

2010

Particle-in-cell simulation of astrophysical plasmas: probing the origin of cosmic rays

Thomas Alan Stroman
Iowa State University

Follow this and additional works at: <http://lib.dr.iastate.edu/etd>

 Part of the [Physics Commons](#)

Recommended Citation

Stroman, Thomas Alan, "Particle-in-cell simulation of astrophysical plasmas: probing the origin of cosmic rays" (2010). *Graduate Theses and Dissertations*. 11780.
<http://lib.dr.iastate.edu/etd/11780>

This Dissertation is brought to you for free and open access by the Graduate College at Iowa State University Digital Repository. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

**Particle-in-cell simulation of astrophysical plasmas:
probing the origin of cosmic rays**

by

Thomas Alan Stroman

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Major: Astrophysics

Program of Study Committee:
Martin Pohl, Major Professor
Frank Krennrich
Steve Kawaler
James Cochran
Elgin Johnston

Iowa State University

Ames, Iowa

2010

Copyright © Thomas Alan Stroman, 2010. All rights reserved.

TABLE OF CONTENTS

| | |
|-------------------------------------------------------------------------------------------------------------|----|
| LIST OF TABLES | iv |
| LIST OF FIGURES | v |
| CHAPTER 1. Introduction | 1 |
| CHAPTER 2. Historical and theoretical background | 4 |
| 2.1 The discovery of cosmic rays | 4 |
| 2.2 Cosmic rays: what we know | 5 |
| 2.3 Supernovae: the cosmic-ray connection | 9 |
| 2.4 Diffusive shock acceleration (DSA) | 11 |
| 2.5 The limits of DSA | 15 |
| 2.6 A special case: Bell’s non-resonant streaming instability | 17 |
| CHAPTER 3. Particle-in-cell plasma simulation | 22 |
| 3.1 The particle-in-cell method with TRISTAN | 22 |
| 3.1.1 The necessity of compromise | 25 |
| 3.1.2 PIC versus alternative approaches | 27 |
| 3.2 Parallel simulation with THISMPI | 29 |
| CHAPTER 4. Amplification of magnetic fields by Bell’s non-resonant streaming instability | 31 |
| 4.1 Bell’s non-resonant streaming instability | 31 |
| 4.2 First results of PIC simulations | 35 |
| CHAPTER 5. Kinetic simulations of turbulent magnetic-field growth by streaming cosmic rays | 38 |

| | | |
|----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------|-----------|
| 5.1 | Abstract | 38 |
| 5.2 | Introduction | 38 |
| 5.3 | Simulation setup | 41 |
| 5.3.1 | Model | 41 |
| 5.3.2 | Implementation | 42 |
| 5.3.3 | Run-specific parameters | 45 |
| 5.4 | Results and discussion | 46 |
| 5.5 | Summary and conclusions | 54 |
| CHAPTER 6. Do cosmic rays affect instabilities in the transition layer of nonrelativistic collisionless shocks? | | 58 |
| 6.1 | Abstract | 58 |
| 6.2 | Introduction | 59 |
| 6.3 | Objectives and approach | 61 |
| 6.3.1 | Parameter-space configurations | 61 |
| 6.3.2 | Simulation setup | 63 |
| 6.3.3 | Behavioral benchmarks | 64 |
| 6.4 | Comparison with analytical beam-plasma predictions | 65 |
| 6.5 | Results | 69 |
| 6.5.1 | Behavior including cosmic rays | 70 |
| 6.6 | Discussion and conclusions | 76 |
| CHAPTER 7. Conclusion | | 80 |
| BIBLIOGRAPHY | | 83 |

LIST OF TABLES

| | | |
|-----------|----------------------------------------------------------|----|
| Table 5.1 | Parameters and selected results of simulations | 57 |
|-----------|----------------------------------------------------------|----|

LIST OF FIGURES

| | | |
|------------|------------------------------------------------------------------------------|----|
| Figure 2.1 | Cosmic-ray energy spectrum | 7 |
| Figure 2.2 | HST image of Supernova 1994D | 10 |
| Figure 2.3 | X-ray image of SN1006 | 12 |
| Figure 2.4 | Sketch of first-order Fermi acceleration | 21 |
| Figure 4.1 | Drift velocities in Niemiec et al. (2008) | 36 |
| Figure 5.1 | Evolution of magnetic field components | 44 |
| Figure 5.2 | Magnetic field and plasma density, $t_{\gamma_{max}} = 7.0$ | 47 |
| Figure 5.3 | Magnetic field and plasma density, $t_{\gamma_{max}} = 14.0$ | 48 |
| Figure 5.4 | Transverse and drift velocity of plasma | 50 |
| Figure 5.5 | The phase-space distribution of cosmic rays | 55 |
| Figure 6.1 | Growth rates of early unstable modes | 68 |
| Figure 6.2 | Selected results for the “negligible-absent” family of simulations | 71 |
| Figure 6.3 | Selected results for the “negligible-weak” family of simulations | 73 |
| Figure 6.4 | Selected results for the “negligible-strong” family of simulations | 74 |
| Figure 6.5 | Cosmic-ray representation effects | 75 |

CHAPTER 1. Introduction

Humans have observed astronomical phenomena for millennia, but it is only in the last few centuries that astronomers have begun to offer truly satisfactory explanations for many of these. But often the improvements in our understanding of the universe, and in the tools we employ to study it, bring with them the discoveries of new phenomena to explain. Consequently, the science of astrophysics is as active as ever, with each newly uncovered mystery more challenging than its predecessors. One such mystery concerns the origin of the so-called “cosmic rays,” which remains an open question despite several decades of scientific inquiry. Cosmic rays are electrically charged subatomic particles moving nearly at the speed of light, constantly hitting the Earth’s atmosphere from all directions.

Having spent the past few years participating in the ongoing research efforts related to cosmic rays, I can look back and say truthfully that it has been a fascinating and rewarding journey. I thus have three goals as I write this dissertation: first, the traditional goal: to convince my audience that I have performed research at a level worthy of the doctoral degree for which I have done this work; second, to equip my audience with at least a basic understanding of the present state of cosmic-ray physics; and finally, to convