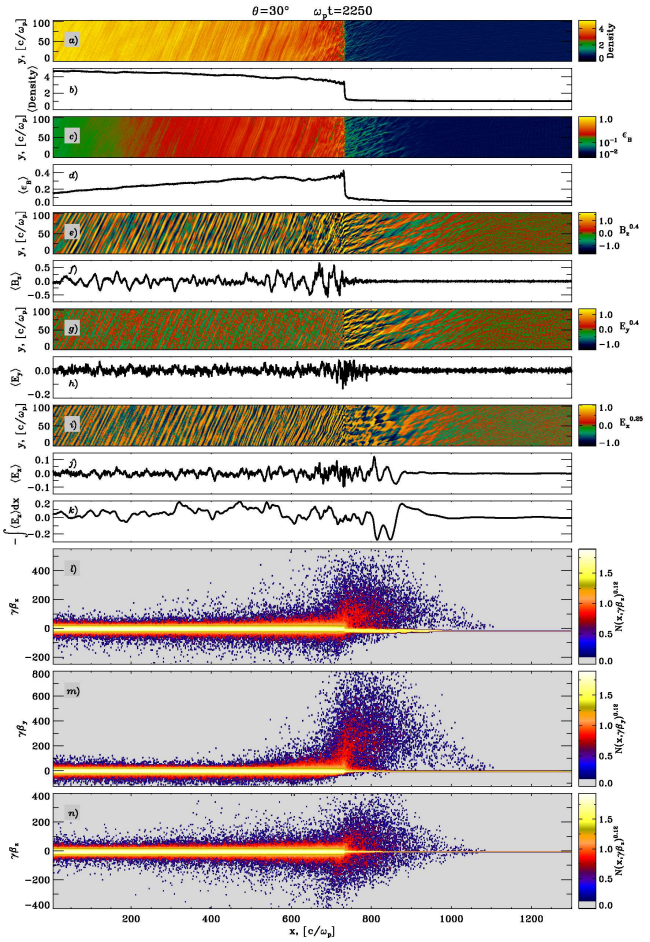


Fig. 2.20.— Shock structure and positron phase space at time $\omega_p t = 2250$ for $\theta = 30°$, with upstream background magnetic field lying in the simulation plane. The fluid quantities are normalized as in Fig. 2.19.

transverse to the upstream flow ($\beta_{\rm y,d} < 0$ and $\beta_{\rm z,d} < 0$ for in-plane and out-of-plane configurations respectively). The transient electromagnetic precursor wave generated by coherent Larmor bunching at the shock, which showed up in $B_{\rm z}$ and $E_{\rm y}$ for out-of-plane oblique configurations, is now seen in the corresponding in-plane fields $B_{\rm y}$ and $E_{\rm z}$ with comparable wavelengths and amplitude (further upstream than Figs. 2.19 and 2.20 show).

The phase space plots in panels *l*-*n* of Figs. 2.19 and 2.20 are similar to their counterparts in Figs. 2.5 and 2.7 respectively, provided that out-of-plane $\gamma\beta_{\rm z}$ and $\gamma\beta_{\rm y}$ are compared with in-plane $\gamma\beta_{\rm y}$ and $-\gamma\beta_{\rm z}$ respectively. As observed for out-of-plane configurations, the returning high-energy particles slide along the upstream magnetic field, which means for the in-plane case that they show large and positive $\gamma\beta_{\rm x}$ and $\gamma\beta_{\rm y}$ (panels l and m in Figs. 2.19, 2.20). The in-plane excess of positrons with $\gamma\beta_{\rm z} > 0$ in the upstream region close to the shock (panel n in Figs. 2.19 and 2.20) is due to ∇B drift, and it corresponds to the asymmetry towards negative values of $\gamma\beta_{\rm y}$ (panel m in Figs. 2.5 and 2.7 respectively) observed for out-of-plane configurations.

Aside from the similarities between in-plane and out-of-plane results, Figs. 2.19 and 2.20 improve our knowledge of the internal structure of oblique shocks. The two-dimensional plots of magnetic energy density (panel c) and $B_{\rm z}$ (panel e) of Figs. 2.5 and 2.7 show islands of low magnetic field in the downstream region of the strong shock. Fewer islands are produced at the shock at later times; when formed, each island slowly decays with time at constant x-location, and it disappears when caught up by the slow shock. The corresponding panels in Figs. 2.19 and 2.20 clarify that such islands are actually tubes of low magnetic energy stretching along the downstream magnetic field, as confirmed also by 3D simulations with 256 cells along each transverse dimension.

Most importantly, the electromagnetic component of the upstream oblique waves, which was mostly seen in $B_{\rm z}$ and $E_{\rm y}$ for the out-of-plane case (panels *e, g* in Figs. 2.5, 2.7), shows up mainly in the same field components also for in-plane configurations. In other words, the wave magnetic field is still mostly perpendicular to the simulation plane, regardless of the orientation of the upstream background magnetic field with respect to the simulation plane. A further study of such modes will therefore have to be fully three-dimensional to capture their true physical nature. What Figs. 2.19 and 2.20 (panels *e, g, i*) add to our understanding of such waves is that, in the region far upstream where they are triggered, the surface of constant phase is roughly a cone of constant opening angle around the upstream background field.

7.2. Particle Energy Spectra

Fig. 2.21 shows at $\omega_{\rm p}t = 9000$ a comparison of downstream particle spectra between in-plane (black line) and out-of-plane (red line) configurations: for $\theta = 15^\circ$ (panel *a*), the spectrum is remarkably similar between the two cases, whereas a significant difference is seen for $\theta = 30^\circ$ (panel *b*), with the high-energy tail for the out-of-plane configuration being much more pronounced than in its in-plane counterpart (and the peak of the low-energy Maxwellian being correspondingly cooler). However, the out-of-plane spectrum for $\theta = 30^\circ$ is in good agreement with a 3D simulation having 64 cells along each transverse dimension (green line in panel *b* of Fig. 2.21), which suggests that results from our 2.5D out-of-plane simulations may be confidently applied to realistic three-dimensional scenarios.

Concerning the fact that the $\theta = 30^\circ$ spectrum at $\omega_{\rm p}t = 9000$ shows a significant difference at high energies between in-plane and out-of-plane configurations, we report that the spectrum computed in the shock transition layer (not shown) is remarkably similar between the two cases, suggesting that it is not the intrinsic acceleration efficiency to differ but rather the transmission probability of shock-accelerated particles into the downstream medium.

8. Appendix (B): The Sharp Transition in Acceleration Efficiency between Subluminal and Superluminal Shocks

In Fig. 2.11 we have shown that only subluminal shocks present a pronounced suprathermal tail in downstream particle spectra, whereas the particle distribution function behind superluminal shocks does not significantly deviate from a Maxwellian. Here, we show that the drop in acceleration efficiency between subluminal and superluminal configurations happens remarkably close to the boundary $\theta_{\rm crit} \approx 34^\circ$ predicted by MHD calculations, and it is rather abrupt.

In Fig. 2.22, downstream spectra for $\theta = 28^\circ$ (violet), 30° (red), 31° (black), 32° (blue), 35° (green) and 45° (yellow) are compared for different times. Within the subluminal range, lower obliquities ($\theta = 28^\circ$) show the most significant suprathermal tail at early times, whereas shocks with magnetic inclinations closer to $\theta_{\rm crit}$ ($\theta = 30^\circ$ and 31°) are the most efficient accelerators at later times. The closer to the superluminality threshold, the faster the time evolution of the spectrum, in agreement with the fact that particle acceleration in shocks with $\theta \lesssim \theta_{\rm crit}$ is mostly mediated by SDA, which is a very fast process (Jokipii 1982). More importantly, the onset of efficient SDA

is extremely rapid, since it is driven by a positive feedback, as discussed in §5.3: as the number of accelerated particles which escape upstream gets large enough to trigger the generation of the upstream oblique waves, a larger fraction of the incoming particles will be deflected from their straight path and their pitch angle at the shock will more likely satisfy the relation in eq. (2-3). More particles will then be reflected upstream and injected into the acceleration process, in a runaway cascade towards larger acceleration efficiency. The dramatic difference observed in the high-energy tail of $\theta = 31^\circ$ between $\omega_{\rm p}t = 2250$ and $\omega_{\rm p}t = 4500$ and of $\theta = 32^\circ$ between $\omega_{\rm p}t = 6750$ and $\omega_{\rm p}t = 13500$ (thin blue line in the bottom right panel of Fig. 2.22) corresponds to the onset of efficient SDA. In response to the increased acceleration efficiency, the peak of the low-energy Maxwellian shifts to lower temperatures.

Since the positive feedback required for the onset of efficient SDA relies on upstream waves triggered by the returning particles, we expect that in superluminal shocks, where particles are not able to return upstream along the magnetic field, acceleration should be strongly suppressed. Indeed, the downstream spectra of the superluminal shocks $\theta = 35^\circ$ and $\theta = 45^\circ$ are mostly thermal, although a minor high-energy tail, further suppressed as θ increases, is observed at late times. Since spectra computed ahead of the shock (not plotted) do not show any sign of high-energy particles, we believe that the residual suprathermal tail observed downstream is probably populated by particles shock-drift accelerated during a single shock crossing from upstream to downstream and then advected downstream. We are confident that these conclusions hold true for times longer than the timespan of our simulations, for the following reasons: *i*) between $\omega_{\rm p}t = 6750$ and $\omega_{\rm p}t = 9000$, the downstream spectra of superluminal shocks do not significantly evolve (see also Fig. 2.15), so that they are presumably close to a steady-state; *ii*) in the whole timespan explored by our simulations, there is no evidence for energetic particles escaping ahead of superluminal shocks, so that no upstream waves should be generated and the onset of efficient SDA should never occur.

The subpanels in the last panel of Fig. 2.22 show at time $\omega_{\rm p}t = 9000$ the power-law slope of the suprathermal tail and the fraction of particles and energy stored in the tail. We see that, with increasing magnetic obliquity within the subluminal range, the high-energy tail becomes flatter, more populated and more energetic. We remark that $\theta = 32^\circ$, which is still subluminal, does not accelerate very efficiently yet at $\omega_{\rm p}t = 9000$, but it will finally show a suprathermal tail similar to $\theta = 30^\circ$ and $\theta = 31^\circ$ (see the thin blue line in the bottom right panel of Fig. 2.22, for $\omega_{\rm p}t = 13500$). We also report that the downstream spectra for $\theta = 30^\circ$ and $\theta = 31^\circ$ do not significantly evolve between $\omega_{\rm p}t = 9000$ and $\omega_{\rm p}t = 13500$ (not shown), apart from the linear increase of the high-energy

cutoff. This suggests that the timespan of our simulations can reasonably capture the evolution of $\theta \lesssim \theta_{\rm crit}$ shocks until they reach a steady-state, and that the values we quote in the subpanels of Fig. 2.22 (and Fig. 2.11 as well) are close to the saturation values.

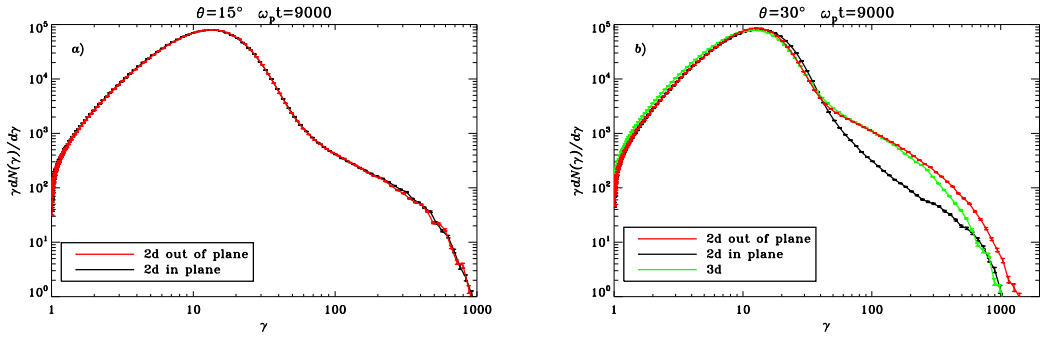


Fig. 2.21.— Comparison of downstream particle spectra at time $\omega_{\rm p}t = 9000$ among different geometries of the upstream background magnetic field: *a*) For fixed magnetic obliquity $\theta = 15°$, magnetic field lying either in the simulation plane (black) or in a plane perpendicular to the simulation plane (red); *b*) For fixed magnetic obliquity $\theta = 30°$, two-dimensional simulations with either in-plane (black) or out-of-plane (red) magnetic field and a three-dimensional run (green).

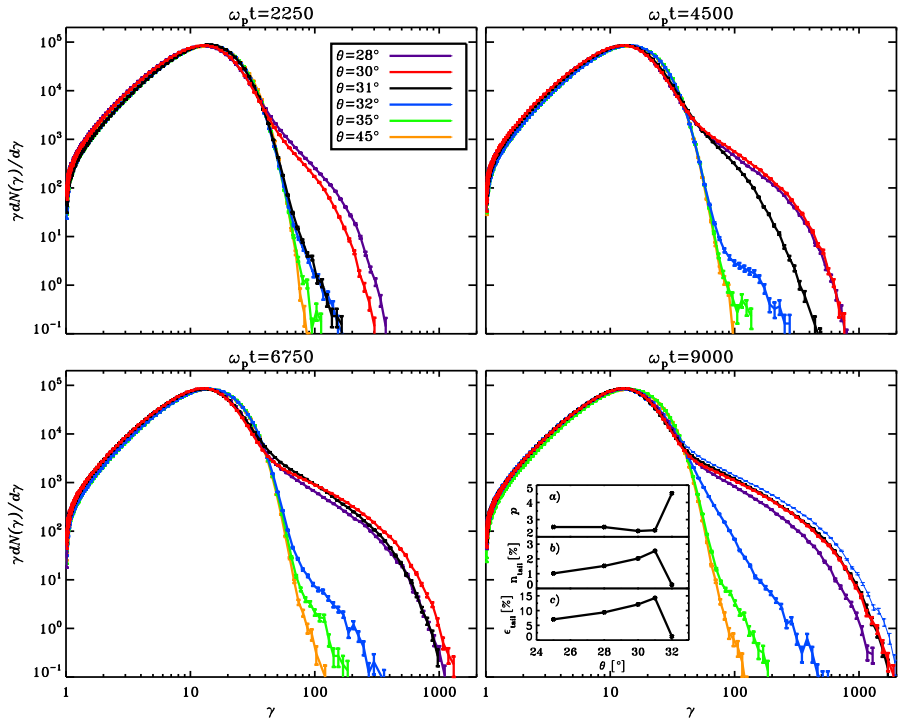


Fig. 2.22.— Time evolution of downstream particle spectra for different obliquities, across the boundary $\theta_{\rm crit} \approx 34°$ between subluminal and superluminal shocks: $\theta = 28°$ (violet), 30° (red), 31° (black), 32° (blue), 35° (green) and 45° (yellow). The subpanels *a*)-*c*) in the last panel show at time $\omega_{\rm p}t = 9000$ the power-law slope of the suprathermal tail and the fraction of particles and energy stored in the tail, as a function of the obliquity angle θ. The thin blue line in the bottom right panel shows the particle spectrum for $\theta = 32°$ at $\omega_{\rm p}t = 13500$.

Bibliography

Achterberg, A., Gallant, Y. A., Kirk, J. G., & Guthmann, A. W. 2001, MNRAS, 328, 393

Alsop, D. & Arons, J. 1988, Physics of Fluids, 31, 839

Amato, E. & Arons, J. 2006, ApJ, 653, 325

Appl, S. & Camenzind, M. 1988, A&A, 206, 258

Ballard, K. R. & Heavens, A. F. 1991, MNRAS, 251, 438

Bednarz, J. & Ostrowski, M. 1998, Physical Review Letters, 80, 3911

Begelman, M. C. & Kirk, J. G. 1990, ApJ, 353, 66

Bell, A. R. 1978, MNRAS, 182, 147

Blandford, R. & Eichler, D. 1987, Phys. Rep., 154, 1

Blandford, R. D. & Ostriker, J. P. 1978, ApJ, 221, L29

Bret, A. & Dieckmann, M. E. 2008, Physics of Plasmas, 15, 062102

Buneman, O. 1993, in "Computer Space Plasma Physics", Terra Scientific, Tokyo, 67

Chang, P., Spitkovsky, A., & Arons, J. 2008, ApJ, 674, 378

Chen, G. & Armstrong, T. P. 1975, in International Cosmic Ray Conference, Vol. 5, 1814–1819

Coroniti, F. V. 1990, ApJ, 349, 538

de Hoffmann, F. & Teller, E. 1950, Physical Review, 80, 692

Del Zanna, L., Amato, E., & Bucciantini, N. 2004, A&A, 421, 1063

Drury, L. O. 1983, Reports on Progress in Physics, 46, 973

Ellison, D. C. & Double, G. P. 2004, Astroparticle Physics, 22, 323

Gabuzda, D. C., Murray, É., & Cronin, P. 2004, MNRAS, 351, L89

Gallant, Y. A., Hoshino, M., Langdon, A. B., Arons, J., & Max, C. E. 1992, ApJ, 391, 73

Greenwood, A. D., Cartwright, K. L., Luginsland, J. W., & Baca, E. A. 2004, Journal of Computational Physics, 201, 665

Gruzinov, A. & Waxman, E. 1999, ApJ, 511, 852

Hoshino, M. 2008, ApJ, 672, 940

Hoshino, M. & Arons, J. 1991, Physics of Fluids B, 3, 818

Hoshino, M., Arons, J., Gallant, Y. A., & Langdon, A. B. 1992, ApJ, 390, 454

Jokipii, J. R. 1982, ApJ, 255, 716

Jones, F. C., Jokipii, J. R., & Baring, M. G. 1998, ApJ, 509, 238

Keshet, U., Katz, B., Spitkovsky, A., & Waxman, E. 2008, ArXiv:astro-ph/0802.3217

Keshet, U. & Waxman, E. 2005, Physical Review Letters, 94, 111102

Kirk, J. G., Guthmann, A. W., Gallant, Y. A., & Achterberg, A. 2000, ApJ, 542, 235

Kirk, J. G. & Heavens, A. F. 1989, MNRAS, 239, 995

Komissarov, S. S. & Lyubarsky, Y. E. 2003, MNRAS, 344, L93

—. 2004, MNRAS, 349, 779

Landau, L. D., Pitaevskii, L. P., & Lifshitz, E. M. 1984, Electrodynamics of continuous media (Elsevier Science & Technology Books)

Langdon, A. B., Arons, J., & Max, C. E. 1988, Physical Review Letters, 61, 779

Lee, M. A., Shapiro, V. D., & Sagdeev, R. Z. 1996, J. Geophys. Res., 101, 4777

Lyubarsky, Y. 2006, ApJ, 652, 1297

Lyubarsky, Y. & Kirk, J. G. 2001, ApJ, 547, 437

Lyubarsky, Y. & Liverts, M. 2008, ApJ, 682, 1436

Lyutikov, M. & Blandford, R. 2003, ArXiv:astro-ph/0312347

Medvedev, M. V. & Loeb, A. 1999, ApJ, 526, 697

Niemiec, J. & Ostrowski, M. 2004, ApJ, 610, 851

Ostrowski, M. & Bednarz, J. 2002, A&A, 394, 1141

Pétri, J. & Lyubarsky, Y. 2007, A&A, 473, 683

Pushkarev, A. B., Gabuzda, D. C., Vetukhnovskaya, Y. N., & Yakimov, V. E. 2005, MNRAS, 356, 859

Sikora, M., Begelman, M. C., Madejski, G. M., & Lasota, J.-P. 2005, ApJ, 625, 72

Sironi, L. & Goodman, J. 2007, ApJ, 671, 1858

Spitkovsky, A. 2005, in AIP Conference Series, Vol. 801, 345-350, astro-ph/0603211

—. 2008a, ApJ, 673, L39

—. 2008b, ApJ, 682, L5

Synge, J. L. 1957, The Relativistic Gas (North-Holland, Amsterdam), 33

Webb, G. M., Axford, W. I., & Terasawa, T. 1983, ApJ, 270, 537

Weibel, E. S. 1959, Physical Review Letters, 2, 83

Chapter 3

Particle Acceleration in Relativistic Magnetized Collisionless Electron-Ion Shocks

Abstract

We investigate shock structure and particle acceleration in relativistic magnetized collisionless electron-ion shocks by means of 2.5D particle-in-cell simulations with ion-to-electron mass ratios (m_i/m_e) ranging from 16 to 1000. We explore a range of inclination angles between the pre-shock magnetic field and the shock normal. In "subluminal" shocks, where relativistic particles can escape ahead of the shock along the magnetic field lines, ions are efficiently accelerated via the first-order Fermi process. The downstream ion spectrum consists of a relativistic Maxwellian and a high-energy power-law tail, which contains $\sim 5\%$ of ions and $\sim 30\%$ of ion energy. Its slope is -2.1 ± 0.1. The scattering is provided by short-wavelength non-resonant modes produced by Bell's instability (Bell 2004), whose growth is seeded by the current of shock-accelerated ions that propagate ahead of the shock. Upstream electrons enter the shock with lower energy than ions (albeit by only a factor of $\sim 5 \ll m_i/m_e$), so they are more strongly tied to the field. As a result, only $\sim 1\%$ of the incoming electrons are accelerated at the shock before being advected downstream, where they populate a steep power-law tail (with slope -3.5 ± 0.1). For "superluminal" shocks, where relativistic particles cannot outrun the shock along the field, the self-generated turbulence is not strong enough to permit efficient Fermi acceleration, and the ion and electron downstream spectra are consistent

with thermal distributions. The incoming electrons are heated up to equipartition with ions, due to strong electromagnetic waves emitted by the shock into the upstream. Thus, efficient electron heating ($\gtrsim 15\%$ of the upstream ion energy) is the universal property of relativistic electron-ion shocks, but significant nonthermal acceleration of electrons ($\gtrsim 2\%$ by number, $\gtrsim 10\%$ by energy, with slope flatter than -2.5) is hard to achieve in magnetized flows and requires weakly magnetized shocks (magnetization $\sigma \lesssim 10^{-3}$), where magnetic fields self-generated via the Weibel instability are stronger than the background field. These findings place important constraints on the models of AGN jets and Gamma Ray Bursts that invoke particle acceleration in relativistic magnetized electron-ion shocks.

1. Introduction

Nonthermal emission from Pulsar Wind Nebulae (PWNe), jets from Active Galactic Nuclei (AGNs), Gamma-Ray Bursts (GRBs) and supernova remnants (SNRs) is usually modeled as synchrotron or inverse Compton radiation from a power-law population of electrons accelerated at collisionless shocks. The slope of the power-law tail and the acceleration efficiency, i.e., the fraction of particles and energy stored in the tail, are usually parameterized *ad hoc* to fit the observations, due to the lack of a fully self-consistent theory of particle acceleration in shocks.

When encountering a collisionless shock, charged particles may be accelerated by means of two basic mechanisms. In first-order Fermi acceleration (or Diffusive Shock Acceleration, DSA), particles diffuse back and forth across the shock front and gain energy by scattering from magnetic turbulence embedded in the converging flows (e.g., Blandford & Ostriker 1978; Bell 1978; Drury 1983; Blandford & Eichler 1987).[1] Alternatively, in magnetized flows, particles may also gain energy directly from the background motional electric field $\mathbf{E} = -\boldsymbol{\beta} \times \mathbf{B}$ while they gyrate around the shock. The latter process is given different names depending on the barrier that reflects the particles back upstream, thereby allowing for multiple energizations. If the reflecting barrier is magnetic, due to the shock compression of the upstream field, the acceleration mechanism is called Shock Drift Acceleration (SDA; e.g., Chen & Armstrong 1975; Webb et al. 1983; Begelman & Kirk 1990). If the barrier is electrostatic, caused by electron-ion charge separation at the shock, the

[1]The second-order Fermi process, where acceleration of particles is due to stochastic encounters with randomly-propagating waves, does not involve a shock. In this chapter, we only discuss the first-order Fermi mechanism, which we simply call "the Fermi process".

process is named Shock Surfing Acceleration (SSA; e.g., Lee et al. 1996; Hoshino & Shimada 2002; Shapiro & Üçer 2003). The SSA mechanism can only operate in electron-ion shocks (and not in electron-positron shocks), since no electrostatic barrier appears if the incoming species have the same rigidity. In general, not only the elemental composition, but also the magnetization and bulk Lorentz factor of the upstream flow may be important in determining the mechanism responsible for particle energization and the resulting acceleration efficiency. If the upstream medium is magnetized, an additional parameter is the obliquity angle that the upstream magnetic field makes with the shock direction of propagation.

In the absence of turbulence, cold particles are constrained to slide along the ordered magnetic field lines, which are advected downstream from the shock. For high magnetic inclinations, in order to return upstream the particles should be moving along the field faster than the speed of light (Begelman & Kirk 1990). If efficient acceleration requires repeated crossings of the shock, such "superluminal" geometries should be very inefficient for particle acceleration (Kirk & Heavens 1989; Ballard & Heavens 1991). This conclusion can be questioned in the following cases: (*i*) a fraction of the incoming particles is pre-heated such that their gyro-radius is larger than the thickness of the shock, and they can experience multiple shock crossings before being advected downstream; (*ii*) there are appreciable fluctuations of the magnetic field at the shock, that may create local "subluminal" configurations where particle acceleration can occur. In both cases, the final outcome is difficult to predict. In support of the pre-heating scenario, a candidate mechanism has been proposed for electron heating in magnetized oblique shocks. The synchrotron maser instability at relativistic shock fronts (Langdon et al. 1988) generates a coherent train of intense electromagnetic "precursor" waves propagating upstream, which push on electrons and make them lag behind ions. The difference in bulk velocities between electrons and ions generates longitudinal electrostatic oscillations (Lyubarsky 2006), that eventually dissipate by heating the upstream electrons (Hoshino 2008). However, it is not clear if this will result in electron acceleration, besides heating. Moreover, if the incoming electrons are too hot, the synchrotron maser instability may be suppressed, thereby self-limiting the whole process.

Alternatively, as mentioned above in (*ii*), particle acceleration in superluminal shocks may be permitted if a sufficiently strong turbulence reorients the field at the shock, opening subluminal channels where particles can return upstream. If the pre-shock medium is not turbulent by itself, such magnetic field fluctuations need to be generated by shock-accelerated particles that escape upstream across field lines. However, the existence of these particles is not obvious *a priori*, since in

turn it requires sufficient magnetic turbulence to mediate their acceleration. In any case, the highly nonlinear problem of wave generation and particle heating and acceleration needs to be addressed with a self-consistent approach.

Fully kinetic particle-in-cell (PIC) simulations provide a powerful tool for exploration of the structure of collisionless shocks from first principles, thus determining self-consistently the interplay between shock-generated waves and accelerated particles. As opposed to semi-analytic kinetic theory methods (e.g., Kirk & Heavens 1989; Ballard & Heavens 1991; Achterberg et al. 2001; Keshet & Waxman 2005) or Monte Carlo test-particle simulations (e.g., Bednarz & Ostrowski 1998; Niemiec & Ostrowski 2004; Ellison & Double 2004; Baring 2010), PIC simulations can tackle the problem of particle acceleration in shocks without the need for simplifying assumptions about the nature of the magnetic turbulence or the details of wave-particle interactions.

Multi-dimensional PIC simulations of relativistic *unmagnetized* shocks have been presented by Spitkovsky (2005) for electron-positron flows (see also Chang et al. 2008; Keshet et al. 2009; Haugboelle 2010), and by Spitkovsky (2008a) for electron-ion flows. Spitkovsky (2008b) and Martins et al. (2009) have shown, respectively for electron-positron and electron-ion plasmas, that unmagnetized shocks naturally produce accelerated particles as part of the shock evolution. A comprehensive study of particle acceleration in relativistic *magnetized electron-positron* shocks has been performed by Sironi & Spitkovsky (2009a, hereafter SS09) by means of multi-dimensional PIC simulations. In agreement with earlier one-dimensional (1D) results (Langdon et al. 1988; Gallant et al. 1992), they found that particle acceleration in superluminal pair shocks is extremely inefficient, meaning that self-generated turbulence is not strong enough to allow for significant particle diffusion across field lines.

In this chapter, we extend the analysis of SS09 by investigating via 2.5D PIC simulations the properties of relativistic *magnetized electron-ion* shocks. In particular, we explore how the level of electron heating (namely, the fraction of pre-shock bulk kinetic energy transferred to post-shock electrons) and the efficiency of particle acceleration depend on the upstream bulk Lorentz factor, magnetization, and magnetic obliquity. We find that electrons are heated up to a few tens of percent of the upstream ion energy, regardless of the upstream conditions, and they can even reach equipartition with ions in high-obliquity magnetized shocks. With regards to particle acceleration, efficient nonthermal energization of ions (and electrons, to a lesser degree) via the first-order Fermi process is observed for *subluminal* configurations. The pressure of shock-accelerated ions

propagating ahead of the shock can substantially perturb the incoming flow and alter the shock structure. In contrast, the downstream spectrum of both ions and electrons in *superluminal* shocks does not show any signature of acceleration to nonthermal energies, in agreement with the results of SS09 for electron-positron shocks.

This chapter is organized as follows. In §2 we discuss the setup of our simulations and the magnetic field geometry. In §3 the shock structure and internal physics is investigated for one representative subluminal and one superluminal obliquity; for the same angles, in §4 we explore the mechanisms responsible for particle heating and acceleration. The main results of our work, concerning the efficiency of electron heating and particle acceleration as a function of bulk Lorentz factor, magnetization and magnetic obliquity, are presented in §5. We summarize our findings in §6 and comment on the application of our results to astrophysical scenarios.

2. Simulation Setup

We use the 3D electromagnetic PIC code TRISTAN-MP (Spitkovsky 2005), which is a parallel version of the publicly available code TRISTAN (Buneman 1993) that was optimized for studying collisionless shocks. Our simulation setup parallels SS09 very closely, which we repeat here for completeness.

The shock is set up by reflecting a cold electron-ion "upstream" flow from a conducting wall located at $x = 0$ (Fig. 3.1). The interaction between the incoming beam (that propagates along $-\hat{\boldsymbol{x}}$) and the reflected beam triggers the formation of a shock, which moves away from the wall along $+\hat{\boldsymbol{x}}$. This setup is equivalent to the head-on collision of two identical plasma shells, which would form a forward and reverse shock and a contact discontinuity. Here, we follow only one of these shocks, and replace the contact discontinuity with the conducting wall. The simulation is performed in the "wall" frame, where the "downstream" plasma behind the shock has zero x-velocity. Since the downstream plasma is at rest, no extra boosts are needed to study the spectra.

We perform simulations in both 2D and 3D computational domains, and we find that most of the shock physics is well captured by 2D simulations. Therefore, to follow the shock evolution for longer times with fixed computational resources, we mainly utilize 2D runs. All three components of particle velocities and electromagnetic fields are tracked, however. So, our simulations are effectively "2.5D", i.e., 2D in physical space but 3D in momentum space.

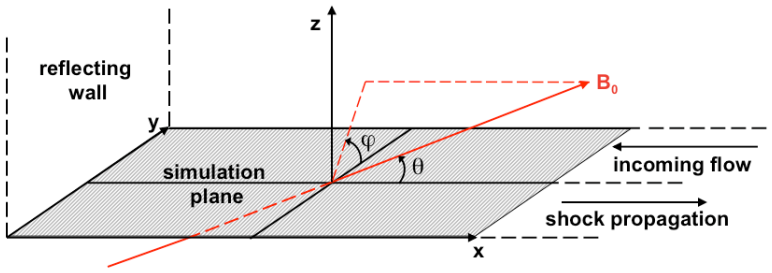


Fig. 3.1.— Simulation geometry. The upstream magnetic field (red arrow), oriented at an angle θ to the shock direction of propagation, can be either in the xy plane of the simulation (for $\varphi = 0°$) or in the xz plane (for $\varphi = 90°$).

We use a rectangular simulation box in the xy plane, with periodic boundary conditions in the y direction (Fig. 3.1). Each computational cell is initialized with two electrons and two ions, but we also performed limited experiments with a larger number of particles per cell (up to 8 per species), obtaining essentially the same results. The relativistic *electron* skin depth for the incoming plasma ($c/\omega_{\rm pe}$) is resolved with 10 computational cells and the simulation timestep is $\Delta t = 0.045\,\omega_{\rm pe}^{-1}$. Here, $\omega_{\rm pe} \equiv (4\pi e^2 n_e/\gamma_0 m_e)^{1/2}$ is the relativistic electron plasma frequency for the upstream flow, with electron number density n_e (measured in the wall frame) and bulk Lorentz factor γ_0. We employ a reduced mass ratio $m_i/m_e = 16$, which allows to follow the shock evolution for longer times (in units of $\omega_{\rm pi}^{-1} = (m_i/m_e)^{1/2}\omega_{\rm pe}^{-1}$), while still clearly separating the ion and electron dynamical scales. As we show in Appendix 7, we obtain essentially the same results when using higher mass ratios (we tried $m_i/m_e = 100$, and up to $m_i/m_e = 1000$ in some cases), which suggests that a mass ratio $m_i/m_e = 16$ is already "large" enough to capture the correct acceleration physics in our shocks. In units of the relativistic *ion* skin depth $c/\omega_{\rm pi} = (m_i/m_e)^{1/2} c/\omega_{\rm pe}$, the computational domain is $\sim 25\,c/\omega_{\rm pi}$ wide (along y), which corresponds to 1024 cells for $m_i/m_e = 16$ and 3072 cells for $m_i/m_e = 100$. We also tried larger domains, up to 4096 cells wide (or $\sim 100\,c/\omega_{\rm pi}$ for $m_i/m_e = 16$), obtaining similar results. Our computational box expands along x at the speed of light (see below), reaching the length of $\sim 10,000\ c/\omega_{\rm pi}$ (or $\sim 400,000$ cells for $m_i/m_e = 16$) at the final time of our longest simulation.

The incoming electron-ion stream is injected along $-\hat{\mathbf{x}}$ with bulk Lorentz factor $\gamma_0 = 15$ and a small thermal spread $\Delta\gamma = 10^{-4}$. As we comment in §5, our results can be rescaled to account for a different γ_0 (we explore a wide range in γ_0, from 3 to 50). The upstream flow is seeded

with a background magnetic field B_0 such that the ratio of magnetic to kinetic energy density is $\sigma \equiv B_0^2/4\pi\gamma_0 m_i n_i c^2 = 0.1$, where n_i $(=n_e)$ is the number density of incoming ions.[2] In §5 we show how our results change for lower magnetizations, down to $\sigma = 10^{-5}$, at which point the shock properties approach the unmagnetized case discussed by Spitkovsky (2008a). The magnetized regime we explore here may be relevant for internal shocks in GRBs and AGN jets, as we discuss in §6.

For each value of the magnetization, we explore a range of magnetic inclinations by varying the angle θ between the shock direction of propagation $+\hat{\mathbf{x}}$ and the upstream magnetic field $\mathbf{B}_0$ (Fig. 3.1). The magnetic obliquity angle θ is measured in the wall frame. We vary θ from $\theta = 0°$, which corresponds to a "parallel" shock, with magnetic field aligned with the shock normal, up to $\theta = 90°$, i.e., a "perpendicular" shock, with magnetic field along the shock front. For $\theta \neq 0°$, in the upstream medium we also initialize a motional electric field $\mathbf{E}_0 = -\boldsymbol{\beta}_0 \times \mathbf{B}_0$, where $\boldsymbol{\beta}_0 = -\beta_0\,\hat{\mathbf{x}}$ is the three-velocity of the injected plasma. As we show in Appendix 8, our results do not depend on the orientation of the magnetic field with respect to the simulation plane, as parameterized by the azimuthal angle φ shown in Fig. 3.1 ($\varphi = 0°$ for field in the xy plane of the simulations, $\varphi = 90°$ for field in the xz plane).

As a byproduct of the shock evolution, particles and electromagnetic waves may propagate upstream from the shock at the speed of light. To ensure that the long-term evolution of the shock is captured correctly, they should not be removed from the computational domain, otherwise we would artificially suppress any feedback they may have on the shock. For this reason, we employ a "moving injector" (receding from the wall at the speed of light) and an expanding simulation box (see SS09 for details). This permits us to follow the shock evolution as far as the computational resources allow, preserving all the particles and waves generated by the shock.

However, since the shock velocity is smaller than the speed of light, the distance between the injector and the shock increases linearly with time. In all PIC codes, a numerical heating instability arises when cold relativistic plasma propagates for large distances over the numerical grid (Dieckmann et al. 2006). For *electron-positron* shocks, in order to suppress this instability we let our moving injector periodically jump backward (i.e., towards the wall), so that the distance

[2]The magnetization parameter as defined above pertains to ions. For electrons, $\sigma_{e0} = (m_i/m_e)\,\sigma$ at injection, where the electron kinetic energy is $\gamma_0 m_e c^2$. So, it seems that the electron magnetization would depend on m_i/m_e, for fixed σ. However, on their way to the shock, electrons increase their average energy up to a fraction α of the initial ion energy $\gamma_0 m_i c^2$, where α is independent of m_i/m_e (see Appendix 7). So, the "effective" electron magnetization $\sigma_{e,\mathrm{eff}} = (1/\alpha)\,\sigma$ does not depend on the mass ratio.

between the shock and the injector stayed roughly constant (see the "jumping injector" technique in SS09). However, for *electron-ion* shocks this is not a viable solution, since the incoming flow is heavily influenced by particles and waves generated by the shock, as we show in §3. By resetting the position of the injector, we would discard such particles and waves, thus dramatically interfering with the long-term evolution of the shock. Instead, we notice that the onset of the numerical instability is significantly delayed if we reduce the transverse size of our simulations. For long ($\gtrsim 2000\,\omega_{\rm pi}^{-1}$) simulations of strongly magnetized high-obliquity shocks, which are more affected by the numerical instability, we then employ a computational box with only 256 transverse cells (down to 64, in some cases). At $\omega_{\rm pi}t \lesssim 2000$ (i.e., before the growth of the instability), the results from such small simulations are identical to our "fiducial" runs with 1024 transverse cells, showing that the relevant 2D properties of the shock are still preserved. For $m_i/m_e = 16$, a transition to the 1D regime is observed only for computational boxes with less than 32 transverse cells.

3. Shock Structure

We now describe the structure and internal physics of relativistic magnetized electron-ion shocks. In this section, as well as in §4, we fix the upstream bulk Lorentz factor ($\gamma_0 = 15$) and magnetization ($\sigma = 0.1$), and we investigate two representative magnetic inclination angles: $\theta = 15^\circ$ and $\theta = 75^\circ$. The dependence of our findings on the bulk Lorentz factor, magnetization and magnetic obliquity of the upstream flow is discussed in §5.

According to the criterion anticipated in §1, a shock is "superluminal" if particles cannot escape ahead of the shock by sliding along the magnetic field. In the upstream frame, this corresponds to magnetic inclinations θ' such that $\cos\theta' > \cos\theta'_{\rm crit} = \beta'_{\rm sh}$, where $\beta'_{\rm sh}$ is the shock speed in the upstream frame. In the simulation frame, the critical obliquity angle that separates subluminal and superluminal configurations will be $\theta_{\rm crit} = {\rm arccot}[\gamma_{\rm sh}\,(\beta_0 + \beta_{\rm sh})]$, where $\beta_{\rm sh}$ and $\gamma_{\rm sh}$ are the shock velocity and Lorentz factor in the simulation frame (see SS09 for a detailed discussion). For $\gamma_0 = 15$ and $\sigma = 0.1$, the critical obliquity angle is $\theta_{\rm crit} \simeq 34^\circ$, and it stays confined within a relatively narrow range (between $\sim 26^\circ$ and $\sim 42^\circ$) for relativistic ($\gamma_0 \gtrsim 2$) flows with moderate magnetization ($\sigma \lesssim 1.0$), as shown in Fig. 2 of SS09. It follows that $\theta = 15^\circ$ is a subluminal shock, and $\theta = 75^\circ$ is superluminal.

Figs. 3.2-3.6 present the internal structure of the shock as a function of the longitudinal coordinate x. For the subluminal angle $\theta = 15^\circ$, Fig. 3.2 covers the whole longitudinal extent of the

simulation domain, whereas Fig. 3.3 focuses on a smaller region around the shock, as delimited by the vertical dashed red lines in the first panel of Fig. 3.2. Both figures refer to $\omega_{\rm pi}t = 2250$, when the shock is already fully developed. The temporal evolution of the shock, from the early stages up to the self-similar state approached at $\omega_{\rm pi}t \gtrsim 2000$, is shown in Fig. 3.4. Figs. 3.5 and 3.6 correspond respectively to Figs. 3.2 and 3.3 for the superluminal angle $\theta = 75^\circ$, at time $\omega_{\rm pi}t = 1350$. Here, we choose an earlier time (with respect to $\omega_{\rm pi}t = 2250$ for $\theta = 15^\circ$), to avoid the numerical instability discussed in §2, which grows at $\omega_{\rm pi}t \sim 2000$ in high-obliquity shocks. Yet, time $\omega_{\rm pi}t = 1350$ already captures the main properties of the long-term evolution of the shock.

This section is organized as follows. In §3.1 we present a general description of the structure of relativistic magnetized shocks in electron-ion plasmas. We focus on the properties they share with electron-positron shocks, and emphasize the differences. Then, we separately discuss the cases of subluminal (in §3.2, for $\theta = 15^\circ$) and superluminal (in §3.3, for $\theta = 75^\circ$) shocks, describing how the obliquity of the field affects the shock structure.

3.1. Magnetized Relativistic Electron-Ion Shocks: General Overview

For all oblique magnetic configurations (i.e, with the exception of strictly parallel and perpendicular shocks), a fluid structure with two shocks (a strong "fast" shock and a weak "slow" shock) is seen in our PIC simulations of both electron-positron (SS09) and electron-ion flows, as expected from analytic theory (Majorana & Anile 1987) and MHD calculations (Komissarov 2003). The incoming fluid does not stop completely at the fast shock. The residual bulk velocity is larger than the local slow magnetosonic speed, and a slow shock is formed behind the fast shock (i.e., closer to the wall, in our simulations). The slow shock transition in our PIC simulations is not as sharp as MHD would predict, due to insufficient dissipation along the magnetic field. The slow shock becomes slower and weaker as θ approaches either 0° or 90°, and relatively stronger and faster for intermediate obliquities. For oblique configurations, we will neglect the slow shock and refer to the fast shock simply as "the shock".

The jump in density and electromagnetic fields at the fast shock, as well as the shock velocity, are in agreement with MHD calculations, and with PIC simulations of electron-positron shocks (see Table 1 in SS09). As observed in pair shocks, at low obliquities ($\theta \lesssim 45^\circ$) the effective adiabatic index of the downstream plasma is 4/3, as in a 3D relativistic gas, whereas for $\theta \gtrsim 45^\circ$ it tends to 3/2, as in a 2D fluid. In fact, right behind the shock the particle motion is mostly confined to

the plane orthogonal to the field, which for large inclination angles is nearly degenerate with the direction of propagation of the incoming flow. This prevents efficient isotropization along the field and results in a 2D adiabatic index.[3] Isotropization proceeds with distance behind the shock, and the particle distribution far downstream is roughly isotropic.

The mechanism that mediates randomization of particles at the shock varies depending on the field obliquity. In *electron-positron* magnetized plasmas, SS09 found that low-obliquity shocks are mediated by Weibel-like filamentation instabilities (Weibel 1959; Medvedev & Loeb 1999; Gruzinov & Waxman 1999), as happens for unmagnetized shocks (e.g., Spitkovsky 2005).[4] The free energy for the instability comes from the counter-streaming between the incoming flow and the shock-accelerated particles propagating upstream.

Low-obliquity shocks in *electron-ion* magnetized plasmas are still mediated by counter-streaming instabilities, but the nature of the instability is different. For our fiducial values $\gamma_0 = 15$ and $\sigma = 0.1$, the shock evolution at early times ($\omega_{\rm pi} t \lesssim 300$) is governed by the *electron* Weibel instability, which has the largest growth rate ($\sim \omega_{\rm pe} \gg \omega_{\rm pi}$). Later on, the so-called Bell's current-driven instability develops (Bell & Lucek 2001; Bell 2004; Reville et al. 2006), and the shock results from nonlinear steepening of the circularly-polarized Alfvénic-type waves associated with the instability (see §3.2 for details). The source for the instability is the electric current of shock-accelerated ions that propagate ahead of the shock, and the polarization of the resulting modes is *non-resonant* with respect to the ion gyro-motion. As we argue in Appendix 9, Bell's instability governs the evolution of the shock only for a limited range of magnetizations ($10^{-2} \lesssim \sigma \lesssim 0.3$, for fixed $\gamma_0 = 15$). If $\sigma \gtrsim 0.3$, the polarization of the dominant mode changes, and the waves are now generated via gyro-frequency *resonance* with the high-energy ions heading upstream (Kulsrud & Pearce 1969). For $\sigma \lesssim 10^{-2}$, the magnetic field is dynamically unimportant, and the shock is mediated by the *ion* Weibel instability (as in unmagnetized shocks, see Spitkovsky 2008a).

---

[3] We remark that this effect is not an artificial consequence of the reduced dimensionality of our computational domain, but it holds also for 3D simulations. Yet, it is more severe for 2D simulations with out-of-plane fields, since in this case the plane perpendicular to the field almost coincides with the simulation plane, and particle isotropization is even less efficient (see Appendix 8).

[4] More precisely, the instability results from the coupling of the electromagnetic filamentation (Weibel) instability with the electrostatic two-stream instability, as explained by Bret (2009). The resulting modes are oblique with respect to the streaming direction.

In contrast, high-obliquity shocks are mediated by magnetic reflection of the incoming flow off the shock-compressed magnetic field (Alsop & Arons 1988, SS09), for both electron-positron and electron-ion plasmas. The coherent gyration of the incoming particles at the shock triggers the so-called synchrotron maser instability (Hoshino & Arons 1991), where the ring of gyrating particles in the shock transition breaks up into bunches of charge. The bunches radiate a coherent train of transverse electromagnetic waves propagating into the upstream (Gallant et al. 1992; Hoshino et al. 1992; Spitkovsky 2005), with wavelength comparable to the Larmor radius of incoming *electrons* in the shock-compressed fields. For *electron-positron* shocks, this "precursor wave" does not have a significant impact on the upstream plasma (SS09). In *electron-ion* flows, the guiding-center velocity of the incoming electrons decreases, since they experience relativistic transverse oscillations in the strong field of the precursor wave. Due to their high mass, ions are less affected by the wave, and proceed at close to their initial velocity. The resulting difference in bulk velocity between the two species will generate a longitudinal electric field (Lyubarsky 2006), so that electrons are boosted toward the shock whereas ions are deboosted.[5] As a result, energy equipartition between electrons and ions may be achieved even before the flow arrives at the shock front, as we discuss in §3.3.

On the other hand, if the efficiency of ion-to-electron energy transfer *in the upstream* is moderate or negligible, the incoming ions will enter the shock with bulk energy much larger than electrons, so that they will penetrate deeper into the shock-compressed field (roughly, in proportion to the ratio of ion to electron Larmor radii). The resulting charge separation establishes a net electrostatic field *at the shock* pointing toward the upstream. The associated cross-shock electric potential attracts the incoming electrons into the downstream and decelerates the incoming ions, so that electrons may still end up with a significant fraction of the ion energy.

3.2. Subluminal Shocks: $0^\circ \leq \theta < \theta_{\rm crit}$

In this section, we describe the internal structure of subluminal magnetized electron-ion shocks. We take $\theta = 15^\circ$ as a representative case, but our results will apply to the whole range of subluminal angles, unless otherwise noted. In both electron-positron and electron-ion plasmas, subluminal shocks are characterized by the presence of a diffuse stream of shock-reflected particles

[5] A clarification is required here regarding our terminology: by "boost" or "deboost", we mean changes in the bulk x-momentum of the fluid; "heating" is used when the comoving particle distribution broadens; finally, "acceleration" is something pertaining only to a subsample of the particles, extracted from the bulk of the fluid and energized up to suprathermal velocities.

that propagate ahead of the shock. The counter-streaming between these "returning" particles and the incoming flow triggers the generation of waves in the upstream medium, which in turn affect the structure of the shock and the process of particle acceleration.

The longitudinal ion phase space in Fig. 3.2(e) shows the injected ions as a cold dense beam with $\gamma_i\beta_{\rm xi} \simeq -15$. In the shock transition layer (located at $\sim 730\,c/\omega_{\rm pi}$ for $\omega_{\rm pi}t = 2250$), the incoming ions are isotropized and thermalized. A diffuse population of high-energy ions moves ahead of the shock with $\gamma_i\beta_{\rm xi} > 0$, following the upstream magnetic field; as we show in §4.1, these "returning" ions have been accelerated at the shock. In the upstream ion spectrum (red lines in Fig. 3.2(h) and (i)), they appear as a high-energy broad bump, whereas the low-energy narrow peak is populated by the incoming particles. If transmitted downstream, the shock-accelerated ions will populate a power-law nonthermal tail, which is seen in the energy spectrum of Fig. 3.2(g) (red line) beyond the thermal distribution.

In contrast, very few electrons propagate back upstream from the shock (see the longitudinal phase space of electrons in Fig. 3.2(f)). This anticipates that in subluminal magnetized shocks electrons are accelerated with lower efficiency than ions.[6] Despite the lack of returning electrons, the electron spectrum downstream from the shock (blue line in panel (g)) deviates from a purely thermal distribution, which suggests that some electrons may be accelerated to suprathermal energies as they cross the shock from upstream to downstream. We refer to §4.1 for a detailed discussion of electron acceleration.

The fact that returning ions greatly outnumber electrons has multiple effects on the structure of the shock, which were absent in electron-positron flows. First, the positively-charged cloud of returning ions perturbs the incoming flow, causing significant transfer of energy from ions to electrons. At the leading edge of the population of returning ions ($x \sim 1800\,c/\omega_{\rm pi}$), the incoming particles cross a kind of double layer, where ions are boosted toward the shock and electrons are deboosted (see the average particle energy in Fig. 3.2(d), at $x \sim 1700\,c/\omega_{\rm pi}$; red for ions, blue for electrons). Beyond that point, electrons are significantly heated (see Fig. 3.2(f) at $800\,c/\omega_{\rm pi} \lesssim x \lesssim 1500\,c/\omega_{\rm pi}$), possibly via small-scale plasma oscillations (Spitkovsky et al., in prep). The average electron energy increases toward the shock at the expense of the ion bulk energy, reaching up to $\sim 15\%$ of the kinetic energy of injected ions (blue line in Fig. 3.2(d)). However,

[6]The very few high-energy electrons seen with $\gamma_e\beta_{\rm xe} > 0$ ahead of the shock (see Fig. 3.2(f), and the high-energy bump in the blue line of Fig. 3.2(i)) have been accelerated in the first stages of evolution, when the shock was mediated by the electron Weibel instability. When ion-driven instabilities take over, electron acceleration shuts off.

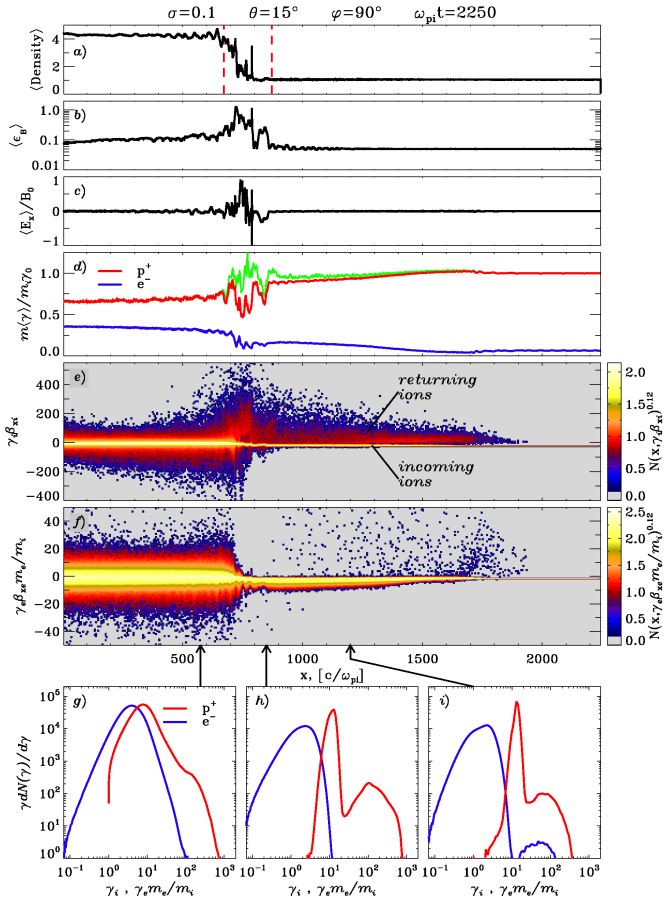


Fig. 3.2.— Internal structure of a $\theta = 15°$ subluminal shock at $\omega_{\rm p}t = 2250$. As a function of the longitudinal coordinate x, the following quantities are plotted: (a) y-averaged particle number density, in units of the upstream value; (b) y-averaged magnetic energy fraction $\epsilon_{\rm B} \equiv B^2/8\pi\gamma_0 m_i n_i c^2$; (c) y-averaged longitudinal electric field $E_{\rm x}$, normalized to the upstream magnetic field B_0; (d) mean particle energy (red for ions, blue for electrons) for particles moving toward the shock, in units of the bulk energy of injected ions (green line includes also the reflected ions; the equivalent for electrons overlaps with the blue line); longitudinal phase space of ions (e) and electrons (f); (g)-(i) particle energy spectra (red for ions, blue for electrons), at three locations across the flow, as marked by arrows at the bottom of panel (f).

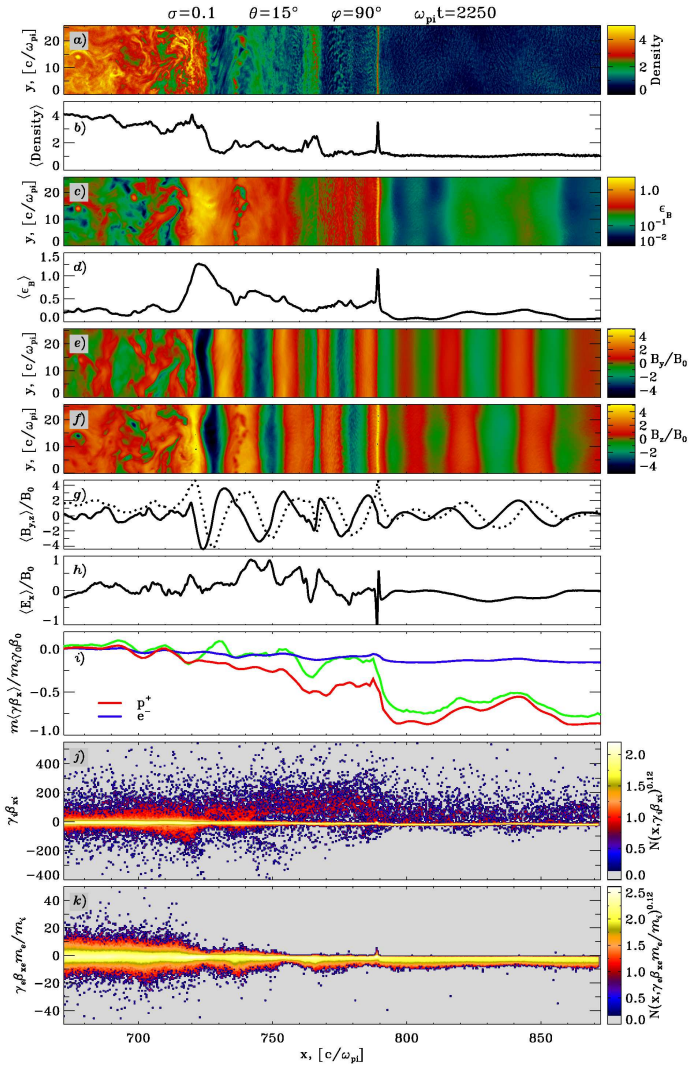


Fig. 3.3.— Internal structure of a $\theta = 15°$ subluminal shock at $\omega_p t = 2250$, zooming in on a region around the shock, as delimited by the vertical dashed red lines in Fig. 3.2(a). The shock is located at $x \sim 730\, c/\omega_p$. As a function of the longitudinal coordinate x, the following quantities are plotted: particle number density, in units of the upstream value (2D plot in the xy simulation plane, panel (a); 1D y-averaged profile, panel (b)); magnetic energy fraction $\epsilon_B \equiv B^2/8\pi\gamma_0 m_i n_i c^2$ (2D in (c), 1D in (d)); 2D plots of B_y (e) and B_z (f), normalized to the upstream magnetic field B_0; (g) corresponding 1D profiles, for B_y (solid) and B_z (dashed); (h) 1D profile of the longitudinal electric field E_x, in units of B_0; (i) mean particle x-momentum (red for ions, blue for electrons) for particles in the bulk (i.e., neglecting accelerated particles), in units of the momentum of injected ions (green line includes also the accelerated ions; the equivalent for electrons overlaps with the blue line); longitudinal phase space of ions (j) and electrons (k).

electrons still enter the shock with lower average energy than ions. The resulting charge separation in the shock-compressed magnetic field generates a net cross-shock electrostatic field ($E_{\rm x} > 0$, seen at $x \sim 730\,c/\omega_{\rm pi}$ in Fig. 3.2(c)). The associated cross-shock potential further increases the average electron energy; in the downstream, electrons have $\sim 30\%$ of the initial ion energy.

Furthermore, the imbalance between returning ions and electrons establishes a net (ion-driven) current in the upstream medium. The compensating current by the incoming electrons triggers the growth of the so-called Bell's instability, which is expected in this regime of shock parameters (Bret 2009; Lemoine & Pelletier 2010). As shown in panel (e) (for $B_{\rm y}$) and (f) (for $B_{\rm z}$) of Fig. 3.3, the instability generates circularly-polarized Alfvénic-type waves in the upstream medium, with wavevector along the background field.[7] The polarization of the waves is such that they are non-resonant with respect to the returning ions, and their wavelength ($\sim 25\,c/\omega_{\rm pi}$) is much smaller than the Larmor radius of the highest-energy particles. As shown in Fig. 3.3(g) (solid for $B_{\rm y}$, dotted for $B_{\rm z}$), the amplitude of the waves increases as they approach the shock, reaching up to ~ 4 times the strength of the initial field (measured in the simulation frame). Right before the shock ($x \sim 730\,c/\omega_{\rm pi}$), more than 50% of the bulk kinetic energy of the upstream flow has been converted into wave magnetic energy (see the magnetic energy fraction $\epsilon_{\rm B} \equiv B^2/8\pi\gamma_0 m_i n_i c^2$ in Fig. 3.3(d)). After the shock, the flow isotropizes, the electric current required to drive the instability vanishes, and the waves dissipate their energy into particle energy. The importance of Bell's waves for ion acceleration will be discussed in §4.1.

In the vicinity of the shock ($730\,c/\omega_{\rm pi} \lesssim x \lesssim 790\,c/\omega_{\rm pi}$), the pressure of returning ions is so large that they heavily impact the structure of the transition region and the dynamics of the incoming flow. Due to the push of the returning ions, the incoming plasma slows down significantly, by more than 50%, well before the shock front (see the jump in the average x-momentum of incoming ions at $x \sim 790\,c/\omega_{\rm pi}$, red line in Fig. 3.3(i)). The deceleration and consequent compression of the incoming fluid produces spikes in the number density and the magnetic energy, as observed at $x \sim 790\,c/\omega_{\rm pi}$ in Fig. 3.3(a)-(d). Such spikes are transient quasi-periodic structures, which get advected downstream after $\sim 100\,\omega_{\rm pi}^{-1}$ (or $\sim 30\,\omega_{\rm ci}^{-1}$, with the relativistic ion Larmor frequency $\omega_{\rm ci} \equiv \sqrt{\sigma}\,\omega_{\rm pi}$), and then reform at a slightly larger distance from the shock.[8] This is the characteristic signature of the

[7] This is most clear in simulations with in-plane $\mathbf{B}_0$ (i.e., $\varphi = 0^\circ$). Here, $\varphi = 90^\circ$, and one can only appreciate that the projection of the wavevector onto the simulation plane lies along $\hat{\mathbf{x}}$ (see Figs. 3.3(e) and (f)).

[8] They resemble the so-called Short Large-Amplitude Magnetic Structures, or SLAMS, observed at the Earth's bow shock (e.g., Schwartz et al. 1992).

so-called "shock reformation" process. Evidence of shock modification by the accelerated particles may also be seen in the ion energy spectrum, as we discuss in §4.1.

Time Evolution

The picture outlined above refers to $\omega_{\rm pi}t = 2250$, when the number of shock-accelerated ions ahead of the shock is already large enough to affect the structure of the transition region. Fig. 3.4 follows the shock evolution from $\omega_{\rm pi}t = 562$ to $\omega_{\rm pi}t = 2250$, showing how the y-averaged profiles of number density (panel (a)), magnetic energy fraction $\epsilon_{\rm B} \equiv B^2/8\pi\gamma_0 m_i n_i c^2$ (panel (b)), and transverse magnetic field $B_{\rm z}$ (panel (c)) change with time.

At early times ($\omega_{\rm pi}t = 562$, blue line), when the pressure of returning ions is negligible with respect to the ram pressure of the incoming flow, the density jump occurs on the scale of a few ion Larmor radii, like an ideal MHD shock. The magnetic field at the shock ($\epsilon_{\rm B} \sim 0.15$, blue line in panel (b)) mostly results from the compression of the upstream field, with a minor contribution from counter-streaming instabilities.

Starting from $\omega_{\rm pi}t = 1125$ (yellow line), the shock structure changes. For $\theta = 15^\circ$, the leading edge of the population of shock-accelerated ions recedes from the shock with x-velocity $\simeq 0.86\,c$ (corresponding to particles moving along the background field at the speed of light). Since the shock speed is $\simeq 0.32\,c$, the distance between the shock and the head of returning ions increases with time. Bell's instability, which at earlier times was suppressed by advection into the shock, can now amplify the field to highly nonlinear values ($B_{\rm z}$ up to a factor of ~ 4 larger than B_0, see green and red lines in panel (c); similar values for $B_{\rm y}$). The magnetic energy at the shock saturates at equipartition with the upstream kinetic energy ($\epsilon_{\rm B} \sim 0.5$, green and red lines in panel (b)).

The increase in the strength of Bell's waves is driven by the growth of the population of returning ions. At late times (green for $\omega_{\rm pi}t = 1687$, red for $\omega_{\rm pi}t = 2250$), the shock transition region becomes much wider (panel (a)), due to the increasing push of the returning ions on the incoming flow. The time-averaged density profile (see the thin black line in panel (a) for $\omega_{\rm pi}t = 2250$) displays the typical structure of "cosmic-ray modified" shocks (e.g., Berezhko & Ellison 1999), with a smooth density increase (the "cosmic-ray precursor") ahead of the main shock (here properly called "subshock"), now located at $x \sim 730\,c/\omega_{\rm pi}$. The magnetic energy stays large ($\epsilon_{\rm B} \sim 0.5$) across the whole transition region, with two characteristic peaks, one at the head of the cosmic-ray

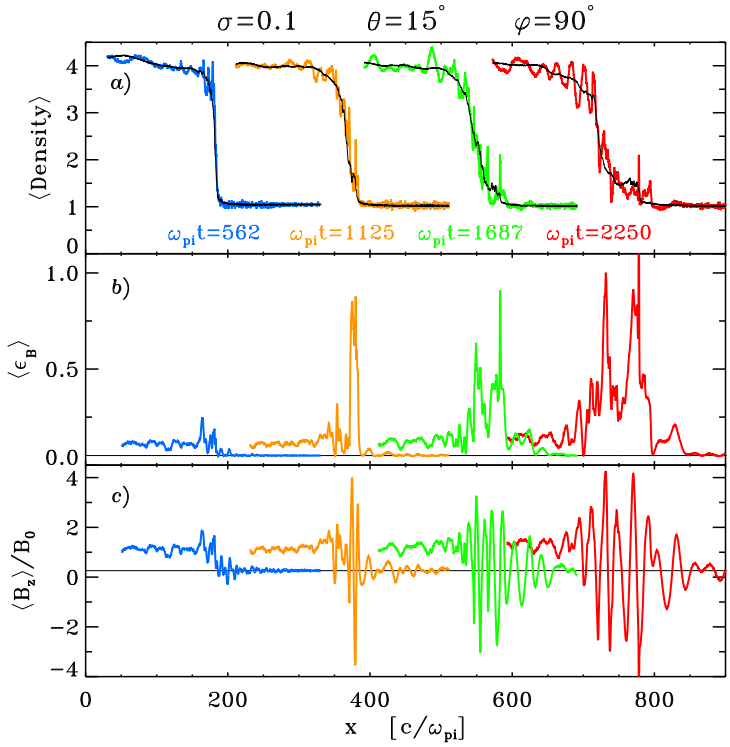


Fig. 3.4.— Time evolution of the internal structure of a $\theta = 15^\circ$ subluminal shock: $\omega_{\rm p}t = 562$ (blue), $\omega_{\rm p}t = 1125$ (yellow), $\omega_{\rm p}t = 1687$ (green), and $\omega_{\rm p}t = 2250$ (red). (a) y-averaged profiles of the particle number density; for each time, the superimposed thin black line is obtained by averaging over a narrow temporal window ($100\,\omega_{\rm pi}^{-1}$ wide), centered on that time. (b) y-averaged profiles of the magnetic energy fraction $\epsilon_{\rm B}$, with the horizontal black line showing the value at injection ($= 0.5\,\sigma = 0.05$). (c) y-averaged profiles of the transverse magnetic field $B_{\rm z}$ (in units of B_0), with the horizontal black line showing the value at injection ($= \sin 15^\circ \simeq 0.26$). For each time, we only plot a limited region ($\sim 300\,c/\omega_{\rm p}$ wide) around the instantaneous location of the shock.

precursor, where the incoming plasma is first decelerated and compressed, and the other at the subshock, where the flow comes to rest (see green and red lines in panel (b)). As clarified by panel (c), most of the magnetic energy in the region between the peaks results from the compression of Bell's waves.

At later times (we follow the shock evolution up to $\omega_{\rm pi}t = 5062$), the shock structure does not qualitatively change, but the thickness of the transition region (where $\epsilon_{\rm B} \sim 0.5$) still expands at $\simeq 0.1\,c$. A detailed study of the relation between shock width and mean free path of diffusively-accelerated ions will be presented elsewhere.

3.3. Superluminal Shocks: $\theta_{\rm crit} < \theta \leq 90^\circ$

We now turn to the structure of a shock with obliquity $\theta = 75^\circ$, as a representative example of superluminal shocks. In such shocks, particles following the magnetic field cannot propagate back from the shock into the upstream. In fact, in Fig. 3.5 no particles are seen ahead of the shock with positive x-momentum (neither ions nor electrons, phase spaces in Fig. 3.5(e) and (f) respectively). The shock transition region (located at $x \sim 600\,c/\omega_{\rm pi}$ for $\omega_{\rm pi}t = 1350$), which is not perturbed by the pressure of energetic particles, is much thinner than for $\theta = 15^\circ$, on the scale of a few Larmor radii of the incoming ions (compare the density profile in panel (b) between Fig. 3.3 and Fig. 3.6). Despite the absence of returning particles, the incoming flow can still be perturbed by the presence of the shock, e.g., via electromagnetic waves emitted by the shock into the upstream.

As discussed in §3.1, coherent Larmor gyration of the incoming electrons in the shock-compressed field triggers the synchrotron maser instability (Hoshino & Arons 1991), which creates a train of transverse electromagnetic waves propagating into the upstream. These "precursor" waves can be barely seen as short-scale ripples stretched along y in the upstream region of Fig. 3.6(c) (for $\epsilon_{\rm B}$) and Fig. 3.6(e) (for $B_{\rm y}$), especially in the vicinity of the shock.

The incoming electrons oscillate in the transverse field of the precursor waves, and their guiding-center velocity decreases. Since ions have larger inertia, their speed is not appreciably altered by the waves. Behind the head of the precursor ($x \sim 1200\,c/\omega_{\rm pi}$ in Fig. 3.5), the difference between the bulk velocities of ions and electrons generates a longitudinal "wakefield" $E_{\rm x} > 0$ (Lyubarsky 2006), so that electrons are initially boosted toward the shock, whereas ions are deboosted (see the average particle energy in Fig. 3.5(d) at $x \gtrsim 1150\,c/\omega_{\rm pi}$; red for ions, blue for

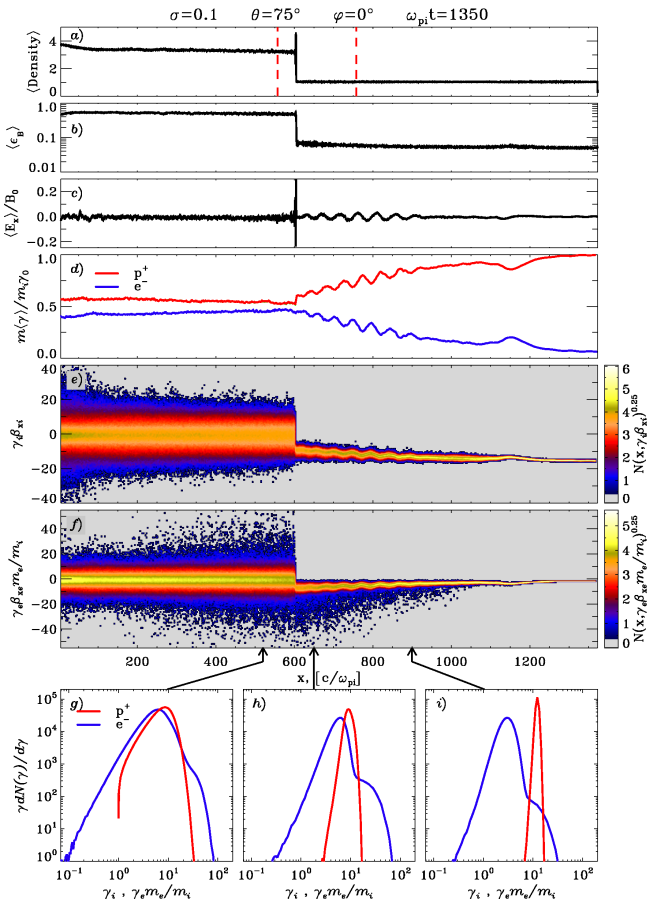


Fig. 3.5.— Internal structure of a $\theta = 75°$ superluminal shock at $\omega_p t = 1350$. See the caption of Fig. 3.2 for details. Here, the slow shock can be seen at $x \lesssim 50\, c/\omega_p$ as an increase in density (panel (a)) and decrease in magnetic energy (panel (b)).

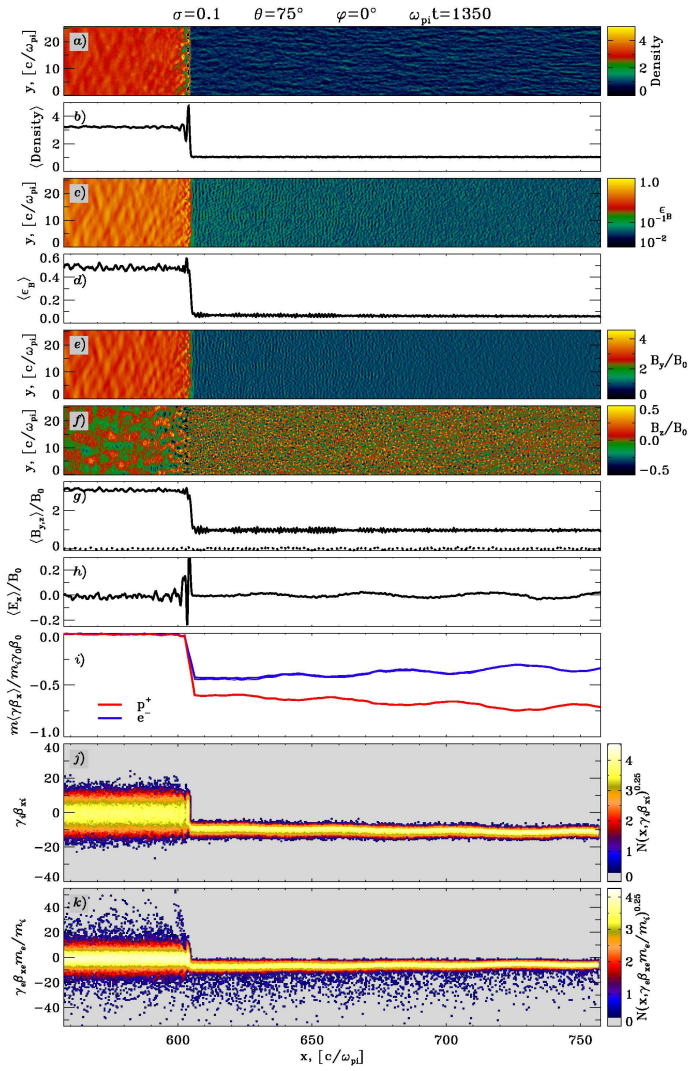


Fig. 3.6.— Internal structure of a $\theta = 75°$ superluminal shock at $\omega_p t = 1350$, zooming in on a region around the shock, as delimited by the vertical dashed red lines in Fig. 3.5(a). See the caption of Fig. 3.3 for details.

electrons). This initiates electrostatic wakefield oscillations in the incoming plasma (see E_x wiggles in Fig. 3.5(c) at $600\,c/\omega_{pi} \lesssim x \lesssim 900\,c/\omega_{pi}$), that mediate a *quasi-periodic* exchange of energy between ions and electrons in the upstream (see analogous wiggles in Fig. 3.5(d)).

If the strength of the precursor were uniform throughout the upstream region, after each wakefield oscillation the electron Lorentz factor would come back to the value γ_0 at injection (Lyubarsky 2006). Instead, we observe a *secular* increase in the average energy of electrons (and decrease for ions) as the incoming flow approaches the shock (Fig. 3.5(d) at $600\,c/\omega_{pi} \lesssim x \lesssim 1100\,c/\omega_{pi}$). This corresponds to a gradient in the amplitude of the precursor waves, which are stronger in the vicinity of the shock. Their radiative push on the incoming electrons will then be larger right in front of the shock, and weaker farther upstream. This imbalance establishes a net wakefield $E_x > 0$ throughout the upstream region (in addition to the oscillations discussed above), which results in a systematic increase of the average electron energy toward the shock, as seen in Fig. 3.5(d). Alternatively, one could define a "ponderomotive" electric potential (proportional to the radiative push of the precursor), whose increase toward the shock would explain the profile of average electron energy in Fig. 3.5(d) (Hoshino 2008).

As a result, electrons and ions enter the shock with roughly the same energy. Given the resulting absence of electron-ion charge separation at the shock, no significant cross-shock electric field is generated, so that the electron and ion average energies do not appreciably change across the shock (Fig. 3.5(d) at $x \sim 600\,c/\omega_{pi}$). So, in high-obliquity shocks the upstream ponderomotive potential turns out to be more important than the cross-shock potential, yet it is not included in most models of relativistic magnetized shocks (e.g., Gedalin et al. 2008).

The increase in electron bulk energy toward the shock is accompanied by substantial electron heating, i.e., broadening of the electron distribution. In the upstream, the electron spectrum shows two components (see blue line in Fig. 3.5(h) and (i)). The low-energy "thermal" part, that contains most of the particles, broadens and shifts to higher energies as the flow propagates toward the shock (the shift corresponds to the secular increase in electron energy of Fig. 3.5(d)). The high-energy component, which gets more populated in the vicinity of the shock, extends a factor of ~ 10 higher in energy than the thermal peak. It is seen in the electron phase space of Fig. 3.5(f) as a diffuse beam moving toward the shock, with much higher x-momentum than the bulk flow. In the downstream spectrum, this extra component survives as a tail at high energies (blue line in Fig. 3.5(g)).

As we discuss in §4.2, this high-energy component should be interpreted as a separate hotter (but still thermal) population, rather than as a nonthermal tail of accelerated particles. It is populated by electrons that, on their way to the shock, happen to receive a substantial kick toward the upstream by the wakefield oscillations mentioned above. At first, the energy of these electrons will decrease, because the kick is opposite to their x-momentum. As explained in 4.2, the electron energy will then oscillate with the gyro-period, periodically reaching a maximum that, for strong kicks, may be much larger than the characteristic energy of electrons in the bulk (see Fig. 3.5(f)). In support of this interpretation, we see in the phase space of Fig. 3.5(f) that injection of electrons into the high-energy component starts at $x \sim 1050\, c/\omega_{\rm pi}$, right behind the location where the average electron energy is driven to a minimum (Fig. 3.5(d) at $x \sim 1100\, c/\omega_{\rm pi}$) by the wakefield oscillations in Fig. 3.5(c). Closer to the shock, where the wakefield oscillations are stronger (Fig. 3.5(c) at $600\, c/\omega_{\rm pi} \lesssim x \sim 900\, c/\omega_{\rm pi}$), more electrons are kicked back upstream and injected into the high-energy spectral component. This explains why the electron high-energy tail gets more populated toward the shock (compare Fig. 3.5(h) and (i)). In the following, since the high-energy component in the electron spectrum is ultimately powered by the dissipation of wakefield oscillations, we shall refer to this process as "wakefield heating".

In contrast, ions are not significantly affected by the wakefield oscillations, due to their larger mass. The upstream ion spectrum is then consistent with a single component (red line in Fig. 3.5(h) and (i)), which results in a purely Maxwellian distribution in the downstream region (red line in Fig. 3.5(g)).

We remark that the main features discussed here for $\theta = 75°$, and in particular the emergence of a separate component of hot electrons in the upstream, are common to all superluminal configurations, up to $\theta = 90°$.

4. Particle Acceleration

In this section, we discuss particle acceleration in subluminal and superluminal shocks. We fix the magnetization ($\sigma = 0.1$) and bulk Lorentz factor ($\gamma_0 = 15$) of the upstream flow, and we investigate one representative subluminal ($\theta = 15°$, in §4.1) and one superluminal ($\theta = 75°$, in §4.2) magnetic obliquity. For each case, we describe the time evolution of the downstream energy spectrum and the mechanism responsible for particle energization. All particle spectra presented in

this section, as well as in §5, are extracted from a $100\,c/\omega_{\rm pi}$-wide slab centered at a fixed distance of $150\,c/\omega_{\rm pi}$ behind the shock.

4.1. Subluminal Shocks: $0^\circ \leq \theta < \theta_{\rm crit}$

Time Evolution

Fig. 3.7 shows the time evolution of the ion (upper panel) and electron (lower panel) energy spectrum for $\theta = 15^\circ$. As shown in panel (c), the spectrum of both ions (red) and electrons (blue) clearly deviates from a Maxwellian (dashed), with a substantial population of nonthermal particles arranged in a power-law tail (dot-dashed).

The high-energy tail in the ion spectrum (upper panel) grows with time as more and more particles are shock-accelerated. At late times (red curve for $\omega_{\rm pi}t = 5062$), the slope approaches -2, which corresponds to equal energy contributions by each decade in Lorentz factor. It follows that, as the upper cutoff of the spectrum increases linearly in time (red line in panel (b)), the energy content of the ion nonthermal tail steadily grows, reaching $\sim 30\%$ of the total ion energy at $\omega_{\rm pi}t = 5062$. By number, the tail is still dominated by the lower energies, and its fractional contribution to the ion census saturates at $\sim 5\%$. At late times (e.g., $\omega_{\rm pi}t = 5062$), the tail does not resemble a simple power-law with constant slope, but it seems to be somewhat concave, which may be due to the modification of the shock by accelerated ions (see §3.2), as discussed for non-relativistic shocks in SNRs by Amato & Blasi (2005, 2006).

As more ion energy is stored in the nonthermal tail, the peak of the ion thermal distribution shifts to lower energies. A similar trend is observed for electrons (lower panel in Fig. 3.7), with the thermal bump moving to lower Lorentz factors as the acceleration efficiency increases.[9] However, the tail in the electron spectrum contains less particles ($\sim 2\%$) and energy ($\sim 10\%$) than the ion tail, it is much steeper (with slope -3.5 ± 0.1 at $\omega_{\rm pi}t = 5062$), and its upper energy cutoff does not grow very fast in time (blue line in panel (b), multiplied by 5 for clarity). Most importantly, as discussed in §3.2, no electrons are seen to propagate back upstream from the shock, as opposed to the large population of returning ions (compare the phase spaces in panels (e) and (f) of Fig. 3.2).

[9]We note that, whereas the partition of energy between the thermal peak and the nonthermal tail changes with time, the average electron energy (that includes both components) stays relatively constant, at $\sim 30\%$ of the initial ion energy (blue line in Fig. 3.7(a)).

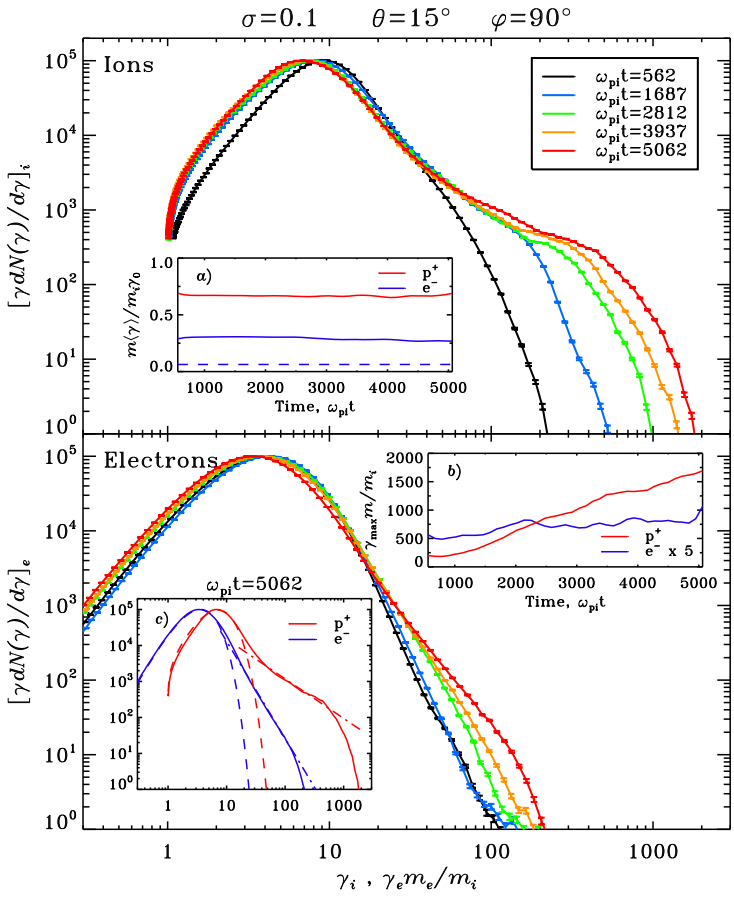


Fig. 3.7.— Time evolution of the downstream energy spectrum in a $\theta = 15°$ subluminal shock, for ions (upper panel) and electrons (lower panel): $\omega_{\rm p}t = 562$ (black), $\omega_{\rm p}t = 1687$ (blue), $\omega_{\rm p}t = 2812$ (green), $\omega_{\rm p}t = 3937$ (yellow), and $\omega_{\rm p}t = 5062$ (red). Subpanels: (a) time evolution of the downstream *mean* particle energy (red for ions, blue for electrons), in units of the bulk energy of injected ions, with the horizontal dashed blue line showing the value expected for electrons in the absence of any ion-to-electron energy transfer ($= m_e/m_i \simeq 0.06$); (b) time evolution of the downstream *maximum* particle energy (red for ions; blue for electrons, multiplied by 5 for clarity); (c) fit to the electron (blue) and ion (red) spectrum at $\omega_{\rm p}t = 5062$, with a low-energy Maxwellian (dashed) plus a high-energy power law (dot-dashed). In panel (c), axes are the same as in the main plot.

This suggests that there should be a substantial difference between the acceleration process of ions and electrons, as we now describe.

Energization Mechanism

The orbit of a representative high-energy ion from the simulation of a subluminal shock ($\theta = 15^\circ$) is plotted in Fig. 3.8. For $\omega_{\rm pi}t \lesssim 1200$, the particle is approaching the shock from upstream with bulk x-momentum $\sim -\gamma_0\beta_0 \simeq -15$.[10] From $\omega_{\rm pi}t \sim 1200$ to $\omega_{\rm pi}t \sim 2100$, it stays at the shock gyrating around the mean magnetic field (see the ion x-location relative to the shock in panel (b)), and its Lorentz factor increases up to $\gamma_i \sim 600$ (panel (a)). Energy gain occurs primarily when the ion is in the upstream region, whereas its Lorentz factor does not significantly change while downstream (e.g., from $\omega_{\rm pi}t \sim 1700$ to $\omega_{\rm pi}t \sim 1850$).[11] As the particle energy grows, its Larmor radius also proportionally increases (see the particle trajectory in panel (d)). The selected ion is finally transmitted downstream at $\omega_{\rm pi}t \sim 2100$, where it will populate the nonthermal tail seen in the ion spectrum of Fig. 3.7. In contrast, the shock-accelerated ions that are eventually reflected back upstream will contribute to the beam of returning ions seen in Fig. 3.2(e), which are responsible for triggering the growth of Bell's instability (see §3.2). In turn, the waves generated by the instability mediate the process of ion acceleration, as we now describe.

In panel (a), different colors show which component of the electric field governs the energy change. We see that energy gain is primarily associated with $E_{\rm y}$ (green) and $E_{\rm z}$ (red), which are indeed the two components associated with Bell's circularly-polarized waves (see Fig. 3.3(g)). From the 1D profiles of $E_{\rm z}$ time-stacked in panel (b), we see that the wavelength of Bell's modes is typically smaller than the ion Larmor radius (compare with blue line in panel (b)). It follows that the interaction between the selected ion and the waves will be *non-resonant*, i.e., during the half Larmor period when the ion is in the upstream, it will encounter multiple wave fronts, resulting in repeated accelerations and decelerations on a sub-Larmor scale. This is clearly shown in the energy

[10] We point out that, for subluminal configurations, by "shock" we actually mean the subshock, located at $x \sim 730\, c/\omega_{\rm pi}$ in Fig. 3.3. Strictly speaking, "upstream" and "downstream" are then defined with respect to the subshock.

[11] This is a consequence of our choice for the simulation frame, which coincides with the downstream plasma frame. It follows that, in the downstream region, no motional electric field is present, and the particle energy does not appreciably change. In the shock frame, energy change would be seen on both sides of the shock.

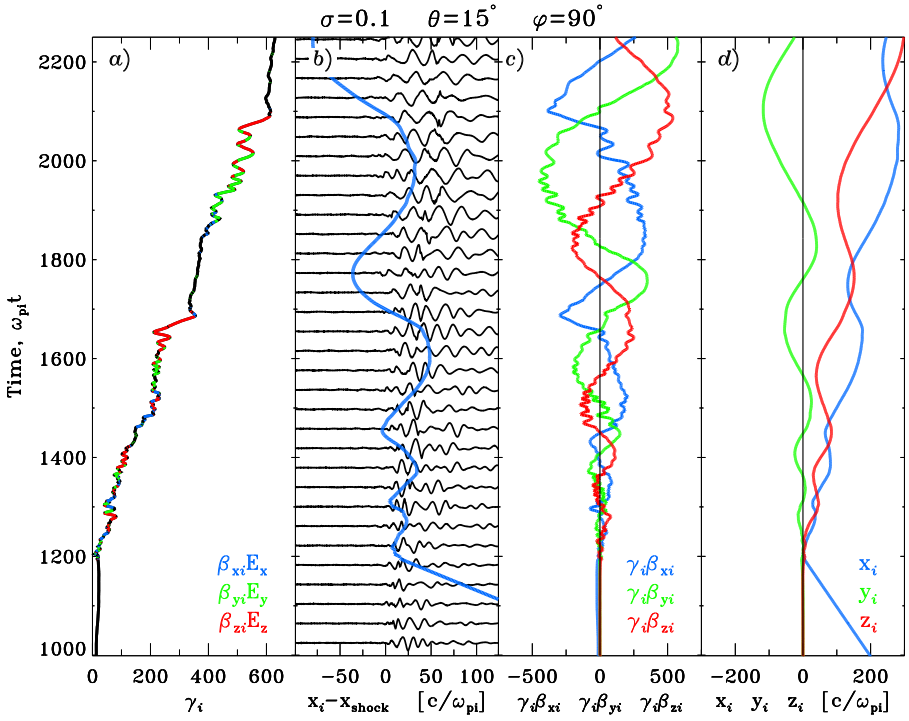


Fig. 3.8.— Trajectory of a representative high-energy ion extracted from the simulation of a $\theta = 15°$ subluminal shock. Panel (a): time evolution of the particle energy; different colors show which component of the electric field governs the energy change (blue for $E_{\rm x}$, green for $E_{\rm y}$, red for $E_{\rm z}$), for the portions of the trajectory with the largest rate of acceleration or deceleration. Panel (b): x-location of the particle relative to the shock, superimposed over y-averaged profiles of $E_{\rm z}$ stacked in time. The electric field is still measured in the wall frame, but shifted along x so that the shock appears stationary. Panel (c): time evolution of the particle 4-velocity, with the same color coding as in panel (a). Panel (d): time evolution of the particle position (x-location is relative to the first encounter with the shock), with the same color coding as in panel (a). We remind that in this case ($\varphi = 90°$) the upstream background electric field $\mathbf{E}_0$ is along $-\hat{\mathbf{y}}$.

evolution of panel (a) and in the 4-velocities of panel (c), where Bell's waves cause short-scale wiggles superimposed over the (otherwise smooth) Larmor gyration.

The overall energy gain results from favorable encounters (i.e., with the ion velocity locally aligned with the wave electric field) being more frequent than unfavorable ones. In particular, significant energization may occur when the ion is crossing the shock from upstream to downstream, as seen at $\omega_{\rm pi}t \sim 1675$ and $\omega_{\rm pi}t \sim 2075$ (compare panels (a) and (b)). In this case, the ion can remain in phase with the waves for a larger fraction of its orbit, since both the ion and the waves are moving toward the shock. Efficient acceleration results if the wave electric field is aligned with the ion velocity (mostly oriented along $+\hat{\mathbf{z}}$ in this phase, for a positively-charged particle gyrating around the mean magnetic field). Overall, the character of the energization process points towards a Fermi-like acceleration mechanism (or DSA), driven by non-resonant interactions of the ion with the self-generated Bell's waves.[12] Strictly speaking, the ion in Fig. 3.8 only experiences the early stages of DSA, since it never penetrates farther than one Larmor radius into the upstream, before being advected downstream. Yet, a fraction of the high-energy ions extracted from our simulation do diffuse far upstream before being scattered back towards the shock, as expected for DSA.

During the acceleration process shown in Fig. 3.8, the ion gyrocenter moves predominantly along the mean magnetic field (in the xz plane, for $\varphi = 90^\circ$), as shown by the time evolution of x_i (blue) and z_i (red) in panel (d). In addition, a clear drift can be seen in y_i (green) towards negative values, i.e., in the direction of the background $\mathbf{E}_0$ (see SS09 for an analogous effect for positrons in pair shocks).

As explained in §1, particles reflected by the shock can be accelerated by the background motional electric field $\mathbf{E}_0$ while drifting along the shock surface. Here, we do not differentiate between SDA and SSA, but we generically term this energization mechanism as "direct acceleration" (as opposed to "diffusive acceleration") or "$\mathbf{E}_0$-driven acceleration". We remind that $\mathbf{E}_0 = -\boldsymbol{\beta}_0 \times \mathbf{B}_0$ is the motional field at injection, and it does not include any contribution from self-generated turbulence.

In Fig. 3.9 we quantify the relative importance of this $\mathbf{E}_0$-driven acceleration mechanism with respect to the DSA process discussed above. For each of the high-energy ions extracted from the

[12] We remark that the detailed nature of the waves is not instrumental in the process of ion acceleration, since the Fermi mechanism only requires a sufficient power in turbulent fluctuations, regardless of their origin. Yet, the details of wave-particle interactions (e.g., resonant versus non-resonant) may be important for the rate of acceleration.

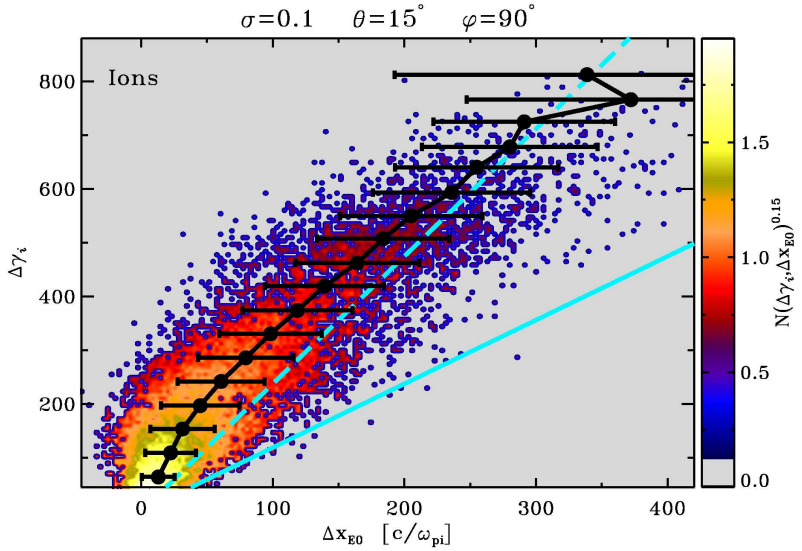


Fig. 3.9.— Relative contribution of different acceleration mechanisms, for high-energy ions (such that $\gamma_i/\gamma_0 > 4$ at $\omega_p t = 2250$) in a $\theta = 15°$ subluminal shock. Here, Δx_{E0} is the ion displacement along $\mathbf{E}_0$, and $\Delta\gamma_i$ is the overall change in Lorentz factor. The 2D histogram shows the ion phase space density $N(\Delta\gamma_i, \Delta x_{E0})$. For each bin in $\Delta\gamma_i$, the value of Δx_{E0} averaged over the ion distribution is shown as a black dot with 1σ error bars. The contribution $\Delta\gamma_{E0}$ expected from direct acceleration by $\mathbf{E}_0$ is the solid cyan line. The dashed cyan line is the predicted $\Delta\gamma_{E0}$ if the field $\mathbf{E}_0$ were twice as large.

simulation, we measure the total energy gain $\Delta\gamma_i$ and the drift $\Delta x_{\rm E0}$ along the upstream field $\mathbf{E}_0$. For each bin in $\Delta\gamma_i$, the average value of $\Delta x_{\rm E0}$ is shown as a black dot with error bars. From $\Delta x_{\rm E0}$, we can compute the expected energy gain $\Delta\gamma_{\rm E0} = (q/mc^2)\, E_0\, \Delta x_{\rm E0}$ due to direct acceleration by the $\mathbf{E}_0$ field.[13] We see that the resulting $\Delta\gamma_{\rm E0}$ (solid cyan line) is much lower than the actual energy gain (black dots), which suggests that DSA in Bell's waves plays a major role in the energization of ions. In fact, $\mathbf{E}_0$-driven acceleration contributes at most $\sim 50\%$ of the total energy gain (see the dashed cyan line, which is the predicted $\Delta\gamma_{\rm E0}$ if the field $\mathbf{E}_0$ were twice as large).

A separate argument needs to be made for electrons, whose acceleration path cannot be exactly the same as for ions, given the different spectrum (Fig. 3.7), and above all the absence of returning electrons (compare phase spaces for ions and electrons in panels (j) and (k) of Fig. 3.3, respectively). In Fig. 3.10, we plot the orbit of a representative high-energy electron in a $\theta = 15^\circ$ subluminal shock. The upstream trajectory of the electron gets perturbed (see panels (b) and (d) at $\omega_{\rm pi}t \sim 2000$) even before entering the shock, by the density spike formed $\sim 50\, c/\omega_{\rm pi}$ ahead of the shock by the pressure of returning ions. The electron energy (panel (a)) does not appreciably change before reaching the shock at $\omega_{\rm pi}t \sim 2050$. Hereafter, the electron bounces several times across the shock transition layer (from $\omega_{\rm pi}t \sim 2050$ to $\omega_{\rm pi}t \sim 2175$), gaining energy by scattering off the Bell's waves (whose wavelength imprints the short-scale modulation seen in the 4-velocities of panel (c) for $2050 \lesssim \omega_{\rm pi}t \lesssim 2125$), similarly to what happens for ions.

The final energy of the selected electron ($\gamma_e m_e/m_i \sim 120$) is much smaller than the downstream energy of the representative ion followed in Fig. 3.8 ($\gamma_i \sim 600$). This reflects the lower bulk energy (smaller by a factor of ~ 5, see Fig. 3.2(d)) with which electrons enter the shock, with respect to ions. Due to their smaller Larmor radius, they are confined closer to the shock, and they do not penetrate very far upstream, where the electric field of Bell's waves would be stronger (or, in the standard Fermi picture, they do not sample the full velocity difference of the converging flows). Also, since they are more strongly tied to the magnetic field, they leave the acceleration region earlier by advection into the downstream (the electron in Fig. 3.10 stays at the shock for $\omega_{\rm pi}\Delta t \sim 125$, as opposed to $\omega_{\rm pi}\Delta t \sim 900$ for the ion in Fig. 3.8), which justifies the scarcity of returning electrons.[14] In summary, electrons display a lower energy gain per acceleration cycle

[13] This formula assumes that most of the energy gain occurs in the upstream medium (where the background field is $\mathbf{E}_0$), which is correct since in the simulation frame no motional electric fields should persist in the downstream region.

[14] Electrons are also attracted into the downstream by the cross-shock electric field $E_{\rm x} > 0$, created by electron-ion charge separation at the shock. We clearly see this effect for some electrons, but not for the particle in Fig. 3.10.

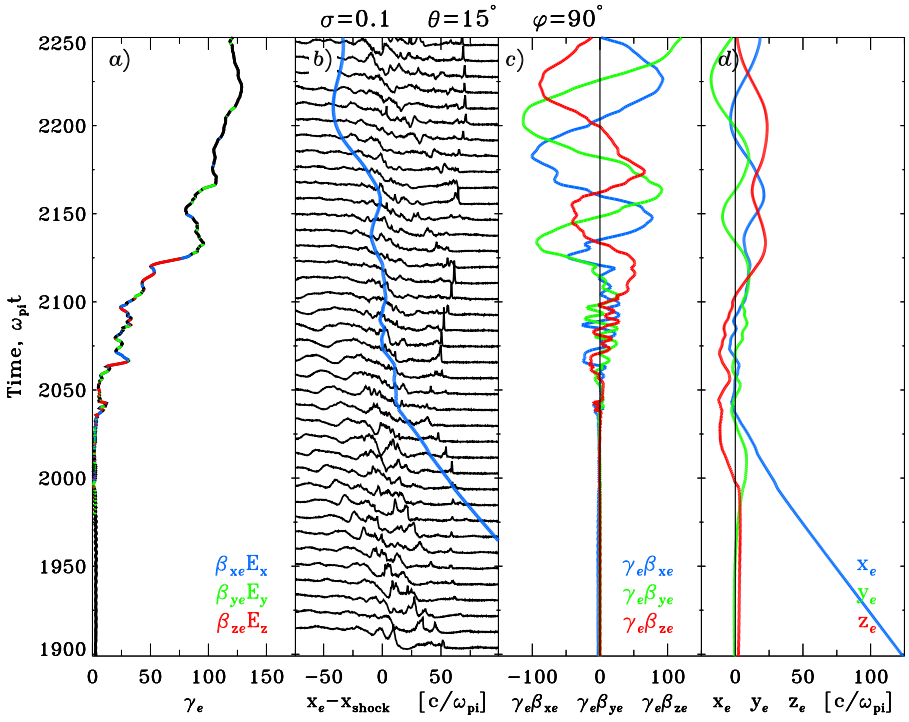


Fig. 3.10.— Trajectory of a representative high-energy electron extracted from the simulation of a $\theta = 15°$ subluminal shock. Panels as in Fig. 3.8, but here the fluid quantity plotted in panel (b) is the particle number density. Also, the electron Lorentz factor and dimensionless momentum have been divided by the mass ratio, although not indicated in the plot labels.

(relative to ions), and a higher escape probability from the acceleration region. The combination of these two effects explains why in Fig. 3.7 electrons present a power-law tail much steeper than ions.

As a side note, we point out that the density profiles stacked in panel (b) of Fig. 3.10 clearly show the process of shock reformation discussed in §3.2. We see a density spike emerging ahead of the shock and then being advected into the shock on a typical timescale of $\sim 100\,\omega_{\rm pi}^{-1}$. Then, a new spike forms, at a slightly larger distance (in fact, the shock transition region widens with time), and the whole process starts again.

4.2. Superluminal Shocks: $\theta_{\rm crit} < \theta \leq 90^\circ$

Time Evolution

Fig. 3.11 follows the time evolution of the ion (upper panel) and electron (lower panel) downstream spectrum in a $\theta = 75^\circ$ superluminal shock. At all times, the ion spectrum is consistent with a purely Maxwellian distribution, i.e., nonthermal acceleration of ions is suppressed in superluminal shocks. The location of the thermal peak varies depending on the amount of ion energy transferred to electrons before the shock, via the mechanism described in §3.3.

The electron spectrum at early times ($\omega_{\rm pi}t = 562$, black curve) is also compatible with a single Maxwellian distribution, but at later times (blue for $\omega_{\rm pi}t = 2250$, green for $\omega_{\rm pi}t = 3937$), a separate component emerges at high energies. We interpret this component as the result of extra *heating* of some electrons by the wakefield oscillations described in §3.3, rather than as nonthermal *acceleration*, for two main reasons. First, in any model of electron acceleration, we would expect a steady increase in the maximum electron Lorentz factor, at odds with what we observe in the blue line of panel (b). Also, we find that the downstream electron spectrum can be well fitted with two Maxwellians (panel (c)), with the hotter Maxwellian (dot-dashed line, with a temperature ~ 3 times higher than the low-energy Maxwellian) accounting for the high-energy component. A fit that employs a low-energy Maxwellian plus a power law with exponential cutoff is less satisfactory.

The appearance of the high-energy component is accompanied by an increase in the average electron energy, at the expense of ions (panel (a) for $\omega_{\rm pi}t \lesssim 3000$; red for ions, blue for electrons). For $\omega_{\rm pi}t \gtrsim 3000$, the mean electron energy decreases (panel (a)), and the high-energy component recedes ($\omega_{\rm pi}t = 5625$, yellow curve), until it finally disappears ($\omega_{\rm pi}t = 7312$, red curve). The similarity between the time evolution of the average electron energy in panel (a) and the maximum

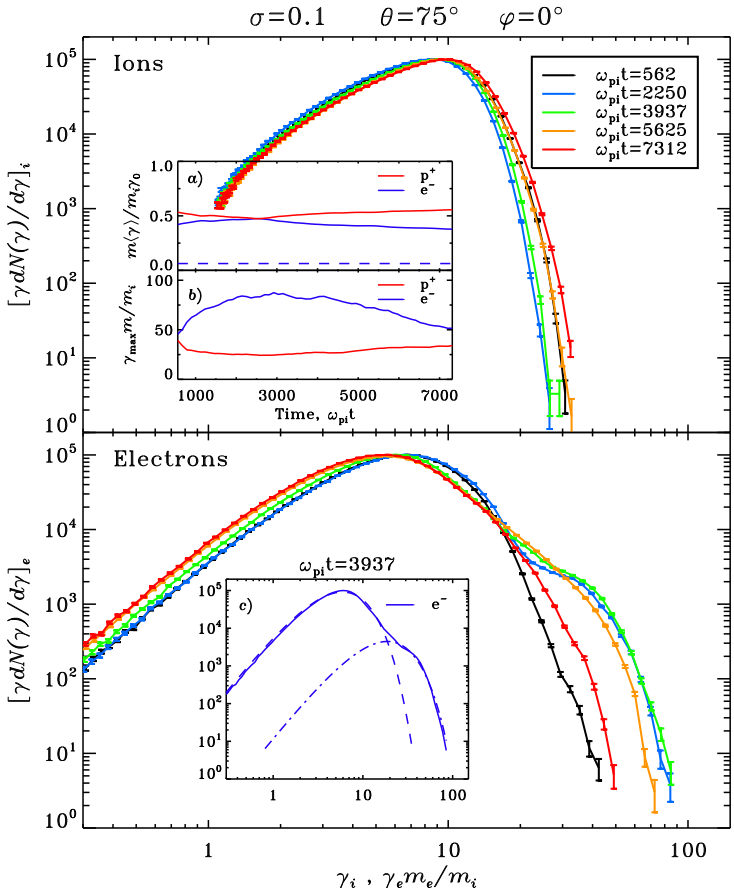


Fig. 3.11.— Time evolution of the downstream energy spectrum in a $\theta = 75°$ superluminal shock, for ions (upper panel) and electrons (lower panel): $\omega_{\rm p}t = 562$ (black), $\omega_{\rm p}t = 2250$ (blue), $\omega_{\rm p}t = 3937$ (green), $\omega_{\rm p}t = 5625$ (yellow), and $\omega_{\rm p}t = 7312$ (red). Subpanels as in the caption of Fig. 3.7, but here the electron spectrum in panel (c) is fitted with two Maxwellians.

electron Lorentz factor in panel (b) (a good proxy for the importance of the high-energy component) suggests that a common mechanism should regulate both the low-energy "thermal" part and the high-energy tail, as we now describe.

At the beginning, electrons enter the shock with x-momentum $= -\gamma_0\beta_0 \simeq -15$ and a small thermal spread. The synchrotron maser instability at the shock front produces powerful electromagnetic waves (Poynting flux in excess of $c\,E_0^2/4\pi$) that illuminate the upstream flow, causing efficient transfer of bulk momentum from ions to electrons (see §3.3). As electrons with larger bulk energy enter the shock, the synchrotron maser instability is more efficient in producing strong precursor waves, and the transfer of energy from ions to electrons in the upstream becomes even more significant. This self-reinforcing cycle repeats until electrons reach energy equipartition with ions (panel (a), at $\omega_{\rm pi}t \sim 2500$). As the strength of the electromagnetic precursor increases, wakefield oscillations of larger amplitude are generated in the upstream (Lyubarsky 2006). As anticipated in §3.3, their dissipation then results in a more pronounced high-energy component in the electron spectrum. This explains why the maximum electron Lorentz factor increases for $\omega_{\rm pi}t \lesssim 3000$ (blue line in panel (b)).

Heated by the upstream wakefield oscillations, electrons now enter the shock much hotter than before, i.e., their comoving momentum spread is larger. In such conditions, the growth of the synchrotron maser instability is quenched (Hoshino & Arons 1991), and the electromagnetic precursor becomes less powerful. This reduces the transfer of bulk energy from ions to electrons in the upstream, causing the decrease in average electron energy observed at $\omega_{\rm pi}t \gtrsim 3000$ in panel (a).[15] Also, the upstream wakefield oscillations now become weaker, so that fewer electrons are injected into the high-energy spectral component, and the maximum electron Lorentz factor decreases (panel (b), at $\omega_{\rm pi}t \gtrsim 3000$).

In our simulation, we find that the electron spectrum approaches a steady state for $\omega_{\rm pi}t \gtrsim 7000$ (we followed the shock evolution up to $\omega_{\rm pi}t \sim 10,000$). However, since the incoming electrons are now colder, we cannot exclude that at later times the whole process may repeat, with efficient transfer of energy to electrons accompanied by the emergence of a new high-energy spectral component. Instead, we do not expect the ion spectrum to deviate from a Maxwellian distribution,

[15] This effect has been verified with controlled numerical experiments, in which we initialized the incoming particles with a larger thermal spread than the usual $\Delta\gamma = 10^{-4}$. For $\Delta\gamma \gtrsim 10^{-1}$, both the power of the electromagnetic precursor and the upstream ion-to-electron energy transfer are gradually suppressed, as well as the high-energy component in the electron spectrum.

since ions are not significantly affected by the upstream wakefield oscillations, due to their larger inertia.

Finally, we comment on the difference between our 2D simulations and the 1D results presented by Hoshino (2008), that reported nonthermal *acceleration* of electrons in perpendicular magnetized electron-ion shocks. First, we find that 1D simulations over-estimate the power of the electromagnetic precursor, since they cannot resolve variations along the shock surface (Fig. 3.6(c) shows that the strength of the electromagnetic precursor is not uniform in y). As a result, 1D runs tend to over-estimate the efficiency of ion-to-electron energy transfer, and the importance of the high-energy component in the electron spectrum. For $m_i/m_e = 16$, computational domains with more than 64 transverse cells are required to obtain consistent 2D results.[16] Second, Hoshino (2008) followed the shock evolution only up to $\omega_{\rm pi}t \sim 850$, a timespan too short to appreciate the saturation and decrease in the maximum electron Lorentz factor shown in panel (b). As mentioned above, this piece of evidence is essential for our interpretation of the high-energy spectral component as a result of wakefield *heating*, rather than *acceleration*.

Energization Mechanism

The trajectory of a representative high-energy electron in a superluminal shock ($\theta = 75^\circ$) is presented in Fig. 3.12. Panel (a) shows that all of the energy gain, up to $\gamma_e m_e/m_i \sim 70$, occurs in the upstream region *before* the electron encounters the shock at $\omega_{\rm pi}t \sim 1050$ (see the electron x-location relative to the shock in panel (b)). After that, the electron is advected downstream, with no appreciable change in energy.

For $\omega_{\rm pi}t \lesssim 850$, the electron is first boosted and then deboosted (see energy, panel (a); or x-momentum, blue line in panel (c)), in response to the wakefield wiggle seen at $\sim 350\, c/\omega_{\rm pi}$ ahead of the shock ($E_{\rm x}$ in panel (b)). So far, the trajectory of the electron in Fig. 3.12 presents no qualitative difference from the bulk of incoming electrons. But at $\omega_{\rm pi}t \sim 850$, presumably due to a strong upstream-oriented kick imparted by the wakefields, the energy of the selected electron becomes so small that it can easily decouple from the bulk flow. As anticipated in §3.3, this is the

[16] In fact, to follow the spectral evolution shown in Fig. 3.11, we have used a simulation box with 64 transverse cells, which is small enough to suppress the growth of the numerical instability mentioned in §2, yet large enough to preserve the relevant 2D properties of the shock.

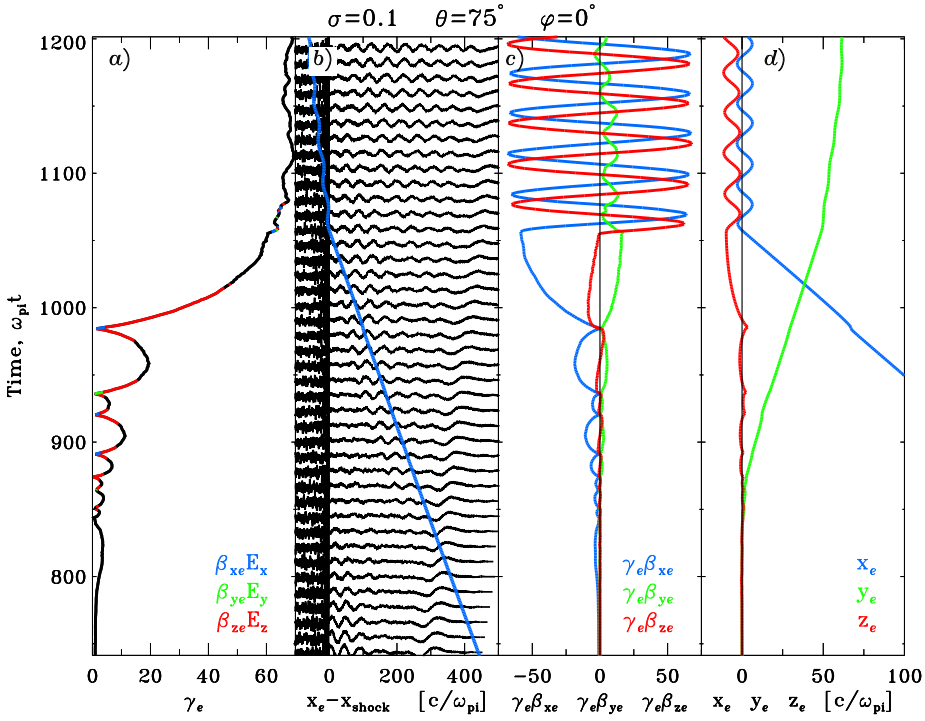


Fig. 3.12.— Trajectory of a representative high-energy electron extracted from the simulation of a $\theta = 75°$ superluminal shock. Panels as in Fig. 3.8, but here the fluid quantity plotted in panel (b) is $E_{\rm x}$. Also, the electron Lorentz factor and dimensionless momentum have been divided by the mass ratio, although not indicated in the plot labels. We remind that in this case ($\varphi = 0°$), the upstream background electric field $\mathbf{E}_0$ is along $+\hat{\mathbf{z}}$.

"injection" step for all the electrons that will end up in the high-energy spectral component seen in Fig. 3.11.

After decoupling from the bulk flow, the electron gyrocenter starts moving with drift velocity $\mathbf{E}_0 \times \mathbf{B}_0/B_0^2$, which for the field configuration employed here ($\varphi = 0°$) corresponds to a straight path in the xy plane (see y_e in panel (d) for $\omega_{\rm pi}t \gtrsim 850$). At the same time, the electron gets pulled by the upstream motional electric field $\mathbf{E}_0$, and it performs a series of Larmor cycles (see energy in panel (a) and 4-velocity in panel (c), for $\omega_{\rm pi}t \gtrsim 850$). When the electron velocity is aligned with $-\mathbf{E}_0$ the particle energy grows, but it decreases by the same amount when the two are oppositely directed (see red segments in panel (a)). Overall, no *net* energy gain is expected from $\mathbf{E}_0$ during a *complete* Larmor cycle. In fact, as seen from the upstream frame, the electron motion is just a Larmor gyration around the background magnetic field, with energy determined by the magnitude of the kick imparted by the wakefields.

The final electron energy is determined by the combination of two factors: (*i*) the amplitude of the energy oscillations seen in panel (a); and (*ii*) the gyro-phase with which the particle enters the shock. Regarding (*i*), secular variations (i.e., averaged over gyro-phase) in the electron energy are determined by the kicks imparted by the wakefields. In particular, the electron will be most sensitive to such kicks when its energy is smaller. At this point, as seen from the upstream frame, the electron is moving away from the shock. If $E_{\rm x} < 0$, the electron will receive a kick in the same direction of its momentum, and its upstream-frame energy will increase. As seen from the downstream frame, this corresponds to a larger amplitude of the resulting energy oscillation (e.g., at $x \sim 960\, c/\omega_{\rm pi}$ in panel (a)). The opposite happens if the phase of wakefield oscillations is such that $E_{\rm x} > 0$, so that the following energy oscillation will have a smaller amplitude (e.g., at $x \sim 900\, c/\omega_{\rm pi}$ in panel (a)). As a confirmation of the importance of wakefields, we see that the amplitude of the energy oscillations in panel (a) increases dramatically for $\omega_{\rm pi}t \gtrsim 925$, when the selected electron enters a region with strong wakefields ($E_{\rm x}$ in panel (b), following the particle track for $\omega_{\rm pi}t \gtrsim 925$).

On the other hand, depending on the gyro-phase with which the particle encounters the shock (point (*ii*) above), the electron will end up in a different energy bin of the downstream high-energy Maxwellian (dot-dashed line in Fig. 3.11(c)).

Finally, we remark that, despite the significant wakefield heating experienced in the upstream, which could possibly serve as pre-injection for a Fermi-like process, the electron in Fig. 3.12 does not bounce back from the shock into the upstream, due to insufficient turbulence on the

downstream side. In fact, in our simulations of superluminal shocks, we do not find any evidence for returning electrons (see the electron phase space in Fig. 3.5(f)), confirming that Fermi acceleration is suppressed for superluminal configurations.

5. Dependence on Upstream Conditions

Armed with a better understanding of the energization mechanisms that operate in relativistic magnetized electron-ion shocks, we now explore how the efficiency of ion-to-electron energy transfer and of particle acceleration depends on the conditions of the upstream flow. First, keeping the bulk Lorentz factor ($\gamma_0 = 15$) and magnetization ($\sigma = 0.1$) fixed, as we did so far, we investigate the transition between subluminal and superluminal shocks by exploring the full range of magnetic obliquities, from $\theta = 0^\circ$ to $\theta = 90^\circ$ (in §5.1). Then, for the two representative obliquities discussed above (one subluminal, $\theta = 15^\circ$; one superluminal, $\theta = 75^\circ$), we extend our analysis to different values of the magnetization (in §5.2) and bulk Lorentz factor (in §5.3).

For different upstream parameters, we compare the particle energy spectrum computed at a distance of $150\,c/\omega_{\rm pi}$ behind the shock (averaged over a $100\,c/\omega_{\rm pi}$-wide slab) at the fixed time $\omega_{\rm pi}t = 2250$. Yet, as shown in Figs. 3.7 and 3.11, the downstream particle spectrum is not in steady state. In subluminal shocks (Fig. 3.7), the nonthermal tail grows with time, so that the acceleration efficiency we compute at $\omega_{\rm pi}t = 2250$ should be taken as a lower limit. In superluminal shocks (Fig. 3.11), the electron high-energy component is most pronounced between $\omega_{\rm pi}t \sim 2000$ and $\omega_{\rm pi}t \sim 3000$, and then it disappears, although a cyclical behavior is likely, yet undemonstrated (see §4.2). At $\omega_{\rm pi}t = 2250$, we are capturing this high-energy component when it is close to its maximum, and more likely to result in an observational signature.

In the following, we fit the downstream particle spectrum with a three-dimensional Maxwellian plus a power-law tail with an exponential cutoff. This is not strictly appropriate for superluminal magnetized shocks, where the high-energy component of the electron spectrum is better fitted with a second (hotter) Maxwellian, as discussed in §4.2. Although this component actually results from *heating* (rather than *acceleration*), in this section we treat it as a nonthermal tail, for easier comparison with observations.[17]

[17]Observationally, a suprathermal component in the radiation spectrum will most likely be ascribed to nonthermal acceleration of particles, and fitted as a power-law.

The form chosen for our fitting function allows to split the particle population into a *thermal* and a *nonthermal* component, and to follow how the relative importance of the two sub-populations changes with the upstream parameters. In Figs. 3.13-3.17, panel (a) presents the mean downstream particle energy (red for ions, blue for electrons) in units of the bulk energy of injected ions, including *both* thermal *and* nonthermal particles. Panels (b)-(d) provide a better characterization of the nonthermal component: panel (b) shows the slope p of the best-fitting power law, such that $dN_{\rm PL}(\gamma)/d\gamma \propto \gamma^{-p}$; panels (c) and (d) present respectively the fraction of particles and energy contained in the nonthermal tail (red for ions and blue for electrons, relative to the total number or energy *of that species*).[18] In subpanels (b)-(d), dotted lines (instead of the usual solid lines) mark the cases in which a nonthermal power-law tail is *not* a physically-motivated model for the high-energy component (i.e., for electrons in superluminal magnetized shocks, as discussed above).

5.1. Dependence on Magnetic Obliquity θ

Fig. 3.13 presents the downstream particle spectrum at $\omega_{\rm pi}t = 2250$ (upper panel for ions, lower panel for electrons), for a sample of magnetic obliquities covering the full range from $\theta = 0^\circ$ to 90°. In this section, we describe how the efficiency of particle acceleration and electron heating depends on magnetic inclination, emphasizing the importance of the critical boundary $\theta_{\rm crit} \simeq 34^\circ$ between subluminal and superluminal configurations.

The ion spectrum in subluminal shocks (blue curve for $\theta = 0^\circ$, green for $\theta = 15^\circ$, red for $\theta = 30^\circ$) shows a prominent nonthermal tail at high energies, together with a thermal bump at low energies. Instead, for superluminal shocks (black for $\theta = 45^\circ$, yellow for $\theta = 75^\circ$, purple for $\theta = 90^\circ$), the ion spectrum does not appreciably deviate from a Maxwellian distribution. As shown in panels (c) and (d), the efficiency of ion acceleration (red lines) abruptly drops between $\theta = 30^\circ$ and $\theta = 35^\circ$, as the magnetic obliquity passes the critical value $\theta_{\rm crit} \simeq 34^\circ$.

For subluminal configurations, the ion nonthermal tail shows a non-monotonic behavior with angle. Ion acceleration is most efficient for $\theta = 15^\circ$, with a fraction $\sim 4\%$ of particles and $\sim 22\%$ of energy contained in the tail at $\omega_{\rm pi}t = 2250$ (red dots for $\theta = 15^\circ$ in panels (c) and (d), respectively). As discussed in §4.1, at $\omega_{\rm pi}t = 5062$ the ion acceleration efficiency reaches as much as $\sim 5\%$ by

[18] For instance, to obtain the fraction of energy in nonthermal *electrons* relative to the bulk energy of injected *ions* (which is often required in phenomenological models of GRBs and AGN jets), the value read from the blue line in panel (d) should be multiplied by the value of the blue line in panel (a).

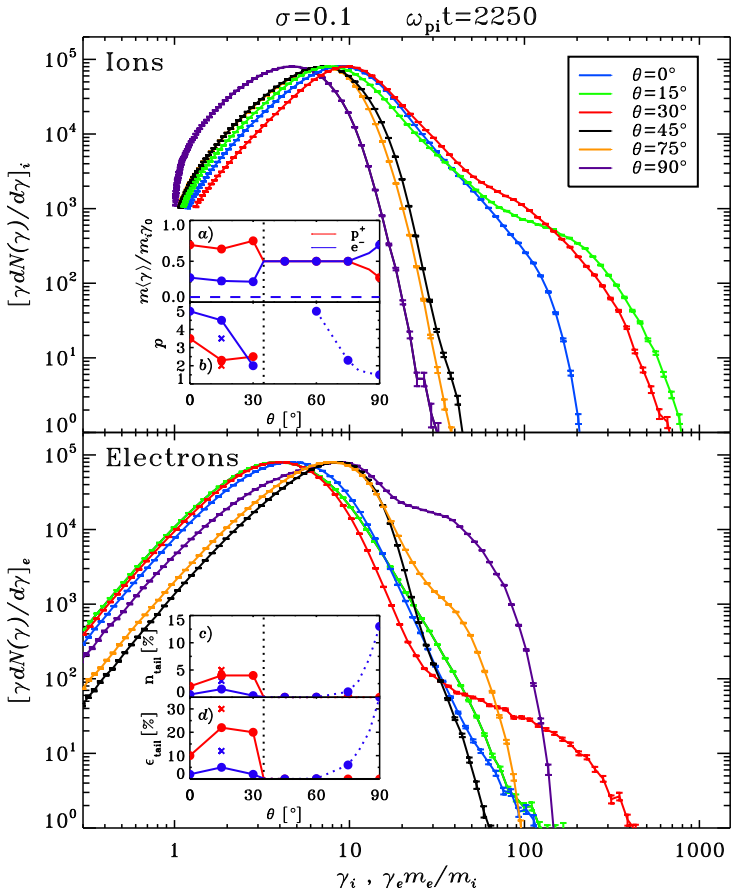


Fig. 3.13.— Downstream particle spectrum at $\omega_{\rm p}t = 2250$ (upper panel for ions, lower panel for electrons) for different magnetic obliquities, from $\theta = 0°$ (parallel shock) to $\theta = 90°$ (perpendicular shock). In subpanels, subluminal and superluminal geometries are separated by a vertical dotted black line, corresponding to the critical angle $\theta_{\rm crit} \simeq 34°$. Subpanels: (a) downstream ion (red) and electron (blue) average energy, in units of the bulk energy of injected ions, with the horizontal dashed blue line showing the value expected for electrons in the absence of any ion-to-electron energy transfer ($= m_e/m_i \simeq 0.06$); (b) power-law slope of the nonthermal tail (red for ions, blue for electrons); (c) fraction of ions (red) or electrons (blue) in the nonthermal tail; (d) fraction of energy in the ion (red) or electron (blue) nonthermal tail, with respect to the total kinetic energy of that species. In subpanels (b)-(d), dotted blue lines mark the cases in which a nonthermal power-law tail is not a physically-motivated model for the electron high-energy component (see text). Crosses for $\theta = 15°$ (red for ions, blue for electrons) correspond to values measured at $\omega_{\rm p}t = 5062$ (instead of $\omega_{\rm p}t = 2250$), to show how the nonthermal tail changes in time.

number and $\sim 30\%$ by energy (red crosses in panels (c) and (d)). By comparison, at $\omega_{\rm pi}t = 2250$ the ion nonthermal tail for $\theta = 0^\circ$ only accounts for $\sim 2\%$ of particles and $\sim 10\%$ of ion energy, whereas the acceleration efficiency for $\theta = 30^\circ$ ($\sim 4\%$ by number, $\sim 20\%$ by energy) is comparable to $\theta = 15^\circ$, but the tail extends to smaller Lorentz factors. Also, the ion tail is flatter for $\theta = 15^\circ$, with a slope $p \sim 2.2$, compared to $p \sim 3.5$ for $\theta = 0^\circ$ and $p \sim 2.5$ for $\theta = 30^\circ$ (red line in panel (b)).

At earlier times (e.g., $\omega_{\rm pi}t = 1125$, not shown here), the trend was different, with larger angles (yet still subluminal) showing flatter ion tails that extended to higher Lorentz factors and contained more particles and energy. In fact, this was the hierarchy that SS09 observed in electron-positron shocks (see Fig. 11 in SS09, same colors as here). In SS09, such trend was attributed to a change in the dominant acceleration mechanism, with DSA prevailing in quasi-parallel shocks and SDA most efficient for $15^\circ \lesssim \theta \lesssim \theta_{\rm crit}$. In electron-ion shocks, as we have shown in §4.1, magnetic turbulence created by the returning ions (in the form of Bell's waves) can dramatically affect the acceleration process, much more than in electron-positron shocks, where the amplitude of self-generated upstream turbulence was much lower (see SS09). The hierarchy seen in the subluminal ion spectra of Fig. 3.13 will then reflect the influence of magnetic inclination on the strength of Bell's waves.

Since the gyrocenters of returning particles tend to move along the upstream field, for higher (yet subluminal) obliquities the returning ions will be confined closer to the shock. This has two opposite consequences. First, since the returning ions are responsible for triggering Bell's instability, at larger obliquities its growth is more likely to be suppressed by advection into the shock. In fact, at $\omega_{\rm pi}t = 2250$, Bell's modes are already fully developed for $\theta = 15^\circ$, whereas for $\theta = 30^\circ$ they are just starting to grow, and the acceleration process is still governed by SDA, as in electron-positron flows (SS09). On the other hand, Bell's instability saturates when the self-generated magnetic energy approaches the energy density of returning ions (Riquelme & Spitkovsky 2009), which is larger for higher obliquities since the returning particles accumulate in a smaller region ahead of the shock, as discussed above. It follows that Bell's waves are stronger for $\theta = 15^\circ$ than for $\theta = 0^\circ$. Overall, the best compromise at $\omega_{\rm pi}t = 2250$ is achieved for $\theta = 15^\circ$, where the waves are both powerful and rapidly-growing (compared to advection).

At later times, the situation may change. The growth of Bell's waves for $\theta = 30^\circ$ will not be limited by advection any longer, and the turbulence will saturate at higher amplitudes than for $\theta = 15^\circ$, giving efficient ion acceleration via DSA. Therefore, the trend in acceleration efficiency observed for subluminal electron-positron shocks by SS09 may eventually be restored for ions

in electron-ion shocks as well. However, here it will be driven by the increase with angle in the self-generated turbulence that mediates DSA, whereas in electron-positron flows it follows the amplitude of the background motional electric field responsible for SDA.

As discussed in §4.1, acceleration of electrons in subluminal shocks is depressed with respect to ions. For nearly-parallel shocks (blue curve for $\theta = 0°$, green for $\theta = 15°$), a minor fraction ($\sim 1\%$) of the incoming electrons are accelerated during their passage through the shock, as explained in §4.1. In the downstream, they populate a steep ($p \sim 5$) power-law tail which only contains $\sim 3\%$ of the total electron energy. The tail may get somewhat flatter with time, thus accounting for more particles and energy (see blue crosses in panels (b)-(d), for $\theta = 15°$), but the acceleration efficiency for electrons remains smaller than for ions by at least a factor of three. Close to the critical obliquity $\theta_{\rm crit} \simeq 34°$ (see $\theta = 30°$, red curve), a more pronounced tail appears in the electron spectrum, populated by particles accelerated at the shock via the SDA mechanism, in analogy to what happens in electron-positron flows (SS09). However, despite extending to high energies (up to $\gamma_e m_e/m_i \sim 400$ at $\omega_{\rm pi}t = 2250$), the electron tail for $\theta = 30°$ only contains $\sim 1\%$ of particles and $\sim 2\%$ of electron energy.

Instead, a prominent high-energy tail is observed in the electron spectrum of superluminal shocks, beyond the thermal bump. As the obliquity angle increases from $\theta = 75°$ (yellow curve) to $\theta = 90°$ (purple), the tail stretches to higher Lorentz factors (maximum $\gamma_e m_e/m_i$ from ~ 95 to ~ 150), and it contributes a larger fraction of electrons (from $\sim 2\%$ to $\sim 13\%$) and electron energy (from $\sim 6\%$ to as much as $\sim 35\%$). If this high-energy component is fitted with a power law (but a fit with a second Maxwellian is actually more physically motivated, see §4.2), we find that it gets flatter for higher obliquities (from $p \sim 2.3$ for $\theta = 75°$ to $p \sim 1.5$ for $\theta = 90°$).

As explained in §4.2, this high-energy component results from extra heating of some electrons by the upstream wakefield oscillations described in §3.3, that are ultimately generated by the radiative push of the precursor wave on the incoming flow. Since the amplitude of wakefield oscillations increases with the strength of the precursor wave (Lyubarsky 2006), which is larger for higher obliquities (roughly $\propto E_0^2 \propto \sin^2\theta$), the electron high-energy component will be more pronounced for quasi-perpendicular shocks, as seen in Fig. 3.13. Being due to *heating*, rather than *acceleration*, we do not expect the electron tail in superluminal shocks to extend with time to higher Lorentz factors. Rather, it may recede to lower energies, and possibly be hidden by the low-energy thermal bump, as shown in §4.2 for $\theta = 75°$.

In superluminal shocks, the electromagnetic precursor wave also causes efficient transfer of energy from ions to electrons ahead of the shock (see §3.3). As a result, in the downstream region the electron and ion average energies are comparable (see Fig. 3.13(a), for $35^\circ \lesssim \theta \lesssim 75^\circ$), or electrons may be even hotter than ions (for $75^\circ \lesssim \theta \lesssim 90^\circ$). In subluminal shocks ($0^\circ \lesssim \theta \lesssim 30^\circ$), electrons are boosted and heated ahead of the shock, as a result of the perturbation induced on the incoming flow by the cloud of returning ions (see §3.2). With respect to superluminal angles, here the downstream electron energy is a smaller fraction ($\sim 20\% - 30\%$) of the initial ion bulk energy, but still much in excess of $m_e/m_i \simeq 6\%$ (for $m_i/m_e = 16$), the value we would expect in the absence of any ion-to-electron energy transfer.

In Appendix 7, we show that the values quoted here for the efficiency of electron heating and of particle acceleration are essentially unchanged for larger (and more realistic) mass ratios.

5.2. Dependence on Magnetization σ

In this section, we explore the dependence of shock thermalization and acceleration upon the magnetization of the upstream flow. We keep the magnetic obliquity fixed and vary the magnetization from $\sigma = 10^{-5}$, a virtually unmagnetized shock, up to $\sigma = 1.0$, where the upstream field is in equipartition with the kinetic energy of injected particles. We discuss separately the cases of subluminal (in §5.2, for $\theta = 15^\circ$) and superluminal (in §5.2, for $\theta = 75^\circ$) shocks.

Subluminal Shocks: $0^\circ \leq \theta < \theta_{\rm crit}$

Fig. 3.14 presents the downstream particle spectrum of a $\theta = 15^\circ$ subluminal shock, for different magnetizations. The ion spectrum (upper panel) shows in all cases a pronounced nonthermal tail, whereas a significant high-energy component appears for electrons (lower panel) only at $\sigma \lesssim 10^{-2}$. This corresponds to a change in the mechanism that mediates the shock.

For $\sigma \lesssim 10^{-2}$ (blue curve for $\sigma = 10^{-2}$, red for $\sigma = 10^{-3}$), the shock is mediated by the filamentation (Weibel) instability, as in strictly unmagnetized flows (Spitkovsky 2005, 2008a). The free energy for the instability comes from the counter-streaming between the incoming plasma and the shock-accelerated particles that propagate ahead of the shock. In turn, magnetic turbulence

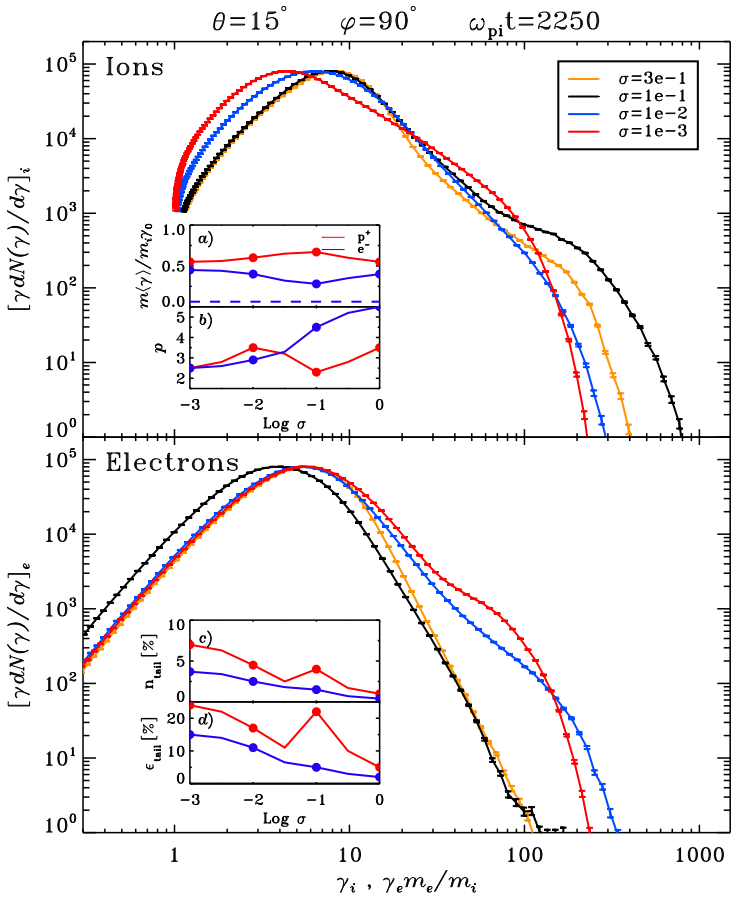


Fig. 3.14.— Downstream particle spectrum at $\omega_{\rm p}t = 2250$ (upper panel for ions, lower panel for electrons) for different magnetizations. We fix the magnetic obliquity $\theta = 15°$ (subluminal shock). Subpanels as in Fig. 3.13.

generated by the instability mediates nonthermal acceleration of particles at the shock front, via a Fermi-like mechanism (Spitkovsky 2008b; Martins et al. 2009).

In fact, the ion spectrum in such shocks shows a prominent nonthermal tail. As the magnetization decreases below $\sigma = 10^{-2}$, the tail becomes flatter (from $p \sim 3.5$ at $\sigma = 10^{-2}$ to $p \sim 2.5$ at $\sigma = 10^{-3}$, red line in Fig. 3.14(b)) and it contains a larger fraction of ions (from $\sim 5\%$ to $\sim 7\%$, Fig. 3.14(c)) and ion energy (from $\sim 17\%$ to $\sim 24\%$, Fig. 3.14(d)). Similarly, the electron spectrum for $\sigma \lesssim 10^{-2}$ presents a significant high-energy tail, comparable to the ion tail in slope (compare red and blue lines in panel (b), for $\sigma \lesssim 10^{-2}$) and maximum Lorentz factor. The acceleration efficiency for electrons is not much smaller than for ions, at most a factor of two (compare red and blue lines in panels (c) and (d), for $\sigma \lesssim 10^{-2}$).[19]

The similarity between the acceleration properties of electrons and ions in weakly magnetized flows ($\sigma \lesssim 10^{-2}$) comes from the fact that the two species enter the shock with comparable bulk energy (compare red and blue lines in panel (a), for $\sigma \lesssim 10^{-2}$), due to efficient transfer of energy from ions to electrons ahead of the shock (Spitkovsky et al., in prep). The Fermi mechanism, which only depends on the particle Larmor radius (i.e., on particle energy, rather than mass), will then proceed in a similar way for ions and electrons.

The situation changes for $\sigma \gtrsim 10^{-2}$. Here, the magnetic field is strong enough so that the upstream plasma cannot filament, and the growth of Weibel instability is suppressed. In these conditions, the current of shock-accelerated ions moving ahead of the shock can trigger Bell's instability, as discussed in §3.2. For $10^{-2} \lesssim \sigma \lesssim 0.3$, the waves generated by Bell's mechanism are responsible for mediating the shock and driving the acceleration of ions. They are most powerful for $\sigma \sim 0.1$ (black curve), which then corresponds to a peak in the ion acceleration efficiency ($\sim 4\%$ by number, $\sim 22\%$ by energy) and a flatter nonthermal tail ($p \sim 2.2$).

In contrast, acceleration of electrons is depressed in highly magnetized flows (black curve for $\sigma = 0.1$), as discussed in §4.1. As the magnetization increases, returning particles are confined closer to the shock (in units of $c/\omega_{\rm pi}$, but same distance in Larmor radii), so their perturbation to the incoming flow extends over a smaller length. It follows that the transfer of energy from ions to electrons in the upstream has less time to proceed, and it is therefore less efficient (see panel (a),

[19]For both electrons and ions, the downstream spectrum at even lower magnetizations (we explored down to $\sigma = 10^{-7}$) is identical to the case $\sigma = 10^{-3}$ shown here, which is therefore a good representation of weakly magnetized (and unmagnetized) shocks.

from $\sigma = 10^{-3}$ to $\sigma = 0.1$). For $\sigma = 0.1$, electrons enter the shock with $\sim 15\%$ of the initial ion energy (see §3.2), so their Larmor radius is small compared to ions. Then, only a small fraction of electrons ($\sim 1\%$, panel (c)) are accelerated at the shock before being advected downstream, where they populate a steep nonthermal tail ($p \sim 4.5$, panel (b)).

At even higher magnetizations ($\sigma \gtrsim 0.3$, yellow curve for $\sigma = 0.3$), the instability that mediates the shock changes again, and the dominant mode is now generated via gyro-frequency resonance with the returning ions (see Appendix 9). However, the overall shape of the electron and ion spectrum does not significantly change (compare black, for $\sigma = 0.1$, and yellow, for $\sigma = 0.3$), showing that, despite the different polarization of the dominant mode (now resonant, as opposed to Bell's non-resonant waves), the physics of the acceleration is not significantly altered. However, the level of magnetic turbulence generated by the instability becomes smaller for $\sigma \gtrsim 0.1$, which explains the observed decrease in acceleration efficiency, especially for ions (Fig. 3.14(c) and (d), for $\sigma \gtrsim 0.1$).

The results presented in Fig. 3.14 refer to $\omega_{\rm pi}t = 2250$. At later times, the nonthermal tail will extend to higher Lorentz factors in all cases, yet at a different rate for different magnetizations. Since ion acceleration is basically regulated by the degree of magnetic turbulence, which is the largest for $\sigma \sim 0.1$, the fastest evolution in the ion spectral tail will be observed for $\sigma \sim 0.1$ (black curve), which in fact extends to the highest energy. Instead, the electron tail evolves more rapidly for moderate magnetizations ($\sigma = 10^{-2}$ in blue), since electron acceleration in high-σ flows is impeded by advection into the downstream (see §4.1). In general, the hierarchy in maximum energy observed in Fig. 3.14 as a function of σ corresponds to an analogous trend in the rates of acceleration.

Finally, we remark that the results reported here for $\theta = 15^\circ$ are representative of the general behavior of low-obliquity shocks.

Superluminal Shocks: $\theta_{\rm crit} < \theta \leq 90^\circ$

The downstream particle spectrum at different magnetizations for a superluminal shock with $\theta = 75^\circ$ is shown in Fig. 3.15. The ion spectrum (upper panel) is always consistent with a purely Maxwellian distribution, with no evidence for nonthermal particles, except for the lowest magnetization ($\sigma = 10^{-5}$, green curve). The non-Maxwellian shape of the ion spectrum for

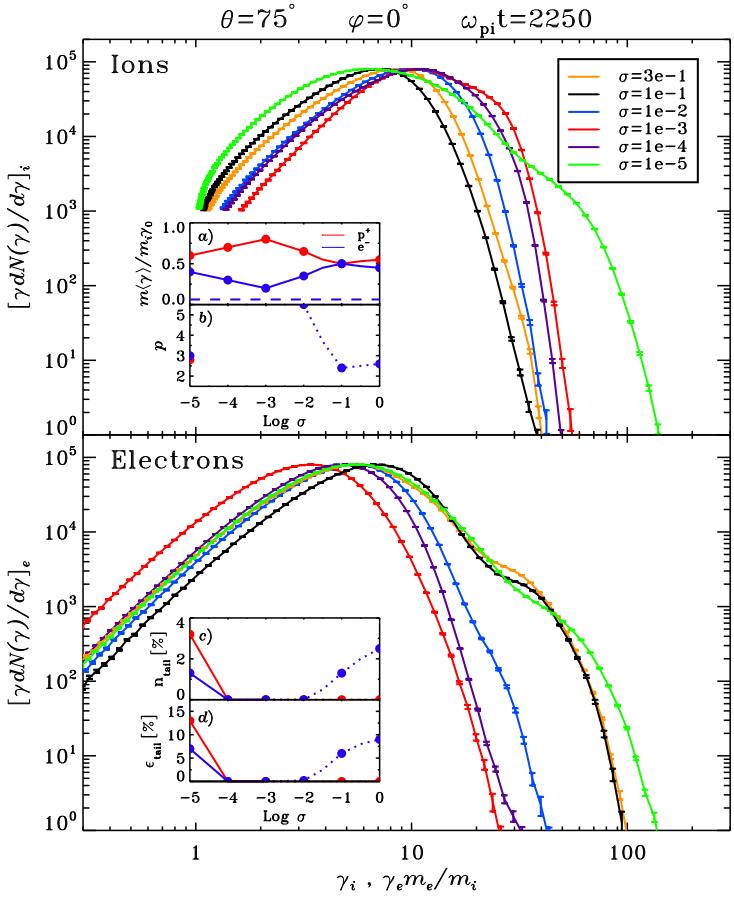


Fig. 3.15.— Downstream particle spectrum at $\omega_{\rm p}t = 2250$ (upper panel for ions, lower panel for electrons) for different magnetizations. We fix the magnetic obliquity $\theta = 75^\circ$ (superluminal shock). Subpanels as in Fig. 3.13.

$\sigma = 10^{-3}$ (red curve) results from incomplete ion thermalization at the shock front, and it will eventually relax to a Maxwellian distribution further downstream from the shock. In contrast, a prominent high-energy component appears in the electron spectrum (lower panel) for both very low (green for $\sigma = 10^{-5}$) and very high (black for $\sigma = 0.1$, yellow for $\sigma = 0.3$) magnetizations, but not in between. However, the mechanism that regulates the electron high-energy tail is markedly different between low and high magnetizations, as we now describe.

In high-σ flows ($\sigma \gtrsim 0.1$; black curve for $\sigma = 0.1$, yellow for $\sigma = 0.3$), the electromagnetic precursor wave generated by the synchrotron maser instability ensures efficient transfer of energy

from ions to electrons ahead of the shock, as discussed in §3.3. For $\sigma \gtrsim 0.1$, the precursor is strong enough to drive electrons up to energy equipartition with ions (compare red and blue lines in panel (a) for $\sigma \gtrsim 0.1$). As described in §4.2, powerful wakefield oscillations excited by the precursor wave in the upstream region dissipate their energy by producing a pronounced high-energy tail in the electron spectrum (black for $\sigma = 0.1$, yellow for $\sigma = 0.3$). The tail contains $\sim 2\%$ of electrons and $\sim 8\%$ of electron energy (blue line in panels (c) and (d), respectively). Again, we remind that in superluminal magnetized shocks ($\sigma \gtrsim 0.1$), the electron high-energy tail is due to wakefield *heating* (rather than *acceleration*), and it does *not* extend in time to higher Lorentz factors (see Fig. 3.11, for $\sigma = 0.1$).

At intermediate magnetizations ($10^{-3} \lesssim \sigma \lesssim 0.1$; blue curve for $\sigma = 10^{-2}$, red for $\sigma = 10^{-3}$), the synchrotron maser instability is gradually suppressed, since it is harder for the electron motion at the shock to stay coherent on the longer gyration time. It follows that the electromagnetic precursor is weaker (Gallant et al. 1992), and ion-to-electron energy transfer in the upstream is suppressed. As a result, the electron spectrum shifts to lower energies, whereas ions get hotter (see panel (a), from $\sigma = 0.1$ down to $\sigma = 10^{-3}$). A minor high-energy component, still due to wakefield heating, survives in the electron spectrum at $\sigma = 10^{-2}$ (blue curve), but it disappears for $\sigma = 10^{-3}$ (red curve). Despite the inefficiency of electron heating upstream of the shock, for $\sigma = 10^{-3}$ the downstream electron population still contains as much as $\sim 20\%$ of the initial ion kinetic energy (blue line in panel (a)). This is considerably in excess of $m_e/m_i \simeq 6\%$ (dashed blue line in panel (a)), the value we would expect in the absence of any ion-to-electron energy transfer.[20] In this regime, electrons gain energy not in the upstream region, but directly at the shock front by the cross-shock potential in the ion foot.

At lower magnetizations ($10^{-5} \lesssim \sigma \lesssim 10^{-3}$; purple for $\sigma = 10^{-4}$), the ion foot, whose scale is set by the Larmor radius of incoming ions in the shock-compressed field, becomes wider, so that electron heating by the cross-shock potential is more efficient (panel (a), for $\sigma \lesssim 10^{-3}$). Finally, for $\sigma \lesssim 10^{-5}$ (we have tested down to $\sigma = 10^{-7}$), the ion foot is so wide that the counter-streaming between incoming and reflected ions can trigger the Weibel instability, before the upstream flow gets advected into the shock. As observed for unmagnetized shocks (Spitkovsky 2008a), this is accompanied by strong electron heating in the Weibel filaments and electron and ion acceleration via a Fermi-like mechanism. This explains the nonthermal tail seen in the electron and ion spectra

[20]This result has been confirmed with larger mass ratios, up to $m_i/m_e = 1000$.

of low-σ flows (green curve for $\sigma = 10^{-5}$). For ions, the tail is relatively flat ($p \sim 2.7$, panel (b)) and it contains $\sim 3\%$ of ions and $\sim 13\%$ of ion energy (panels (c) and (d), respectively). The acceleration efficiency for electrons is smaller by roughly a factor of two (compare red and blue lines in panels (c) and (d), for $\sigma = 10^{-5}$). Unlike in high-σ shocks, here the ion and electron nonthermal tails extend with time to higher energies.

For $\sigma = 10^{-5}$, the ion and electron spectra in Fig. 3.15 approach the results obtained for $\theta = 15^\circ$ at low magnetizations (e.g., $\sigma = 10^{-3}$, red curve in Fig. 3.14). Independently of magnetic obliquity, the shock now behaves as the unmagnetized case discussed by Spitkovsky (2008a). We find that the transition to the unmagnetized regime happens at lower magnetizations for higher obliquities ($\sigma \sim 10^{-2}$ for $\theta = 15^\circ$, $\sigma \sim 10^{-4}$ for $\theta = 75^\circ$), when the self-generated Weibel fields are comparable to the transverse component of the background field ($\propto \sqrt{\sigma}\sin\theta$). Neglecting the obliquity dependence, $\sigma \sim 10^{-3}$ may be taken as a representative value for the transition point.

We remark that the results presented here for $\theta = 75^\circ$ are a good representation of the overall behavior of superluminal shocks, up to the perpendicular case $\theta = 90^\circ$.

5.3. Dependence on Bulk Lorentz Factor γ_0

In this section, we keep the magnetization fixed at $\sigma = 0.1$ and explore the dependence of our results upon the upstream bulk Lorentz factor, from $\gamma_0 = 3$ to $\gamma_0 = 50$. We analyze separately the case of subluminal ($\theta = 15^\circ$, in §5.3) and superluminal ($\theta = 75^\circ$, in §5.3) shocks.

Subluminal Shocks: $0^\circ \leq \theta < \theta_{\rm crit}$

Fig. 3.16 shows the downstream particle spectrum of a $\theta = 15^\circ$ subluminal shock, for different upstream bulk Lorentz factors. Spectra are shifted along the x-axis by $15/\gamma_0$, to facilitate comparison with the reference case $\gamma_0 = 15$ discussed in the previous sections. Once normalized in this way, the spectra for $\gamma_0 \gtrsim 15$ overlap almost perfectly (black for $\gamma_0 = 15$, blue for $\gamma_0 = 30$, red for $\gamma_0 = 50$), suggesting that the physics that regulates the acceleration of particles (both ions and electrons) is insensitive to the upstream bulk Lorentz factor, for highly relativistic flows. In fact, in our simulations we find that for $\gamma_0 \gtrsim 15$ the waves generated by Bell's instability have the same power and wavelength, once their power is normalized to the upstream bulk kinetic flux

and the wavelength is measured in units of the upstream ion skin depth (or Larmor radius). The scaling seen in Fig. 3.16 for $\gamma_0 \gtrsim 15$ then follows from the fact that Bell's modes are of paramount importance for particle acceleration in subluminal shocks (see §4.1).

Spectra for lower Lorentz factors (green for $\gamma_0 = 5$, yellow for $\gamma_0 = 3$) still share the same overall properties of $\gamma_0 = 15$, albeit with a tendency for lower high-energy cutoffs and cooler electron distributions. Overall, the properties of $\gamma_0 \gtrsim 3$ subluminal shocks seem to be well captured by our representative case $\gamma_0 = 15$.

In mildly relativistic magnetized flows ($\gamma_0 \lesssim 2$, for $\sigma = 0.1$), the shock is mediated by a different process, such that the dominant mode is now generated by gyro-frequency resonance with the returning ions, rather than by Bell's non-resonant instability (see Appendix 9). This may change the efficiency of particle acceleration. Such trans-relativistic regime is not the subject of this chapter and will be addressed elsewhere.

Superluminal Shocks: $\theta_{\rm crit} < \theta \leq 90^\circ$

As shown in Fig. 3.17, the spectrum of $\theta = 75^\circ$ superluminal shocks presents a much more complicated dependence on the upstream bulk Lorentz factor. The ion spectrum (upper panel) is consistent with a thermal distribution in the whole range of γ_0 we explore (from $\gamma_0 = 3$ to $\gamma_0 = 50$), but for $\gamma_0 \lesssim 5$ (green for $\gamma_0 = 5$, yellow for $\gamma_0 = 3$) the ion Maxwellian is much hotter than for $\gamma_0 \gtrsim 10$ (black for $\gamma_0 = 15$, blue for $\gamma_0 = 30$, red for $\gamma_0 = 50$). Correspondingly, the electron distribution (lower panel), which peaks at roughly the same energy as ions for $\gamma_0 \gtrsim 10$, shifts to much lower energies for $\gamma_0 \lesssim 5$.

Lyubarsky (2006) argues that the synchrotron maser instability should only operate for $\gamma_0 \gtrsim (m_i/m_e)^{1/3}$. In this regime, the shock-compressed field grows on a scale smaller than the electron Larmor radius, and the non-adiabatic motion of incoming electrons into the shock will result in a ring-like distribution in phase space, which is synchrotron-maser unstable. It follows that a powerful electromagnetic precursor will only be generated in high-γ_0 shocks. Our results show that, for $m_i/m_e = 16$, the transition occurs at $\gamma_0 \sim 10$. For $\gamma_0 \gtrsim 10$, when the precursor wave is strong, electrons reach energy equipartition with ions ahead of the shock. As discussed in §4.2, powerful electrostatic oscillations generated in the wake of the precursor dissipate their energy by heating a fraction of the incoming electrons, that eventually populate the high-energy component

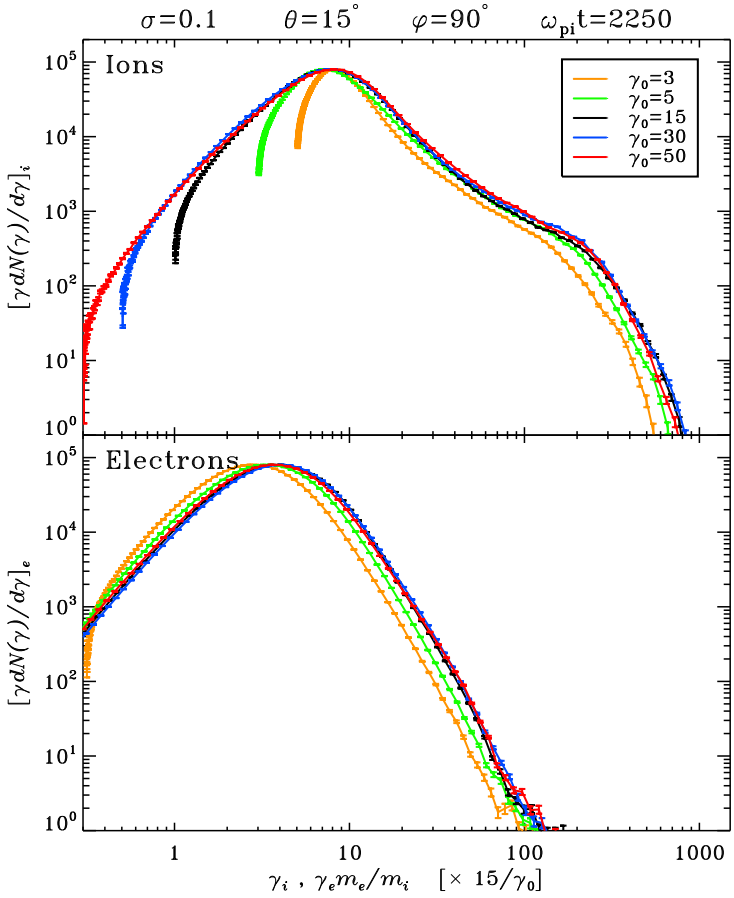


Fig. 3.16.— Downstream particle spectrum at $\omega_{\rm p}t = 2250$ (upper panel for ions, lower panel for electrons) for different upstream bulk Lorentz factors. We fix the magnetic obliquity $\theta = 15^\circ$ (subluminal shock) and the magnetization $\sigma = 0.1$. Spectra are shifted along the x-axis by $15/\gamma_0$ to facilitate comparison with the reference case $\gamma_0 = 15$.

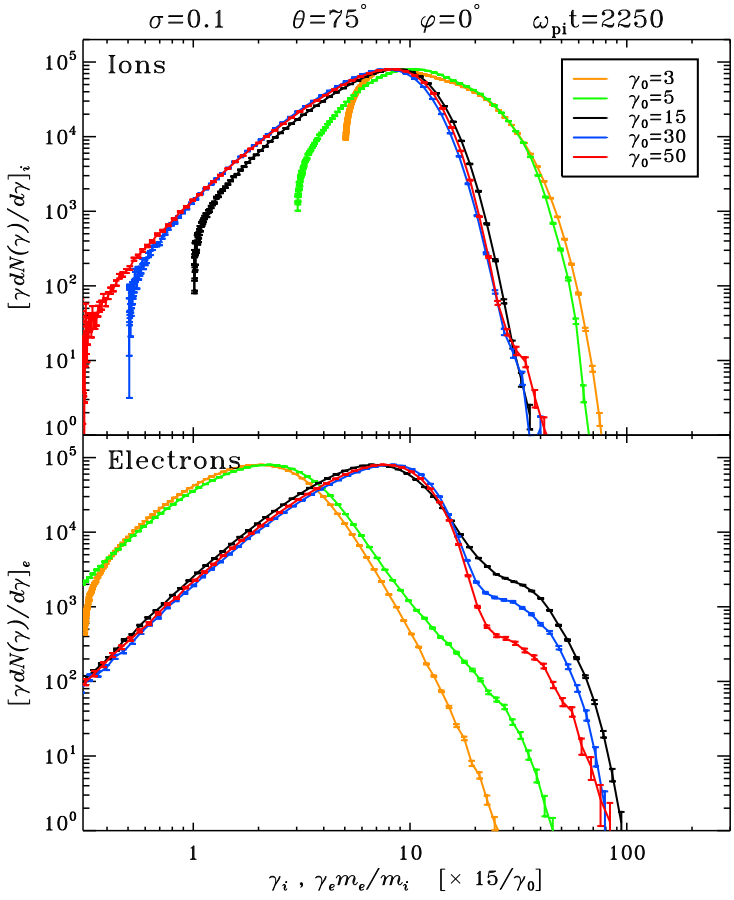


Fig. 3.17.— Downstream particle spectrum at $\omega_{\rm p}t = 2250$ (upper panel for ions, lower panel for electrons) for different upstream bulk Lorentz factors. We fix the magnetic obliquity $\theta = 75°$ (superluminal shock) and the magnetization $\sigma = 0.1$. Spectra are shifted along the x-axis by $15/\gamma_0$ to facilitate comparison with the reference case $\gamma_0 = 15$.

seen for $\gamma_0 \gtrsim 10$ in the electron spectrum of Fig. 3.17. Instead, for $\gamma_0 \lesssim 10$, wakefield oscillations are much weaker, so that the electron high-energy component tends to disappear (green for $\gamma_0 = 5$, yellow for $\gamma_0 = 3$). Also, since the precursor is less powerful in low-γ_0 flows, ion-to-electron energy transfer in the upstream is suppressed, and the downstream electron spectrum shifts to lower energies. Yet, the downstream electron population still contains an appreciable fraction ($\sim 15\%$) of the initial ion energy, due to significant electron energization by the cross-shock potential.

A note of caution is required concerning the electron high-energy tail for $\gamma_0 \gtrsim 10$. As shown in Fig. 3.11, this component grows in time up to a maximum (reached between $\omega_{\rm pi}t \sim 2000$ and $\omega_{\rm pi}t \sim 3000$), and then it recedes (although we cannot exclude a cyclical behavior). The growth time scales with the wavelength of wakefield oscillations, which is longer (in units of $c/\omega_{\rm pi}$) for larger bulk Lorentz factors (Lyubarsky 2006). It follows that at fixed $\omega_{\rm pi}t = 2250$, while $\gamma_0 = 15$ (black curve in Fig. 3.17) is already at maximum, the electron high-energy population is still building up for $\gamma_0 = 30$ (blue curve) and $\gamma_0 = 50$ (red curve). At later times (not shown in Fig. 3.17), the electron spectrum for $\gamma_0 = 30$ and $\gamma_0 = 50$ will roughly resemble the case $\gamma_0 = 15$ presented here.

In summary, our results for $\gamma_0 = 15$ are a good proxy for highly relativistic flows ($\gamma_0 \gtrsim 10$), but they do not apply to mildly relativistic shocks.

6. Summary and Discussion

We have explored by means of 2.5D PIC simulations the internal structure and acceleration properties of relativistic collisionless shocks propagating in a magnetized electron-ion plasma. This chapter complements the study of SS09, that analyzed relativistic shocks in magnetized electron-positron plasmas. To investigate how the shock properties depend on the conditions of the upstream flow, we vary the bulk Lorentz factor γ_0, the magnetization σ and the angle θ between the upstream magnetic field and the shock normal. Our results, confirmed by extensive convergence studies, can be summarized as follows:

- For $\sigma \sim 0.1$, the magnetic field is so strong that charged particles are constrained to move along the field lines, which are advected downstream from the shock. If the magnetic obliquity $\theta < \theta_{\rm crit} \simeq 34^\circ$ ("subluminal" shocks), downstream particles can return upstream by following the magnetic field. In this case, the electric current of "returning" ions can trigger the so-called

Bell's instability (Bell 2004), which generates powerful circularly-polarized Alfvénic-type waves in the upstream. The shock results from the nonlinear steepening of such waves, which happens quasi-periodically every few ion gyro-times (the so-called "shock reformation" process). Instead, for $\theta > \theta_{\rm crit}$ ("superluminal" shocks), no particle can propagate ahead of the shock along the field, and counter-streaming instabilities are suppressed. In this case, the shock is mediated by magnetic reflection of the incoming flow off the shock-compressed field (e.g., Alsop & Arons 1988). In short, the dichotomy between subluminal and superluminal configurations, presented by SS09 for electron-positron flows, still holds in electron-ion shocks.

- If particle acceleration requires repeated bounces back and forth across the shock, as in the standard first-order Fermi picture, only *subluminal* shocks should be efficient accelerators. Or, as measured in the upstream frame, the magnetic field needs to be within an angle $\sim \theta_{\rm crit}/\gamma_0 \simeq 34^\circ/\gamma_0$ from the shock normal. In subluminal shocks, particles are accelerated mostly via non-resonant interactions with Bell's waves (a Fermi-like process), although for $30^\circ \lesssim \theta \lesssim \theta_{\rm crit}$ acceleration by the upstream motional electric field may be important as well (Shock Drift/Surfing Acceleration; see SS09 for similar conclusions). In the downstream spectrum, the shock-accelerated particles populate a power-law nonthermal tail. For our representative case $\theta = 15^\circ$, the ion tail contains $\sim 5\%$ of ions and $\sim 30\%$ of ion energy, and its slope is -2.1 ± 0.1. The tail stretches linearly with time to higher Lorentz factors. Electrons enter the shock with a significant fraction ($\sim 15\%$) of the initial ion energy, yet their Larmor radius is still small compared to ions, and they are rapidly advected downstream with the magnetic field. It follows that the acceleration of electrons is less efficient ($\sim 2\%$ by number and $\sim 10\%$ by energy, with a steep tail of slope -3.5 ± 0.1). In the upstream, the pressure of shock-accelerated ions (in fact, few electrons propagate back upstream) can significantly alter the structure of the shock, forming a smooth density "precursor" ahead of the actual "subshock" (this is the characteristic profile of "cosmic-ray modified" shocks). The overall width of the shock layer increases with time at $\simeq 0.1\,c$.

- In *superluminal* shocks, the synchrotron maser instability (Hoshino & Arons 1991) generates a train of transverse electromagnetic "precusor" waves propagating into the upstream, that push on the incoming electrons, making them lag behind ions. This initiates longitudinal "wakefield" oscillations in the upstream, that transfer energy from ions to electrons (Lyubarsky 2006). In high-obliquity shocks, electrons enter the shock with bulk energy comparable to ions. So, electrons acquire most of their energy in the *upstream* rather than at the *shock*, in contrast

to what is assumed in models that only consider the cross-shock potential (e.g., Gedalin et al. 2008). A few percent of the incoming electrons receive an extra amount of heating from the dissipation of wakefield oscillations. They eventually populate a prominent high-energy feature in the downstream electron spectrum, which may resemble a (short) power-law tail. However, we remark that it results from *heating*, and not *acceleration* (as argued by Hoshino 2008), and it does *not* extend in time to higher energies. Rather, as the incoming electrons get hotter, the generation of electromagnetic precursor waves is suppressed (Hoshino & Arons 1991), and the electron high-energy component recedes to lower energies. Hereafter, a cyclical behavior is likely, though undemonstrated. Despite the significant wakefield heating experienced in the upstream, which may facilitate injection into a diffusive acceleration process, no signature of Fermi-accelerated electrons is seen in magnetized superluminal shocks (see also SS09).

- The properties of magnetized subluminal shocks do not significantly depend on the upstream bulk Lorentz factor, in the relativistic regime $\gamma_0 \gtrsim 3$ we explore. For superluminal angles, the synchrotron maser instability is suppressed for $\gamma_0 \lesssim 10$ (Lyubarsky 2006). So, for low-γ_0 shocks the electromagnetic precursor is weaker, and transfer of energy from ions to electrons in the upstream is much less efficient. Yet, downstream electrons still account for as much as $\sim 15\%$ of the initial ion energy, due to significant electron energization by the cross-shock potential in the ion foot.

- For lower magnetizations ($\sigma \lesssim 10^{-3}$), the incoming plasma can filament, and the shock is mediated by the ion Weibel instability (Weibel 1959; Medvedev & Loeb 1999; Gruzinov & Waxman 1999), which ensures strong electron heating and efficient particle acceleration, for both ions and electrons (Spitkovsky 2008a; Martins et al. 2009). The nonthermal tail in the downstream ion spectrum contains $\sim 5\%$ of ions and $\sim 20\%$ of ion energy, and its slope is -2.5 ± 0.1. A similar tail is present in the electron spectrum, with comparable slope but a smaller particle ($\sim 2\%$) and energy ($\sim 10\%$) content. For magnetizations $\sigma \lesssim 10^{-3}$, the fields generated by the Weibel instability are stronger than the background field, and no difference persists between subluminal and superluminal configurations.

- An important question is whether our 2.5D simulations can capture the relevant three-dimensional physics of electron-ion shocks. We performed a limited number of 3D experiments with relatively small computational grids (~ 3 ion skin depths along each transverse dimension), obtaining

essentially the same results as in our 2.5D simulations. Yet, large 3D runs are certainly desirable to confirm the overall picture presented here.

Our findings may constrain the composition, Lorentz factor and magnetization of jets in Active Galactic Nuclei (AGNs) and Gamma-Ray Bursts (GRBs).[21] Synchrotron and inverse Compton emission from the core of blazar jets is usually attributed to high-energy electrons accelerated in mildly-relativistic internal shocks ($\gamma_0 \sim 2$ in the jet comoving frame). Such shocks should be quasi-perpendicular, since polarization measurements of radio knots (Gabuzda et al. 2004; Pushkarev et al. 2005) indicate that the magnetic field is mostly transverse to the jet axis. Electron acceleration in internal shocks is also invoked to explain the prompt emission of GRBs (e.g., Piran 2004). If the jet flow carries a substantial toroidal field, internal shocks in GRBs will be magnetized and quasi-perpendicular. For both AGNs and GRBs, it is inferred that the emitting electrons should contain a substantial fraction ($\epsilon_e \sim 10\%$) of the flow energy, and their spectrum is usually modeled as a power law extending over several decades in energy.

We find that significant heating of electrons is a universal property of relativistic magnetized shocks. The mechanism involved may vary depending on magnetization and magnetic obliquity, but post-shock electrons account for at least $\sim 20\%$ of the bulk pre-shock energy, regardless of the upstream conditions. Although the agreement between this lower limit and the observational requirement mentioned above is surely encouraging, we remark that most of our downstream electrons belong to the *thermal* component, whereas the emitting particles from AGNs and GRBs are usually thought to be *nonthermal*. Yet, our findings are important to constrain the characteristic downstream electron energy to be $\gtrsim 0.2\,\gamma_0 m_i c^2$, with the following two implications. First, this rules out models of GRB emission (e.g., Bykov & Meszaros 1996) which assume that only a small fraction of electrons are accelerated up to high energies (and share the whole ϵ_e), whereas the bulk of electrons remain cold (at $\sim \gamma_0 m_e c^2$). Second, the minimum Lorentz factor of a hypothetical power-law tail will be $\gtrsim 0.2\,\gamma_0 m_i/m_e$. For models of GRB and AGN emission that require a power-law distribution starting from smaller Lorentz factors, this would suggest that electron-positron pairs may be a major component of the flow (see SS09).

[21] Pulsar Winds are thought to be dominated by electron-positron pairs, and they will not be discussed here (we refer to SS09). If a significant ionic component is present, electron (and positron) acceleration may proceed via the Resonant Cyclotron Absorption mechanism discussed by Hoshino et al. (1992) and Amato & Arons (2006).

Regarding nonthermal acceleration of electrons, we find that quasi-perpendicular magnetized shocks are inefficient particle accelerators, in contrast to what is required in models of GRBs and AGN jets. However, the high-energy component that we find in the electron spectrum of high-obliquity shocks, which actually results from heating (and not acceleration), may be (mis)interpreted as a short power-law tail. In fact, this component might reasonably explain the emission from AGN jets, which often requires an electron distribution extending only for one decade in energy (Celotti & Ghisellini 2008). Otherwise, our findings imply that the jet pre-shock flow should be weakly magnetized ($\sigma \lesssim 10^{-3}$), or maybe seeded with strong small-scale magnetic turbulence, that could facilitate the Fermi process even in superluminal shocks (Sironi & Goodman 2007). Alternatively, electron acceleration may occur not in shocks but in reconnection layers of a Poynting-dominated jet (e.g., Lyutikov & Blandford 2003; Giannios & Spruit 2006).

Finally, if internal shocks in AGNs and GRBs are the sources of High Energy Cosmic Ray ions (e.g., Sironi & Socrates 2010), then the pre-shock magnetic field needs to be nearly aligned with the shock normal, within an angle $\sim \theta_{\rm crit}/\gamma_0 \simeq 34^\circ/\gamma_0$ (as measured in the upstream fluid frame), for the shock to be "subluminal". Alternatively, the acceleration of cosmic ray protons may occur at weakly magnetized shocks, like GRB external shocks or the termination shock ("hot spot") of AGN jets, or in the jet "sheath", where interaction of the jet with the surrounding medium may create subluminal configurations. In any case, we find that ion acceleration is always accompanied by substantial magnetic field amplification, especially ahead of the shock. In the presence of accelerated electrons, the upstream region may then contribute considerably to the observed emission. Synthetic spectra extracted from PIC simulations of shocks (as done in unmagnetized pair shocks by Sironi & Spitkovsky 2009b) may possibly reveal characteristic signatures that could observationally distinguish such *upstream* emission from the *downstream*-based models that are usually employed.

7. Appendix (A): Dependence on the Mass Ratio

The reliability of our results may be questioned based on the artificially small mass ratio ($m_i/m_e = 16$) that we employ in most of our simulations. However, as we show here, for the parameters considered in the present chapter, a mass ratio of $m_i/m_e = 16$ is already large enough to separate the electron and ion scales, and capture the relevant physics of electron-ion shocks.

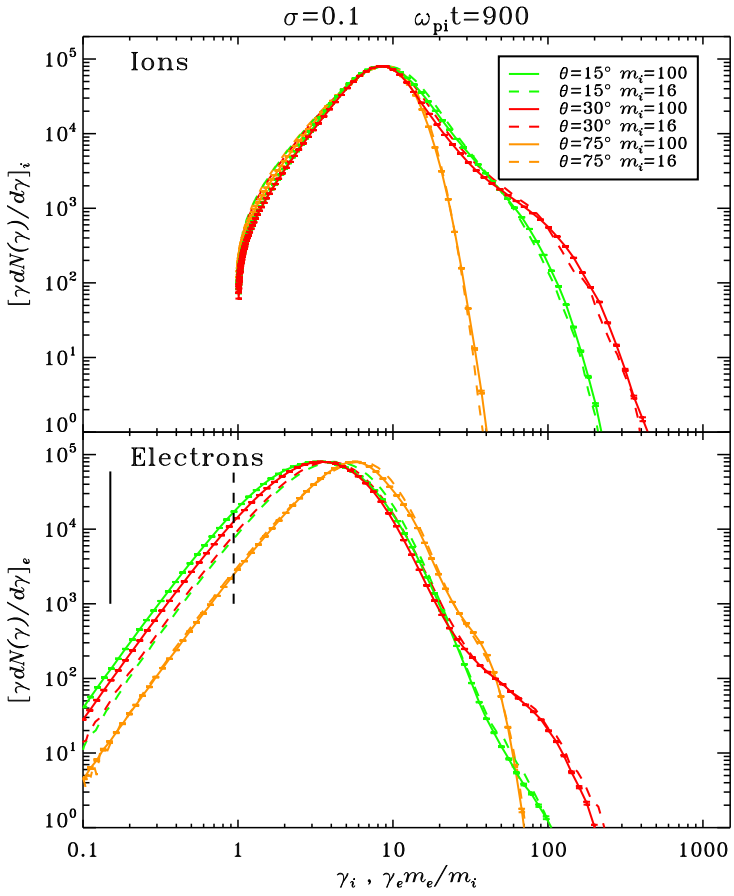


Fig. 3.18.— Comparison of downstream particle spectra (upper panel for ions, lower panel for electrons) between different mass ratios: $m_i/m_e = 16$ (dashed lines) and $m_i/m_e = 100$ (solid lines). The black vertical lines (dashed for $m_i/m_e = 16$, solid for $m_i/m_e = 100$) correspond to $\gamma_0 m_e/m_i$, i.e., the location expected for the peak of the electron distribution in the absence of any ion-to-electron energy transfer.

Fig. 3.18 compares the downstream particle spectra for three representative magnetic obliquities ($\theta = 15°$, $30°$ and $75°$, with $\gamma_0 = 15$ and $\sigma = 0.1$) between two values of the mass ratio: $m_i/m_e = 16$ (dashed lines) and $m_i/m_e = 100$ (solid lines). The spectra are computed at the same time (in units of $\omega_{\rm pi}^{-1}$) and at the same distance behind the shock (in units of $c/\omega_{\rm pi}$). Very good agreement is obtained in all cases, for both ions and electrons.

The downstream electron spectrum peaks at much higher energies than expected in the absence of any ion-to-electron energy transfer (compare with black vertical lines; dashed for $m_i/m_e = 16$, solid for $m_i/m_e = 100$), and the location of the peak is roughly independent of the mass ratio. It follows that the "effective" electron mass, including the average electron Lorentz factor, is a substantial fraction of $\gamma_0 m_i$, independent of m_i/m_e (for $m_i/m_e \gtrsim 16$). The electron physics will then evolve on the ion time and length scales, which explains the similarity in Fig. 3.18 between electron spectra with different m_i/m_e, in terms of both the low-energy thermal part and the high-energy component.

Various tests, for mass ratios up to $m_i/m_e = 1000$ and for different upstream parameters (we changed both σ and γ_0), confirm the overall picture presented here.

8. Appendix (B): Dependence on the Field Orientation with respect to the Simulation Plane

We now assess the robustness of our results upon the orientation of the upstream magnetic field $\mathbf{B}_0$ with respect to the 2D simulation plane. In Fig. 3.19 (upper panel for ions, lower panel for electrons), we plot with red lines the downstream spectrum if $\mathbf{B}_0$ is in the simulation plane (i.e., $\varphi = 0°$), whereas black lines are for $\mathbf{B}_0$ lying in a plane perpendicular to the simulation plane (i.e., $\varphi = 90°$).

For low-obliquity shocks (e.g., $\theta = 15°$, left panel), our results are independent on the angle φ. However, for shocks in which the obliquity is close to the critical boundary $\theta_{\rm crit} \simeq 34°$ between subluminal and superluminal configurations, we observe a substantial suppression of SDA for in-plane magnetic fields, both in electron-positron and electron-ion flows (see $\theta = 30°$ in Fig. 21(b) of SS09, for pair shocks). As discussed in SS09, the case with $\varphi = 90°$ is in better agreement with 3D simulations. For this reason, in the whole range of subluminal angles presented in the current chapter, we employ $\varphi = 90°$.

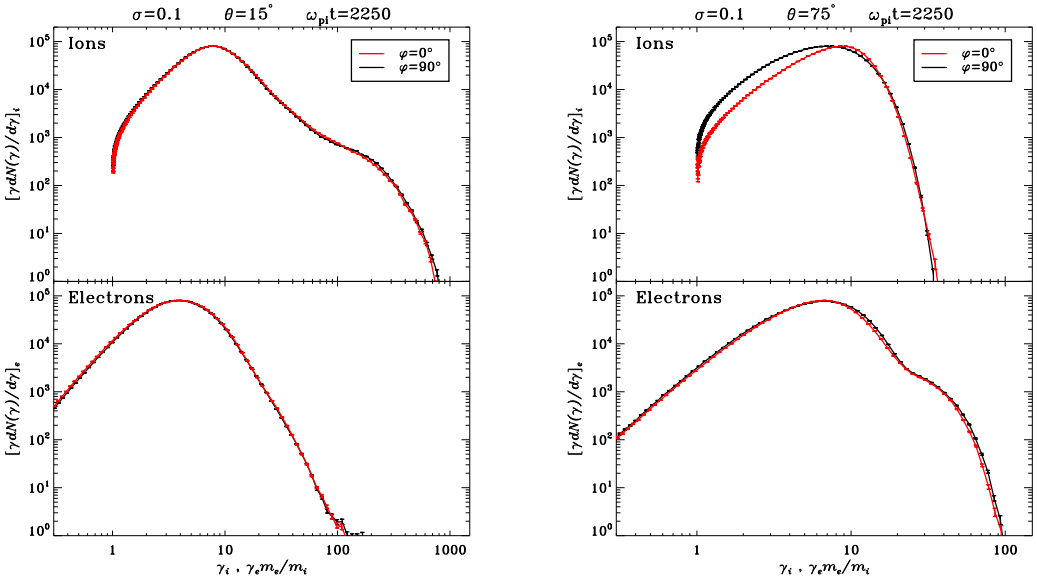


Fig. 3.19.— For fixed magnetic obliquity ($\theta = 15°$ to the left, $\theta = 75°$ to the right), downstream particle spectra at $\omega_{\rm p}t = 2250$ for different orientations of the upstream magnetic field $\mathbf{B}_0$ with respect to the simulation plane: magnetic field lying either in the simulation plane ($\varphi = 0°$, in red) or in a plane perpendicular to the simulation plane ($\varphi = 90°$, in black).

For high-obliquity shocks (e.g., $\theta = 75°$, right panel), the agreement between in-plane and out-of-plane results is also remarkably good. The only appreciable difference is in the low-energy part of the ion spectrum (upper right panel). For $\varphi = 90°$, the plane perpendicular to the magnetic field (where downstream particle orbits would tend to lie) is almost degenerate with the simulation plane, resulting in a downstream ion distribution which is closer to a 2D (rather than 3D) Maxwellian. Instead, for $\varphi = 0°$ the two planes are well separated, and ion isotropization in three dimensions is favored. For this reason, we preferentially employ $\varphi = 0°$ for superluminal shocks.

9. Appendix (C): Ion Counter-Streaming Instabilities

In both unmagnetized and subluminal magnetized shocks, the interaction between the incoming flow and the shock-accelerated particles that propagate back upstream is a potential source of counter-streaming plasma instabilities. In this section, we determine how the nature of the instability depends on the upstream parameters (bulk Lorentz factor γ_0 and magnetization σ, choosing $\theta = 0°$ as a representative subluminal obliquity). Primed quantities are measured in the upstream fluid frame, unprimed quantities in the downstream (simulation) frame.

Let $\zeta_{\rm CR} \equiv n_{\rm CR}/n_i$ be the ratio of cosmic-ray number density to number density of incoming ions. By "cosmic rays", we mean the shock-accelerated ions moving ahead of the shock (we neglect the returning electrons, since their contribution is usually smaller). In the upstream frame, $\zeta'_{\rm CR} = \gamma_0^2\,\zeta_{\rm CR}$, given that the population of returning ions is roughly isotropic in the simulation frame. Similarly, we define $\epsilon_{\rm CR} \equiv U_{\rm CR}/(\gamma_0 m_i n_i c^2)$ as the ratio of cosmic ray energy density to kinetic energy density of the injected ions, such that the average cosmic ray Lorentz factor is $\gamma_{\rm CR} = \gamma_0(\epsilon_{\rm CR}/\zeta_{\rm CR})$. In the upstream frame, $\epsilon'_{\rm CR} = \gamma_0^4\epsilon_{\rm CR}$ and $\gamma'_{\rm CR} = \epsilon'_{\rm CR}/\zeta'_{\rm CR} = \gamma_0\,\gamma_{\rm CR}$. Informed by Figs. 3.14 and 3.16, we can assume to a first approximation that $\zeta_{\rm CR}$ and $\epsilon_{\rm CR}$ (as measured in the downstream frame) do not significantly depend on γ_0 or σ.

The maximum growth rate of Bell's instability ($\omega'_{\rm Bell} \simeq 1/2\,\zeta'_{\rm CR}\,\omega_{\rm pi}$) is achieved for a characteristic wavelength $\lambda'_{\rm Bell} \simeq 4\pi(\sqrt{\sigma'}/\zeta_{\rm CR})\,c/\omega_{\rm pi}$ (Reville et al. 2006). Here, σ' is the magnetization parameter as defined in the upstream frame ($\sigma' = \gamma_0^2\,\sigma$, for $\theta = 0°$). The corresponding Larmor frequency of background ions is $\omega'_{\rm ci} = \gamma_0\,\omega_{\rm ci}$.

Bell's instability governs the evolution of the shock if the following two conditions are simultaneously satisfied: (*i*) its characteristic wavelength $\lambda'_{\rm Bell}$ is smaller than $r'_{\rm L,CR}$, the Larmor radius of cosmic ray ions; (*ii*) its growth rate is smaller than the Larmor frequency of background ions ($\omega'_{\rm Bell} \lesssim \omega'_{\rm ci}$), i.e., the upstream plasma is magnetized. Since $r'_{\rm L,CR} = \gamma'_{\rm CR}/\sqrt{\sigma'}\ c/\omega_{\rm pi}$, the first condition can be rewritten as $2\,\sigma' \lesssim \epsilon'_{\rm CR}$ in the upstream fluid frame, which corresponds, in terms of quantities in the simulation frame, to

$$2\,\sigma \lesssim \gamma_0^2\,\epsilon_{\rm CR} \quad . \tag{3-1}$$

When the above inequality is broken (either for large σ or small γ_0, at fixed $\epsilon_{\rm CR}$), the physics of the instability changes. The dominant mode is now generated via gyro-frequency resonance with the cosmic rays (as opposed to Bell's instability, which produces non-resonant modes). We observe this change of polarization in the dominant mode for magnetizations $\sigma \gtrsim 0.3$ (with fixed $\gamma_0 = 15$) and for bulk Lorentz factors $\gamma_0 \lesssim 2$ (with fixed $\sigma = 0.1$).

The condition (*ii*), that the upstream ions stay magnetized on the growth time of Bell's instability, can be rewritten as $\zeta'_{\rm CR} \lesssim \sqrt{\sigma'}$ in the upstream frame. In the simulation frame,

$$\gamma_0\,\zeta_{\rm CR} \lesssim \sqrt{\sigma} \quad . \tag{3-2}$$

For lower magnetizations, or higher bulk Lorentz factors, the upstream plasma can filament, and the ion Weibel instability dominates. For $\gamma_0 = 15$, this transition happens in our simulations around $\sigma \sim 10^{-2}$.

Bibliography

Achterberg, A., Gallant, Y. A., Kirk, J. G., & Guthmann, A. W. 2001, MNRAS, 328, 393

Alsop, D. & Arons, J. 1988, Physics of Fluids, 31, 839

Amato, E. & Arons, J. 2006, ApJ, 653, 325

Amato, E. & Blasi, P. 2005, MNRAS, 364, L76

—. 2006, MNRAS, 371, 1251

Ballard, K. R. & Heavens, A. F. 1991, MNRAS, 251, 438

Baring, M. G. 2010, ArXiv:astro-ph/1002.3848

Bednarz, J. & Ostrowski, M. 1998, Physical Review Letters, 80, 3911

Begelman, M. C. & Kirk, J. G. 1990, ApJ, 353, 66

Bell, A. R. 1978, MNRAS, 182, 147

—. 2004, MNRAS, 353, 550

Bell, A. R. & Lucek, S. G. 2001, MNRAS, 321, 433

Berezhko, E. G. & Ellison, D. C. 1999, ApJ, 526, 385

Blandford, R. & Eichler, D. 1987, Phys. Rep., 154, 1

Blandford, R. D. & Ostriker, J. P. 1978, ApJ, 221, L29

Bret, A. 2009, ApJ, 699, 990

Buneman, O. 1993, in "Computer Space Plasma Physics", Terra Scientific, Tokyo, 67

Bykov, A. M. & Meszaros, P. 1996, ApJ, 461, L37+

Celotti, A. & Ghisellini, G. 2008, MNRAS, 385, 283

Chang, P., Spitkovsky, A., & Arons, J. 2008, ApJ, 674, 378

Chen, G. & Armstrong, T. P. 1975, in International Cosmic Ray Conference, Vol. 5, 1814–1819

Dieckmann, M. E., Frederiksen, J. T., Bret, A., & Shukla, P. K. 2006, Physics of Plasmas, 13, 112110

Drury, L. O. 1983, Reports on Progress in Physics, 46, 973

Ellison, D. C. & Double, G. P. 2004, Astroparticle Physics, 22, 323

Gabuzda, D. C., Murray, É., & Cronin, P. 2004, MNRAS, 351, L89

Gallant, Y. A., Hoshino, M., Langdon, A. B., Arons, J., & Max, C. E. 1992, ApJ, 391, 73

Gedalin, M., Balikhin, M. A., & Eichler, D. 2008, Phys. Rev. E, 77, 026403

Giannios, D. & Spruit, H. C. 2006, A&A, 450, 887

Gruzinov, A. & Waxman, E. 1999, ApJ, 511, 852

Haugboelle, T. 2010, ArXiv:astro-ph/1007.5082

Hoshino, M. 2008, ApJ, 672, 940

Hoshino, M. & Arons, J. 1991, Physics of Fluids B, 3, 818

Hoshino, M., Arons, J., Gallant, Y. A., & Langdon, A. B. 1992, ApJ, 390, 454

Hoshino, M. & Shimada, N. 2002, ApJ, 572, 880

Keshet, U., Katz, B., Spitkovsky, A., & Waxman, E. 2009, ApJ, 693, L127

Keshet, U. & Waxman, E. 2005, Physical Review Letters, 94, 111102

Kirk, J. G. & Heavens, A. F. 1989, MNRAS, 239, 995

Komissarov, S. S. 2003, MNRAS, 341, 717

Kulsrud, R. & Pearce, W. P. 1969, ApJ, 156, 445

Langdon, A. B., Arons, J., & Max, C. E. 1988, Physical Review Letters, 61, 779

Lee, M. A., Shapiro, V. D., & Sagdeev, R. Z. 1996, J. Geophys. Res., 101, 4777

Lemoine, M. & Pelletier, G. 2010, MNRAS, 402, 321

Lyubarsky, Y. 2006, ApJ, 652, 1297

Lyutikov, M. & Blandford, R. 2003, ArXiv:astro-ph/0312347

Majorana, A. & Anile, A. M. 1987, Physics of Fluids, 30, 3045

Martins, S. F., Fonseca, R. A., Silva, L. O., & Mori, W. B. 2009, ApJ, 695, L189

Medvedev, M. V. & Loeb, A. 1999, ApJ, 526, 697

Niemiec, J. & Ostrowski, M. 2004, ApJ, 610, 851

Piran, T. 2004, Reviews of Modern Physics, 76, 1143

Pushkarev, A. B., Gabuzda, D. C., Vetukhnovskaya, Y. N., & Yakimov, V. E. 2005, MNRAS, 356, 859

Reville, B., Kirk, J. G., & Duffy, P. 2006, Plasma Physics and Controlled Fusion, 48, 1741

Riquelme, M. A. & Spitkovsky, A. 2009, ApJ, 694, 626

Schwartz, S. J., Burgess, D., Wilkinson, W. P., Kessel, R. L., Dunlop, M., & Luehr, H. 1992, J. Geophys. Res., 97, 4209

Shapiro, V. D. & Üçer, D. 2003, Planet. Space Sci., 51, 665

Sironi, L. & Goodman, J. 2007, ApJ, 671, 1858

Sironi, L. & Socrates, A. 2010, ApJ, 710, 891

Sironi, L. & Spitkovsky, A. 2009a, ApJ, 698, 1523 (SS09)

—. 2009b, ApJ, 707, L92

Spitkovsky, A. 2005, in AIP Conf. Ser., Vol. 801, 345

—. 2008a, ApJ, 673, L39

—. 2008b, ApJ, 682, L5

Webb, G. M., Axford, W. I., & Terasawa, T. 1983, ApJ, 270, 537

Weibel, E. S. 1959, Physical Review Letters, 2, 83

Chapter 4

Acceleration of Particles at the Termination Shock of a Relativistic Striped Wind

Abstract

The relativistic wind of obliquely-rotating pulsars consists of toroidal stripes of opposite magnetic field polarity, separated by current sheets of hot plasma. By means of two- and three-dimensional particle-in-cell simulations, we investigate particle acceleration and magnetic field dissipation at the termination shock of a relativistic striped wind. At the shock, the flow compresses and the alternating fields annihilate by driven magnetic reconnection. Irrespective of the stripe wavelength λ or the wind magnetization σ (in the regime $\sigma \gg 1$ of magnetically-dominated flows), shock-driven reconnection transfers all the magnetic energy of alternating fields to the particles, whose average Lorentz factor increases by a factor of σ with respect to the pre-shock value. The shape of the post-shock spectrum depends primarily on the ratio $\lambda/(r_L\sigma)$, where r_L is the relativistic Larmor radius in the wind. The spectrum becomes broader as the value of $\lambda/(r_L\sigma)$ increases, passing from a relativistic Maxwellian to a flat power-law tail with slope around -1.5, populated by particles accelerated by the reconnection electric field. Close to the equatorial plane of the wind, where the stripes are symmetric, the highest energy particles resulting from magnetic reconnection can escape ahead of the shock, and be injected into a Fermi-like acceleration process. In the post-shock spectrum, they populate a power-law tail with slope around -2.5, that extends beyond the flat

component produced by reconnection. Our study suggests that the spectral break between the radio and the optical band in Pulsar Wind Nebulae can be a natural consequence of particle acceleration at the termination shock of striped pulsar winds.

1. Introduction

The broadband spectrum of Pulsar Wind Nebulae (PWNe), extending from the radio up to the gamma-ray band, is usually modeled as synchrotron and inverse Compton radiation from a nonthermal population of electron-positron pairs. It is assumed that the emitting particles are accelerated to a power-law distribution at the so-called termination shock, where the momentum flux of the relativistic pulsar wind is balanced by the confining pressure of the nebula. However, it is still unclear how the wind termination shock can accelerate electrons and positrons (hereafter, simply "electrons") to the relativistic energies required to power the nebular emission.

A flat radio spectrum is the common observational feature of PWNe, whose spectral flux in the radio band scales as $F_\nu \propto \nu^{-\alpha_r}$, with $\alpha_r = 0.0 - 0.3$. For the Crab Nebula, the best studied example of a PWN, the synchrotron cooling time of radio-emitting electrons largely exceeds the age of the nebula, so one cannot exclude that such particles may be "primordial," i.e., accelerated at a very early stage of the PWN evolution (Atoyan 1999). However, observations of wisps in the radio band (Bietenholz et al. 2001) seem to suggest that radio-emitting electrons are being accelerated now, in the same region as those responsible for the optical and X-ray emission. In this case, the observed radio spectral index ($\alpha_r \sim 0.25$ for the Crab; Bietenholz et al. 1997) would imply a distribution of emitting electrons of the form $dN/dE \propto E^{-p}$, with slope $p = 2\alpha_r + 1 \sim 1.5$. To explain the radio through optical emission of the Crab, the power law of shock-accelerated electrons should span at least three decades in particle energy (e.g., Lyubarsky 2003). Flat electron spectra with slope $p < 2$ below GeV energies are also required to model the radio emission of hotspots in radio galaxies (Stawarz et al. 2007) and X-ray observations of luminous blazar sources (Sikora et al. 2009).

For a power-law particle spectrum with $1 < p < 2$, most of the particles lie at the low-energy end of the distribution, but the plasma energy content is dominated by particles at the high-energy cutoff. This contrasts with the energy spectrum expected from Fermi acceleration in relativistic shocks, which normally yields slopes $p > 2$ (e.g., Achterberg et al. 2001; Keshet & Waxman 2005), in which case low-energy particles dominate both by number and by energy. To produce broad spectra with $p < 2$, a mechanism is required that raises the mean electron energy, but leaves most

of the particles at relatively low energies. In the context of pulsar winds, Hoshino et al. (1992) and Amato & Arons (2006) have proposed that, in a wind loaded with ions, resonant absorption of ion cyclotron waves by electrons and positrons could generate flat distributions downstream from the termination shock. However, the required ion injection rate largely exceeds what is expected from standard models of pulsar magnetospheres (e.g., Arons 2007).

An alternative possibility was discussed by Lyubarsky (2003), under the assumption that the flow upstream of the termination shock consists of alternating stripes of opposite magnetic polarity, separated by current sheets of hot plasma (from now on, a "striped wind"). For obliquely-rotating pulsars, this is the configuration expected around the equatorial plane of the wind, where the sign of the toroidal field alternates with the pulsar period. If the wind remains dominated by Poynting flux till the termination shock (which is still an open question, see Lyubarsky & Kirk 2001; Kirk & Skjæraasen 2003), shock-compression of the stripes may drive annihilation of the alternating fields, a process known as "driven magnetic reconnection." The resulting transfer of energy from the fields to the particles could potentially generate flat electron distributions, with slope $1 < p < 2$.

The physics of magnetic reconnection can be captured self-consistently only by means of multi-dimensional particle-in-cell (PIC) simulations. Fully-kinetic PIC simulations provide a powerful tool to explore the microphysics of collisionless plasmas, since they can capture from first principles the fundamental interplay between charged particles and electromagnetic fields. In the context of relativistic magnetic reconnection in pair plasmas, most studies have explored the so-called "undriven reconnection" process (Zenitani & Hoshino 2001, 2005a, 2007, 2008; Jaroschek et al. 2004), where field annihilation is initiated by a transient seed perturbation to an otherwise stable current sheet. As discussed above, this is not the setup expected at the wind termination shock. Here, it is the shock-compression of the flow that steadily drives regions of opposite magnetic field polarity toward each other, causing reconnection.

Shock-driven reconnection in magnetically-dominated flows has been studied by Pétri & Lyubarsky (2007) with analytical methods and one-dimensional (1D) PIC simulations.[1] They find a condition on the wind magnetization and the stripe wavelength needed for full dissipation of the alternating fields at the termination shock (roughly speaking, the post-shock Larmor radius should be larger than the separation between current sheets). However, in 1D all particles gain energy

[1] A similar 1D study, but for winds dominated by kinetic (rather than Poynting) flux, was performed by Nagata et al. (2008).

at the same rate, which results in a Maxwellian distribution downstream from the shock. The generation of nonthermal tails is possible only if different particles experience different energy gains, which requires multi-dimensional simulations. As a preliminary step toward multi-dimensional studies of the termination shock in a striped wind, Lyubarsky & Liverts (2008) have performed 2D simulations of driven magnetic reconnection. Two stripes of opposite magnetic polarity were compressed by means of an external force, which would imitate the effect of a shock. They found that driven magnetic reconnection can produce flat nonthermal tails, with $p \sim 1$. Yet, their study was limited to a single current sheet, and most importantly it did not self-consistently account for the presence of a shock.

In this chapter, we explore via multi-dimensional PIC simulations the acceleration of particles at the termination shock of a striped relativistic electron-positron wind. The physics of a shock interacting with multiple current sheets is captured self-consistently in 2D and 3D, and the transfer of energy from the field to the particles via shock-driven reconnection is investigated from first principles. We find that the alternating fields are completely dissipated upon compression by the shock, and their energy is transferred to the particles, regardless of the properties of the flow. This contrasts with the conclusions of Pétri & Lyubarsky (2007), who found efficient field dissipation only for high magnetizations and short stripe wavelengths. We show that the reduced dimensionality of their 1D model could not correctly capture the growth of the so-called "tearing-mode instability," which is of paramount importance for field annihilation in 2D and 3D configurations. We find that broad particle spectra with flat slopes ($1 < p < 2$) are a common by-product of shock-driven reconnection, but the extent of the power-law tail depends on the wind magnetization and the stripe wavelength.

This chapter is organized as follows. In §2, we discuss the setup of our simulations and the magnetic field geometry. In §3, the properties of the termination shock are investigated for one representative choice of wind parameters. We discuss both the shock structure, with focus on the physics of shock-driven magnetic reconnection, and the process of particle acceleration to nonthermal energies. In §4, we study how the particle spectrum downstream from the shock depends on the physical conditions in the wind, like magnetization and stripe wavelength. Although we mostly employ 2D simulations, in §5 we compare 2D and 3D results, finding good agreement. We summarize our findings in §6 and comment on the application of our results to astrophysical scenarios.

2. Simulation Setup

We use the 3D electromagnetic PIC code TRISTAN-MP (Buneman 1993; Spitkovsky 2005) to study the termination shock of a relativistic striped wind. The shock is set up by reflecting a magnetized electron-positron flow from a conducting wall located at $x = 0$ (Fig. 4.1). The interaction between the incoming beam (that propagates along $-\hat{\boldsymbol{x}}$) and the reflected beam triggers the formation of a shock, which moves away from the wall along $+\hat{\boldsymbol{x}}$ (Fig. 4.1). The simulation is performed in the "wall" frame, where the downstream plasma is at rest.

To follow the shock evolution for longer times with fixed computational resources, we mainly utilize 2D computational domains in the xy plane, with periodic boundary conditions in the y direction. In §5 we compare 2D and 3D simulations and show that 2D runs can capture most of the relevant physics. For both 2D and 3D domains, all three components of particle velocities and electromagnetic fields are tracked.

The incoming electron-positron stream is injected along $-\hat{\boldsymbol{x}}$ with bulk Lorentz factor γ_0. The flow carries a strong magnetic field, oriented along the y direction. It follows that the shock will be "perpendicular," i.e., with magnetic field orthogonal to the shock normal. The choice to initialize the field in the simulation plane (i.e., oriented along y, as opposed to z) is motivated by the agreement between 2D simulations with in-plane fields and 3D experiments, as we show in §5. The spatial profile of the injected magnetic field, sketched in Fig. 4.1, is

$$\begin{aligned} B_y(x,t) &= B_0 \tanh\left\{\frac{1}{\delta}\left[\alpha + \cos\left(\frac{2\pi(x+\beta_0 ct)}{\lambda}\right)\right]\right\} \\ &\equiv B_0 \tanh\zeta \end{aligned} \tag{4-1}$$

and its associated motional electric field is $E_z = \beta_0 B_y$. Here, β_0 is the three-velocity of the injected plasma. This particular choice for the magnetic field is motivated by analytical studies of the structure of pulsar winds (e.g., Pétri & Kirk 2005). At each given time, eq. (4-1) shows that the field profile is periodic with wavelength λ (see Fig. 4.1). Current sheets, where the magnetic field vanishes, are identified by the condition $\zeta \simeq 0$, and their half-thickness is $\Delta = \lambda\delta/2\pi \ll \lambda$. Across each current sheet, the field reverses its sign from $+B_0$ to $-B_0$ (or vice versa), where B_0 is the field strength in the region outside the current sheets, which we shall call "cold wind." In the following, we parameterize the field strength via the so-called magnetization parameter $\sigma \equiv B_0^2/4\pi\gamma_0 m n_{c0} c^2$, taken to be much larger than unity. Here, m is the electron (or positron) mass and n_{c0} is the density of particles in the cold wind. By defining the relativistic skin depth in the cold wind

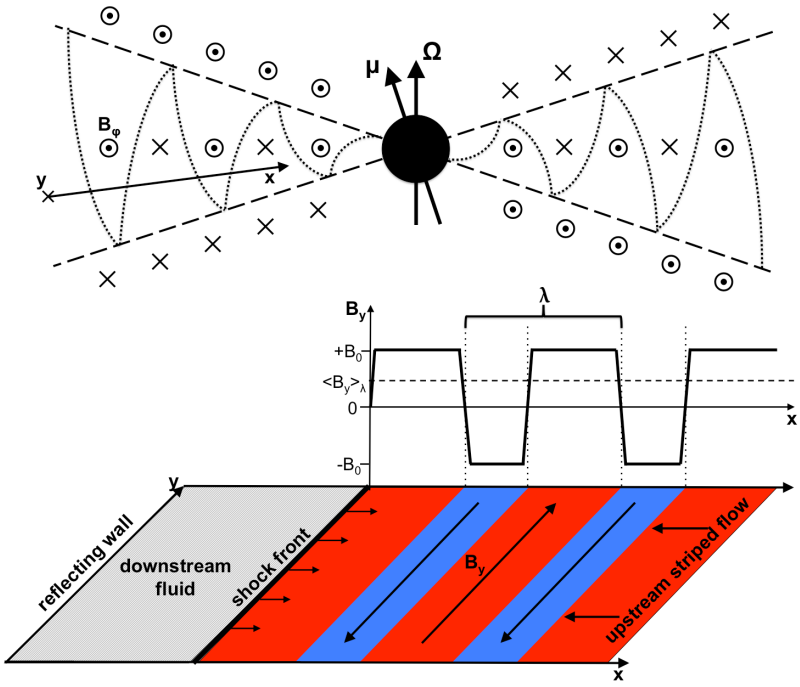


Fig. 4.1.— Upper panel: poloidal structure of the striped pulsar wind, according to the solution by Bogovalov (1999). The arrows denote the pulsar rotational axis (along $\boldsymbol{\Omega}$, vertical) and magnetic axis (along $\boldsymbol{\mu}$, inclined). Within the equatorial wedge bounded by the dashed lines, the wind consists of toroidal stripes of alternating polarity (see the reversals of B_φ), separated by current sheets (dotted lines). At latitudes higher than the inclination angle between $\boldsymbol{\Omega}$ and $\boldsymbol{\mu}$ (i.e., beyond the dashed lines), the field does not alternate. Lower panel: simulation geometry. For 2D runs, the simulation domain is in the xy plane, oriented as shown in the upper panel (so, $\hat{\boldsymbol{x}} = -\hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{y}} = -\hat{\boldsymbol{\varphi}}$). The incoming flow propagates along $-\hat{\boldsymbol{x}}$, and the shock moves away from the reflecting wall (located at $x = 0$) toward $+\hat{\boldsymbol{x}}$. The magnetic field lies in the simulation plane along the y direction, and its polarity alternates with wavelength λ (red stripe for $B_y = +B_0$, blue for $B_y = -B_0$). A net stripe-averaged field $\langle B_y \rangle_\lambda > 0$ is set up by choosing red stripes wider than the blue stripes. For the pulsar wind sketched above, $\langle B_y \rangle_\lambda > 0$ is realized below the equator.

$c/\omega_{\rm p} \equiv \sqrt{\gamma_0 mc^2/4\pi e^2 n_{c0}}$, and the relativistic Larmor radius $r_L \equiv \gamma_0 mc^2/eB_0$, we can rewrite the magnetization parameter as $\sigma = (c/\omega_{\rm p})^2/r_L^2$. Finally, the parameter α ($-1 < \alpha < 1$) is a measure of the net field (i.e., averaged over one wavelength λ), such that $\langle B_y \rangle_\lambda / B_0 = \alpha/(2-|\alpha|)$. Although the magnetic field intensity in the cold wind is always B_0, the wavelength-averaged field is not necessarily zero, because two neighboring stripes generally have different widths (see Fig. 4.1). In the pulsar wind, one expects $\langle B_y \rangle_\lambda = 0$ (which corresponds to $\alpha = 0$) only in the equatorial plane (see upper panel in Fig. 4.1). We remark that a net field $\langle B_y \rangle_\lambda \neq 0$ does not play the role of a guide field, which in our geometry would correspond to a uniform component of magnetic field in the z direction. No significant guide field is expected at the termination shock of pulsar winds, and the physics of shock-driven reconnection in the presence of a guide field will be discussed elsewhere.

In the cold wind, each computational cell is initialized with two electrons and two positrons, with a small thermal spread $\Theta_c \equiv kT_c/mc^2 = 10^{-4}$. We have also performed limited experiments with a larger number of particles per cell (up to 32 per species), obtaining essentially the same results. Such a background of cold particles is initialized also within the current sheets, but with the addition of a hot component whose density profile is $n_h = n_{h0}/\cosh^2\zeta$. Here $n_{h0} \equiv \eta\, n_{c0}$ is the density of hot particles at the center of the current sheet. The temperature of this hot component is determined by pressure balance between the gaseous part (inside the current sheet) and the magnetic part (in the cold wind), which yields a thermal spread $\Theta_h \equiv kT_h/mc^2 = \sigma/2\eta$ for a 3D relativistic Maxwellian. Finally, the variation of magnetic field across a current sheet needs to be sustained by a current flowing along $\hat{\boldsymbol{z}}$. In the wind frame, electrons and positrons in the current sheet should be drifting in opposite directions with the same speed $\beta_h = \sqrt{\sigma}/(\eta\gamma_0)\,(c/\omega_{\rm p})/\Delta$. In summary, to initialize the distribution of hot electrons (and positrons) in the frame of the simulations, we need to make two successive Lorentz boosts: first, from their proper frame to the wind frame, where electrons and positrons counter-stream along $\hat{\boldsymbol{z}}$ with velocity $\pm\beta_h$; then, from the wind frame to the simulation frame, where the wind propagates along $-\hat{\boldsymbol{x}}$ with Lorentz factor γ_0.[2]

In our study, we adopt $\gamma_0 = 15$ as our fiducial bulk Lorentz factor. As discussed in §4, we actually explore a wide range of Lorentz factors, from $\gamma_0 = 3$ to $\gamma_0 = 375$, and find basically the same results, modulo an overall shift in the energy scale. We vary the stripe wavelength by almost two orders of magnitude, from $\lambda = 20\,c/\omega_{\rm p}$ to $\lambda = 1280\,c/\omega_{\rm p}$, and we discuss its effects on the

[2]In the range of parameters we explore, we find that $\beta_h \ll 1$, so the first Lorentz boost can be safely neglected.

structure of the shock and the process of shock-driven reconnection in §4.1. The dependence on the wind magnetization, though always in the regime of highly magnetized flows ($\sigma \gg 1$), is presented in §4.2, where we vary σ from 10 to 100. Different values of the stripe-averaged magnetic field are investigated in §4.3, where we vary the parameter $\alpha = 2\langle B_y \rangle_\lambda / (B_0 + |\langle B_y \rangle_\lambda|)$ between 0.0 and 0.95. For pulsar winds, large values of $|\alpha|$ would be expected at high latitudes above the midplane, whereas $\alpha = 0$ along the equator (see upper panel in Fig. 4.1).

The structure of the current sheet is completely determined once we fix the value of the density excess $\eta = n_{h0}/n_{c0}$, and of the sheet half-thickness $\Delta = \lambda\delta/2\pi$. We choose $\eta = 3$, motivated by analytical studies of pulsar winds (Lyubarsky & Kirk 2001). However, we find that our results do not appreciably depend on the value of η (we have also tried with $\eta = 1.5$ and $\eta = 6$), provided that the contribution of current sheets to the upstream particle flux stays negligible (i.e., $\eta\delta \ll 1$). Regarding the sheet half-thickness, we typically employ $\Delta = c/\omega_{\rm p}$ (for each choice of λ, this determines the value of $\delta = 2\pi\Delta/\lambda$), but we have verified that our results are not very sensitive to this parameter (we have tried up to $\Delta = 8\,c/\omega_{\rm p}$, for $\lambda = 320\,c/\omega_{\rm p}$). The choice for the sheet thickness is guided, on the one hand, by the requirement that the contribution of current sheets to the upstream plasma energy flux should be small (i.e., $\sigma\delta \ll 1$). On the other hand, we enforce $\Delta \gtrsim 0.5\,c/\omega_{\rm p}$, to ensure that the incoming current sheets are stable to the process of *undriven* reconnection. With this choice, we are confident that the trigger for magnetic dissipation will come entirely from the *shock-driven* compression of the flow.

We choose the relativistic electron (or positron) skin depth in the cold wind $c/\omega_{\rm p}$ such that the smallest scale in the system, which for $\sigma > 1$ is the relativistic Larmor radius $r_L = (c/\omega_{\rm p})/\sqrt{\sigma}$, is resolved with at least a few computational cells.[3] We resolve $c/\omega_{\rm p}$ with 7.5 computational cells for $\sigma = 10$, 10 cells for $\sigma = 20$, 13 cells for $\sigma = 30$, 17 cells for $\sigma = 50$, and 24 cells for $\sigma = 100$. For each value of the magnetization, we have verified that our results do not substantially change when doubling the resolution (for $\sigma = 10$, we have tried with up to $c/\omega_{\rm p} = 30$ cells). Our computational domain is typically $\sim 400\,c/\omega_{\rm p}$ wide (along the y direction). In particular, we choose a transverse size of 3072 cells for $\sigma = 10$, of 4096 cells for $\sigma = 20$, of 5120 cells for $\sigma = 30$, of 7168 cells for $\sigma = 50$, and of 9216 cells for $\sigma = 100$. A large box is of paramount importance for the consistency of our findings. As we show in Appendix 7, the results of our simulations converge only when the

[3]The skin depth and Larmor radius for the hot particles in the current sheet are larger than their counterparts in the cold wind by a factor of $\sqrt{\sigma/2\eta^2}$ and $\sigma/2\eta$, respectively.

computational box is larger than the half-wavelength of the stripes. It follows that 1D simulations, as the ones presented by Pétri & Lyubarsky (2007), cannot capture all the relevant physics, especially for long wavelengths.

The results presented in the following sections refer to shocks that have already reached a steady state. For the typical parameters we explore, this happens before $\omega_{\rm p}t = 3000$, but in some cases (especially for long stripe wavelengths), we need to evolve our simulations up to longer times. In Appendix 8, we follow the entire time evolution of a representative shock, from the earliest stages until it approaches a steady state.

3. Shock Structure and Particle Acceleration

In this section, we present the properties of a shock that propagates in a relativistic electron-positron flow with *alternating* fields (a "striped wind"). The case of *uniform* pre-shock fields has been investigated by Sironi & Spitkovsky (2009) and Sironi & Spitkovsky (2011) via multi-dimensional PIC simulations (for electron-positron and electron-ion plasmas, respectively). They studied how the magnetic field strength and obliquity can affect the efficiency of the Fermi process, where particles gain energy by bouncing back and forth across the shock (e.g., Blandford & Eichler 1987). They found that in perpendicular shocks i.e., with field orthogonal to the shock normal, the Fermi mechanism can operate only if the flow is weakly magnetized. For magnetizations $\sigma \gtrsim 10^{-3}$, the field is so strong that charged particles are forced to slide along the field lines, whose orientation prohibits repeated crossings of the shock, thus inhibiting the Fermi process.

In a medium with alternating fields, we find that the efficiency of Fermi acceleration is controlled by the magnitude of the stripe-averaged field $\langle B_y \rangle_\lambda$. Striped flows with $\langle B_y \rangle_\lambda / B_0 \ll 1$ are equivalent to weakly magnetized shocks in uniform fields, with respect to particle acceleration via the Fermi process. In this section, we choose a value for the stripe-averaged field $\langle B_y \rangle_\lambda / B_0 \simeq 0.05$ (corresponding to $\alpha = 0.1$) that is large enough to inhibit the Fermi process. This allows us to focus solely on the process of shock-driven reconnection, and to isolate its effects on the shock structure and the particle spectrum. In pulsar winds, the special case $\langle B_y \rangle_\lambda = 0$ should be realized only along the equatorial plane of the wind. By choosing $\alpha = 0.1$, we address the generic properties of the flow above (or below) the midplane.

In this section, we adopt $\gamma_0 = 15$ and $\sigma = 10$ as our fiducial values for the upstream bulk Lorentz factor and magnetization. To clarify the physics of shock-driven reconnection, we employ relatively large values for the stripe wavelength ($\lambda = 320\, c/\omega_{\rm p}$ or $\lambda = 640\, c/\omega_{\rm p}$), and we choose $\alpha = 0.1$, corresponding to a stripe-averaged field $\langle B_y \rangle_\lambda / B_0 \simeq 0.05$. The dependence of the shock properties on γ_0, σ, λ, and α is investigated in §4. There, we also explore the special case $\alpha = 0$, where shock-driven reconnection coexists with Fermi acceleration.

3.1. Shock Structure

The typical structure of a relativistic shock propagating in a striped wind is shown in Figs. 4.2 and 4.3, which refer to $\omega_{\rm p} t = 3750$, after the flow structure has reached a steady state. The longitudinal profile of the transition region, along the direction x of shock propagation, is presented in Fig. 4.2, whereas Fig. 4.3 shows the flow properties in the xy plane of the simulation, zooming in on a region around the shock (as delimited by the vertical dashed red lines in Fig. 4.2(a)).

The incoming flow ($x \gtrsim 2600\, c/\omega_{\rm p}$) carries an alternating magnetic field of wavelength $\lambda = 640\, c/\omega_{\rm p}$ (Fig. 4.2(c); black line for B_y, red for E_z). The field reverses its polarity across current sheets filled with hot overdense plasma (see the spikes in density and average particle energy of Fig. 4.2(a) and (d), respectively), where the magnetic energy is much lower than in the cold wind (see the troughs in the magnetic energy fraction $\epsilon_B = B^2/8\pi\gamma_0 m n_{c0} c^2$, black line of Fig. 4.2(b)). For $x \gtrsim 2600\, c/\omega_{\rm p}$, the only perturbation to the incoming flow comes from the so-called electromagnetic precursor, a coherent train of transverse electromagnetic waves propagating into the upstream at the speed of light (see the short-scale ripples at $x \simeq 3850\, c/\omega_{\rm p}$ in Fig. 4.2(b) and (c)). The electromagnetic precursor wave is generated in the initial stages of shock evolution, before shock-driven reconnection appreciably changes the structure of the shock. At such early times, the flow resembles a magnetized shock propagating in uniform fields (see Appendix 8), which is well known to radiate electromagnetic waves into the upstream (Gallant et al. 1992; Hoshino 2008; Sironi & Spitkovsky 2009, 2011).

The steady-state structure of the flow in Fig. 4.2 shows the presence of two shocks. The main shock (which we shall call the "hydrodynamic shock" from now on, for reasons that will become clear below) corresponds to the jump in density occurring at $x \simeq 1000\, c/\omega_{\rm p}$. At $x \simeq 2600\, c/\omega_{\rm p}$, well ahead of the hydrodynamic shock, the incoming flow crosses a fast magnetohydrodynamic (MHD) shock, whose location is highlighted in Fig. 4.2 by a vertical dotted blue line. At the

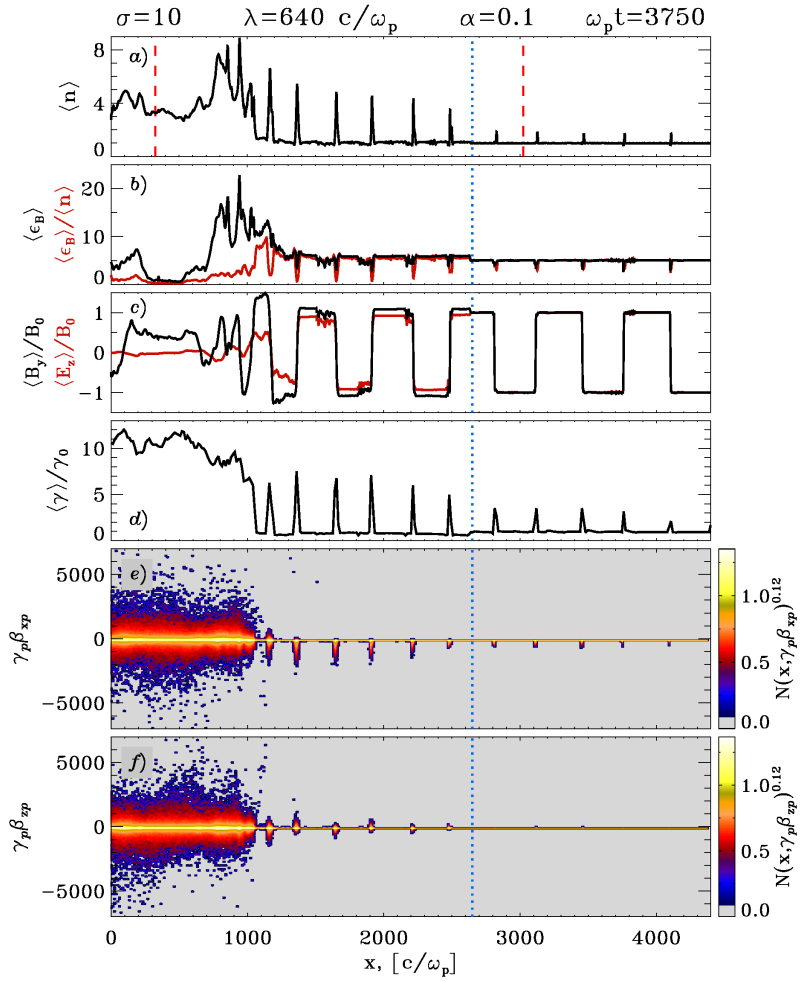


Fig. 4.2.— Internal structure of the flow at $\omega_{\rm p}t = 3750$, for stripe wavelength $\lambda = 640\,c/\omega_{\rm p}$ and magnetization $\sigma = 10$. The stripe-averaged field is $\langle B_y\rangle_\lambda/B_0 \simeq 0.05$, corresponding to $\alpha = 0.1$. The hydrodynamic shock is located at $x \simeq 1000\,c/\omega_{\rm p}$. In all panels, the location of the fast MHD shock ($x \simeq 2600\,c/\omega_{\rm p}$) is indicated with a vertical dotted blue line. As a function of the longitudinal coordinate x, the following quantities are plotted: (a) y-averaged particle number density n, in units of the upstream value; (b) y-averaged magnetic energy fraction $\epsilon_B \equiv B^2/8\pi\gamma_0 m n_{c0} c^2$ (black line), and y-averaged magnetic energy per particle ϵ_B/n (red line); (c) y-averaged magnetic field B_y (black line) and electric field E_z (red line), normalized to the upstream magnetic field B_0; (d) mean kinetic energy per particle, in units of the bulk energy at injection; (e) $x - \gamma\beta_{\rm x}$ positron phase space; (f) $x - \gamma\beta_{\rm z}$ positron phase space.

fast shock, the wind is decelerated and compressed, as shown by the modest increase in density (Fig. 4.2(a)), magnetic energy (black line in Fig. 4.2(b)), and transverse magnetic field B_y (black line in Fig. 4.2(c)). The deceleration of the flow appears most clearly in panel (c), as a mismatch between the motional electric field E_z and the magnetic field B_y behind the fast shock (red and black line, respectively). Irrespective of the initial bulk Lorentz factor of the flow, the plasma drift velocity ($= -E_z/B_y\,\hat{\boldsymbol{x}}$) downstream from the fast shock is only moderately relativistic, since $E_z/B_y \simeq 0.8$ there (Fig. 4.2(c)). The flow deceleration at the fast shock is accompanied by a partial isotropization of the particle distribution, mostly in the xz plane orthogonal to the field. Overall, the flow kinetic energy decreases across the fast shock from $\gamma_0 = 15$ to $\langle\gamma\rangle \simeq 10$ (Fig. 4.2(d)), to compensate for the growth in magnetic energy (Fig. 4.2(b)).

The fast MHD shock has a dramatic impact on the structure of the striped wind. Across the fast shock, the stripe wavelength shrinks from $\lambda = 640\,c/\omega_{\rm p}$ to $\lambda \simeq 550\,c/\omega_{\rm p}$ (see Fig. 4.2(c)). Most importantly, the compression induced by the passage of the fast shock through a given current sheet triggers the onset of magnetic reconnection. Small-scale islands develop inside the current sheet as a result of the tearing-mode instability. The growth and evolution of reconnection islands is clearly shown in the 2D plots of density and magnetic energy of Fig. 4.3(a) and (b), respectively. The current sheet at $x \simeq 2500\,c/\omega_{\rm p}$, just downstream from the fast shock (located at $x \simeq 2600\,c/\omega_{\rm p}$), appears partitioned into a number of overdense islands with strong magnetic fields (red point-like regions in Fig. 4.3(b), at $x \simeq 2500\,c/\omega_{\rm p}$). A magnetic X-point exists in between each pair of neighboring islands, where field lines of opposite polarity break and reconnect. As the flow recedes from the fast shock, the small-scale islands created inside each current sheet coalesce to form bigger islands. This temporal evolution corresponds, in the snapshot taken at $\omega_{\rm p}t = 3750$ in Fig. 4.3, to a pattern of bigger (and so, fewer) islands for current sheets that are farther behind the fast shock.

As reconnection proceeds with distance downstream from the fast shock, more and more magnetic energy is released and transferred to particles. In fact, the spikes in average particle energy of Fig. 4.2(d) become stronger for current sheets farther from the MHD shock (see also the longitudinal phase space of positrons in Fig. 4.2(e), for $1000\,c/\omega_{\rm p} \lesssim x \lesssim 2600\,c/\omega_{\rm p}$). The transfer of magnetic energy to the particles continues up to the point where reconnection islands fill the entire region between neighboring current sheets. Now, the striped structure of the flow is erased, and a hydrodynamic shock forms (located at $x \simeq 1000\,c/\omega_{\rm p}$ in Figs. 4.2 and 4.3).

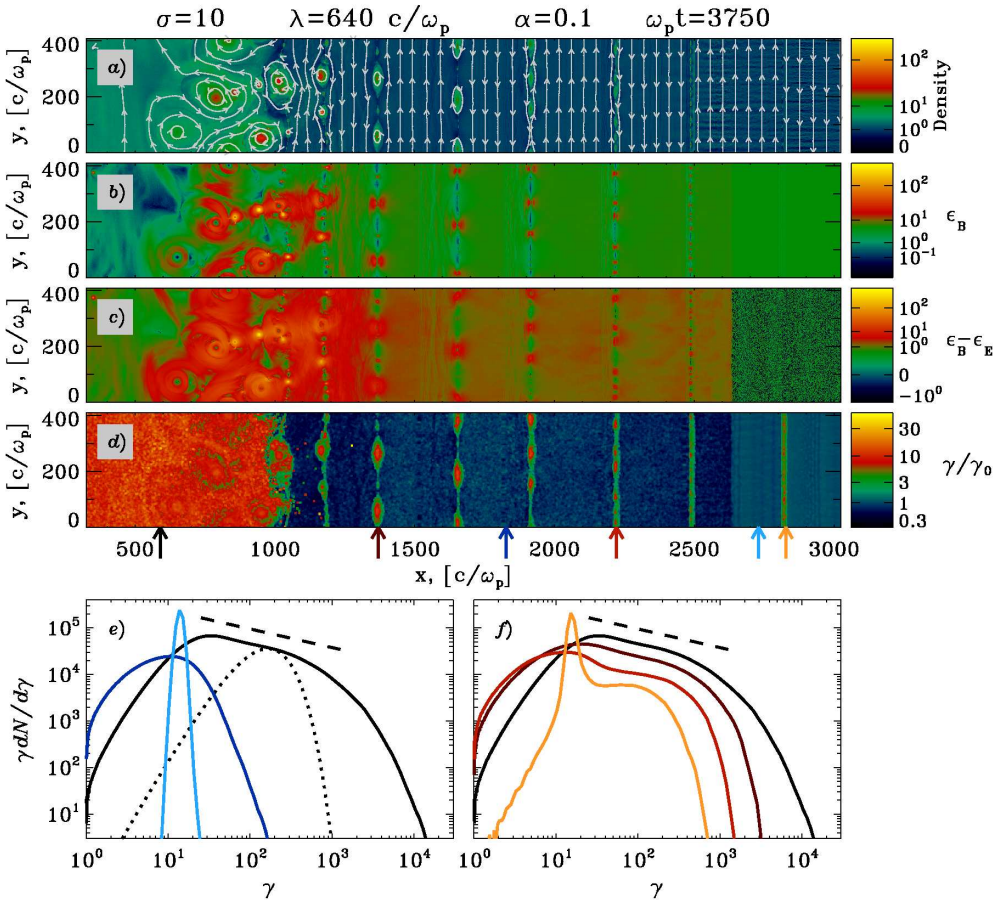


Fig. 4.3.— Internal structure of the flow at $\omega_{\rm p}t = 3750$, for $\lambda = 640\,c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$, zooming in on a region around the shock, as delimited by the vertical dashed red lines in Fig. 4.2(a). The hydrodynamic shock is located at $x \simeq 1000\,c/\omega_{\rm p}$, and the fast MHD shock at $x \simeq 2600\,c/\omega_{\rm p}$. The following quantities are plotted: (a) 2D plot in the simulation plane of the particle number density, in units of the upstream value, with contours showing the magnetic field lines; (b) 2D plot of the magnetic energy fraction $\epsilon_B \equiv B^2/8\pi\gamma_0 m n_{c0} c^2$; (c) 2D plot of the difference $\epsilon_B - \epsilon_E$, where the electric energy fraction is $\epsilon_E \equiv E^2/8\pi\gamma_0 m n_{c0} c^2$; (d) 2D plot of the mean kinetic energy per particle, in units of the bulk energy at injection; (e)-(f) particle energy spectra, at different locations across the flow, respectively outside (panel (e)) or inside (panel (f)) of current sheets. The color of each spectrum matches the color of the corresponding arrow at the bottom of panel (d), showing where the spectrum is computed. The dotted line in panel (e) is a Maxwellian with the same average energy as the downstream particles; the dashed lines in panels (e) and (f) indicate a power-law distribution with slope $p = 1.4$.

Behind the hydrodynamic shock, the flow comes to rest (in the wall frame of our simulations), and the particle distribution is isotropic in three dimensions (compare longitudinal and transverse phase spaces in Fig. 4.2(e) and (f), respectively). The post-shock number density approaches $n_{\rm d} \simeq 4\, n_{\rm u}$ (Fig. 4.2(a), where $n_{\rm u}$ is the density ahead of the fast shock), and the shock velocity is $\beta_{\rm sh} \simeq 1/3$ (see Appendix 8), in agreement with the jump conditions of a relativistic *unmagnetized* 3D plasma. In fact, most of the energy per particle downstream from the hydrodynamic shock is in kinetic form (as opposed to the dominant electromagnetic component of the incoming striped flow). This becomes clear when comparing the profile of the mean *kinetic* energy per particle (Fig. 4.2(d)) with the 1D plot of the mean *magnetic* energy per particle (red line in Fig. 4.2(b)).[4] Behind the hydrodynamic shock, the average kinetic energy per particle increases from $\langle\gamma\rangle/\gamma_0 \simeq 1$ up to $\langle\gamma\rangle/\gamma_0 \simeq \sigma + 1 \simeq 11$, as expected in the case of full dissipation of magnetic fields. Correspondingly, the mean magnetic energy per particle drops to zero (red line in Fig. 4.2(b)). As best seen in experiments with a smaller stripe wavelength, in the case $\alpha \neq 0$ a net magnetic field survives downstream from the shock, as predicted by Pétri & Lyubarsky (2007), and shown by the 1D profile of B_y in Fig. 4.2(c) (black line). This residual magnetic field results from shock-compression of the stripe-averaged upstream field (we remind that $\alpha = 0.1$ here, corresponding to $\langle B_{\rm y}\rangle_\lambda/B_0 \simeq 0.05$).

We point out that the picture described above represents accurately the long-term behavior of the shock. In particular, the distance between the fast and the hydrodynamic shock ($\simeq 1600\, c/\omega_{\rm p}$ in Figs. 4.2 and 4.3) stays approximately constant in time (see Appendix 8). The hydrodynamic shock moves with $\beta_{\rm sh} \simeq 1/3$, and the fast MHD shock remains far enough ahead of the main shock such that to guarantee that reconnection islands will fill the space between neighboring current sheets, by the time the flow enters the hydrodynamic shock. For a fixed reconnection rate, this implies that the distance between the fast and the hydrodynamic shock will scale linearly with the stripe wavelength (as will the size of reconnection islands just ahead of the main shock), a prediction that we have directly verified in our simulations.

Finally, we refer to Appendix 7 for a detailed discussion of the differences between our 2D results and the 1D simulations presented by Pétri & Lyubarsky (2007). As apparent in Fig. 4.3, the process of island coalescence, which is essential for the formation of the hydrodynamic shock, can be captured correctly only with multi-dimensional simulations. Indeed, for the parameters

[4] We point out that in the incoming flow electric and magnetic fields equally contribute to the energy balance, whereas the red line in Fig. 4.2(b) only includes magnetic fields. However, downstream from the hydrodynamic shock, the flow is at rest (in the simulation frame), and electric fields vanish (see red line in Fig. 4.2(c)).

employed in this section, the 1D model by Pétri & Lyubarsky (2007) would predict negligible field dissipation, in sharp contrast with our findings. This clearly emphasizes the importance of multi-dimensional physics for our understanding of shock-driven reconnection.

3.2. Particle Spectrum and Acceleration

The particle energy spectrum at different locations through the flow, as marked by arrows at the bottom of Fig. 4.3(d), is shown in Fig. 4.3(e) and (f). Panel (e) shows the particle distribution in the cold wind, whereas panel (f) focuses on the hot plasma within current sheets. Downstream from the hydrodynamic shock, the striped structure is completely erased, and no difference persists between current sheets and cold wind. Here, the particle spectrum (black line in both panels) is in the form of a flat power-law tail (with spectral index $p \simeq 1.4$, dashed line) extending from $\gamma_{\rm min} \simeq 30$ to $\gamma_{\rm max} \simeq 500$, where it cuts off exponentially. Since $1 < p < 2$, most of the particles are found at $\gamma \simeq \gamma_{\rm min}$, but most of the energy is contributed by particles at $\gamma \simeq \gamma_{\rm max}$. For comparison, a 3D Maxwellian with the same average Lorentz factor $\langle\gamma\rangle = \gamma_0(\sigma + 1) \simeq 165$ is plotted in panel (e) as a dotted line, to show that the actual particle spectrum is much broader than a thermal distribution.[5] By following the flow from the fast to the hydrodynamic shock, we now clarify how such a particle spectrum is generated.

Most of the incoming plasma is contained in the cold striped wind exterior to current sheets, and its spectrum is initially a cold Maxwellian drifting with bulk Lorentz factor $\gamma_0 = 15$ (light blue line in Fig. 4.3(e)). After crossing the fast shock, the flow decelerates and it partly isotropizes, which explains why the particle spectrum is now broader and it peaks at lower energies (dark blue line in Fig. 4.3(e), but see also the 2D plot of average particle energy in Fig. 4.3(d), and its 1D projection in Fig. 4.2(d)).

The evolution of the particle population within current sheets is much more dramatic, as shown in Fig. 4.3(f).[6] Ahead of the fast shock, the spectrum at $x \simeq 2850\, c/\omega_{\rm p}$ (yellow line in panel (f)) shows, in addition to the cold plasma in the striped wind, the hot particles in the current sheet. Their distribution is initialized as a drifting 3D Maxwellian with thermal spread $\Theta_h = \sigma/2\eta$,

[5] We remark that the particle spectrum behind the hydrodynamic shock does not significantly evolve with time, after the shock has reached a steady state (see Appendix 8).

[6] We remind that the spatial pattern shown in Fig. 4.3, as a function of distance behind the MHD shock, corresponds to the time history of a given fluid element after its crossing of the fast shock.

so that in the simulation frame the spectrum should peak at $\langle\gamma\rangle = 3\gamma_0\Theta_h \simeq 75$, as observed (we remind that $\sigma = 10$ and $\eta = 3$). Behind the fast shock (red line for $x \simeq 2200\,c/\omega_{\rm p}$, brown line for $x \simeq 1350\,c/\omega_{\rm p}$), the particle spectrum consists of two components: a low-energy peak, reminiscent of the particle distribution far from current sheets (compare with the dark blue line in Fig. 4.3(e)); and a high-energy component, that grows in number and extends to higher and higher energies, as the flow propagates toward the hydrodynamic shock. This component is populated by particles that were initially outside the current sheet, and have entered the sheet in the course of the reconnection process. They mostly end up in the reconnection islands of Fig. 4.3, that act as reservoirs of particles (Fig. 4.3(a)) and particle energy (the high-energy regions of panel (d), colored in red, match very well the overdense islands of panel (a)). In fact, the main contributions to the high-energy end of the spectrum at $x \simeq 1350\,c/\omega_{\rm p}$ (brown line in Fig. 4.3(f)) come from the two major islands of that current sheet.

As magnetic reconnection proceeds with distance behind the fast shock, particles are accelerated to higher and higher energies, as we discuss below in more detail. This continues until reconnection islands grow as big as the distance between neighboring sheets, at which point the current sheets occupy the entire flow (and the hydrodynamic shock forms). This explains the smooth transition from the spectrum within current sheets to the distribution downstream of the hydrodynamic shock (yellow through black lines in Fig. 4.3(f)). Also, since most of the particles that end up with high energies were initially part of the cold wind, the resulting spectrum behind the hydrodynamic shock is almost independent of the initial setup of the current sheet (most notably, sheet thickness Δ and overdensity η).

We now investigate in more detail the physics of particle acceleration within current sheets. In Fig. 4.4, we follow the trajectories of a representative sample of positrons, extracted from the simulation of a striped wind with $\lambda = 320\,c/\omega_{\rm p}$ (but our conclusions hold for all wavelengths). All the selected positrons are initially in the cold wind (i.e., far from current sheets), at roughly the same x-location (within a range of $5\,c/\omega_{\rm p}$). At first they propagate toward the fast shock, which they cross at $\omega_{\rm p}t \simeq 700$, as shown by the spreading of their energies in Fig. 4.4(a). They encounter the closest current sheet (in this case, the one that was to the right of their initial position) at $\omega_{\rm p}t = 1008$, as shown in Fig. 4.4(b) together with the corresponding 2D plot of particle number density. Starting from this time, their histories can diverge significantly, depending on their y-location at the moment of interaction with the current sheet. The particles that will end up with relatively low energies ($\gamma < 3\gamma_0$, black lines in panel (a)) are found at $\omega_{\rm p}t = 1008$ exclusively around

magnetic islands (their locations are indicated with filled circles in panel (b)). At later times, they will get trapped on the outskirts of the growing islands, without appreciable changes in energy.

On the contrary, the positrons that will eventually reach high energies ($\gamma > 6\gamma_0$, red and blue lines in Fig. 4.4(a)) are concentrated at $\omega_{\rm p}t = 1008$ in the vicinity of magnetic X-points (the particle locations are shown with open circles in Fig. 4.4(b)). It is at X-points that magnetic field lines of opposite polarity are converging and tearing (white contours in Fig. 4.4(b) show the magnetic field lines), and magnetic energy is transferred to particles. Starting from $\omega_{\rm p}t = 1008$ (vertical dashed line in panel (a)), the energies of these positrons grow simultaneously and explosively (red and blue lines in Fig. 4.4(a)), due to the reconnection electric field parallel to the X-line (Zenitani & Hoshino 2001; Jaroschek et al. 2004; Lyubarsky & Liverts 2008). We find that X-point acceleration represents the first stage of energization for most of the particles that will end up with high energies. Particles pre-accelerated at X-points continue to gain energy by compression (Lyubarsky & Liverts 2008; Giannios 2010), as they are advected from the X-point into the closest island. No major systematic changes of energy are observed after the particles get trapped within magnetic islands. Different X-points may provide different energy gains, as shown by the red (respectively, blue) curves in panel (a), which refer to high-energy positrons that get accelerated at the X-point above (respectively, below) the small magnetic island in the middle of the current sheet. In summary, the energy gain of a given particle is determined on the one hand by the strength of the reconnection electric field at the X-point, and on the other hand by the time available for acceleration, as the particle flows from the X-point into the closest island.

In retrospect, the importance of X-points for the generation of nonthermal particles is not surprising. In magnetically-dominated flows, particles can untie from magnetic field lines and be injected into the acceleration process only in regions where the electric field is comparable or exceeds the magnetic field, i.e., green or blue regions in the current sheets of Fig. 4.3(c) (where we plot the difference $\epsilon_B - \epsilon_E$ between magnetic and electric energy). Based on this criterion, X-points are natural sites of particle energization, since the magnetic field vanishes there. As discussed above, the main agent of particle acceleration at X-points is the reconnection electric field, which for the magnetic geometry employed in this chapter is oriented along z (more precisely, along $+\hat{\boldsymbol{z}}$ if the magnetic field B_y is negative to the left of the current sheet and positive to the right; along $-\hat{\boldsymbol{z}}$ in the opposite case). The sign of the reconnection electric field leaves an imprint on the particle z-momentum. As shown by the positron transverse phase space in Fig. 4.2(f), current sheets where the reconnection field is expected to be along $+\hat{\boldsymbol{z}}$ are mostly populated by high-energy positrons

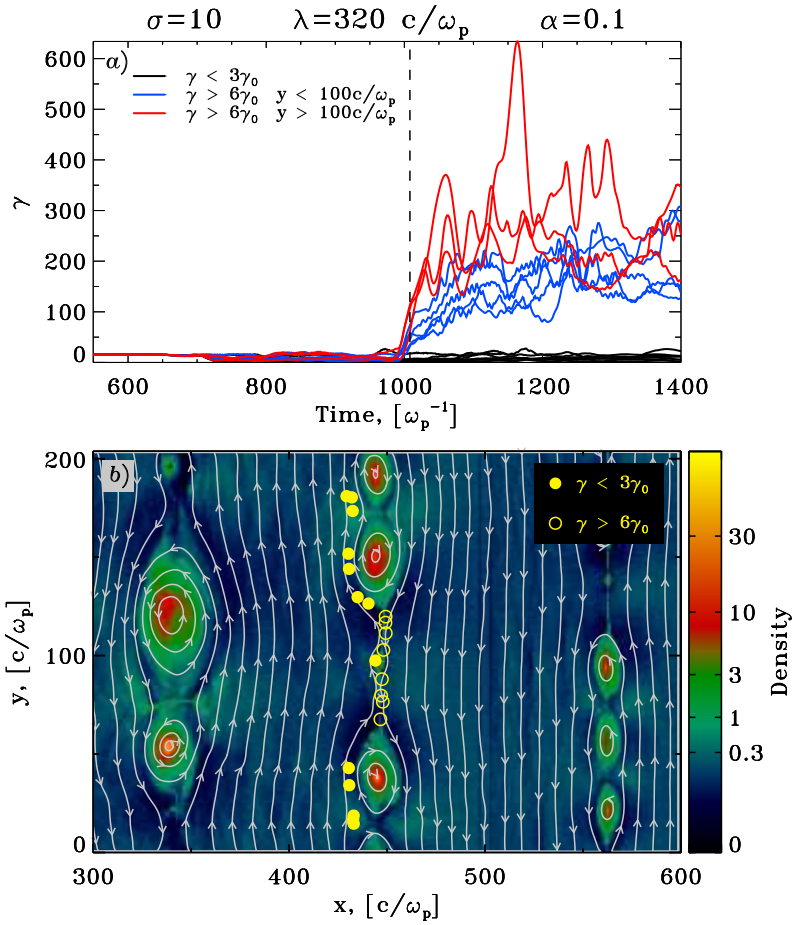


Fig. 4.4.— Energy evolution of a sample of representative positrons interacting with a current sheet, in a striped flow with $\lambda = 320\,c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$. All the selected positrons have been initialized as part of the cold wind, at roughly the same x-location. Panel (a): Energy histories of the selected positrons, with lines of different color depending on their final energy (black if $\gamma < 3\gamma_0$, blue or red if $\gamma > 6\gamma_0$) and their y-location at $\omega_{\rm p}t = 1008$ (blue if $y < 100\,c/\omega_{\rm p}$, red if $y > 100\,c/\omega_{\rm p}$). Panel (b): Magnetic field lines (white contours) superimposed on the 2D plot of particle number density, at time $\omega_{\rm p}t = 1008$ (marked as a vertical dashed line in panel (a)); the locations of the selected positrons are shown as yellow filled or open circles, depending on their final energy (filled circles for particles with $\gamma < 3\gamma_0$, open circles for $\gamma > 6\gamma_0$).

with $\gamma\beta_z > 0$ (and electrons with $\gamma\beta_z < 0$), and the opposite holds where the reconnection field is along $-\hat{\boldsymbol{z}}$.

Accelerated by the reconnection electric field, high-energy particles stream at first along the z direction, and then, by effect of the reconnected magnetic field B_x, they are advected along y into the closest island. In our 2D simulations, where the z dimension is not resolved, the relevant timescale for acceleration is the time to drift from an X-point into the closest magnetic island. As the flow propagates downstream from the fast shock, magnetic islands grow and merge (see Fig. 4.3(a)), so that the typical distance between an X-point and the closest island increases, allowing particles to be accelerated to higher energies. This explains why the upper cutoff in the particle spectra of Fig. 4.3(f) shifts to higher energies for current sheets farther from the fast shock.[7]

Since our 2D simulations cannot resolve potential perturbations to the particle motion along z, the validity of our conclusions may be questioned, if the structure of X-points in a realistic 3D scenario is much different from the 2D picture analyzed here. In particular, in 3D we expect the current sheet to be folded, by effect of the so-called drift-kink instability (Zenitani & Hoshino 2005a, 2007). The characteristic wavelength of the kink mode in the xz plane will introduce a different length scale, that could potentially compete with the distance between X-points and islands (in the xy plane) in constraining the maximum energy of accelerated particles. In §5 we address this important point, and show that the latter constraint is more restrictive than the former, at least for the parameters explored in this chapter. It follows that our 2D simulations with in-plane fields can capture very well the main 3D properties of the shock, and in particular the physics of particle acceleration.

Although we find that the reconnection electric field at X-points is the main agent of particle acceleration, we have verified, by tracing the orbits of a large number of particles extracted directly from the simulation, that a variety of other mechanisms can play a role in particle energization. During the coalescence of magnetic islands, particles may be accelerated by the anti-reconnection electric field at the X-point located in between the two merging islands (Jaroschek et al. 2004; Oka et al. 2010). Alternatively, particle energization may be due to a kind of first-order Fermi process within magnetic islands, when particles reflect from the ends of contracting islands

[7] Another factor that drives the trend in the upper spectral cutoff of Fig. 4.3(f), though less important than the timescale argument mentioned above, is the increase in the strength of the reconnection field as the flow is compressed toward the hydrodynamic shock.

following coalescence (Drake et al. 2006). Also, the second-order Fermi process can operate between randomly-moving islands just downstream from the hydrodynamic shock. The two latter mechanisms change primarily the components of particle momentum along the plane where the magnetic field lies (the xy plane, in our simulations). On the contrary, for reconnection and anti-reconnection at X-points, it is the momentum component along the reconnection electric field (along z, in our case) that receives the greatest boost.

The energization mechanisms discussed so far are peculiar to the process of magnetic reconnection, being seeded by the field dissipation itself, or by the resulting fluid motions. In shock-dominated systems, two more acceleration processes can be relevant, whose free energy comes from the converging motion of the pre-shock and post-shock flows. In first-order Fermi acceleration (or diffusive shock acceleration, DSA), particles diffuse back and forth across the shock front and gain energy by scattering from magnetic turbulence embedded in the converging flows (e.g., Blandford & Ostriker 1978; Bell 1978; Blandford & Eichler 1987). In magnetized flows, particles may also gain energy directly from the large-scale motional electric field while they gyrate around the shock front, a process known as shock-drift acceleration (SDA; e.g., Chen & Armstrong 1975; Webb et al. 1983; Begelman & Kirk 1990). As we describe in the following section, the SDA process is most efficient in shocks where $\langle B_y \rangle_\lambda / B_0 \gtrsim 0.015$, whereas the DSA mechanism can only operate if the stripe-averaged field is small, $\langle B_y \rangle_\lambda / B_0 \lesssim 0.015$.

4. Dependence on the Wind Properties

In this section, we investigate how the physical conditions in the wind can affect the structure of the shock and the physics of magnetic reconnection, with particular focus on the particle spectrum downstream from the hydrodynamic shock. We will explore the dependence of our results on the stripe wavelength (in §4.1), the wind magnetization (in §4.2), the strength of the stripe-averaged field (in §4.3), and the wind bulk Lorentz factor (in §4.4).

The 1D model presented by Pétri & Lyubarsky (2007) predicts that magnetic dissipation should not appreciably affect the dynamics of the flow, provided that $\lambda_{c/\omega_p}/\sigma \gtrsim 8\,\xi_1/5$. Here, $\lambda_{c/\omega_p} = \lambda/(c/\omega_p)$ is the stripe wavelength in units of the relativistic skin depth in the cold wind, and $\xi_1 = 6 - 10$ is the ratio of downstream sheet thickness to downstream Larmor radius (its value is calibrated with 1D PIC simulations). In this case, the striped structure of the wind will be preserved downstream, just compressed, and the flow will satisfy the MHD jump conditions for a

strongly magnetized fluid. On the contrary, full dissipation of the alternating fields should occur if $\lambda_{c/\omega_p}/\sqrt{\sigma} \lesssim 4\,\xi_1/5$, when the thickness of downstream current sheets becomes comparable to the distance between neighboring sheets. In this case, Pétri & Lyubarsky (2007) find that the post-shock flow is purely hydrodynamical, since all the magnetic field energy has been transferred to the particles.

As we discuss below, the results of our 2D and 3D simulations suggest a different picture. We find that complete dissipation of the alternating fields by magnetic reconnection is a common by-product of the shock evolution, even in the cases when the 1D model of Pétri & Lyubarsky (2007) predicts that the stripes should be preserved downstream from the shock. As suggested by Fig. 4.3, multi-dimensional simulations are of key importance to correctly capture the evolution of the tearing-mode instability, which is ultimately responsible for field dissipation and particle acceleration.

Even though complete field dissipation (and transfer of field energy to the particles) occurs irrespective of the wind properties, the shape of the downstream particle spectrum is sensitive to the stripe wavelength and the wind magnetization. In most cases, the particle spectrum can be described as a power law in energy, with a flat slope of index $1 < p < 2$. The slope is moderately sensitive to the pre-shock parameters, but it is the width of the spectrum that depends most dramatically on the properties of the wind. If the spectrum can be modeled as a power law of index p between $\gamma_{\rm min}$ and $\gamma_{\rm max}$, we can derive a relation between the lower and upper cutoffs by requiring that the mean particle Lorentz factor be $\langle\gamma\rangle = \gamma_0(\sigma+1)$, as expected for full dissipation of magnetic energy. We obtain

$$\left(\frac{\gamma_{\rm min}}{\gamma_0}\right)^{p-1}\left(\frac{\gamma_{\rm max}}{\gamma_0}\right)^{2-p} = \frac{2-p}{p-1}\,(\sigma+1) \quad . \tag{4-2}$$

For $1 < p < 2$ and fixed σ, we see that this expression includes both the case of narrow distributions with $\gamma_{\rm min} \simeq \gamma_{\rm max} \simeq \gamma_0\sigma$, and the case of wide spectra with $\gamma_{\rm min} \simeq \gamma_0$ and $\gamma_{\rm max} \simeq \gamma_0\sigma^{1/(2-p)}$. As we show in §4.1, the stripe wavelength is the main factor, at fixed wind magnetization, that controls the transition from narrow Maxwellian-like spectra to broad power-law tails.

In Figs. 4.5-4.7, 4.9 and 4.12, we show the dependence of the downstream particle spectrum on the conditions in the wind. In most cases, the spectrum is measured between $700\,c/\omega_p$ and $300\,c/\omega_p$ downstream from the hydrodynamic shock.[8] In the subpanel of each figure, we estimate the mean

[8]The width of the slab where we compute the spectrum is calibrated such that to average out all potential variations along x, and to provide an estimate of the characteristic downstream spectrum.

particle kinetic energy $\langle\gamma\rangle/\gamma_0$ (black line) and the width of the particle spectrum, measured as the ratio $\gamma_{\rm max}/\gamma_{\rm min}$ of upper to lower energy cutoff (red line). More precisely, $\gamma_{\rm min}$ will be defined as the Lorentz factor where most of the particles lie (i.e., where $\gamma\, dN/d\gamma$ peaks), whereas $\gamma_{\rm max}$ will correspond to the location where most of the energy resides (i.e., where $\gamma^2 dN/d\gamma$ is at maximum).

4.1. Dependence on the Stripe Wavelength

Fig. 4.5 shows how the downstream particle spectrum changes for different stripe wavelengths λ, keeping fixed the magnetization $\sigma = 10$, the stripe-averaged field $\langle B_y\rangle_\lambda/B_0 \simeq 0.05$ (corresponding to $\alpha = 0.1$), and the bulk Lorentz factor $\gamma_0 = 15$. We explore a wide range of wavelengths, from $\lambda = 20\, c/\omega_{\rm p}$ up to $\lambda = 1280\, c/\omega_{\rm p}$.

Regardless of the value of λ, the structure of the flow resembles the picture presented in §3.1. A fast MHD shock propagates into the striped wind, compressing the incoming current sheets and triggering magnetic reconnection. Due to reconnection, the striped structure is eventually erased, and a hydrodynamic shock forms. Downstream from the shock, the flow is nearly unmagnetized, apart from the field component associated with the shock-compression of the stripe-averaged field $\langle B_y\rangle_\lambda$. It follows that the downstream mean kinetic energy per particle should be close to the value $\langle\gamma\rangle/\gamma_0 = \sigma + 1$ expected in the case of complete field dissipation. The value of $\langle\gamma\rangle$ measured in our simulations (black line in the subpanel of Fig. 4.5) is approximately constant with respect to λ, and is consistent with such expectation.[9] Clearly, this is in contrast with the 1D model of Pétri & Lyubarsky (2007), which would predict negligible field dissipation if $\lambda \gtrsim 130\, c/\omega_{\rm p}$ (for $\xi_1 = 8$ and $\sigma = 10$).

Even though the mean kinetic energy per particle does not appreciably vary with λ, the shape of the spectrum does change, with a clear tendency for broader spectra at longer stripe wavelengths. As shown by the red line in the subpanel of Fig. 4.5, the ratio $\gamma_{\rm max}/\gamma_{\rm min}$ increases with wavelength from $\gamma_{\rm max}/\gamma_{\rm min} \simeq 1$ (at $\lambda = 20\, c/\omega_{\rm p}$) up to $\gamma_{\rm max}/\gamma_{\rm min} \simeq 20$ (at $\lambda = 1280\, c/\omega_{\rm p}$). Longer wavelengths yield smaller values of $\gamma_{\rm min}$ (see main plot), which then imply larger values of $\gamma_{\rm max}$, given the relation in eq. (4-2).

---

[9] The trend of smaller values of $\langle\gamma\rangle/\gamma_0$ with decreasing λ (black line in the subpanel of Fig. 4.5) is due to the fact that the fractional contribution of current sheets (with fixed thickness $2\Delta \simeq 2\, c/\omega_{\rm p}$) to the upstream flow is larger for smaller λ. When averaged over one stripe wavelength, this results in a lower upstream energy per particle for smaller λ (including both magnetic and kinetic contributions), which is then reflected in the post-shock value of $\langle\gamma\rangle/\gamma_0$.

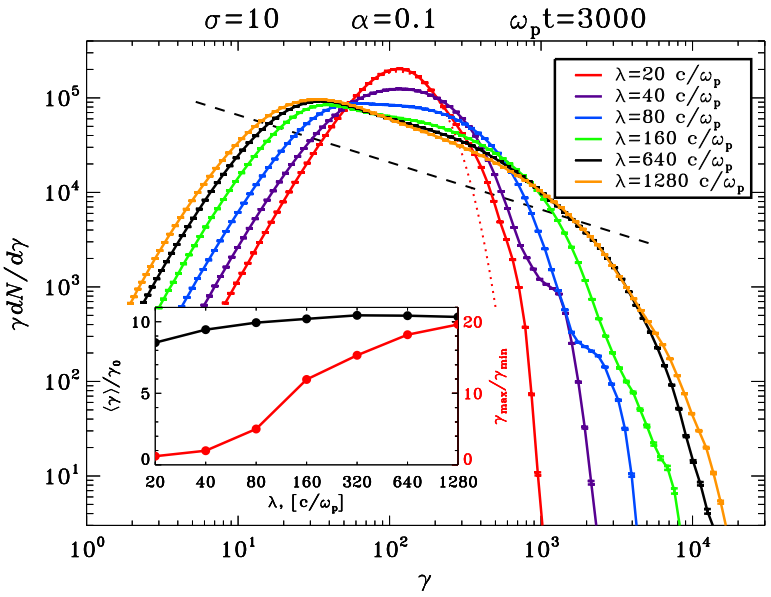


Fig. 4.5.— Downstream particle spectrum at $\omega_{\rm p}t = 3000$ for different stripe wavelengths λ, in a flow with $\sigma = 10$ and $\alpha = 0.1$. We vary the stripe wavelength from $\lambda = 20\,c/\omega_{\rm p}$ up to $\lambda = 1280\,c/\omega_{\rm p}$. The dotted red line is a Maxwellian with the same average energy as the spectrum colored in red (which refers to $\lambda = 20\,c/\omega_{\rm p}$); the dashed black line indicates a power-law distribution with slope $p = 1.5$. In the subpanel, the black line shows the average Lorentz factor of downstream particles $\langle\gamma\rangle$ as a function of the stripe wavelength λ (axis on the left side). In the same subpanel, the red line presents the dependence of $\gamma_{\rm max}/\gamma_{\rm min}$ on λ (axis on the right side). Here, $\gamma_{\rm min}$ is defined as the location where $\gamma dN/d\gamma$ peaks (i.e., where most of the particles reside), whereas $\gamma_{\rm max}$ is the Lorentz factor where $\gamma^2 dN/d\gamma$ peaks (i.e., where most of the energy lies).

This trend can be easily understood by considering the structure of the flow for different stripe wavelengths. As described in §3.1, the hydrodynamic shock is located at the point where reconnection islands grow large enough to fill the entire space in between neighboring current sheets. Since the distance between sheets is proportional to λ, the same scaling should hold for the size of reconnection islands, just upstream of the hydrodynamic shock.[10] Shorter wavelengths will then result in more numerous islands of smaller size, whereas fewer but bigger islands will be present for longer λ (see Fig. 4.3). Since an X-point exists in between each pair of neighboring islands (belonging to the same current sheet), the number of X-points per unit length (along the current sheet) will be larger for smaller λ. For short wavelengths, most of the incoming particles will likely pass in the vicinity of one of the numerous X-points, as the flow crosses the hydrodynamic shock. The energy evolution of one particle will then be similar to that of any other particle, with its Lorentz factor increasing from $\gamma \simeq \gamma_0$ up to $\gamma \simeq \gamma_0(\sigma + 1)$, as the particle gains energy from the annihilating fields. This explains why for short stripe wavelengths ($\lambda = 20\, c/\omega_{\rm p}$, red curve in Fig. 4.5) the particle spectrum is so narrow, and similar to a Maxwellian distribution (red dotted curve). In summary, for short stripe wavelengths, all particles are equally close to an X-point, so they gain comparable amounts of energy.

On the other hand, for long wavelengths, the energy evolution of different particles can be extremely diverse, as discussed in §3.2. Particles that stay away from X-points are likely to remain cold, and retain the energy $\gamma \simeq \gamma_0$ they started with. On the contrary, particles that pass through an X-point can gain a significant amount of energy. With increasing stripe wavelength, more particles will belong to the former group, and fewer to the latter, since the number of X-points decreases. This explains why the peaks in the spectra of Fig. 4.5, which track the downstream Lorentz factor of the "typical" particle, shift to lower energies with increasing stripe wavelength. For the same reason, the best-fitting power law becomes softer, with spectral slope increasing from $p \simeq 1.0$ to $p \simeq 1.5$ (indicated as a black dashed line in Fig. 4.5). In response to a lower $\gamma_{\rm min}$ and a steeper slope, the upper cutoff of the spectrum shifts to higher energies, as predicted by eq. (4-2). This is just a consequence of the longer time available for acceleration, at larger stripe wavelengths. As discussed in §3.2, particles injected at X-points are continuously accelerated by the reconnection electric field, while they drift from the X-point into the closest island. Since the spacing between

[10] In the remaining of this section, by "reconnection islands" we will be referring only to the islands in the current sheet just ahead of the hydrodynamic shock, where the shape of the downstream particle spectrum is established, as demonstrated by Fig. 4.3(f).

neighboring islands (in the same current sheet) increases with stripe wavelength, the injected particles will be able to stay in the acceleration region for longer times (and so, to reach higher energies), if λ is larger.

As we discuss in §4.2, the threshold between short and long wavelengths depends on the wind magnetization. Here, by "short λ" we generically mean all cases in which the spectrum is narrow, similar to a Maxwellian. In contrast, the cases with "long λ" have broad power-law spectra with slope $1 < p < 2$. We find that the condition proposed for full dissipation by Pétri & Lyubarsky (2007), namely that $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma} \lesssim 4\,\xi_1/5$, works fairly well to discriminate between short and long wavelengths. For $\xi_1 = 8$, such criterion would predict Maxwellian-like spectra for $\lambda \lesssim 20\,c/\omega_{\rm p}$ (if $\sigma = 10$), $\lambda \lesssim 45\,c/\omega_{\rm p}$ (if $\sigma = 50$), and $\lambda \lesssim 65\,c/\omega_{\rm p}$ (if $\sigma = 100$), which is indeed what we observe in our simulations. In retrospect, this is not surprising. In fact, the condition $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma} \lesssim 4\,\xi_1/5$ comes from requiring that the downstream sheet thickness be larger than the separation of neighboring sheets. In this limit, no distinction should persist between the plasma within (versus outside) the current sheets, i.e., all particles should share the same energy evolution. This obviously results in a Maxwellian distribution.

Finally, we point out that the minor bump emerging at high energies in the spectra of Fig. 4.5 (most notably for wavelengths $\lambda \lesssim 160\,c/\omega_{\rm p}$) is populated by particles that are energized via the shock-drift mechanism (SDA) at the hydrodynamic shock. Such particles gain energy from the stripe-averaged motional electric field $\langle E_z\rangle_\lambda \simeq \langle B_y\rangle_\lambda$, while they gyrate around the shock. Since $\langle E_z\rangle_\lambda \geq 0$ for our simulation setup, shock-drift accelerated positrons should preferentially have $\gamma\beta_{\rm z} > 0$, whereas the opposite should hold for electrons. The few high-$\gamma\beta_{\rm z}$ positrons seen at $x \simeq 1100\,c/\omega_{\rm p}$ in panel (f) of Fig. 4.2 are indeed accelerated via the SDA mechanism. They gyrate around the shock a few times, before being advected downstream by the stripe-averaged magnetic field. Given the limited number of acceleration cycles, the SDA component in the spectra of Fig. 4.5 does not extend in time to higher energies. So, SDA is not a promising candidate to produce broad power-law tails.

The efficiency of injection into the SDA process is higher for shorter stripe wavelengths, as Fig. 4.5 suggests. As seen from the pre-shock frame, a particle coming from the downstream side can be accelerated via SDA in the upstream stripe-averaged field only if it can sample the full wavelength of the striped wind during its upstream residence time. Since particles returning upstream are caught up by the shock after having travelled a distance $r_{L,{\rm u}}/\gamma_0$ (here, $r_{L,{\rm u}}$ is

the particle Larmor radius measured in the upstream frame), we require $\gamma_0(\lambda/2) \lesssim r_{L,\rm u}/\gamma_0$ for efficient SDA. If $\gamma = \xi_2\langle\gamma\rangle \simeq \xi_2\gamma_0\sigma$ is the Lorentz factor of particles at the high-energy end of the distribution ($\xi_2 = 1 - 10$, depending on λ), the previous condition can be rewritten as $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma} \lesssim 2\,\xi_2$. For fixed magnetization, the criterion for efficient SDA is more easily satisfied at short stripe wavelengths, which explains why the normalization of the high-energy bump in Fig. 4.5 increases for smaller λ.

We remark that, apart from the factor ξ_2, which could in principle depend (though weakly) on λ and σ, the criterion for efficient SDA involves the same combination of wind magnetization and stripe wavelength (namely, $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma}$) as the condition discussed at the beginning of this section, which regulates how closely the particle spectrum resembles a Maxwellian distribution (see also Pétri & Lyubarsky 2007). Since $\lambda_{c/\omega_{\rm p}}$ is the wind wavelength normalized to the skin depth of the *pre-shock* fluid, and since the average particle energy increases by a factor of $\simeq \sigma$ across the shock as a result of field dissipation, then $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma}$ is just the wind wavelength measured in units of the *post-shock* plasma skin depth (apart from factors of order unity). In §4.2, we further comment on the importance on the parameter $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma}$ in regulating the physics of the shock and the shape of the downstream particle spectrum.

4.2. Dependence on the Wind Magnetization

In this section, we investigate the dependence of our results on the magnetization of the wind. Fig. 4.6 shows how the downstream spectrum changes with σ, keeping fixed the stripe wavelength $\lambda = 160\,c/\omega_{\rm p}$, the stripe-averaged field $\langle B_y\rangle_\lambda \simeq 0.05$ (corresponding to $\alpha = 0.1$), and the bulk Lorentz factor $\gamma_0 = 15$. We explore a wide range of magnetizations, from $\sigma = 10$ up to $\sigma = 100$.

We find that efficient annihilation of magnetic fields by shock-driven reconnection, and transfer of magnetic energy to the particles, occur irrespective of the wind magnetization. As shown by the black line in the subpanel of Fig. 4.6, the average Lorentz factor of downstream particles approaches in all cases the value $\langle\gamma\rangle \simeq \gamma_0\sigma$ expected for complete field dissipation. As the wind magnetization increases, the spectrum shifts to higher Lorentz factors (just because $\langle\gamma\rangle \propto \sigma$) and it changes in shape, with the part at low energies getting de-populated at the expense of the high-energy component. As a result, the particle spectrum, which could be fitted as a broad power law of index $p \simeq 1.2$ for $\sigma = 10$ (blue curve), approaches a Maxwellian-like distribution for $\sigma = 100$ (red curve).

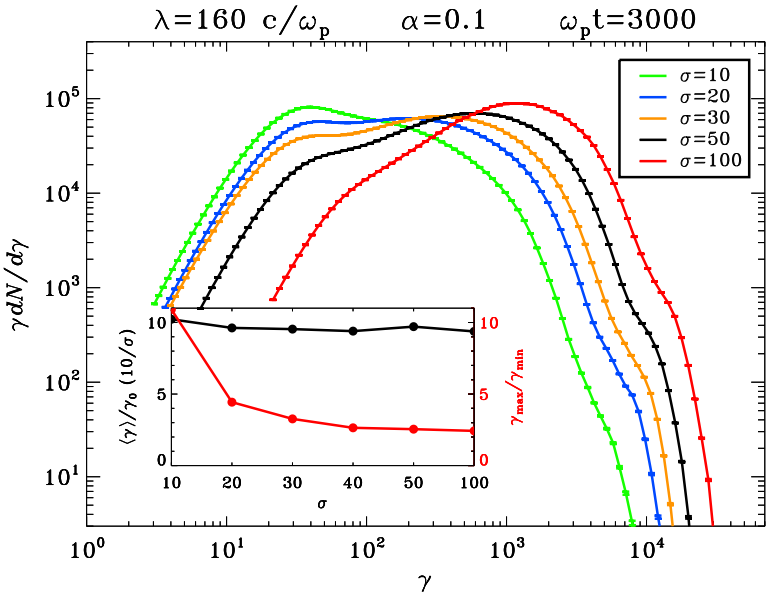


Fig. 4.6.— Downstream particle spectrum at $\omega_{\rm p}t = 3000$ for different magnetizations σ, in a flow with $\lambda = 160\,c/\omega_{\rm p}$ and $\alpha = 0.1$. We vary the wind magnetization from $\sigma = 10$ to $\sigma = 100$. In the subpanel, the black line shows the average Lorentz factor $\langle\gamma\rangle$ of downstream particles as a function of σ (axis on the left side). The value of $\langle\gamma\rangle$ is normalized to the total energy per particle (kinetic + electromagnetic) in the pre-shock flow, i.e., $\gamma_0\sigma$. In the same subpanel, the red line presents the dependence of the spectral width $\gamma_{\rm max}/\gamma_{\rm min}$ on σ (axis on the right side).

As a consequence, the ratio $\gamma_{\rm max}/\gamma_{\rm min}$, a proxy for the width of the spectrum, becomes smaller for increasing σ (red line in the subpanel of Fig. 4.6, with axis on the right).

This is the same trend observed when decreasing λ at fixed σ, as discussed in §4.1. In the limit of very high magnetizations, for fixed wavelength (or very short λ, for fixed σ), the downstream energy spectrum approaches a Maxwellian distribution, when $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma} \lesssim 4\,\xi_1/5$ (here, $\xi_1 = 6 - 10$). For $\lambda = 20\,c/\omega_{\rm p}$, $40\,c/\omega_{\rm p}$, and $80\,c/\omega_{\rm p}$, we have verified that the value of σ above which the spectrum resembles a Maxwellian is in good agreement with this criterion. For $\lambda = 160\,c/\omega_{\rm p}$, we would expect a Maxwellian distribution only for $\sigma \gtrsim 300$ (if $\xi_1 = 8$), beyond the range of magnetizations we explore. However, the trend toward a single-component distribution (as σ increases) is already clear within the limited range of magnetizations investigated in Fig. 4.6.

The dominance of the high-energy component at large magnetizations can be understood in terms of the basic properties of the flow. In experiments with fixed σ and different λ (see

§4.1), we have related the shape of the spectrum to the size of magnetic islands just ahead of the hydrodynamic shock, that scales with stripe wavelength. Here λ is fixed, and the change in spectral shape with magnetization should instead be attributed to the different reconnection rate, namely the velocity $\beta_{\rm rec}$ at which cold plasma flows into the X-points. Zenitani et al. (2009) have measured the reconnection rate in high-σ pair plasmas via 2D two-fluid MHD simulations of undriven reconnection, finding that the inflow velocity scales roughly as $\beta_{\rm rec} \simeq 0.07\sqrt{\sigma}$ in the range $\sigma = 10 - 100$. We have confirmed this scaling in our simulations, by measuring as a function of magnetization (at fixed λ) the distance between the fast and the hydrodynamic shock. Everything else being fixed, this distance should be inversely proportional to the reconnection rate, and in the range of magnetizations explored in our study we find that the reconnection speed scales as $\beta_{\rm rec} \propto \sqrt{\sigma}$.

As the reconnection rate increases with magnetization, a larger number of particles can be advected into a given X-point per unit time, which results in more particles experiencing the reconnection electric field. In other words, at larger σ it is more likely for the "typical" particle to enter an X-point and gain energy from field dissipation. This explains why for high magnetizations the spectra in Fig. 4.6 are dominated by a single component of hot particles accelerated by the reconnection field. A similar argument explains why, for large values of σ, the spectrum does not extend much beyond the average Lorentz factor $\langle\gamma\rangle$. In response to the increase of inflow speed with σ, also the outflow velocity from a given X-point into the closest island increases with magnetization. For fixed λ, which corresponds to a fixed spacing between X-points and magnetic islands, the time available for acceleration by the reconnection field will be shorter, for larger σ. So, the maximum energy at which particles can be accelerated will be lower (relative to the average Lorentz factor $\langle\gamma\rangle$), for higher magnetizations.

Finally, we notice that the minor bump emerging at high energies from the spectra of Fig. 4.6, which is populated by particles accelerated via the SDA mechanism (see §4.1), becomes more prominent for larger magnetizations. This is in good agreement with the criterion $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma} \lesssim 2\,\xi_2$ discussed at the end of §4.1, which dictates that injection into the SDA process should be easier for higher σ (at fixed λ).

In summary, we find that both the physics of shock-driven reconnection, which ultimately governs the shape of the particle spectrum, and the efficiency of SDA are controlled by the same parameter $\lambda_{c/\omega_{\rm p}}/\sqrt{\sigma}$, namely the stripe wavelength measured in units of the *post-shock* skin depth

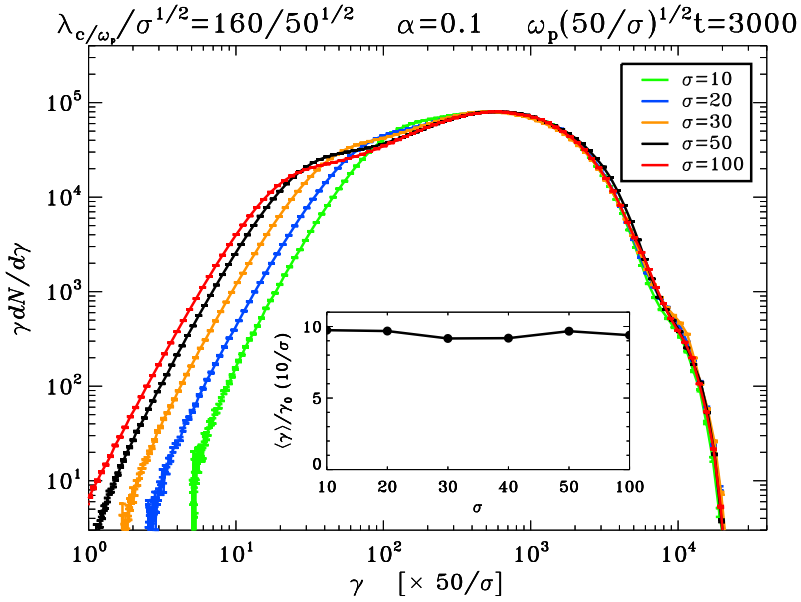


Fig. 4.7.— Downstream particle spectrum for different values of λ and σ, but such that the ratio $\lambda_{c/\omega_p}/\sqrt{\sigma} = 160/\sqrt{50} \simeq 22.5$ is fixed (as well as $\alpha = 0.1$). Here, λ_{c/ω_p} is the wind wavelength in units of the skin depth in the *pre-shock* fluid, so $\lambda_{c/\omega_p}/\sqrt{\sigma}$ is roughly the wavelength normalized to the skin depth of the *post-shock* fluid. The comparison is performed at the same time, in units of the relativistic plasma frequency of the *post-shock* flow ($\simeq \omega_p/\sqrt{\sigma}$). Spectra are shifted along the x-axis by $50/\sigma$ to facilitate comparison with the case $\sigma = 50$. The black line in the subpanel shows the average downstream Lorentz factor, as in Fig. 4.6.

($\simeq \sqrt{\sigma}\, c/\omega_p$). We expect that the main properties of the shock, and in particular the high-energy end of the downstream particle spectrum, should be relatively insensitive to variations in λ or σ, provided that the ratio $\lambda_{c/\omega_p}/\sqrt{\sigma}$ is kept constant.

In Fig. 4.7 we investigate this prediction by comparing the downstream spectrum among flows that have the same value of $\lambda_{c/\omega_p}/\sqrt{\sigma} = 160/\sqrt{50} \simeq 22.5$. We vary the wind magnetization from $\sigma = 10$ up to $\sigma = 100$, or equivalently the stripe wavelength from $\lambda = 72\, c/\omega_p$ up to $\lambda = 226\, c/\omega_p$. The spectra in Fig. 4.7 are shifted along the x-axis by $50/\sigma$, to facilitate comparison with the reference case $\sigma = 50$ (black curve). This compensates for the overall shift to higher Lorentz factors due to the increase in $\langle\gamma\rangle \propto \sigma$, as expected for complete field dissipation.[11] We confirm that the

[11]Strictly speaking, the comparison in Fig. 4.7 is performed at the same $\omega_p t/\sqrt{\sigma}$, namely at the same time in units of the *post-shock* plasma frequency ($\simeq \omega_p/\sqrt{\sigma}$). However, we stress once again that the downstream spectrum does not appreciably evolve in time.

shape of the spectrum, at the high-energy end, is controlled by the single parameter $\lambda_{c/\omega_p}/\sqrt{\sigma}$. For a fixed value of this particular combination of λ and σ, both the high-energy component of the spectrum and the bump of shock-drift accelerated particles are insensitive to variations in either λ or σ.

The shift of the low-energy part to higher Lorentz factors, as σ decreases, is just an artificial consequence of our re-scaling in the x-axis of Fig. 4.7. In reality, the low-energy cutoff stays fixed at some multiple of the pre-shock bulk Lorentz factor γ_0, irrespective of λ or σ, provided $\lambda_{c/\omega_p}/\sqrt{\sigma}$ is kept constant. With decreasing λ, the number of X-points per unit length increases as $\propto \lambda^{-1}$, which would make it easier for the "tyipical" particle to gain energy from the dissipating fields (everything else being equal). However, for lower σ, the inflow velocity into a given X-point is also reduced, as $\propto \sqrt{\sigma}$. Evidently, for constant $\lambda_{c/\omega_p}/\sqrt{\sigma}$, the two opposite effects compensate, resulting in a fixed low-energy spectral cutoff.

4.3. Dependence on the Stripe-Averaged Field

In this section, we explore the dependence of the shock properties on the strength of the stripe-averaged field $\langle B_y \rangle_\lambda$. The magnitude B_0 of the field in the cold wind is fixed by the value of the magnetization σ. A non-zero stripe-averaged field is established if the portion of each stripe with $B_y = +B_0$ has a different width than the region with $B_y = -B_0$. If λ_+ is the length of the former and λ_- of the latter (with the constraint $\lambda_+ + \lambda_- = \lambda$), then we have $\langle B_y \rangle_\lambda / B_0 = (\lambda_+ - \lambda_-)/\lambda$. Alternatively, by defining $\alpha \equiv (\lambda_+ - \lambda_-)/\lambda_+$, we have $\langle B_y \rangle_\lambda / B_0 = \alpha/(2 - |\alpha|)$. In pulsar winds, $|\alpha| \to 1$ at latitudes approaching the inclination angle between magnetic and rotational axes of the pulsar (see the upper panel of Fig. 4.1). The special case $\alpha = 0$, appropriate for the wind midplane, will be discussed in more detail at the end of this section.

The results presented in the previous sections referred to shocks with a moderate value of $\alpha = 0.1$. As a representative case of high-α flows, Fig. 4.8 shows the shock structure for $\alpha = 0.75$, corresponding to $\langle B_y \rangle_\lambda / B_0 \simeq 0.6$. The stripe wavelength $\lambda = 320\, c/\omega_p$ and the wind magnetization $\sigma = 10$ are chosen to facilitate comparison with the results discussed in §3.1. As in the case $\alpha = 0.1$, the pre-shock flow interacts at first with a fast MHD shock (at $x \simeq 2170\, c/\omega_p$ in Fig. 4.8, indicated with a vertical dotted blue line), that triggers magnetic reconnection. With increasing α (and so, with increasing $\langle B_y \rangle_\lambda$), a smaller fraction of the incoming Poynting flux is in the form of alternating fields, and as such available for dissipation. For $\alpha > 0$ (which implies $\lambda_+ > \lambda_-$), reconnection

behind the fast shock will eventually annihilate the field where $B_y = -B_0$, but a finite region where the field stays $B_y = +B_0$ will survive. Dissipation of the alternating component is completed when the size of reconnection islands approaches λ_- (or, in general, the smallest between λ_+ and λ_-), as shown by the 2D plots of density and magnetic energy in Fig. 4.8(a) and (b), respectively. At this point, a second shock forms (at $x \simeq 2050\, c/\omega_{\rm p}$ in Fig. 4.8), which would correspond in the regime of small α (i.e., $\alpha \lesssim 0.1$) to the hydrodynamic shock discussed in §3.1. For larger α, the speed of this second shock increases, from $\beta_{\rm sh} \simeq 1/3$, as appropriate for an unmagnetized post-shock plasma, up to $\beta_{\rm sh} \to 1$, the shock velocity in a $\sigma \gg 1$ flow. In the limit $|\alpha| \to 1$, the shock that in the regime $\alpha \lesssim 0.1$ was identified as "hydrodynamic" becomes eventually degenerate with the fast MHD shock. In the case $\alpha = 0.75$ shown in Fig. 4.8, the fast shock stays at a distance of $\simeq 100\, c/\omega_{\rm p}$ ahead of the main shock, and they both move at $\beta_{\rm sh} \simeq 0.6$.

Downstream from the main shock ($x \lesssim 2050\, c/\omega_{\rm p}$), the residual magnetic flux that has not been annihilated on the way from the fast shock survives within stripes of high magnetic field, separated by regions of relatively low magnetic content (see the 2D and 1D plots of magnetic energy in Fig. 4.8(b) and (d), respectively; and the 1D profile of B_y, black line in Fig. 4.8(e)). The plasma inside the high-field stripes is relatively cold, not having gained much energy from field dissipation. In contrast, the low-field regions are populated with hot particles energized by the reconnection field, as shown in the 1D profile of mean kinetic energy per particle of Fig. 4.8(f) and in the longitudinal positron phase space of Fig. 4.8(g). We find that the post-shock flow retains a striped structure, with magnetically-dominated regions separated by relatively wide sheets of hot plasma, for values of $|\alpha| \gtrsim 0.3$ (or equivalently, for $|\langle B_y \rangle_\lambda|/B_0 \gtrsim 0.15$). As $|\alpha|$ increases, the longitudinal extent of the high-field stripes grows, at the expense of the regions of hot dense plasma, and magnetic pressure dominates over particle pressure in the energetics of the downstream fluid.

The striped structure of the downstream flow leaves a clear imprint on the post-shock particle spectrum, as Fig. 4.9 shows. There, we present how the particle distribution changes when the parameter α is varied between $\alpha = 0.0$ and $\alpha = 0.95$, for fixed wavelength $\lambda = 320\, c/\omega_{\rm p}$, magnetization $\sigma = 10$, and bulk Lorentz factor $\gamma_0 = 15$.

For $\alpha \lesssim 0.3$, the average particle Lorentz factor downstream from the hydrodynamic shock approaches $\langle\gamma\rangle/\gamma_0 \simeq \sigma + 1 \simeq 11$, suggesting full dissipation of magnetic fields (black line in the subpanel of Fig. 4.9). The shape of the spectrum is independent of α, in the regime $\alpha \lesssim 0.1$. The spectra for $\alpha = 0.0$ (red curve) and $\alpha = 0.1$ (black curve) differ only at the upper end, where the

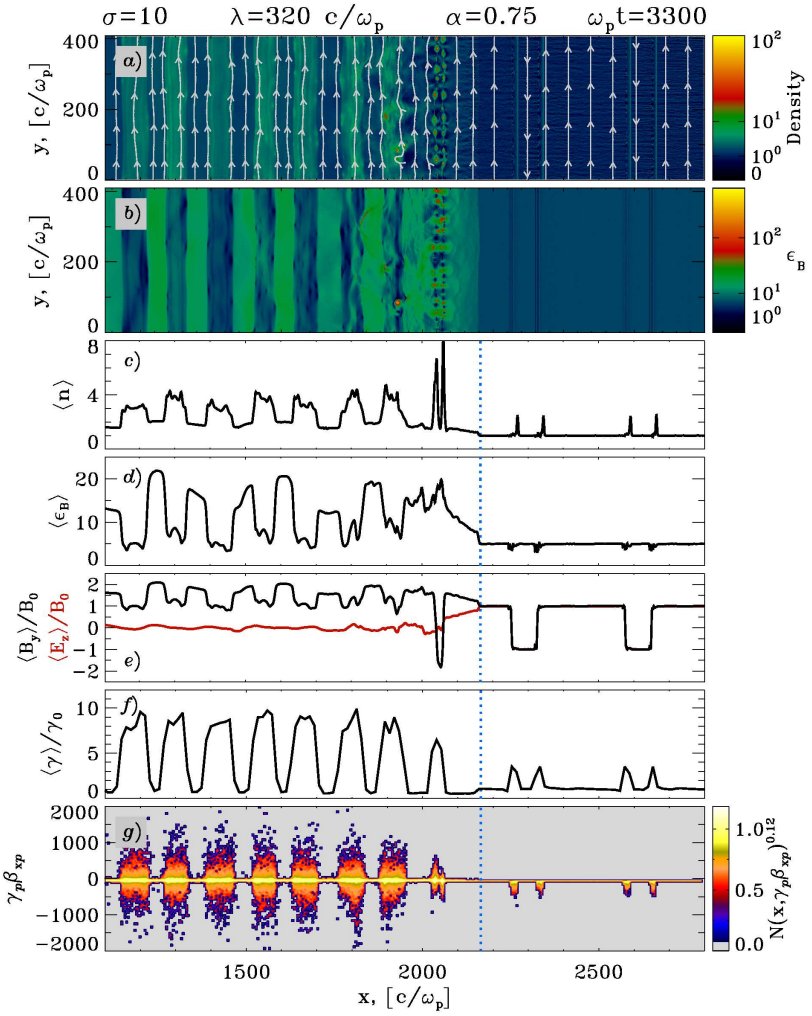


Fig. 4.8.— Internal structure of the flow at $\omega_{\rm p}t = 3300$, for stripe wavelength $\lambda = 320\,c/\omega_{\rm p}$ and magnetization $\sigma = 10$, zooming in on a region around the shock. The stripe-averaged field is $\langle B_y\rangle_\lambda/B_0 \simeq 0.6$, corresponding to $\alpha = 0.75$. The hydrodynamic shock is located at $x \simeq 2050\,c/\omega_{\rm p}$, and the location of the fast MHD shock ($x \simeq 2170\,c/\omega_{\rm p}$) is indicated with a vertical dotted blue line. The following quantities are plotted: (a) 2D plot in the simulation plane of the particle number density, in units of the upstream value, with contours showing the magnetic field lines; (b) 2D plot of the magnetic energy fraction $\epsilon_B \equiv B^2/8\pi\gamma_0 m n_{c0} c^2$; (c) y-averaged particle number density; (d) y-averaged magnetic energy fraction; (e) y-averaged magnetic field B_y (black line) and electric field E_z (red line), normalized to the upstream magnetic field B_0; (f) mean kinetic energy per particle, in units of the bulk energy at injection; (g) $x - \gamma\beta_{\rm x}$ positron phase space.

case $\alpha = 0.1$ shows the high-energy bump of shock-drift accelerated particles. This component disappears for $\alpha = 0.0$, when the stripe-averaged field $\langle E_z \rangle_\lambda = \beta_0 \langle B_y \rangle_\lambda$ powering the SDA process vanishes.

For $\alpha \gtrsim 0.3$, the post-shock particle spectrum consists of two components. The low-energy peak comes from cold plasma residing in the high-field stripes described above, whereas the high-energy part is populated by hot particles that gained energy from field dissipation. As α increases, the fraction of upstream Poynting flux available for dissipation decreases, which explains why the high-energy component in the spectra of Fig. 4.9 gets de-populated, at the expense of the low-energy part. As a consequence, the average particle energy downstream from the shock monotonically drops from $\langle \gamma \rangle \simeq \gamma_0(\sigma + 1)$, corresponding to the case of full dissipation, down to $\langle \gamma \rangle \simeq \gamma_0$, as expected for negligible field annihilation (black line in the subpanel of Fig. 4.9). In the limit $|\alpha| \to 1$ (yellow curve for $\alpha = 0.95$), most of the downstream plasma is contained within the high-field regions, and the low-energy component of the spectrum approaches the result expected for an unstriped wind (purple line).[12]

With increasing α, the high-energy component of the downstream spectrum recedes to lower energies (see the trend in Fig. 4.9, from the blue line for $\alpha = 0.45$, to the yellow line for $\alpha = 0.95$). As described in §4.1, the maximum energy of particles accelerated at X-points depends primarily on the amount of time available for acceleration, which is set by the characteristic distance between a given X-point and the closest island (in the same current sheet). In turn, this scales with the maximum size of reconnection islands, which is determined by the smallest between λ_+ and λ_- (here $\lambda_- < \lambda_+$, since $\alpha > 0$). For larger α at fixed λ, the extent of λ_- decreases, so the maximum energy to which particles can be accelerated by the reconnection electric field decreases. This explains the trend in the upper cutoff of the spectra in Fig. 4.9. Given that the location of the high-energy cutoff is set by the value of λ_- (or equivalently, by the maximum size of reconnection islands), one would expect that simulations with the same λ_-, but different choices of λ and $\alpha > 0$, should yield similar spectra, at the high-energy end. We have tested this, and we find that flows with the same λ_- show magnetic islands of similar sizes, and their downstream spectra almost overlap at high energies. Of course, this scaling does not hold for the low-energy spectral

[12]We remark that the low-energy peak in the spectrum of high-α flows, including the limit of unstriped wind, is not compatible with a Maxwellian (yellow and purple lines in Fig. 4.9). Here, the particle distribution is a ring in momentum space, lying in a plane orthogonal to the field. The location of the low-energy spectral peak scales linearly with γ_0, in agreement with the jump conditions for a highly magnetized fluid.

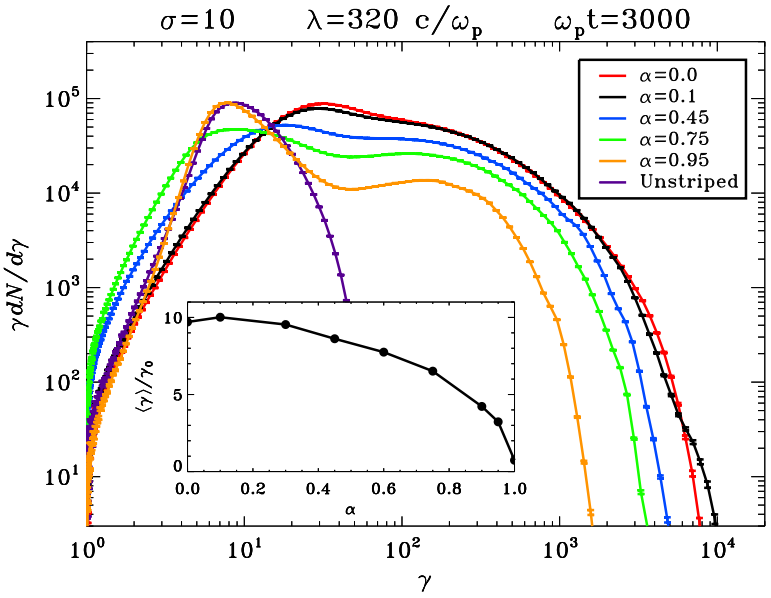


Fig. 4.9.— Downstream particle spectrum at $\omega_{\rm p}t = 3000$ for different values of the stripe-averaged field $\langle B_y \rangle_\lambda$ (or equivalently, of the parameter α), in a flow with $\lambda = 320\, c/\omega_{\rm p}$ and $\sigma = 10$. We vary α from $\alpha = 0$, when the regions with $B_y = +B_0$ and $B_y = -B_0$ have the same width $\lambda_+ = \lambda_- = \lambda/2$, to the limit $\alpha \to 1$, when $\lambda_- \ll \lambda_+ \simeq \lambda$. The limit of an unstriped wind is shown for reference as a purple line. The black line in the subpanel shows the average downstream Lorentz factor as a function of α (with $\alpha = 1.0$ referring to the unstriped wind).

component, since it is populated by particles that were initially in the region with $B_y = +B_0$, and are therefore insensitive to the value of λ_-.

The Case of Zero Stripe-Averaged Field

The special case $\langle B_y \rangle_\lambda = 0$ (or equivalently, $\alpha = 0$), which is relevant for the equatorial plane of pulsar winds, is worth a deeper investigation. As anticipated in §4.1, hot downstream particles can propagate back from the shock into the upstream, and cross at least one wavelength of the striped wind, provided that $\lambda_{c/\omega_p}/\sqrt{\sigma} \lesssim 2\,\xi_2$, where $\xi_2 = 1 - 10$ is a numerical factor that depends weakly on λ and σ. For wind parameters that satisfy such criterion, the motion of hot particles moving into the upstream (the "returning" particles, from now on) will be regulated, to zeroth order, by the stripe-averaged fields $\langle B_y \rangle_\lambda$ and $\langle E_z \rangle_\lambda = \beta_0 \langle B_y \rangle_\lambda$. For $\alpha \gtrsim 0.03$ (or equivalently, $\langle B_y \rangle_\lambda / B_0 \gtrsim 0.015$), the returning particles will be quickly advected downstream by the stripe-averaged fields, after being accelerated via the SDA mechanism discussed in §4.1. On the other hand, for $\alpha \lesssim 0.03$ they can propagate as if the upstream flow were unmagnetized. The special case $\alpha = 0.0$ is taken here as a representative example of the shock structure for $\alpha \lesssim 0.03$.

Fig. 4.10 compares the structure of the flow between the cases $\alpha = 0.1$ (left panels) and $\alpha = 0.0$ (right panels), for the same wavelength $\lambda = 80\,c/\omega_p$ and magnetization $\sigma = 10$. The positron phase space of the case $\alpha = 0.0$ (panel (c), right column) shows a diffuse population of hot particles moving back from the shock into the upstream ($x \gtrsim 1050\,c/\omega_p$), which were absent in the case $\alpha = 0.1$ (panel (c), left column). For both $\alpha = 0.0$ and $\alpha = 0.1$, as the flow propagates from the fast shock (at $x \simeq 1050\,c/\omega_p$) to the hydrodynamic shock (at $x \simeq 1000\,c/\omega_p$), magnetic reconnection transfers energy from the field to the particles, which in the downstream region end up with an average Lorentz factor $\langle \gamma \rangle \simeq \gamma_0(\sigma + 1)$ (see §4.1). In other words, the physics of magnetic reconnection is not sensitive to the value of α, in the regime $\alpha \lesssim 0.1$. However, it is only in the case $\alpha = 0.0$ (or $\alpha \lesssim 0.03$, generally speaking), when the upstream stripe-averaged fields are small, that the hot particles resulting from shock-driven reconnection can stream for large distances ahead of the shock, as observed in the phase space of Fig. 4.10(c) (right panel).

When a sufficient number of returning particles builds up in the pre-shock region, their counter-streaming with respect to the incoming flow can trigger the filamentation (or Weibel) instability (Weibel 1959; Medvedev & Loeb 1999; Gruzinov & Waxman 1999). In fact, the filamentary structures seen ahead of the shock in the 2D plot of B_z (panel (d), right column) are

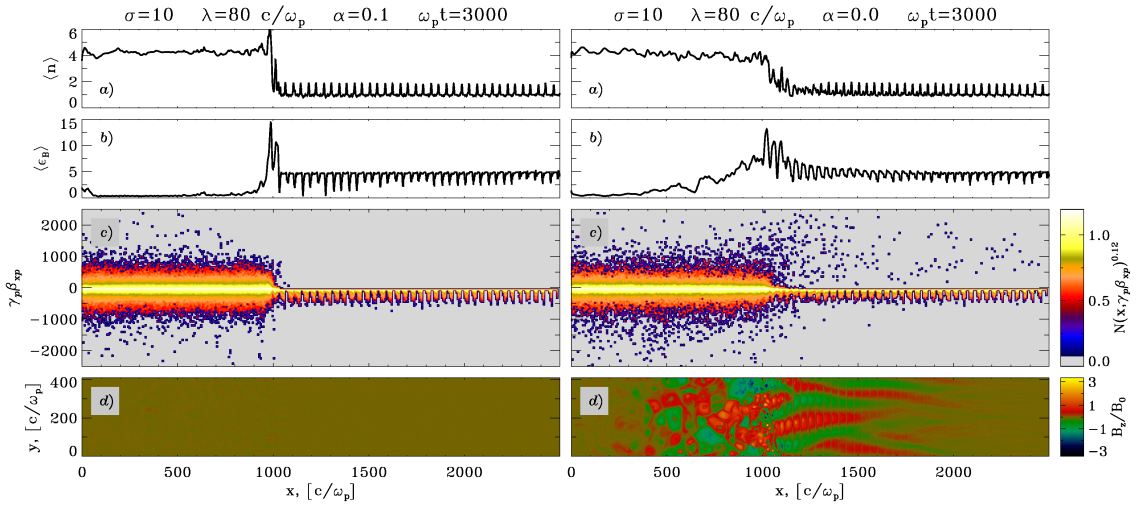


Fig. 4.10.— Comparison of the shock internal structure, for a flow with $\lambda = 80\, c/\omega_{\rm p}$ and $\sigma = 10$, between the case $\alpha = 0.1$ (left panels) and the case $\alpha = 0.0$ (right panels, corresponding to $\langle B_y \rangle_\lambda = 0$). The comparison is performed at $\omega_{\rm p} t = 3000$. In both cases, the shock is located at $x \simeq 1000\, c/\omega_{\rm p}$. As a function of the longitudinal coordinate x, the following quantities are plotted: (a) y-averaged profile of the particle density n, in units of the pre-shock value; (b) y-averaged profile of the magnetic energy fraction ϵ_B; (c) $x - \gamma\beta_{\rm x}$ positron phase space; (d) 2D plot on the simulation plane of the out-of-plane component B_z of the magnetic field.

generated by the Weibel instability. A similar pattern appears in the 2D plots of electric field E_y $(= -\beta_0 B_z)$, particle density, and magnetic energy (not shown). In contrast, no signatures of filamentation are seen for $\alpha = 0.1$ (panel (d), left column), since in this case no particles can propagate far enough ahead of the shock to seed the Weibel instability.

For both $\alpha = 0.0$ and $\alpha = 0.1$, the hydrodynamic shock satisfies the jump conditions appropriate for an unmagnetized 3D relativistic fluid. However, the filamentation instability appreciably changes the structure of the flow for $\alpha = 0.0$, especially at late times. The density profile in the transition region (panel (a)) is much smoother in the case $\alpha = 0.0$ (right column) than it is for $\alpha = 0.1$ (left column). More importantly, the magnetic fields generated ahead of the shock by the Weibel instability survive in the downstream region of the $\alpha = 0.0$ shock, as shown by the 1D profile of magnetic energy in panel (b) (right column, at $x \lesssim 1000\, c/\omega_{\rm p}$). Instead, for $\alpha = 0.1$, the only magnetic field left behind the shock comes from compression of the pre-shock stripe-averaged field, which results in a much lower level of post-shock magnetic energy (see the left panel of Fig. 4.10(b)).

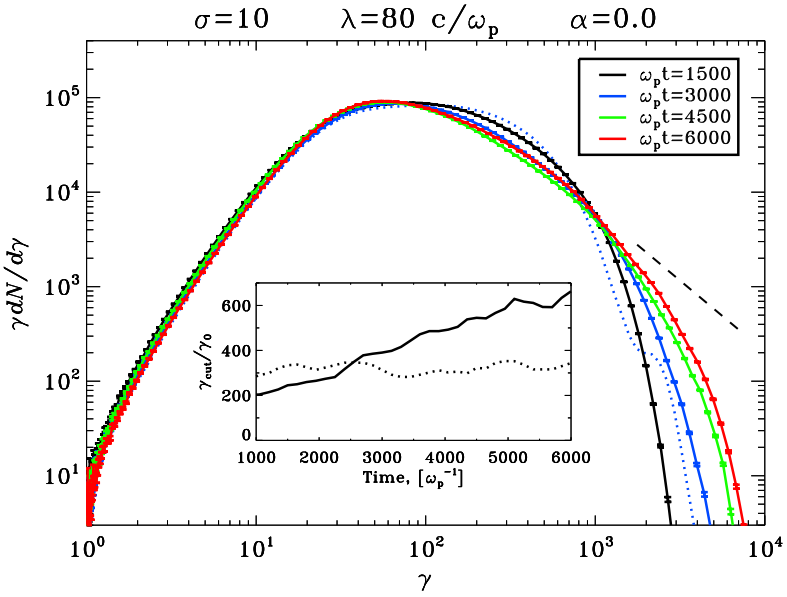


Fig. 4.11.— Time evolution of the downstream spectrum for a shock propagating in a flow with $\lambda = 80\, c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0$ (i.e., zero stripe-averaged field), from $\omega_{\rm p}t = 1500$ up to $\omega_{\rm p}t = 6000$. For comparison, the downstream spectrum for a wind with the same parameters except for $\alpha = 0.1$ is shown as a dotted blue line for $\omega_{\rm p}t = 3000$ (but it does not appreciably evolve in time). In the subpanel, time evolution of the maximum particle Lorentz factor, for $\alpha = 0.0$ (solid line) and $\alpha = 0.1$ (dotted line).

The shape of the downstream energy spectrum is also significantly different between the cases $\alpha = 0.0$ and $\alpha = 0.1$. Fig. 4.11 follows the time evolution of the particle distribution behind a shock with $\alpha = 0.0$, compared with the case $\alpha = 0.1$. As further discussed in Appendix 8, the downstream particle spectrum in a flow with $\alpha \gtrsim 0.03$ does not appreciably change with time, so the dotted blue line in Fig. 4.11, which refers to $\omega_{\rm p}t = 3000$ for the case $\alpha = 0.1$, is a good representation of the downstream energy spectrum at all times (for $\alpha = 0.1$). In contrast, as shown by the solid lines, the evolution of the particle spectrum in the case $\alpha = 0.0$ is much more dramatic. At early times (black line for $\omega_{\rm p}t = 1500$), the spectrum for $\alpha = 0.0$ is remarkably similar to the (time-independent) particle distribution of the case $\alpha = 0.1$ (dotted blue line). This confirms once more that the physics of shock-driven reconnection, which is ultimately responsible for shaping the post-shock spectrum at early times, does not appreciably depend on α, in the regime $\alpha \lesssim 0.1$. As mentioned above, the minor high-energy bump seen in the spectrum of $\alpha = 0.1$ (dotted blue line),

but not for $\alpha = 0.0$ (black line), results from particles accelerated at the hydrodynamic shock via the SDA mechanism.

Starting from $\omega_{\rm p}t = 3000$ (solid blue line), the spectrum for $\alpha = 0.0$ changes significantly. The high-energy part of the broad peak that was present at earlier times (at $\gamma \simeq 300$) gets de-populated, at the expense of a prominent tail of nonthermal particles (at $\gamma \gtrsim 10^3$). The tail contains approximately a few percent of particles and a few percent of the total particle energy. With time, the slope of the nonthermal tail becomes flatter (for comparison, the dashed line in Fig. 4.11 has $p = 2.5$), and its high-energy cutoff extends to higher Lorentz factors, as tracked by the solid line in the subpanel of Fig. 4.11. In contrast, no appreciable evolution in the value of the maximum particle Lorentz factor is observed for the case $\alpha = 0.1$ (dotted line in the subpanel). The particles in the nonthermal tail seen for $\alpha = 0.0$ have been accelerated via a kind of first-order Fermi mechanism (or diffusive shock acceleration, DSA), by scattering off the turbulence generated by the Weibel instability (see Fig. 4.10(d), right column). Their energization proceeds in the same way as for a strictly unmagnetized shock, which was discussed by Spitkovsky (2008). By comparing the black line (for $\omega_{\rm p}t = 1500$) and the blue line (for $\omega_{\rm p}t = 3000$) at $\gamma \gtrsim 10^2$, we infer that injection into the Fermi process is due primarily to the so-called "thermal leakage," i.e., the hottest particles from the downstream region can escape ahead of the shock, where they take part in the Fermi acceleration cycles. Of course, once the fields produced by the Weibel instability grow stronger than the background field B_0 (see the scale in Fig. 4.10(d), right column), the Weibel-generated turbulence in the shock layer may mediate further injection of particles into the Fermi process.

Finally, we remark that, based on the condition $\lambda_{\,c/\omega_{\rm p}}/\sqrt{\sigma} \lesssim 2\,\xi_2$ discussed above, injection into the Fermi process is expected to be more favorable for short stripe wavelengths and highly magnetized winds. In fact, we observe that the efficiency of Fermi acceleration, meaning the fraction of particles participating in the Fermi cycles, is higher for smaller λ (at fixed σ) and for larger σ (at fixed λ). This explains why, for $\lambda = 320\,c/\omega_{\rm p}$ (so, a wavelength longer than the value $\lambda = 80\,c/\omega_{\rm p}$ used here), Fig. 4.9 shows that the spectra for $\alpha = 0.0$ and $\alpha = 0.1$ are almost identical, in contrast with the conclusions of Fig. 4.11. For $\lambda = 320\,c/\omega_{\rm p}$, fewer particles can be injected into the Fermi process, and the structure of the flow at $\omega_{\rm p}t = 3000$ is practically the same for all values of $\alpha \lesssim 0.1$.

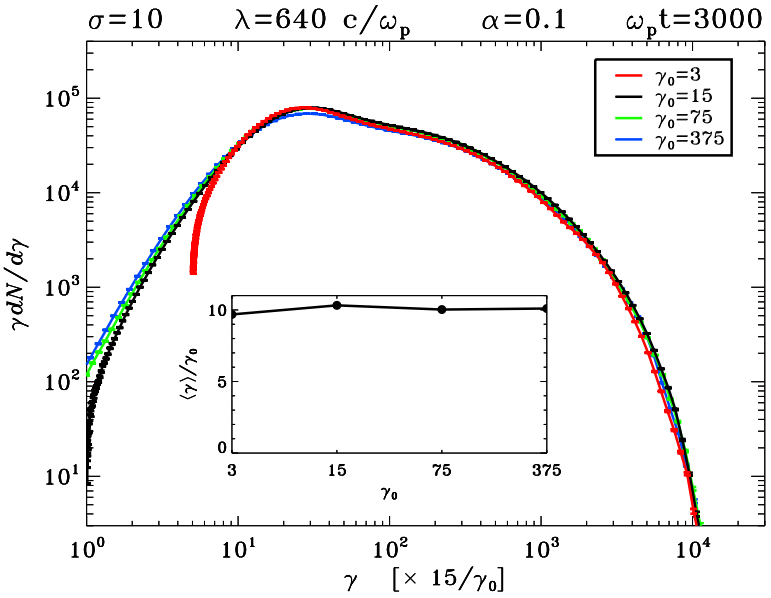


Fig. 4.12.— Downstream particle spectrum at $\omega_{\rm p}t = 3000$ for different values of the pre-shock bulk Lorentz factor γ_0, in a flow with $\lambda = 640\, c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$. We vary γ_0 from $\gamma_0 = 3$ to $\gamma_0 = 375$. Spectra are shifted along the x-axis by $15/\gamma_0$ to facilitate comparison with the reference case $\gamma_0 = 15$. The black line in the subpanel shows the average downstream Lorentz factor as a function of γ_0.

4.4. Dependence on the Bulk Lorentz Factor

Fig. 4.12 shows the particle spectrum downstream from the hydrodynamic shock, for different bulk Lorentz factors of the striped wind (from $\gamma_0 = 3$, in red, up to $\gamma_0 = 375$, in blue). The spectra are shifted along the x-axis by $15/\gamma_0$, to facilitate comparison with the reference case $\gamma_0 = 15$ (black curve) discussed in the previous sections. Once normalized in this way, the spectra overlap almost perfectly, suggesting that the physics of shock-driven reconnection is not sensitive to the upstream bulk Lorentz factor. The value of γ_0 enters the definition of the relativistic plasma skin depth $c/\omega_{\rm p}$ ($\propto \sqrt{\gamma_0}$) and of the relativistic Larmor radius ($r_L = c/\omega_{\rm p}/\sqrt{\sigma} \propto \sqrt{\gamma_0}$), but no residual dependence on γ_0 remains if length and time scales are normalized with respect to $c/\omega_{\rm p}$ and $\omega_{\rm p}^{-1}$, respectively. More specifically, the conditions that control the physics of shock-driven reconnection (see §4.1, and Pétri & Lyubarsky 2007) depend only on the stripe wavelength λ and the magnetization σ, but they do not explicitly depend on the wind Lorentz factor. This explains why our results are nearly the same (modulo an overall shift in energy scale) across a wide range of γ_0.

5. Three-Dimensional Simulations

The results presented so far are based on 2D experiments with magnetic field initialized in the plane of the simulations. This is the best configuration to study the tearing-mode instability, whose wavevector lies in the plane of the field, but it artificially inhibits the growth of cross-plane instabilities, which may also be important for the structure of the shock. In this section, we study the 3D physics of shock-driven reconnection, with particular emphasis on the process of particle energization. The simulation setup parallels very closely what we described in §2, with the magnetic field oriented initially along the y direction. We adopt our fiducial values for $\sigma = 10$ and $\alpha = 0.1$, and we investigate two choices for the stripe wavelength, $\lambda = 80\,c/\omega_{\rm p}$ and $\lambda = 160\,c/\omega_{\rm p}$. For each value of λ, the extent of the simulation domain along y and z is chosen such that our results are not artificially affected by the periodicity of our boundary conditions (see Appendix 7). We adopt $L_y = L_z = 50\,c/\omega_{\rm p}$ for $\lambda = 80\,c/\omega_{\rm p}$, and $L_y = L_z = 100\,c/\omega_{\rm p}$ for $\lambda = 160\,c/\omega_{\rm p}$.

Fig. 4.13 shows the structure of the shock at $\omega_{\rm p}t = 990$, for a striped wind of wavelength $\lambda = 160\,c/\omega_{\rm p}$. In agreement with the results presented in §3.1 for 2D simulations, the structure of the flow in Fig. 4.13 shows the presence of two shocks, a fast MHD shock located at $x \simeq 630\,c/\omega_{\rm p}$ and a hydrodynamic shock at $x \simeq 250\,c/\omega_{\rm p}$. At the fast shock, the flow is decelerated and

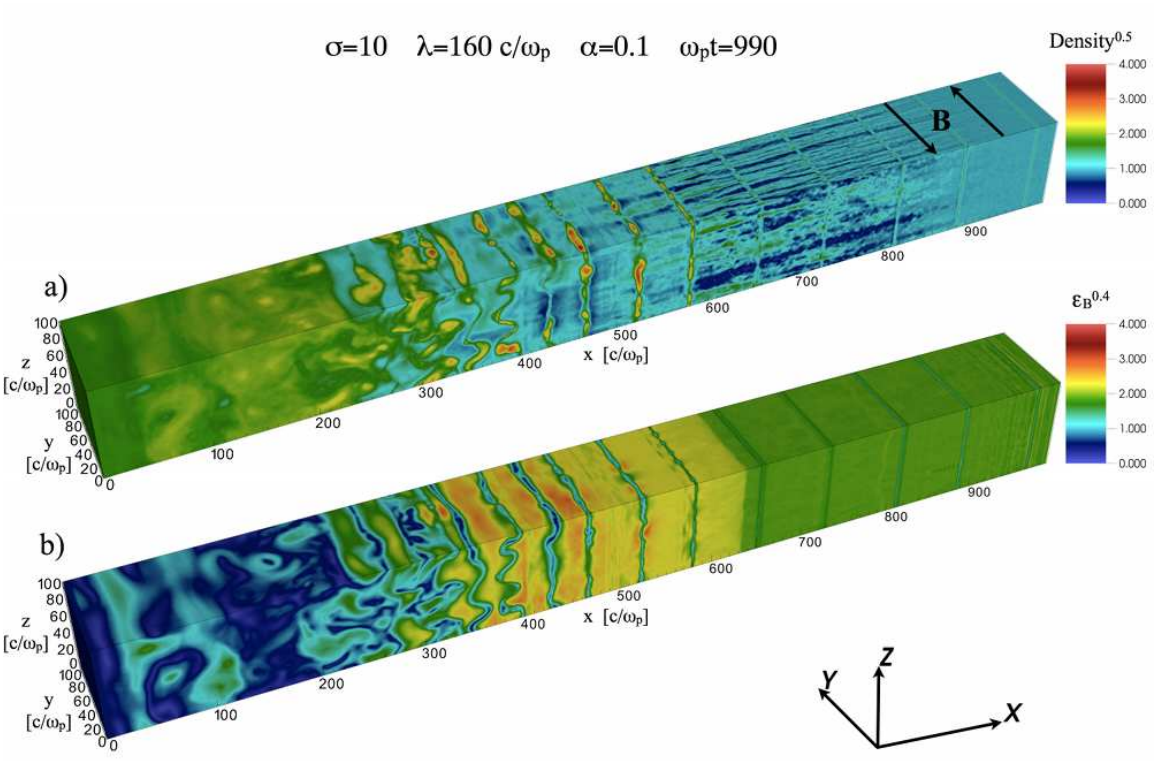


Fig. 4.13.— Internal structure of the flow at $\omega_{\rm p}t = 990$, from the 3D simulation of a striped wind with $\lambda = 160\,c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$. The magnetic field is initialized in the xy plane, as shown by the black arrows in panel (a). The hydrodynamic shock is located at $x \simeq 250\,c/\omega_{\rm p}$, and the fast MHD shock at $x \simeq 630\,c/\omega_{\rm p}$. The following quantities are plotted: (a) 3D plot of the particle number density, in units of the upstream value, with color scale stretched to enhance contrast; (b) 3D plot of the magnetic energy fraction $\epsilon_B \equiv B^2/8\pi\gamma_0 m n_{c0} c^2$, with color scale stretched to enhance contrast.

compressed, as shown by the increase in magnetic energy in Fig. 4.13(b). The passage of the fast shock through the incoming current sheets initiates magnetic field annihilation. The dissipation of alternating fields proceeds as the flow propagates from the fast to the hydrodynamic shock, and little magnetic energy remains downstream from the hydrodynamic shock (see Fig. 4.13(b) at $x \lesssim 250\, c/\omega_{\rm p}$), in agreement with our 2D results. As observed in 2D, the particle density downstream from the hydrodynamic shock (Fig. 4.13(a)) and the shock velocity are in agreement with MHD jump conditions for a relativistic *unmagnetized* fluid. In fact, most of the post-shock energy is in kinetic form, rather than electromagnetic form.[13]

As the flow propagates from the fast to the hydrodynamic shock, the structure of the current sheets is affected by the growth of two competing modes. In the xy plane of the magnetic field, the tearing-mode instability breaks the current sheet into a sequence of high-density islands (see Fig. 4.13(a)), separated by X-points where the field lines tear and reconnect. In the xz plane orthogonal to the field, the folding of the current sheet seen in Fig. 4.13 is governed by the drift-kink instability, driven by the current of fast-drifting plasmas in a thin sheet (Daughton 1998; Zenitani & Hoshino 2005a, 2007). The 2D geometry employed in the previous sections was chosen to select the tearing mode, and suppress the drift-kink mode. In the regime of undriven reconnection, the drift-kink instability is generally thought to grow faster than the tearing mode (Zenitani & Hoshino 2005b, 2007), but for very thin sheets the opposite hierarchy is observed, as shown via 3D PIC simulations by Liu et al. (2011). In our experiments of shock-driven reconnection, the current sheet just behind the fast shock ($x \simeq 580\, c/\omega_{\rm p}$) shows the pattern of magnetic islands typical of the tearing mode, whereas the drift-kink mode appears only farther downstream ($x \lesssim 500\, c/\omega_{\rm p}$).

The nature of the instability that dominates the process of field annihilation leaves an imprint on the resulting particle distribution. Zenitani & Hoshino (2007) have shown with PIC simulations that field dissipation due to the drift-kink instability does not result in nonthermal particle acceleration, the plasma just being heated. This is because the field lines remain straight, so that all particles gain energy at the same rate. Nonthermal particles are produced only by the tearing mode (e.g., Zenitani & Hoshino 2001). In our 2D simulations with in-plane fields, i.e., the geometry required to capture the tearing mode, we have shown that particles are accelerated to nonthermal energies by the reconnection electric field. They initially move along the z direction parallel to

[13] As a side note, we point out that the density filamentation seen in Fig. 4.13(a) ahead of the fast shock is a transient effect of the passage of the electromagnetic precursor discussed in §3.1. The precursor wave, and the resulting density filamentation, do not appreciably affect the shock structure at late times.

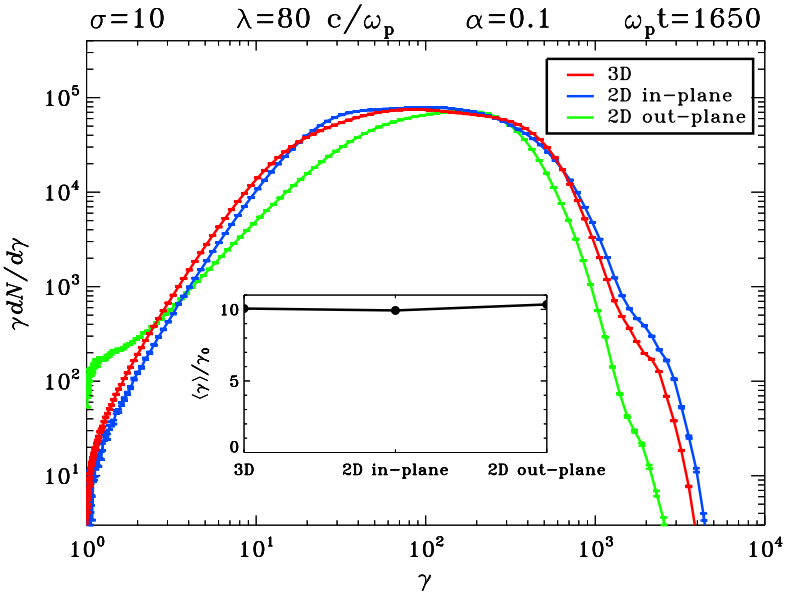


Fig. 4.14.— In a striped wind with $\lambda = 80\,c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$, comparison of the downstream particle spectrum for different magnetic field configurations: 3D simulation (red) and 2D simulations with either in-plane (blue) or out-of-plane (green) magnetic fields. The black line in the subpanel shows the average downstream Lorentz factor $\langle\gamma\rangle$ for the different cases.

the reconnection field, and then they drift along y from a given X-point into the closest magnetic island. In 3D, one could argue that the folding of the current sheet introduced by the drift-kink instability in the xz plane (not resolved in 2D) may deflect the accelerating particles out of the current sheet, thus suppressing their energization.

In Fig. 4.14, we address this important issue, by comparing for $\lambda = 80\,c/\omega_{\rm p}$ the post-shock particle spectrum of a 3D simulation (red line) to the results of 2D experiments with in-plane or out-of-plane magnetic fields (blue and green lines, respectively). The excellent agreement between 3D and 2D in-plane results (red and blue line, respectively) suggests that the physics of particle acceleration by shock-driven reconnection is captured extremely well by our 2D simulations with in-plane fields.[14] In turn, this implies that, at least for $\lambda = 80\,c/\omega_{\rm p}$, the maximum energy of accelerated particles is constrained by the distance between X-points and islands (in the xy plane), rather than by the current sheet folding in the xz plane. Just upstream of the hydrodynamic

[14]The agreement between red and blue lines in the high-energy bump at $\gamma \gtrsim 2000$ also suggests that the physics of SDA is correctly described by our 2D simulations with in-plane fields.

shock, the typical spacing between X-points and islands scales linearly with the stripe wavelength. A similar trend may be expected in the characteristic scale of the drift-kink mode, for the current sheet just ahead of the hydrodynamic shock (see Fig. 4.13(b) at $x \simeq 350\, c/\omega_{\rm p}$). If this is the case, our conclusions will also hold for longer stripe wavelengths, as we have directly verified in the case $\lambda = 160\, c/\omega_{\rm p}$. A detailed 3D study of particle acceleration by shock-driven reconnection will be presented elsewhere.

Finally, we point out that the tension force of a guide field (along z, in our geometry) can easily stabilize the drift-kink mode, with little or no effect on the tearing mode (Zenitani & Hoshino 2008). In the presence of a guide field, the 3D physics of shock-driven reconnection should be described very accurately by 2D simulations with alternating fields lying in the simulation plane. A guide field component, though, is not usually expected in the context of pulsar winds (but see Pétri & Kirk 2005).

6. Summary and Discussion

We have explored by means of 2D and 3D PIC simulations the internal structure and acceleration properties of relativistic shocks that propagate in an electron-positron striped wind, i.e., a flow consisting of stripes of alternating field polarity separated by current sheets of hot plasma. This chapter extends the study of Sironi & Spitkovsky (2009), that investigated relativistic shocks in a pair plasma carrying uniform magnetic fields.

We find that a fast MHD shock propagates into the pre-shock striped flow, compressing the incoming current sheets and initiating the process of *driven magnetic reconnection*, via the tearing-mode instability. Reconnection islands seeded by the passage of the fast shock grow and coalesce, while magnetic energy is dissipated at X-points located in between each pair of islands. When reconnection islands grow so big to occupy the entire region between neighboring current sheets, the striped structure of the flow is erased, and a hydrodynamic shock forms. Downstream from the shock, the average particle energy is larger than in the pre-shock flow by a factor of $\simeq \sigma$, where σ is the wind magnetization. In other words, the energy stored in the alternating fields has been entirely transferred to the particles via shock-driven magnetic reconnection, and the downstream fluid behaves like an unmagnetized plasma. The only field component surviving in the downstream region comes from shock-compression of the stripe-averaged pre-shock field (i.e., the non-alternating component), if non-zero. In our 2D and 3D simulations, we find that

complete annihilation of the alternating fields occurs irrespective of the stripe wavelength λ or the magnetization σ, in contrast with the results of the 1D analysis by Pétri & Lyubarsky (2007). In Appendix 7 we show that multi-dimensional studies are of paramount importance to correctly capture the physics of shock-driven reconnection.

The main agent of particle energization is the reconnection electric field, as particles drift from a given X-point into the closest island. Whether all particles have comparable energy gains, or only a few particles are accelerated to high energies, and the majority stay cold, depends sensitively on the properties of the wind. We have explored the dependence of the downstream particle spectrum on the stripe wavelength λ, the wind magnetization σ (or equivalently, the field strength B_0), and the stripe-averaged magnetic field $\langle B_y \rangle_\lambda$ (or equivalently, the parameter α, with $\alpha = 0$ for $\langle B_y \rangle_\lambda = 0$, and $|\alpha| \to 1$ in the limit $|\langle B_y \rangle_\lambda| \to B_0$). For fixed $\alpha = 0.1$, we find that the shape of the post-shock spectrum depends primarily on the combination $\lambda/(r_L\sigma)$, where r_L is the relativistic Larmor radius in the striped wind. This is just the stripe wavelength measured in units of the post-shock plasma skin depth. For small values of $\lambda/(r_L\sigma)$ ($\lesssim$ a few tens), the spectrum resembles a Maxwellian distribution with mean energy $\langle\gamma\rangle \simeq \gamma_0\sigma$, where γ_0 is the bulk Lorentz factor of the pre-shock flow. In the limit of very large values of $\lambda/(r_L\sigma)$ ($\gtrsim$ a few hundreds), the spectrum approaches a broad power-law tail, with a flat spectral slope $p \simeq 1.5$. In this regime, the tail extends from $\gamma_{\rm min} \simeq \gamma_0$ up to $\gamma_{\rm max} \simeq \gamma_0\sigma^{1/(2-p)}$. In terms of the flow structure, the former case (i.e., a Maxwellian-like spectrum) is realized when all particles pass in the vicinity of an X-point, thus gaining energy from field dissipation. In contrast, the latter case occurs when most of the particles remain far from X-points, thus staying cold, and the energy of annihilating fields is transferred to only a few particles, which end up dominating the energy content of the flow.

As $|\alpha|$ increases from 0.1 up to 1, the fraction of upstream Poynting flux in the form of alternating fields becomes smaller. As the amount of field energy available for dissipation decreases, the average particle Lorentz factor downstream from the shock recedes from $\langle\gamma\rangle \simeq \gamma_0\sigma$, as expected for complete field annihilation, down to $\langle\gamma\rangle \simeq \gamma_0$, the result of an unstriped wind. In the opposite limit $|\alpha| \lesssim 0.1$, i.e., in the case of nearly symmetric stripes, the particles accelerated by the reconnection electric field to the highest energies can escape ahead of the shock, unaffected by the stripe-averaged field that would tend to advect them back downstream. Their counter-streaming with respect to the incoming flow can seed the filamentation (or Weibel) instability, as it happens for shocks propagating in unmagnetized plasmas (Spitkovsky 2005, 2008). In turn, the turbulence generated by the Weibel instability can mediate particle acceleration to even higher energies, via a

Fermi-like diffusive process. The nonthermal tail produced by the Fermi mechanism is steeper than that given by magnetic reconnection, with characteristic slopes $p \gtrsim 2.5$. Typically, a few percent of particles are injected into the Fermi process, and the injection efficiency scales in inverse proportion to the ratio $\lambda/(r_L\sigma)$. With time, the nonthermal tail of Fermi-accelerated particles extends to higher energies. For $|\alpha| \gtrsim 0.1$, the propagation of particles far ahead of the shock is inhibited by the stripe-averaged field, and Fermi acceleration is suppressed. In this case, a few percent of particles can get accelerated by the stripe-averaged electric field $\langle E_z \rangle_\lambda$ while gyrating around the shock front, a process known as shock-drift acceleration (i.e., Begelman & Kirk 1990). The tail resulting from shock-drift acceleration has a limited energy extent, and it does not evolve in time.

We have confirmed our findings with an extensive set of convergence tests. Most importantly, as we have shown in §5, 2D simulations with in-plane magnetic fields – the setup that we have primarily adopted in this chapter – can reproduce very accurately the shock physics and particle spectrum observed in 3D simulations. In other words, the current sheet folding caused by the drift-kink instability in the plane orthogonal to the field (Zenitani & Hoshino 2007) does not represent a serious constraint for the maximum particle energy achieved in the course of the reconnection process. Rather, the maximum energy of accelerated particles is limited by the time to drift from a given X-point into the closest island.

These findings can place important constraints on current models of Pulsar Wind Nebulae (PWNe), bubbles of synchrotron-emitting plasma powered by the relativistic wind of young pulsars. If the magnetic and rotation axes of the pulsar are not aligned, the wind consists of stripes of opposite magnetic field polarity, alternating with the pulsar period. The wind wavelength will be $\lambda = 2\pi R_{LC}$, where $R_{LC} = c/\Omega$ is the so-called light-cylinder radius, and Ω is the angular frequency of the pulsar. The stripe-averaged field vanishes along the equatorial plane (i.e., $\alpha = 0$ there), and the striped structure persists up to a latitude equal to the inclination angle between the pulsar's magnetic and rotational axes. At higher latitudes, the wind carries a uniform (i.e., non-alternating) field, so the physics discussed in the present chapter does not apply (see, instead, Gallant et al. 1992; Hoshino et al. 1992; Amato & Arons 2006; Sironi & Spitkovsky 2009, 2011).

Our results suggest that efficient dissipation of the alternating fields should occur at the termination shock of pulsar winds, irrespective of the properties of the flow. It follows that current observational constraints on the upstream magnetization parameter (Rees & Gunn 1974; Kennel & Coroniti 1984a; Del Zanna et al. 2004) should be attributed not to the total Poynting flux, but to

the Poynting flux associated with the stripe-averaged magnetic field, the only component surviving downstream from the shock. The upstream flow may be Poynting-dominated, provided that most of the Poynting flux is in the form of an alternating field, which is annihilated at the shock. Thus, our results provide a solution to the so-called "sigma problem," i.e., that pulsar winds should be magnetically-dominated at the light-cylinder radius, but kinetically-dominated downstream from the termination shock. The idea of shock-driven reconnection as a solution to the sigma problem was originally proposed by Lyubarsky (2003). Our work demonstrates that complete annihilation of the alternating fields occurs regardless of the properties of the wind, in the regime of relativistic magnetically-dominated flows.

In addition, the particle energy spectra extracted from our simulations can be used directly to interpret the radiative signature of PWNe. The radio spectrum of the Crab Nebula, the prototype of the class of PWNe, requires a population of nonthermal particles with a flat spectral slope ($p \simeq 1.5$), extending at least across three decades in energy, from 10^2 MeV up to 10^5 MeV. The particle spectrum should be steeper at higher energies, with $p \simeq 2.5$, to explain the optical and X-ray flux.

Our results suggest that, for a slope $p \simeq 1.5$, the width of the electron power-law tail resulting from reconnection will be $\simeq \sigma^2$. To produce a particle distribution extending over three decades in energy, the wind magnetization at the termination shock needs to be $\sigma \gtrsim 30$. Most importantly, broad particle spectra are produced by shock-driven reconnection only if $\lambda/(r_L\sigma) \gtrsim$ a few tens. For pulsar winds, the particle number density scales with distance from the pulsar as $n = n_{LC}(R_{LC}/R)^2$, where the density at the light-cylinder radius is usually written as $n_{LC} = \kappa\, n_{GJ}$. Here, $n_{GJ} = \Omega B_{LC}/2\pi ec$ is the Goldreich-Julian density (Goldreich & Julian 1969), and κ is the so-called multiplicity. Conservation of energy along the flow streamlines implies that at the termination shock $\gamma_0(1+\sigma)\kappa = \omega_{LC}/2\,\Omega$, where $\omega_{LC} = eB_{LC}/mc$. It follows that at the termination shock ($R = R_{TS}$)

$$\frac{\lambda}{r_L\sigma} \simeq 4\pi\kappa\frac{R_{LC}}{R_{TS}} \quad . \tag{4-3}$$

The constraint $\lambda/(r_L\sigma) \gtrsim$ a few tens required to produce broad particle spectra is extremely challenging for current models of PWNe and pulsar magnetospheres. The radius of the Crab termination shock in the equatorial plane is inferred from X-ray observations to be $R_{TS} \simeq 0.1$ pc $\simeq 5\times10^8 R_{LC}$ (Hester et al. 2002). Most available models estimate $\kappa \simeq 10^4 - 10^6$ (Bucciantini et al. 2011), depending on whether or not radio-emitting electrons are included in the

analysis. Based on our findings, the resulting value of $\lambda/(r_L\sigma) \lesssim 0.01$ would yield a Maxwellian-like spectrum, at odds with the wide flat spectrum required by observations.

If radio-emitting electrons are produced in the equatorial plane by shock-driven reconnection, a revision of the existing estimates of κ is required. In this respect, we point out that the values of κ quoted above are averages over latitude, and one cannot exclude that particle injection into the pulsar wind is highly anisotropic, with multiplicity as large as 10^8 along the equatorial plane. Of course, any given choice of κ will constrain γ_0 and σ via $\gamma_0(1+\sigma)\kappa = \omega_{LC}/2\,\Omega$, a relation that holds at each latitude. From the equatorial plane, radio-emitting electrons will quickly fill the whole nebula, transported by the strong fluid motions observed in MHD models of PWNe downstream from the termination shock (e.g., Komissarov & Lyubarsky 2004; Del Zanna et al. 2004; Camus et al. 2009). This would provide a natural explanation for the lack of gradients in the radio spectral slope of the Crab (Bietenholz & Kronberg 1992). Alternatively, the radio part of the spectrum may be produced at higher latitudes, where the termination shock is closer to the pulsar, and the ratio in eq. (4-3) becomes larger. However, in this case one should also account for the fact that, at high latitudes, the fraction of upstream Poynting flux available for dissipation is lower (since $|\alpha|$ is larger), which results in a narrower spectrum, everything else being fixed.[15]

On the other hand, the small value of $\lambda/(r_L\sigma)$ expected in the equatorial plane on the basis of current models of the Crab Nebula is the most favorable setup for particle acceleration via the Fermi diffusive process, with efficiency of a few percent by number. In this respect, one could interpret the optical and X-ray signature of the Crab, which requires a particle spectrum with $p \simeq 2.5$, as synchrotron emission from such Fermi-accelerated particles. Based on our results, optical and X-ray emitting electrons should be produced close to the equatorial plane of the wind, where $|\alpha| \lesssim 0.1$. If the same equatorial wedge is responsible for accelerating the radio-emitting electrons, our findings predict that they should outnumber the optical and X-ray emitting particles by a factor of few hundred. This is in agreement with current models of PWNe spectra, which require $\kappa \simeq 10^6$ if the radio band is included in the analysis (Bucciantini et al. 2011), a value which is two orders of magnitude larger than if only higher frequency bands are considered ($\kappa \simeq 10^4$, Kennel & Coroniti 1984b).

[15] Also, the termination shock at high latitudes is oblique with respect to the wind velocity, a configuration not studied here.

Although our work focuses primarily on the physics of PWNe, shocks propagating into striped flows may be present also in blazar jets and gamma-ray bursts (e.g., Thompson 2006), where they may provide an appealing explanation, based on strong physical grounds, for any flat particle distribution inferred from observations. The broad-band emission of hotspots in radio galaxies, usually interpreted as the termination shocks of relativistic jets, indicates that the spectra of accelerated electrons need to be flat ($1 < p < 2$) below GeV energies (Stawarz et al. 2007). In the case of luminous blazar sources, very flat electron spectra below GeV energies are inferred directly from X-ray observations (Celotti & Ghisellini 2008; Sikora et al. 2009). The steeper particle spectrum required at $\gtrsim$GeV energies, with slope $p \simeq 2.5$, could instead result from Fermi acceleration across the shock, if the stripe-averaged field satisfies $|\alpha| \lesssim 0.1$.

7. Appendix (A): Dependence on the Transverse Size of the Simulation Box

In Fig. 4.15 we investigate the dependence of our findings on the transverse size L_y of the computational domain, for $\lambda = 640\,c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$. The agreement between the black curve ($L_y = 400\,c/\omega_{\rm p}$) and the red line ($L_y = 800\,c/\omega_{\rm p}$) suggests that our results become insensitive to the transverse size of the box (i.e., they converge with respect to L_y), for boxes larger than $400\,c/\omega_{\rm p}$. Generally speaking, we find that a good criterion for convergence is $L_y \gtrsim \lambda/2$. When this condition is fulfilled, the growth of magnetic islands behind the fast shock is not artificially inhibited by the periodicity of our boundaries in the y direction. Magnetic reconnection then proceeds up to the point when islands from two neighboring current sheets will merge, which happens when their size is $\simeq \lambda/2$ (for $\alpha \lesssim 0.1$).

On the other hand, for $L_y \lesssim 300\,c/\omega_{\rm p}$ the growth and coalescence of reconnection islands within a given current sheet artificially stops when only one island is left in the sheet, and its size approaches L_y. At this point, a second shock forms, located behind the fast shock. For $L_y \gtrsim 400\,c/\omega_{\rm p}$, this would correspond to the hydrodynamic shock discussed in §3.1, where the striped structure of the flow is erased, and field energy is entirely transferred to the particles. Instead, in the case $L_y \lesssim 300\,c/\omega_{\rm p}$, complete dissipation of the alternating fields is prohibited by the fact that the growth of islands is artificially inhibited by the transverse extent of the box. Here, the post-shock flow retains a striped structure, with regions of hot plasma separated by highly magnetized walls of cold particles. With decreasing L_y, less magnetic energy is transferred to the particles behind the fast shock, and the post-shock flow gets more dominated by magnetic field

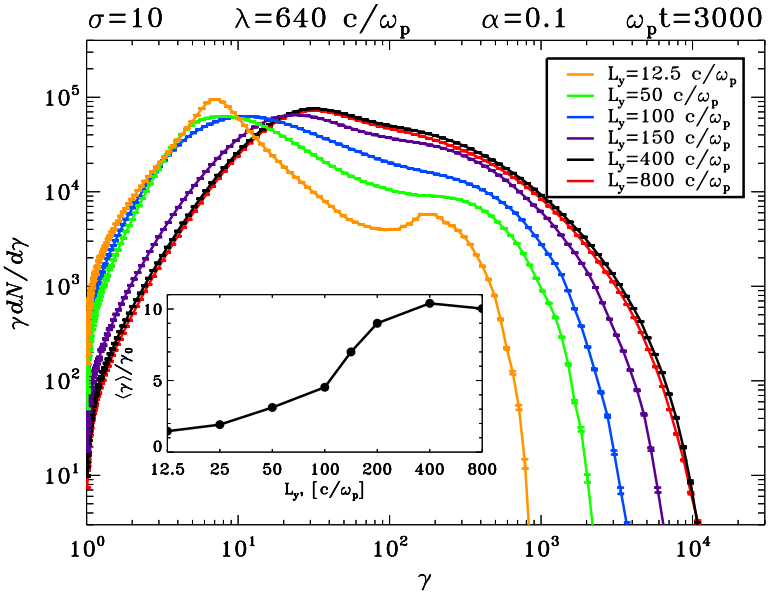


Fig. 4.15.— Downstream particle spectrum at $\omega_{\rm p}t = 3000$ for different values of the transverse size L_y of our simulation box, in a flow with $\lambda = 640\,c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$. We vary L_y from $L_y = 800\,c/\omega_{\rm p}$ down to $L_y = 12.5\,c/\omega_{\rm p}$, which approaches the 1D setup of Pétri & Lyubarsky (2007). The black line in the subpanel shows the average downstream Lorentz factor $\langle\gamma\rangle$ as a function of the box width L_y.

energy, with respect to particle kinetic energy (see the decrease in average particle energy with decreasing L_y, in the subpanel of Fig. 4.15). This also explains why in Fig. 4.15 the normalization of the low-energy peak, populated by the cold particles in the high-field regions, grows with decreasing L_y, at the expense of the high-energy component of hot particles that gained energy from field dissipation.

As discussed in the main body of the chapter (see §4.1), the maximum size of magnetic islands (which here is set by the smallest between $\lambda/2$ and L_y) correlates with the highest energy at which particles can be accelerated by the reconnection electric field. The trend in the upper spectral cutoff of Fig. 4.15 can then be explained by the decrease with L_y in the maximum size of reconnection islands. In this respect, the spectra in Fig. 4.15, interpreted as a monotonic trend in the size of reconnection islands, can be directly compared to the spatial sequence of spectra in Fig. 4.3(f), which follow the growth of magnetic islands as the flow propagates from the fast to the hydrodynamic shock. A similar trend as in Fig. 4.15 is also observed in the large-α limit of Fig. 4.9 (yet for a different wavelength, $\lambda = 320\,c/\omega_{\rm p}$ as opposed to $\lambda = 640\,c/\omega_{\rm p}$). However, in that case the maximum size of magnetic islands was set by a physical constraint, namely the smallest between λ_+ (where $B_y = +B_0$) and λ_- (where $B_y = -B_0$), whereas here it is the transverse size of the box which artificially limits the growth of magnetic islands for $L_y \lesssim 300\,c/\omega_{\rm p}$.

As the downstream region gets more magnetized with decreasing L_y, the main shock, which would be moving at $\beta_{\rm sh} \simeq 1/3$ for $L_y \gtrsim 400\,c/\omega_{\rm p}$, propagates at a faster velocity, eventually catching up with the fast MHD shock, in the limit $L_y \ll \lambda$. In this regime, we recover the 1D results by Pétri & Lyubarsky (2007). The wind parameters employed in this section, $\lambda = 640\,c/\omega_{\rm p}$ and $\sigma = 10$, satisfy the condition $\lambda_{\,c/\omega_{\rm p}}/\sigma \gtrsim 8\,\xi_1/5$ (where $\xi_1 = 6 - 10$). In this case, the 1D model by Pétri & Lyubarsky (2007) would predict negligible field dissipation, in agreement with our results for a very narrow box (yellow line for $L_y = 12.5\,c/\omega_{\rm p}$; see also the subpanel, for small values of L_y). However, as clarified by Fig. 4.15, this is just an artificial consequence of the reduced dimensionality of their model, which cannot correctly capture the development of the tearing-mode instability, and the resulting growth and coalescence of magnetic islands. In a more realistic 2D scenario, complete field dissipation is achieved by the time the flow enters the hydrodynamic shock. This stresses that multi-dimensional simulations are essential for our understanding of shock-driven reconnection.

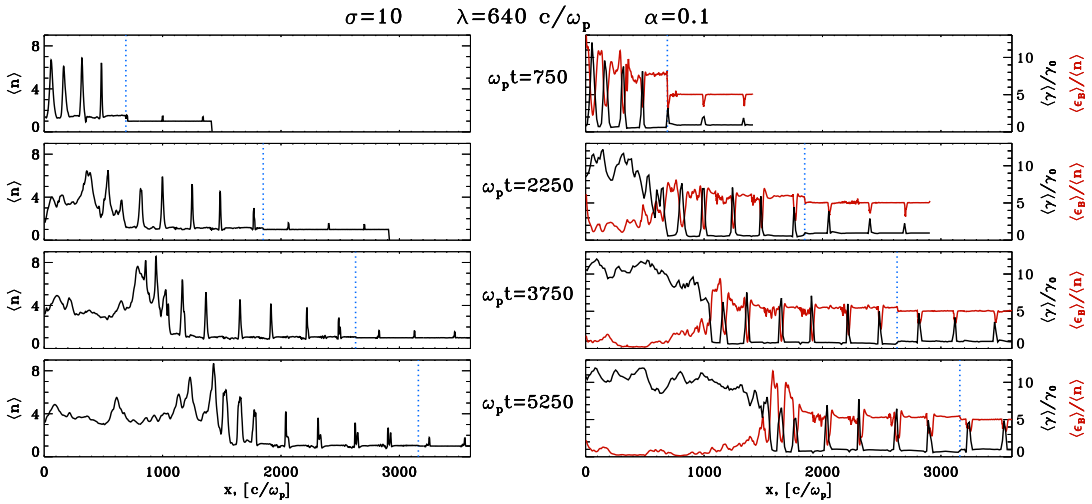


Fig. 4.16.— Time evolution of the internal structure of a shock propagating in a striped flow with $\lambda = 640\, c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$. We follow the evolution of the shock from $\omega_{\rm p}t = 750$ up to $\omega_{\rm p}t = 5250$ (from top to bottom). Left panels show the y-averaged particle number density, in units of the upstream value. Right panels show the mean kinetic energy per particle (black lines) and the mean magnetic energy per particle (red lines). In all panels, the location of the fast MHD shock is indicated as a vertical dotted blue line.

8. Appendix (B): Time Evolution of the Shock

In Fig. 4.16 we follow the time evolution of the shock, from the earliest stages until it reaches a steady state. We choose our fiducial values for $\lambda = 640\,c/\omega_{\rm p}$, $\sigma = 10$, and $\langle B_y \rangle_\lambda \simeq 0.05$ (corresponding to $\alpha = 0.1$), but our main conclusions hold basically for the whole parameter space we have investigated. At very early times (uppermost row, for $\omega_{\rm p} t = 750$), only the fast MHD shock is present (its location at different times is shown as a vertical dotted blue line in Fig. 4.16). The incoming flow decelerates at the fast shock, forming a ring in momentum space lying in the xz plane orthogonal to the magnetic field. Since the spread in y-momentum is much smaller than in the other two directions, the fluid behind the fast shock can be treated as a 2D relativistic plasma. In fact, the increase in density, magnetic energy, and transverse magnetic field across the fast shock, as well as the drop in average kinetic energy per particle, are in good agreement with the MHD jump conditions for a 2D relativistic fluid.

Behind the fast shock, the energy content of the flow at $\omega_{\rm p} t = 750$ is dominated by magnetic fields. However, even at such early times, we see that current sheets are getting wider (left panel) and populated with hotter particles (see the mean kinetic energy per particle, black line in the right panel), as we proceed farther downstream from the fast shock. As described in §3, a larger and larger fraction of the cold wind is being channeled into reconnection islands, seeded by the passage of the fast shock through the upstream current sheets. At X-points located in between each pair of neighboring islands, magnetic energy is being transferred to the particles, which explains why the peaks in mean *kinetic* energy per particle are getting higher with distance behind the fast shock (black line in the right panel), at the expense of the mean *magnetic* energy per particle (red line). Yet, the striped structure, with magnetized regions of cold plasma separated by hot current sheets, still persists downstream from the fast shock.

The shock structure changes significantly at $\omega_{\rm p} t = 2250$ (second row from the top). Behind the fast shock, which is now much weaker (see the modest density jump at $x \simeq 1850\,c/\omega_{\rm p}$), reconnection islands have enough time to grow and fill the entire region in between neighboring current sheets. At $x \simeq 700\,c/\omega_{\rm p}$, the striped structure is erased, and a hydrodynamic shock forms (see the density profile in the left panel). The density plateau behind the hydrodynamic shock is in agreement with the expected jump conditions for a relativistic unmagnetized 3D plasma, namely $n_{\rm d}/n_{\rm u} \simeq 4$. In fact, the energy balance presented in the right panel shows that, by the time the flow crosses

the hydrodynamic shock, most of the field energy (red line) has been converted to particle kinetic energy (black line).

A similar picture holds for later times ($\omega_{\rm p}t = 3750$ in the third row from the top, $\omega_{\rm p}t = 5250$ in the lowermost row). The density behind the hydrodynamic shock stabilizes at $n_{\rm d} \simeq 4\,n_{\rm u}$, and the transfer of energy from magnetic to kinetic form becomes more and more complete. The change of adiabatic index due to the transition from a magnetically-dominated fluid to a kinetically-dominated plasma results in a lower pressure to drive the fast shock, that decelerates (compare the locations of the vertical dotted blue lines at different times). Between $\omega_{\rm p}t = 3750$ and $\omega_{\rm p}t = 5250$, the overall structure of the transition region reaches a steady state. The fast shock propagates ahead of the hydrodynamic shock, roughly at a fixed distance (i.e., the speed of the fast shock approaches the velocity $\beta_{\rm sh} \simeq 1/3$ of the hydrodynamic shock). The separation between the two shocks obeys the requirement that reconnection islands seeded by the fast shock should fill the entire region in between current sheets, by the time the flow arrives at the hydrodynamic shock. When the flow structure has reached the steady state, we find that the particle spectrum downstream from the hydrodynamic shock does not significantly change with time.

9. Appendix (C): Relativistic Striped Winds in the Presence of a Guide Field

The results presented so far are obtained in the absence of a guide field, i.e., the component of magnetic field aligned with the particle current in the sheets. In this section, we investigate the effects of a guide field component on the structure of the termination shock and the process of magnetic dissipation, assuming that the alternating fields are oriented either along or perpendicular to the xy simulation plane (so that the guide field is along z or y, respectively). In addition to the alternating fields of strength B_0, we initialize in the pre-shock flow a uniform guide field of intensity B_g (with $0.0 \leq B_g/B_0 \leq 1.0$), both within and outside the current sheets. As we have anticipated in §5, a guide field component is generally not expected in the context of pulsar winds, but the results of this section could be relevant for gamma-ray bursts and blazar jets.

The role of the guide field in the case of *undriven reconnection* has been investigated by Zenitani & Hoshino (2008) by means of linear analysis and PIC simulations. They found that the tearing mode is relatively insensitive to the strength of the guide field, whereas the relativistic drift-kink instability is stabilized by the tension force of the guide field if $B_g/B_0 \gtrsim 0.3$. In this section, we perform a similar analysis in the context of *shock-driven reconnection* at the termination

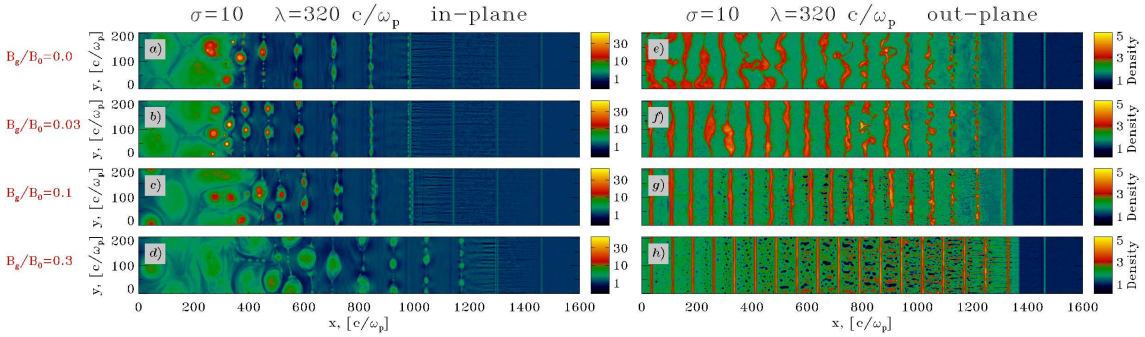


Fig. 4.17.— 2D structure of the particle number density (in units of the upstream value), for alternating fields lying either in the simulation plane (left column) or orthogonal to the plane (right column). We vary the guide field strength B_g/B_0 as indicated on the left. The figure refers to a flow with $\lambda = 320\, c/\omega_{\rm p}$, $\sigma = 10$, and $\alpha = 0.1$, at time $\omega_{\rm p} t = 1500$.

shock of striped winds, focusing on the process of particle acceleration. For alternating fields lying in the plane of the simulations, the drift-kink mode is not resolved, and we can isolate the effects of the guide field on the tearing mode. On the contrary, if the alternating fields are orthogonal to the computational domain, we can study whether the guide field can stabilize the drift-kink mode.

Fig. 4.17 shows the structure of the flow in the two cases (left column for in-plane alternating fields, right column for out-of-plane fields), for different strengths of the guide field, ranging from $B_g/B_0 = 0.0$ up to 0.3. The wavelength and magnetization of the striped wind are $\lambda = 320\, c/\omega_{\rm p}$ and $\sigma = 10$ in all cases. The left column shows that, in the case of in-plane alternating fields (so that the guide field is along z), the development of shock-driven reconnection and the resulting structure of the flow are nearly insensitive to the strength of the guide field. For all the values of B_g we investigate (B_g/B_0 from 0.0 to 0.3, from top to bottom), a fast shock propagates into the striped wind, compressing the incoming current sheets and triggering the tearing mode, which results in the magnetic islands seen in the density structure of Fig. 4.17 (left column). When the size of magnetic islands becomes comparable to the distance between neighboring current sheets, the striped structure of the flow is erased, and the main shock forms (at $x \simeq 400\, c/\omega_{\rm p}$ in the left column of Fig. 4.17), as described in §3. The alternating fields annihilate completely at the shock, whereas the guide field is compressed, in agreement with MHD jump conditions. For stronger guide fields, the magnetic pressure in the post-shock plasma is larger, and the shock propagates at a

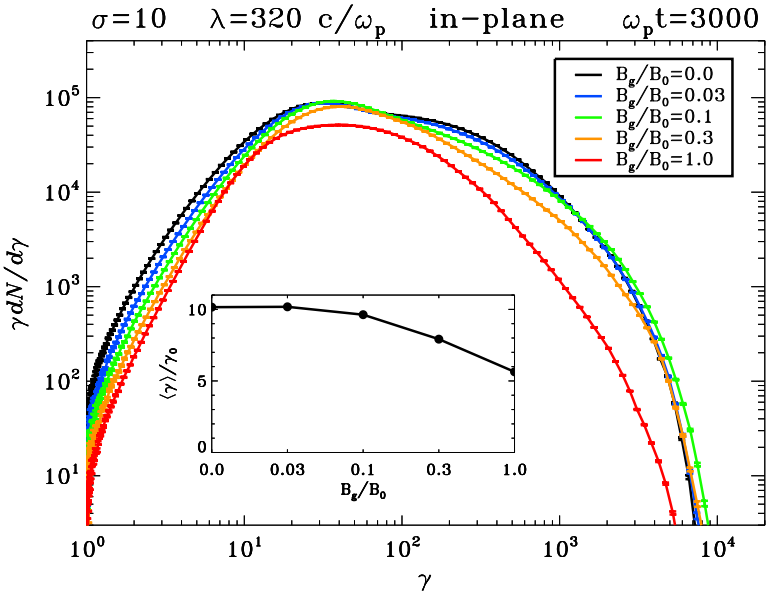


Fig. 4.18.— Downstream particle spectrum at $\omega_{\rm p}t = 3000$ for different values of the guide field strength B_g/B_0 (from 0.0 to 1.0), in a flow with $\lambda = 320\,c/\omega_{\rm p}$, $\sigma = 10$, $\alpha = 0.1$, and alternating fields lying within the simulation plane. In the subpanel, the black line shows the mean Lorentz factor $\langle\gamma\rangle/\gamma_0$ of post-shock particles as a function of B_g/B_0.

higher speed. This explains why the location of the main shock shifts toward $+\hat{\boldsymbol{x}}$ as the guide field strength increases from 0.0 to 0.3 (panel (a) through (d) in the left column).

For $B_g/B_0 = 0.3$ (panel (d)), the fast shock is farther ahead of the main shock, with respect to the case $B_g/B_0 \lesssim 0.1$.[16] As discussed in §3, the distance between the fast and the main shock is set by the reconnection rate (i.e., the rate of dissipation of alternating fields), everything else being fixed, so that our results suggest that the reconnection rate should be lower, for larger values of the guide field. Inflows into the reconnection X-points and outflows from the X-points into the neighboring islands (within the same current sheet) are expected to be impeded by a strong guide field, since the plasma should be moving across the guide field lines. This explains why the rate of field annihilation decreases for large values of B_g/B_0.

[16]The location of the fast shock cannot be easily inferred from the density plot of Fig. 4.17. Using 1D profiles of the magnetic and electric fields (not shown), we see that the fast shock is located at $x \simeq 1100\,c/\omega_{\rm p}$ for $B_g/B_0 \lesssim 0.1$, and at $x \simeq 1300\,c/\omega_{\rm p}$ for $B_g/B_0 = 0.3$.

In summary, apart from minor differences in shock speed and reconnection rate, the overall flow structure in the case of in-plane alternating fields – the configuration required to capture the tearing mode instability – is not appreciably changed by the presence of the guide field, up to $B_g/B_0 = 0.3$. This is confirmed by Fig. 4.18, where we compare the post-shock particle energy spectra for values of B_g/B_0 ranging from 0.0 to 1.0. We find that the particle spectrum departs from the result obtained in the absence of a guide field (black line) only for $B_g/B_0 \gtrsim 0.3$ (yellow line for $B_g/B_0 = 0.3$, red line for $B_g/B_0 = 1.0$). As discussed above, for strong guide fields the rate of particle inflow into the X-points is smaller, which results in fewer particles being accelerated by the reconnection electric field. This explains why the high-energy component in the particle spectra of Fig. 4.18 is suppressed for $B_g/B_0 \gtrsim 0.3$. In turn, this corresponds to the decrease of mean particle Lorentz factor shown in the subpanel of Fig. 4.18, with respect to the value $\langle\gamma\rangle \simeq \gamma_0(\sigma+1)$ expected in the case of complete field annihilation.[17]

So far we have discussed the effects of a guide field component on the development of the tearing instability, the only mode expected for alternating fields lying in the simulation plane. The right column of Fig. 4.17 shows instead the effect of the guide field in the case of alternating fields oriented along z, i.e., perpendicular to the simulation plane. For this geometry, the guide field is oriented along y. The flow structure in the absence of a guide field ($B_g/B_0 = 0.0$ in panel (e) of Fig. 4.17) shows a fast MHD shock located at $x \simeq 1400\, c/\omega_{\rm p}$, which compresses the incoming flow and triggers the drift-kink mode, manifested in the folding of the current sheets downstream from the shock.[18] The right column of Fig. 4.17 clearly shows that the drift-kink mode is stabilized once the guide field exceeds $B_g/B_0 \simeq 0.1$, in agreement with the conclusions by Zenitani & Hoshino (2008). For $B_g/B_0 = 0.3$ in Fig. 4.17(h), the current sheets downstream from the shock do not show any sign of the kink mode. Due to the suppression of the drift-kink mode for $B_g/B_0 \gtrsim 0.1$, most of the post-shock energy stays in the magnetic field, and the plasma remains cold. This is clearly shown in the particle energy spectra of Fig. 4.19, for different strengths of the guide field.

---

[17]The different normalization of the spectral peak for $B_g/B_0 = 1.0$ (red line in Fig. 4.18) results from the fact that the post-shock flow is more magnetized, for extreme values of the guide field. It follows that the shock speed is higher, so the post-shock density is lower.

[18]The speed of the shock is much larger here than in the case of in-plane fields with $B_g/B_0 = 0.0$ of Fig. 4.17(a). This results primarily from the different adiabatic index imposed by the 2D geometry of our simulations. If the alternating fields are in the simulation plane, the particle gyro-motion around the field will involve the third dimension (i.e., orthogonal to the plane), so that the post-shock plasma will be isotropic in 3D. In contrast, for out-of-plane alternating fields, the motion of the plasma in the presence of the field is confined to the 2D plane of the simulations, resulting in a 2D adiabatic index. The fact that the tearing mode is more effective than the drift-kink mode in converting field energy to particle energy further contributes to the difference in shock speed between panel (a) and (e).

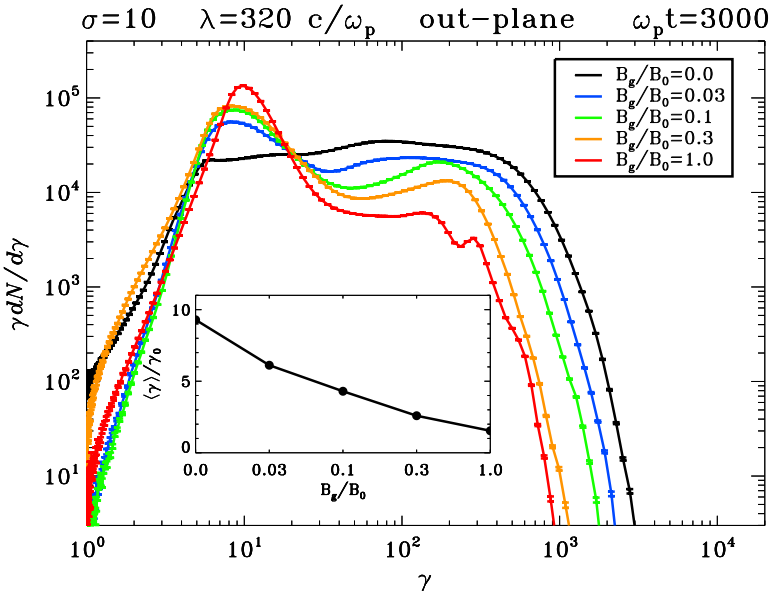


Fig. 4.19.— Downstream particle spectrum at $\omega_{\rm p}t = 3000$ for different values of the guide field strength B_g/B_0 (from 0.0 to 1.0), in a flow with $\lambda = 320\,c/\omega_{\rm p}$, $\sigma = 10$, $\alpha = 0.1$, and alternating fields orthogonal to the simulation plane. In the subpanel, the black line shows the mean Lorentz factor $\langle\gamma\rangle/\gamma_0$ of post-shock particles as a function of B_g/B_0.

As the value of B_g/B_0 increases, the spectrum passes from a flat distribution with mean Lorentz factor $\simeq \gamma_0(\sigma+1)$ (black line for $B_g/B_0 = 0.0$), as expected for efficient field annihilation, to a mono-energetic distribution peaking at $\gamma \simeq \gamma_0$ (red line for $B_g/B_0 = 1.0$), as in the case of negligible field dissipation. The subpanel of Fig. 4.19 shows that the suppression of the drift-kink mode, and the resulting inhibition of particle heating, occurs for guide field strengths as small as $B_g/B_0 \simeq 0.03$.

10. Appendix (D): The Termination Shock of Electron-Ion Striped Winds

So far, we have investigated the physics of driven reconnection at the termination shock of striped electron-positron winds, with focus on pulsar winds. Both theoretical studies of pulsar magnetospheres (Arons 2007) and models of the radiative signature from Pulsar Wind Nebulae (Bucciantini et al. 2011) suggest that the wind is primarily composed of electron-positron pairs, with no evidence for heavier species. On the other hand, the composition of gamma-ray burst and AGN jets is less certain, and current models cannot exclude the presence of a significant hadronic component. In this section, we extend our study of shock-driven reconnection to relativistic striped flows composed of electrons and protons.

Our numerical setup parallels very closely what we described in §2, but the spatial resolution is now set by the electron scale, in terms of the relativistic electron plasma skin depth $c/\omega_{\rm pe} \equiv \sqrt{\gamma_0 m_e c^2/4\pi e^2 n_{c0}}$ and the electron Larmor radius $r_{Le} \equiv \gamma_0 m_e c^2/eB_0$. On the other hand, the flow magnetization is defined using the ion scale as $\sigma \equiv (c/\omega_{\rm pi})^2/r_{Li}^2$, where $c/\omega_{\rm pi} = \sqrt{m_i/m_e}\, c/\omega_{\rm pe}$ is the relativistic ion skin depth (for ions of mass m_i) and r_{Li} is the ion Larmor radius. It follows that the smallest scale we need to resolve in our simulations is the electron Larmor radius $r_{Le} = (\sqrt{m_e/m_i}/\sigma)\, c/\omega_{\rm pe}$, but the temporal and spatial structure of the flow will be regulated by the ion dynamics. We explore the dependence of our results on the choice of the mass ratio m_i/m_e, studying the cases $m_i/m_e = 25$, 100 and 400.

Fig. 4.20 shows that our results do not depend on the value of m_i/m_e. Both the ion (upper panel) and the electron (lower panel) post-shock energy spectra overlap for different m_i/m_e, when computed at the same time (in units of the inverse ion plasma frequency $\omega_{\rm pi}^{-1}$) and for flows with the same stripe wavelength (in units of $c/\omega_{\rm pi}$). Both the main component of the spectrum, resulting from shock-driven magnetic reconnection, and the minor high-energy bump populated by particles accelerated via SDA by the stripe-averaged field $\langle E_z \rangle_\lambda$, are independent of our numerical choice

of m_i/m_e. In other words, our results can be confidently extrapolated to the realistic mass ratio $m_i/m_e = 1836$.

Most importantly, a comparison between the ion and the electron energy spectra suggests that the two species behave essentially in the same way (the spectra of electrons and ions overlap at high energies). In fact, analysis of the flow structure reveals that electrons and ions are already in energy equipartition behind the fast shock. In part, electron heating is due to the effect of the electromagnetic precursor wave radiated by the fast MHD shock, in analogy to what happens for electron-ion shocks propagating in uniform fields (Chapter 3). If electrons enter the fast shock with lower energy than ions, they are further heated by the strong electrostatic potential resulting from charge separation in the shock-compressed magnetic field, again similarly to what we described in Chapter 3. In summary, electrons and ions enter the process of shock-driven magnetic reconnection, and the resulting particle acceleration, with roughly the same energy $\simeq \gamma_0 m_i c^2/2$. Then, the flow behaves as an electron-positron plasma. This is in fact confirmed in the subpanel of Fig. 4.20, which shows that behind the hydrodynamic shock the mean energy of electrons (blue line) and ions (red line) is the same, and their sum (black line) approaches the result expected for complete field dissipation.

Extending the conclusions presented in the previous sections to the case of electron-ion flows, we could then argue that the physics of electron-ion striped winds should be regulated primarily by the combination $\lambda/(r_{Li}\sigma)$. We have confirmed with an extensive set of numerical experiments that $\lambda/(r_{Li}\sigma)$ is indeed the relevant parameter controlling the shape of the spectrum (both of ions and of electrons) in electron-ion striped flows.

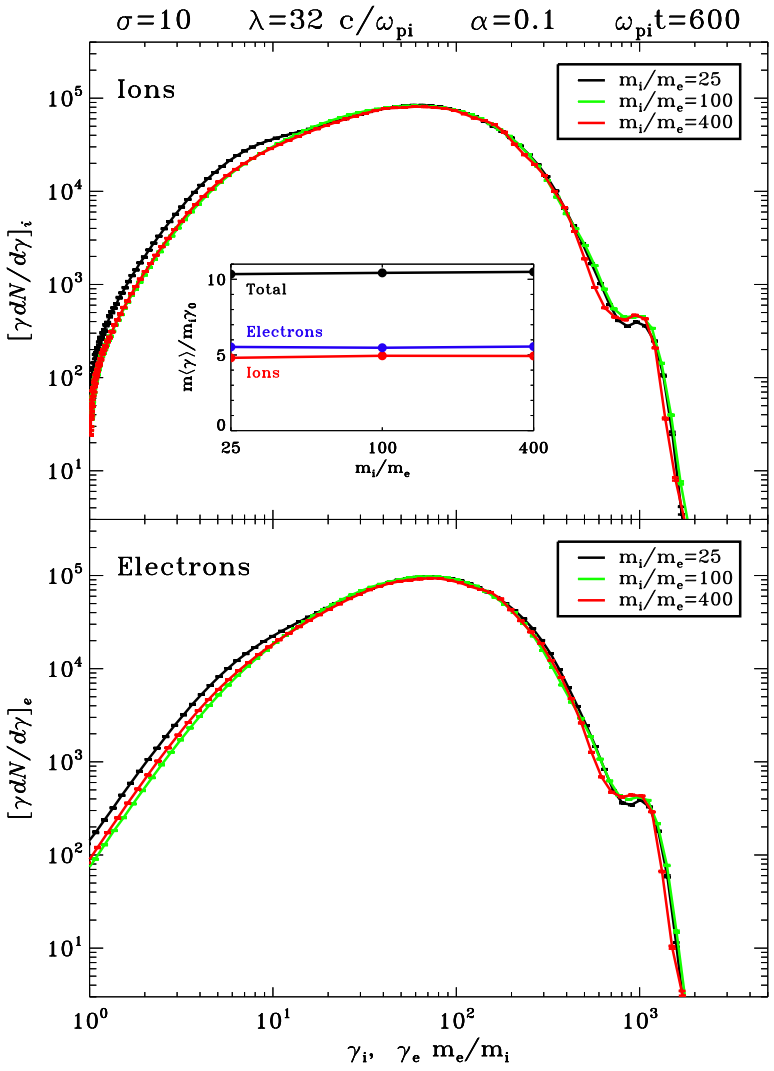


Fig. 4.20.— Particle spectrum behind the termination shock of electron-ion striped winds for different values of the mass ratio: $m_i/m_e = 25$ (black), 100 (green), and 400 (red). The ion spectrum is in the upper panel (x-axis is γ_i), the electron spectrum in the lower panel (x-axis is $\gamma_e m_e/m_i$). In the subpanel, ion (red) and electron (blue) mean energy per particle, in units of $\gamma_0 m_i c^2$; the black line shows the average energy per particle (including both electrons and ions), in units of $\gamma_0 m_i c^2$.

Bibliography

Achterberg, A., Gallant, Y. A., Kirk, J. G., & Guthmann, A. W. 2001, MNRAS, 328, 393

Amato, E. & Arons, J. 2006, ApJ, 653, 325

Arons, J. 2007, ArXiv/astro-ph:0708.1050

Atoyan, A. M. 1999, A&A, 346, L49

Begelman, M. C. & Kirk, J. G. 1990, ApJ, 353, 66

Bell, A. R. 1978, MNRAS, 182, 147

Bietenholz, M. F., Frail, D. A., & Hester, J. J. 2001, ApJ, 560, 254

Bietenholz, M. F., Kassim, N., Frail, D. A., Perley, R. A., Erickson, W. C., & Hajian, A. R. 1997, ApJ, 490, 291

Bietenholz, M. F. & Kronberg, P. P. 1992, ApJ, 393, 206

Blandford, R. & Eichler, D. 1987, Phys. Rep., 154, 1

Blandford, R. D. & Ostriker, J. P. 1978, ApJ, 221, L29

Bogovalov, S. V. 1999, A&A, 349, 1017

Bucciantini, N., Arons, J., & Amato, E. 2011, MNRAS, 410, 381

Buneman, O. 1993, in "Computer Space Plasma Physics", Terra Scientific, Tokyo, 67

Camus, N. F., Komissarov, S. S., Bucciantini, N., & Hughes, P. A. 2009, MNRAS, 400, 1241

Celotti, A. & Ghisellini, G. 2008, MNRAS, 385, 283

Chen, G. & Armstrong, T. P. 1975, in International Cosmic Ray Conference, Vol. 5, 1814–1819

Daughton, W. 1998, J. Geophys. Res., 103, 29429

Del Zanna, L., Amato, E., & Bucciantini, N. 2004, A&A, 421, 1063

Drake, J. F., Swisdak, M., Che, H., & Shay, M. A. 2006, Nature, 443, 553

Gallant, Y. A., Hoshino, M., Langdon, A. B., Arons, J., & Max, C. E. 1992, ApJ, 391, 73

Giannios, D. 2010, MNRAS, 408, L46

Goldreich, P. & Julian, W. H. 1969, ApJ, 157, 869

Gruzinov, A. & Waxman, E. 1999, ApJ, 511, 852

Hester, J. J., Mori, K., Burrows, D., Gallagher, J. S., Graham, J. R., Halverson, M., Kader, A., Michel, F. C., & Scowen, P. 2002, ApJ, 577, L49

Hoshino, M. 2008, ApJ, 672, 940

Hoshino, M., Arons, J., Gallant, Y. A., & Langdon, A. B. 1992, ApJ, 390, 454

Jaroschek, C. H., Lesch, H., & Treumann, R. A. 2004, ApJ, 605, L9

Kennel, C. F. & Coroniti, F. V. 1984a, ApJ, 283, 694

—. 1984b, ApJ, 283, 710

Keshet, U. & Waxman, E. 2005, Physical Review Letters, 94, 111102

Kirk, J. G. & Skjæraasen, O. 2003, ApJ, 591, 366

Komissarov, S. S. & Lyubarsky, Y. E. 2004, MNRAS, 349, 779

Liu, W., Li, H., Yin, L., Albright, B. J., Bowers, K. J., & Liang, E. P. 2011, Physics of Plasmas, 18, 052105

Lyubarsky, Y. & Kirk, J. G. 2001, ApJ, 547, 437

Lyubarsky, Y. & Liverts, M. 2008, ApJ, 682, 1436

Lyubarsky, Y. E. 2003, MNRAS, 345, 153

Medvedev, M. V. & Loeb, A. 1999, ApJ, 526, 697

Nagata, K., Hoshino, M., Jaroschek, C. H., & Takabe, H. 2008, ApJ, 680, 627

Oka, M., Phan, T., Krucker, S., Fujimoto, M., & Shinohara, I. 2010, ApJ, 714, 915

Pétri, J. & Kirk, J. G. 2005, ApJ, 627, L37

Pétri, J. & Lyubarsky, Y. 2007, A&A, 473, 683

Rees, M. J. & Gunn, J. E. 1974, MNRAS, 167, 1

Sikora, M., Stawarz, Ł., Moderski, R., Nalewajko, K., & Madejski, G. M. 2009, ApJ, 704, 38

Sironi, L. & Spitkovsky, A. 2009, ApJ, 698, 1523

—. 2011, ApJ, 726, 75

Spitkovsky, A. 2005, in AIP Conf. Ser., Vol. 801, Astrophysical Sources of High Energy Particles and Radiation, ed. T. Bulik, B. Rudak, & G. Madejski, 345

Spitkovsky, A. 2008, ApJ, 673, L39

Stawarz, Ł., Cheung, C. C., Harris, D. E., & Ostrowski, M. 2007, ApJ, 662, 213

Thompson, C. 2006, ApJ, 651, 333

Webb, G. M., Axford, W. I., & Terasawa, T. 1983, ApJ, 270, 537

Weibel, E. S. 1959, Physical Review Letters, 2, 83

Zenitani, S., Hesse, M., & Klimas, A. 2009, ApJ, 696, 1385

Zenitani, S. & Hoshino, M. 2001, ApJ, 562, L63

—. 2005a, ApJ, 618, L111

—. 2005b, Physical Review Letters, 95, 095001

—. 2007, ApJ, 670, 702

—. 2008, ApJ, 677, 530

Chapter 5

Synthetic Spectra from PIC Simulations of Relativistic Collisionless Shocks

Abstract

We extract synthetic photon spectra from first-principles particle-in-cell simulations of relativistic shocks propagating in unmagnetized pair plasmas. The two basic ingredients for the radiation, namely accelerated particles and magnetic fields, are produced self-consistently as part of the shock evolution. We use the method of Hededal & Nordlund (2005) and compute the photon spectrum via Fourier transform of the electric far-field from a large number of particles, sampled directly from the simulation. We find that the spectrum from relativistic collisionless shocks is entirely consistent with synchrotron radiation in the magnetic fields generated by Weibel instability. We can recover the so-called "jitter" regime only if we artificially reduce the strength of the electromagnetic fields, such that the wiggler parameter $K \equiv qB\lambda/mc^2$ becomes much smaller than unity (B and λ are the strength and scale of the magnetic turbulence, respectively). These findings may place constraints on the origin of non-thermal emission in astrophysics, especially for the interpretation of the hard (harder than synchrotron) low-frequency spectrum of Gamma-Ray Bursts.

1. Introduction

Non-thermal photon spectra from Pulsar Wind Nebulae, jets from Active Galactic Nuclei, Gamma-Ray Bursts and Supernova Remnants are usually explained as synchrotron radiation from a power-law population of particles, presumably accelerated in collisionless shocks. The microphysical details of shock acceleration are still poorly known, however, and are the subject of active research.

Particle-in-cell (PIC) simulations of colliding plasma shells have shown that Weibel instability (Weibel 1959; Medvedev & Loeb 1999; Gruzinov & Waxman 1999) converts the free energy of counter-streaming flows into small scale (skin-depth) magnetic fields (Nishikawa et al. 2003, 2005; Silva et al. 2003; Frederiksen et al. 2004; Hededal et al. 2004). The fields grow to sub-equipartition levels and deflect and randomize the bulk flow, thus creating a shock (Spitkovsky 2005, 2008a; Chang et al. 2008; Keshet et al. 2009). A few percent of the incoming particles repeatedly scatter off the magnetic turbulence created by Weibel instability, and eventually populate a power-law high-energy tail in the particle spectrum behind the shock (Spitkovsky 2008b; Martins et al. 2009b; Sironi & Spitkovsky 2009).

Since most of the magnetic power generated by Weibel instability is concentrated on scales as small as a few plasma skin depths, it has been speculated that the emission mechanism in unmagnetized collisionless shocks may be the so-called "jitter" radiation. Whereas the standard synchrotron emission applies to large-scale fields, the jitter regime is realized if the scale λ of the turbulence is such that the wiggler parameter $K \equiv qB\lambda/mc^2 \ll 1$ (Medvedev 2000, 2006; Fleishman 2006a,b). Jitter radiation has been proposed as a solution for the so-called "line of death" puzzle in Gamma-Ray Burst (GRB) spectra, which below the peak frequency are sometimes harder than expected from synchrotron radiation (e.g., Preece 1998). Since PIC simulations self-consistently provide both the strength and the spatial structure of electromagnetic fields, as well as the particle distribution, it is possible to calculate the photon spectrum from first principles, thus determining whether the emission is more synchrotron-like or jitter-like.

In this chapter, we present synthetic spectra extracted from PIC simulations of relativistic unmagnetized collisionless shocks. In §2 we show the simulated shock structure and particle energy spectrum. The numerical technique that we employ to compute the photon spectrum is described in §3.1, and in §3.2 we show our results. In §4 we discuss the implications of our findings for the interpretation of the non-thermal radiation from astrophysical sources.

2. Shock Structure and Particle Energy Spectrum

We use the three-dimensional (3D) electromagnetic PIC code TRISTAN-MP (Buneman 1993; Spitkovsky 2005) to simulate a relativistic shock propagating into an unmagnetized pair plasma. The shock is triggered by reflecting an incoming cold "upstream" flow off a conducting wall at $x = 0$ (e.g., Sironi & Spitkovsky 2009). The simulation is performed in the "wall" or "downstream" frame. The incoming flow propagates along $-\hat{\boldsymbol{x}}$ with Lorentz factor $\gamma_0 = 15$, and the shock moves along $+\hat{\boldsymbol{x}}$.

To follow the shock evolution for longer times with fixed computational resources, we mostly use a 2D simulation domain in the xy plane. However, as we show in §3.2, our main conclusions still hold for a fully 3D simulation. In the case of an unmagnetized 2D shock, only the in-plane components of the velocity, current and electric field, and only the out-of-plane component of the magnetic field are present. Each computational cell is initialized with 16 particles per species. The relativistic plasma skin depth for the upstream flow ($c/\omega_{\rm p}$) is resolved with 10 cells and the simulation timestep is $\Delta t = 0.045\,\omega_{\rm p}^{-1}$. The simulation box is $100\,c/\omega_{\rm p}$ wide (along y) and, at the final simulation time $\omega_{\rm p}t = 4500$, it is $4500\,c/\omega_{\rm p}$ long (along x).

In Fig. 5.1 we show the internal structure of the shock at time $\omega_{\rm p}t = 2250$. The 2D plots of number density (Fig. 5.1a), magnetic energy (Fig. 5.1c) and electric energy (Fig. 5.1d) show the filaments produced by Weibel instability in the upstream region, with characteristic transverse scale $\sim 10\,c/\omega_{\rm p}$. The magnetic filaments are advected with the upstream flow, so in the simulation frame their magnetic and electric components are comparable (black and red line in Fig. 5.1e, respectively). At the shock ($x_{\rm sh} \simeq 1030\,c/\omega_{\rm p}$ at time $\omega_{\rm p}t = 2250$), the filaments merge and the magnetic energy density peaks at $\sim 10\%$ of the upstream kinetic energy density (Fig. 5.1e). The magnetic field decays farther downstream, where the field is confined within islands of typical scale $\sim 20\,c/\omega_{\rm p}$ (Fig. 5.1c).

The particle energy spectrum behind the shock ($x_{\rm sh} - 400\,c/\omega_{\rm p} < x < x_{\rm sh}$; thin solid line in Fig. 5.1h) consists of a relativistic Maxwellian and a high-energy tail, which can be fitted as a power-law of index $\alpha = 2.5$ (dashed line in Fig. 5.1h) with an exponential cutoff (Spitkovsky 2008b). In the upstream spectrum ($x_{\rm sh} < x < x_{\rm sh} + 400\,c/\omega_{\rm p}$; dotted line in Fig. 5.1h), the unshocked beam populates the low-energy peak at $p \equiv \gamma\beta \simeq 15$, whereas the shock-accelerated returning particles with $\gamma\beta_{\rm x} > 0$ (see the hot, diffuse population in Fig. 5.1f) appear as as a high-energy bump. Since at the highest energies the particles ahead of the shock account for nearly half of the total

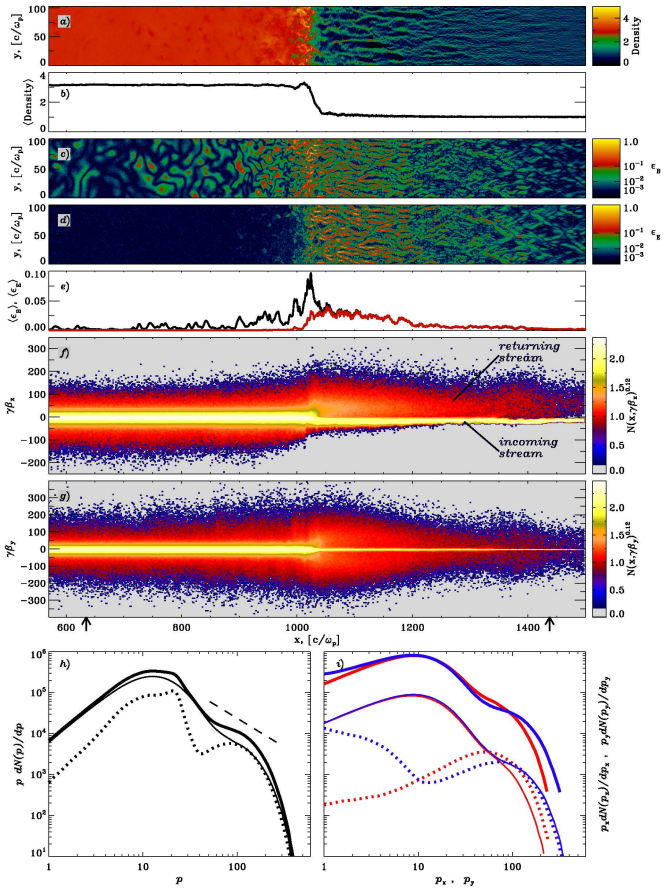


Fig. 5.1.— Shock structure and particle energy spectra at time $\omega_{\rm p}t = 2250$. a) Number density in the simulation plane, normalized to the upstream density. b) Transversely-averaged density. c)-d) Magnetic and electric energy density in the simulation plane, normalized to the upstream kinetic energy density. e) Transversely-averaged magnetic (black) and electric (red) energy density. f) Longitudinal phase space for positrons, shown as a 2D histogram (electron phase space is nearly identical). g) Transverse phase space for positrons. h) Positron spectrum $p\,{\rm d}N/{\rm d}p$ versus p (thick solid line; $p \equiv \gamma\beta = [(\gamma\beta_{\rm x})^2 + (\gamma\beta_{\rm y})^2]^{1/2}$ is the particle 4-velocity) in the region $x_{\rm sh} - 400\,c/\omega_{\rm p} < x < x_{\rm sh} + 400\,c/\omega_{\rm p}$ (delimited by the arrows at the bottom of panel (g)), and the relative contributions of particles behind (thin solid line) and ahead (dotted line) of the shock. The dashed line corresponds to a power-law distribution ${\rm d}N/{\rm d}p \propto p^{-\alpha}$ with $\alpha = 2.5$. i) Positron spectrum $p_{\rm x}\,{\rm d}N/{\rm d}p_{\rm x}$ versus $p_{\rm x}$ (with $p_{\rm x} \equiv \gamma\beta_{\rm x}$) for particles with $p_{\rm x} > 0$ (red lines), and positron spectrum $p_{\rm y}\,{\rm d}N/{\rm d}p_{\rm y}$ versus $p_{\rm y}$ (with $p_{\rm y} \equiv \gamma\beta_{\rm y}$) for particles with $p_{\rm y} > 0$ (blue lines). The line style coding is the same as in panel (h). The total spectrum around the shock (thick solid lines) is shifted upward by a factor of 10 for clarity.

census (compare dotted and thin solid lines in Fig. 5.1h), the high-energy tail in the total spectrum ($x_{\rm sh} - 400\,c/\omega_{\rm p} < x < x_{\rm sh} + 400\,c/\omega_{\rm p}$; thick solid line in Fig. 5.1h) is significantly flatter than in the downstream spectrum.

3. Synthetic Photon Spectrum

3.1. Numerical Technique

We summarize the method introduced by Hededal (2005) and Hededal & Nordlund (2005) (see also Nishikawa 2009; Martins et al. 2009a) to extract synthetic spectra from simulations of collisionless shocks. The electric far-field from a particle with charge q, velocity $\boldsymbol{v} = \boldsymbol{\beta}c$ and acceleration $\dot{\boldsymbol{v}} = \dot{\boldsymbol{\beta}}c$ is (Jackson 1999)

$$\boldsymbol{E}(\boldsymbol{x},t) = \frac{q}{c}\left[\frac{\boldsymbol{n}\times\{(\boldsymbol{n}-\boldsymbol{\beta})\times\dot{\boldsymbol{\beta}}\}}{(1-\boldsymbol{\beta}\cdot\boldsymbol{n})^3 R}\right]_{\rm ret}\,, \tag{5-1}$$

where the unit vector $\boldsymbol{n}$ points toward the observer, at distance R from the emitting particle. Here, the quantity in square brackets is to be evaluated at the retarded time $t' = t - R(t')/c$. The photon spectrum is then computed via the Fourier transform of Poynting flux associated with the field in eq. (5-1). The energy $\mathrm{d}W$ received per unit solid angle $\mathrm{d}\Omega$ (around the direction $\boldsymbol{n}$) and per unit frequency $\mathrm{d}\omega$ can be computed as (Jackson 1999)

$$\frac{\mathrm{d}^2 W}{\mathrm{d}\Omega \mathrm{d}\omega} = \frac{q^2}{4\pi^2 c}\left|\int_{-\infty}^{+\infty}\frac{\boldsymbol{n}\times\{(\boldsymbol{n}-\boldsymbol{\beta})\times\dot{\boldsymbol{\beta}}\}}{(1-\boldsymbol{\beta}\cdot\boldsymbol{n})^2}e^{i\omega(t'-\boldsymbol{n}\cdot\boldsymbol{r}(t')/c)}\mathrm{d}t'\right|^2 \tag{5-2}$$

where $\boldsymbol{r}(t')$ is the particle trajectory. Here, we neglected the Tsytovich-Razin effect due to the dispersive properties of the plasma (Rybicki & Lightman 1979).

In PIC simulations we know the positions, velocities and accelerations of simulation particles with time resolution $\Delta t = 0.045\,\omega_{\rm p}^{-1}$. In order to accurately compute the integral in eq. (5-2), we interpolate the orbit of the selected particles so that to achieve an effective timestep of $0.1\,\Delta t$. For a given choice of $\boldsymbol{n}$, we can then integrate eq. (5-2) to obtain the photon spectrum from each particle. Assuming that the far-fields by different particles are phase-uncorrelated, the total spectrum will be the sum of the spectra of individual particles.

We have implemented eq. (5-2) and tested it for the cases of synchrotron, bremsstrahlung and wiggler/undulator radiation, finding good agreement with analytic solutions. Following Hededal (2005) and Hededal & Nordlund (2005), we have traced test particle orbits in magnetic turbulence

with a *prescribed* wave spectrum, and we have verified the transition from the synchrotron to the jitter regime as the wiggler parameter K becomes smaller than unity.

We now present the photon spectrum resulting from $\sim 10,000$ particles moving in the electromagnetic fields *self-consistently* produced in our PIC simulations. We typically follow the particle trajectories over a time $T = 3000\,\Delta t = 135\,\omega_{\rm p}^{-1}$, which is long enough to reach convergence in the shape of the spectrum, but short compared to the characteristic time of the shock evolution. Thus, the spectrum we obtain may be regarded as "instantaneous." We compute the spectrum for both head-on emission ($\boldsymbol{n} = \hat{\boldsymbol{x}}$) and edge-on emission ($\boldsymbol{n} = \hat{\boldsymbol{y}}$). Our calculations are performed in the downstream fluid frame. An additional Lorentz transformation is required if the downstream medium is moving with respect to the observer.

3.2. Results

We have studied the emission from relativistic shocks with two experiments. First, we have injected a power-law distribution of particles into the downstream region, and we have traced their orbits in a fixed snapshot of the fields of the simulation. The photon spectrum is then calculated with the technique described above. Second, we compute the spectrum from trajectories of particles extracted directly from the simulation. In this case, both the particle distribution and the electromagnetic turbulence are provided by the PIC simulation, and the resulting emission will be the *self-consistent* spectrum from our relativistic shock.

Fig. 5.2 shows the emission spectrum from a power-law population of particles behind the shock ($x_{\rm sh} - 400\,c/\omega_{\rm p} < x < x_{\rm sh}$). We inject a 2D isotropic distribution with $\mathrm{d}N/\mathrm{d}p \propto p^{-\alpha}$ ($\alpha = 2.5$), between $p_{\rm min} = 50$ and $p_{\rm max} = 2000$. The spectral index and the lower cutoff of the distribution are chosen to mimic the high-energy tail of the downstream particle spectrum in the simulation (thin solid line in Fig. 5.1h). The injected particles are evolved in a fixed snapshot of the electromagnetic fields from the PIC simulation, at time $\omega_{\rm p} t = 2250$.

Since electric fields are negligible in the downstream medium (see Fig. 5.1d-e), the resulting photon spectrum will probe the strength and structure of the magnetic fields, and it will clarify which regime – synchrotron or jitter – is appropriate to describe the particle emission. The solid lines in Fig. 5.2 show that the spectrum can be well approximated by two power-law segments. Regardless of the observer's direction $\boldsymbol{n}$ (red line for $\boldsymbol{n} = \hat{\boldsymbol{x}}$, blue line for $\boldsymbol{n} = \hat{\boldsymbol{y}}$), the slope at

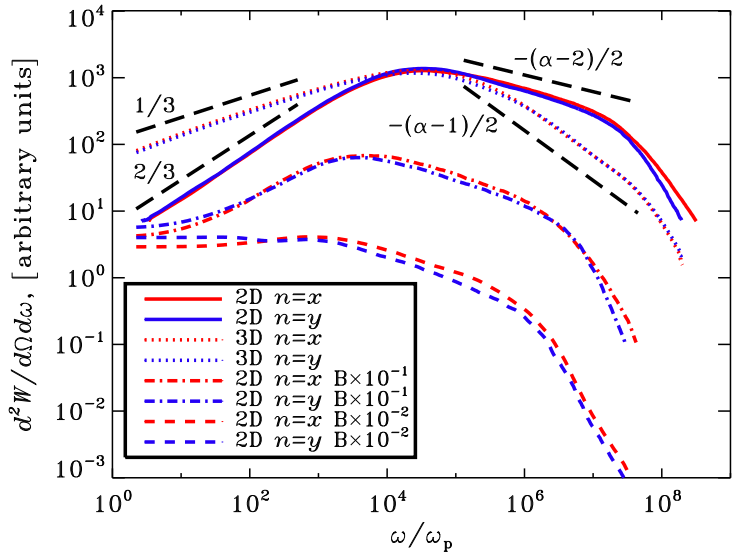


Fig. 5.2.— Photon spectrum from a power-law population of particles ($\mathrm{d}N/\mathrm{d}p \propto p^{-\alpha}$ between $p_{\rm min} = 50$ and $p_{\rm max} = 2000$, with $\alpha = 2.5$) injected behind the shock ($x_{\rm sh} - 400\,c/\omega_{\rm p} < x < x_{\rm sh}$). The particles are traced in the fields of a 2D (solid lines) or 3D (dotted lines) simulation at time $\omega_{\rm p}t = 2250$. The 3D spectra are arbitrarily shifted in normalization to compare with the 2D case. Dot-dashed and dashed lines: photon spectrum from the same particle population, but evolved in 2D electromagnetic fields artificially reduced by a factor of 10 (dot-dashed lines) or 100 (dashed lines). Red lines are for head-on emission ($\boldsymbol{n} = \hat{\boldsymbol{x}}$), blue lines for edge-on emission ($\boldsymbol{n} = \hat{\boldsymbol{y}}$).

low frequencies is remarkably close to 2/3 (black long-dashed line in Fig. 5.2), as expected for synchrotron emission from a 2D particle distribution (e.g., Jackson 1999). Similarly, when we inject a 3D isotropic particle distribution with the same energy spectrum as above in the post-shock region of a 3D simulation, the photon spectrum we obtain (dotted lines in Fig. 5.2) scales as $\propto \omega^{1/3}$ at low frequencies, in agreement with the synchrotron theory. In N dimensions, the high-frequency slope should be $-[\alpha-(4-N)]/2$, which reduces to $-(\alpha-2)/2=-0.25$ in 2D and $-(\alpha-1)/2=-0.75$ in 3D, in agreement with our spectra (see black long-dashed lines). The similarity between the cases $\boldsymbol{n}=\hat{\boldsymbol{x}}$ and $\boldsymbol{n}=\hat{\boldsymbol{y}}$ suggests, given the isotropy of the injected particle distribution, that the downstream magnetic fluctuations are spatially isotropic, as seen in Fig. 5.1c.[1]

A transition to the jitter regime should appear when the wiggler parameter K becomes significantly smaller than unity (Medvedev 2000, 2006; Fleishman 2006a,b). We have tried to artificially lower the value of K by decreasing the strength of the electromagnetic fields, by a factor of 10 (dot-dashed lines in Fig. 5.2, for a 2D simulation) and 100 (dashed lines in Fig. 5.2, for a 2D simulation). The high-frequency power-law decreases in intensity and shifts to lower frequencies.[2] The low-frequency spectrum becomes softer for decreasing B, approaching the flat slope expected in the jitter regime when the shock is viewed edge-on. When we impose a magnetic field spectrum of the form $\boldsymbol{B}(\boldsymbol{k})\propto\delta(\hat{\boldsymbol{k}}-\hat{\boldsymbol{k}}_0)$ (here, $\hat{\boldsymbol{k}}_0$ is a fixed unit vector in k-space) with wiggler parameter $K\ll 1$, we are able to recover the hard low-frequency slope ($\propto\omega^1$) discussed by Medvedev (2000) for head-on emission from shocks. However, we do not observe it for the downstream turbulence self-consistently generated in the simulation, suggesting that the magnetic field fluctuations are not sufficiently ordered.

In the jitter (small deflection) regime, a particle of Lorentz factor γ remains in phase with the radiated photons (of frequency ω) for a time $\sim 4\pi\gamma^2(\omega_{\rm p}/\omega)\,\omega_{\rm p}^{-1}$. By following the particle trajectories over a limited timespan T, we cannot adequately resolve the frequencies below $\omega_{\rm T}\sim 4\pi\gamma_{\rm min}^2/T$.[3] Here, $\gamma_{\rm min}\simeq p_{\rm min}=50$ is the lower cutoff of the power-law distribution of particles. For the jitter spectra in Fig. 5.2 (dashed lines), we have integrated the particle orbits for

[1]The slight difference at high frequencies between $\boldsymbol{n}=\hat{\boldsymbol{x}}$ and $\boldsymbol{n}=\hat{\boldsymbol{y}}$ in the 2D results is due to the residual electric fields at $x\lesssim x_{\rm sh}$ (see Fig. 5.1d-e), and it disappears if electric fields are neglected while computing the spectrum.

[2]We remark that, while in the synchrotron regime both the amplitude and the peak frequency of the spectrum scale linearly with the field strength B, in the jitter regime ($K\ll 1$) the spectral intensity is proportional to $K^2\propto B^2$ and the peak frequency does not depend on B.

[3]We thank J. Kirk for bringing this point to our attention.

$T = 450\,\omega_{\rm p}^{-1}$, so that our results are reliable for $\omega \gtrsim \omega_{\rm T} \simeq 70\,\omega_{\rm p}$.[4] We have also tested that the shape of the spectrum does not appreciably depend on T (we have tried up to $T = 1350\,\omega_{\rm p}^{-1}$).

In Fig. 5.3 we show the photon spectrum resulting from a sample of particles extracted directly from a 2D PIC simulation, followed near $\omega_{\rm p}t = 2250$ in the *time-varying* electromagnetic fields of the simulation. The selected particles start in the region $x_{\rm sh} - 400\,c/\omega_{\rm p} < x < x_{\rm sh} + 400$. The photon spectrum (thick solid lines in Fig. 5.3; red for $\boldsymbol{n} = \hat{\boldsymbol{x}}$, blue for $\boldsymbol{n} = \hat{\boldsymbol{y}}$) confirms that the emission occurs in the synchrotron regime, as the 2/3 slope at low frequencies suggests (black dashed line).

Most of the low-frequency emission is powered by the thermal particles behind the shock, with 4-velocity $p \leq 50$ (dot-dashed lines in Fig. 5.3a). At high frequencies, the emission along $\hat{\boldsymbol{y}}$ (blue solid line in Fig. 5.3a) is more powerful than along $\hat{\boldsymbol{x}}$ (red solid line in Fig. 5.3a). This reflects the fact that the highest energy particles are grazing the shock surface (Spitkovsky 2008b) and they mostly contribute to the emission along $\hat{\boldsymbol{y}}$. The difference between $\boldsymbol{n} = \hat{\boldsymbol{x}}$ and $\boldsymbol{n} = \hat{\boldsymbol{y}}$ in Fig. 5.3a can be directly related to the difference between the particle spectra $p_{\rm x}\,{\rm d}N/{\rm d}p_{\rm x}$ and $p_{\rm y}\,{\rm d}N/{\rm d}p_{\rm y}$ shown in Fig. 5.1i (thick solid lines, respectively red and blue). In fact, due to relativistic beaming, an observer located along $\hat{\boldsymbol{x}}$ will more likely detect the radiation from particles with $p_{\rm x}/p_{\rm y} \gg 1$ (and $p_{\rm x} > 0$), whereas the particles with $p_{\rm y}/p_{\rm x} \gg 1$ (and $p_{\rm y} > 0$) will contribute more to the emission along $\hat{\boldsymbol{y}}$.

As shown in Fig. 5.3b, the spectrum of downstream particles (thin solid lines) dominates the total emission (thick solid lines) in the low-frequency bump, whereas the high frequencies are mostly powered by the upstream particles (dotted lines). This is due to a combination of two effects: first, at the highest particle energies the upstream particles may dominate by number (compare red dotted and thin solid lines in Fig. 5.1i); second, the electromagnetic energy in the upstream region is on average larger than in the downstream, since upstream electric fields are as strong as magnetic fields, and also because the field is built up ahead of the shock on a length scale larger than the decay scale in the downstream (Fig. 5.1e). It is worth pointing out that, at low frequencies, the spectra of both downstream and upstream particles can be fitted with the synchrotron slope 2/3 (black dashed lines).

[4]The spectrum at lower frequencies ($\omega \lesssim \gamma_{\rm min}\,\omega_{\rm p} \simeq 50\,\omega_{\rm p}$) is also likely to be exponentially suppressed by the Tsytovich-Razin effect.

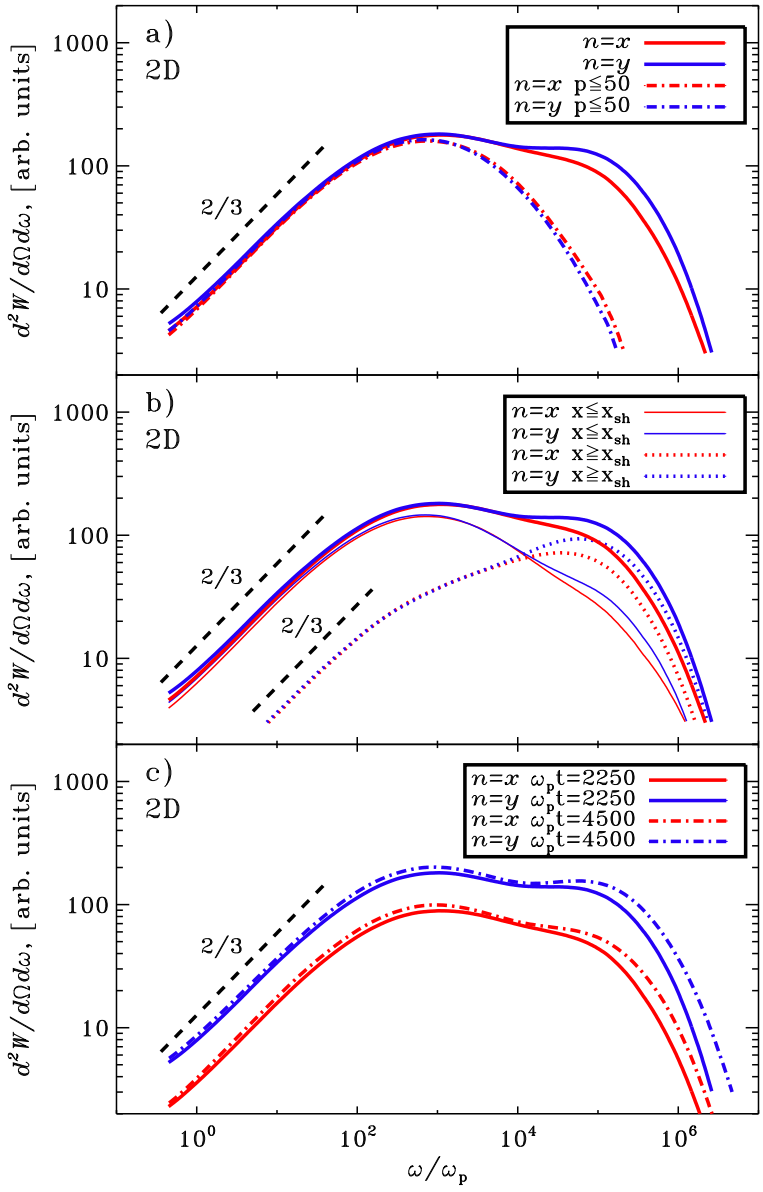


Fig. 5.3.— Photon spectrum (thick solid lines, in all panels) for particles extracted from a 2D PIC simulation, evolved in time-varying electromagnetic fields around $\omega_{\rm p}t = 2250$. Red lines are for head-on emission ($\boldsymbol{n} = \hat{\boldsymbol{x}}$), blue lines for edge-on emission ($\boldsymbol{n} = \hat{\boldsymbol{y}}$). a) Dot-dashed lines only include the contribution of thermal particles, with 4-velocity $p \leq 50$. b) Relative contribution of downstream (thin solid lines) and upstream (dotted lines) particles. c) Time evolution of the photon spectrum, from $\omega_{\rm p}t = 2250$ (thick solid lines) to $\omega_{\rm p}t = 4500$ (dot-dashed lines). Here, the case $\boldsymbol{n} = \hat{\boldsymbol{x}}$ has been shifted downward by a factor of 2 for clarity.

Fig. 5.3c shows the time evolution of the total photon spectrum, from $\omega_{\rm p}t = 2250$ (thick solid lines) to $\omega_{\rm p}t = 4500$ (dot-dashed lines). The upper cutoff in the particle energy spectrum grows with time (Spitkovsky 2008b; Sironi & Spitkovsky 2009), and as a result the photon spectrum extends to higher frequencies. Also, the spectral intensity increases, especially at high frequencies, due to stronger electromagnetic fields and a larger number of accelerated particles.

4. Discussion

We have computed synthetic photon spectra from 2D PIC simulations of relativistic collisionless shocks by following a sample of simulation particles in the time-varying fields of the simulation. The low-frequency part of the spectrum scales as $\propto \omega^{2/3}$, as expected for synchrotron emission from a 2D particle distribution. Although the electromagnetic fluctuations generated by Weibel instability are on small (skin-depth) scales, the particle emission does not occur in the jitter regime.

In retrospect, this is not surprising. The characteristic length scale of the magnetic turbulence is $\lambda \gtrsim 10\, c/\omega_{\rm p}$, and in the shock region, where most of the emission is produced, the magnetic energy reaches a fraction $\epsilon_{\rm B} \simeq 0.1$ of the upstream bulk kinetic energy. This implies $r_{\rm L}/(\, c/\omega_{\rm p}) = \epsilon_{\rm B}^{-1/2} \simeq 3$, where $r_{\rm L}$ is the relativistic Larmor radius of a particle moving with the upstream flow. It follows that $K = \lambda/(r_{\rm L}/\gamma_0) \simeq 3\,\gamma_0$, so that the condition $K \ll 1$ for jitter radiation is unlikely to be satisfied, even for moderately relativistic shocks. Indeed, the photon spectrum we obtain from a $\gamma_0 = 3$ shock is still consistent with synchrotron radiation. For electron-ion shocks, the wiggler parameter will be m_p/m_e times larger, and we are even deeper in the synchrotron regime. Further downstream from the shock, although the magnetic field strength decays, the value of K does not significantly change, since short-wavelength modes are progressively damped (Chang et al. 2008), and the characteristic scale of the turbulence increases.

Although the results presented here mostly employ 2D simulations, we have tested that the main conclusions hold for 3D configurations as well. In particular, we have verified that the photon spectrum from a 3D particle population injected in the downstream turbulence of a fully 3D simulation scales as $\propto \omega^{1/3}$ at low frequencies, in agreement with the synchrotron theory. If the GRB prompt emission results from high-energy particles accelerated in relativistic unmagnetized shocks, it seems that resorting to the jitter regime is not a viable solution for the synchrotron "line of death" puzzle (Preece 1998).

At high frequencies, our results suggest that the contribution of upstream particles to the total emission, which is usually omitted in standard models, is not negligible. It causes the radiation spectrum to be flatter than the corresponding downstream spectrum, thus partly masking the contribution of downstream thermal particles (Giannios & Spitkovsky 2009). This could potentially explain the absence of clear signatures of downstream thermal emission in GRB shocks (Band 1993). Simulations extending to longer times (and higher particle energies) will help to clarify these issues.

We remark that our calculations do not include radiative particle cooling, synchrotron self-absorption and inverse Compton radiation. Still, we have shown that the computation of synthetic spectra from self-consistent PIC simulations provides a powerful tool for studying the origin of astrophysical non-thermal emission.

Bibliography

Band, D. e. a. 1993, ApJ, 413, 281

Buneman, O. 1993, in "Computer Space Plasma Physics", Terra Scientific, Tokyo, 67

Chang, P., Spitkovsky, A., & Arons, J. 2008, ApJ, 674, 378

Fleishman, G. D. 2006a, MNRAS, 365, L11

—. 2006b, ApJ, 638, 348

Frederiksen, J. T., Hededal, C. B., Haugbølle, T., & Nordlund, Å. 2004, ApJ, 608, L13

Giannios, D. & Spitkovsky, A. 2009, ArXiv:astro-ph/0905.1970

Gruzinov, A. & Waxman, E. 1999, ApJ, 511, 852

Hededal, C. 2005, PhD thesis, arXiv:astro-ph/0506559

Hededal, C. & Nordlund, Å. 2005, arXiv:astro-ph/0511662

Hededal, C. B., Haugbølle, T., Frederiksen, J. T., & Nordlund, Å. 2004, ApJ, 617, L107

Jackson, J. D. 1999, Classical Electrodynamics: Third Edition (John Wiley & Sons, Inc.)

Keshet, U., Katz, B., Spitkovsky, A., & Waxman, E. 2009, ApJ, 693, L127

Martins, J. L., Martins, S. F., Fonseca, R. A., & Silva, L. O. 2009a, in SPIE Conf. Ser., Vol. 7359

Martins, S. F., Fonseca, R. A., Silva, L. O., & Mori, W. B. 2009b, ApJ, 695, L189

Medvedev, M. V. 2000, ApJ, 540, 704

—. 2006, ApJ, 637, 869

Medvedev, M. V. & Loeb, A. 1999, ApJ, 526, 697

Nishikawa, K.-I., Hardee, P., Richardson, G., Preece, R., Sol, H., & Fishman, G. J. 2003, ApJ, 595, 555

—. 2005, ApJ, 622, 927

Nishikawa, K.-I. e. a. 2009, in AIP Conf. Ser., Vol. 1133, 235

Preece, R. D. e. a. 1998, ApJ, 506, L23

Rybicki, G. B. & Lightman, A. D. 1979, Radiative Processes in Astrophysics (John Wiley & Sons, Inc.)

Silva, L. O., Fonseca, R. A., Tonge, J. W., Dawson, J. M., Mori, W. B., & Medvedev, M. V. 2003, ApJ, 596, L121

Sironi, L. & Spitkovsky, A. 2009, ApJ, 698, 1523

Spitkovsky, A. 2005, in AIP Conf. Ser., Vol. 801, 345

—. 2008a, ApJ, 673, L39

—. 2008b, ApJ, 682, L5

Weibel, E. S. 1959, Physical Review Letters, 2, 83

Chapter 6

Conclusions and Future Prospects

In this dissertation, I have studied the problem of particle acceleration in relativistic astrophysical shocks by means of *ab initio* particle-in-cell (PIC) plasma simulations. The nonthermal signature of PWNe, GRBs and AGN jets is usually interpreted as synchrotron or inverse Compton emission from a power-law population of relativistic electrons, presumably accelerated at relativistic shocks. Acceleration of protons or heavy ions in the relativistic shocks of GRBs and AGNs is also invoked to explain the flux of Ultra-High Energy Cosmic Rays (UHECRs) detected on Earth. The physical properties of relativistic astrophysical sources can be constrained only by understanding from first principles the micro-physics of particle acceleration in shocks, since this is what ultimately determines the radiative or particle signature observed from these objects.

In the first part of my thesis (Chapters 2 and 3), I have systematically explored by means of PIC simulations the properties of relativistic shocks that propagate in electron-positron and electron-proton plasmas seeded with uniform magnetic fields. I have found that, for highly relativistic flows ($\gamma_0 \gtrsim 5$), efficient particle acceleration occurs only if the pre-shock flow is weakly magnetized (magnetization $\sigma \lesssim 10^{-3}$), or if the upstream field is nearly aligned with the shock direction of propagation (within an angle of $\lesssim 34^\circ/\gamma_0$, as seen from the pre-shock frame). Clearly, the latter constraint is much more serious for highly relativistic shocks (e.g., pulsar winds with $\gamma_0 \gtrsim 10^2$, or the external shocks in early GRB afterglows with $\gamma_0 \gtrsim 10^2$), than for mildly relativistic flows ($\gamma_0 \sim$ a few for internal shocks in GRBs and AGNs, and for the late stages of evolution of GRB afterglows). Trans-relativistic shocks with $\gamma_0 \sim$ a few have not been investigated in this thesis, and are certainly worth of further study.

The simulations I have performed in my thesis have shown that the main properties of the shock are relatively insensitive to the upstream bulk Lorentz factor, provided that the flow is highly relativistic ($\gamma_0 \gtrsim 5$). Preliminary results suggest that the shock physics could change significantly for mildly relativistic flows ($\gamma_0 \lesssim 3$), leading potentially to different conclusions for the efficiency of particle acceleration. Future work on low-γ_0 shocks should address the following questions:

- How does the acceleration efficiency of electrons and ions depend on the upstream conditions, like composition, magnetization, and magnetic obliquity? Recent PIC simulations of non-relativistic quasi-perpendicular shocks (Riquelme & Spitkovsky 2011) have shown signatures of electron injection into a Fermi-like process, but negligible ion acceleration. The opposite seems to hold for non-relativistic quasi-parallel shocks, which are efficient ion accelerators (Gargaté & Spitkovsky 2011). Do the same conclusions apply to trans-relativistic shocks? Any asymmetry found between the acceleration efficiency of electrons and ions may be important for the interpretation of the gamma-ray emission detected by the Fermi telescope, as either inverse Compton by accelerated electrons (leptonic scenario) or π^0 decay from cosmic ray ions (hadronic scenario). Moreover, assessing which shocks can accelerate ions to highly nonthermal energies will help to unveil the sources of UHECRs.

- What is the level of self-generated magnetic turbulence in trans-relativistic shocks? This should be compared with models of nonthermal emission from shock-dominated systems, to give physically-grounded inputs for the phenomenological parameter ϵ_B (defined as the ratio of post-shock magnetic to thermal pressure).

- What are the emission signatures of trans-relativistic shocks? In highly relativistic unmagnetized shocks, I have shown in Chapter 5 that an appreciable contribution to the spectrum comes from high-energy particles streaming ahead of the shock, whose emission is neglected in most models. To understand the emission from GRBs and AGN jets, it is of paramount importance to explore possible observational signatures of the presence of shock-accelerated particles in the upstream medium of low-γ_0 shocks.

In particular, this study of trans-relativistic shocks will clarify whether magnetized ($\sigma \gtrsim 10^{-2}$) perpendicular low-γ_0 shocks, as expected in GRBs and AGN jets, can be efficient particle accelerators. If it turns out that the Fermi process is still suppressed, similarly to what happens in high-γ_0 magnetized perpendicular shocks, then alternative explanations need to be found regarding

the origin of nonthermal particles in GRBs and AGN jets. As I have described in Chapter 4, a valid alternative in the context of pulsar winds is given by shocks propagating in a medium with *alternating* (as opposed to *uniform*) magnetic fields. In a flow with reversing fields, we find that shock-driven reconnection can efficiently transfer the energy of alternating fields to the particles, producing abundant populations of nonthermal particles with flat power-law tails.

The magnetic geometry adopted in Chapter 4, with application to pulsar winds, consists of magnetic stripes of opposite polarity, alternating with the rotation period of the pulsar. In the relativistic jets of GRBs and AGNs, such a regular sequence of periodic field reversals is not to be expected. Rather, due to occasional flips in the field polarity of the accretion flow, a few regions may exist in the jet where the field direction reverses over microscopic plasma scales. In these regions, the field energy is likely to be dissipated, transferred to the particles, and eventually released as radiation. In summary, particle acceleration in relativistic jets is likely to occur within a single current sheet (or a few distant current sheets), in contrast to the sequence of neighboring sheets that characterizes the structure of pulsar winds. Self-consistent PIC simulations will be needed to capture the process of magnetic dissipation and the resulting particle acceleration in this configuration, relevant for GRBs and AGN jets. This work will answer the following questions:

- What are the conditions under which magnetic reconnection can efficiently dissipate the field energy into plasma kinetic energy, in the relativistic jets of GRBs and AGNs? It will be interesting to consider two possible configurations that could lead to fast reconnection: on the one side, reconnection may be driven by the collisions of flows inside the jet; alternatively, the effective gravity force arising in an accelerating jet may distort the pressure equilibrium and facilitate reconnection, as recently proposed by Lyubarsky (2010).

- Can magnetic reconnection accelerate particles up to relativistic energies? In particular, what determines the maximum energy of accelerated particles in the reconnection region, and how does the maximum energy depend on the properties of the jet flow? Thanks to the numerical technique described in Chapter 5, I can extract *ab initio* radiation spectra from PIC simulations, by Fourier transform of the electric far field from a large sample of simulation particles. With this powerful tool, I will be able to assess whether particle acceleration via magnetic reconnection can ultimately explain the prominent nonthermal signatures of astrophysical jets in GRBs and AGNs.

- What is the velocity of plasma inflow into the reconnection region (the so-called reconnection rate) as a function of magnetization, and what are its implications for the rate of field dissipation in astrophysical jets? Also, can relativistic bulk flows in the "exhausts" of reconnection explain the fast time variability observed from GRBs and AGN jets?

In conclusion, I have shown that studies of particle acceleration in relativistic shocks performed by means of first-principle PIC simulations are of paramount importance to provide physically-grounded inputs for astrophysical models of high-energy emission. By systematically exploring the regimes of shock parameters where the Fermi mechanism can operate, I have placed important constraints on the composition and magnetization of the sources of nonthermal radiation. By extracting *ab initio* radiation spectra from my PIC simulations of relativistic shocks, I have been able to interpret from first principles the nonthermal signature of astrophysical sources. The work done in my thesis will help, in concert with dedicated high-energy experiments (Chandra, Fermi, and TeV telescopes), to unveil the nature of astrophysical sources of nonthermal emission.

Bibliography

Gargaté, L. & Spitkovsky, A. 2011, arXiv:astro-ph/1107.0762

Lyubarsky, Y. 2010, ApJ, 725, L234

Riquelme, M. A. & Spitkovsky, A. 2011, ApJ, 733, 63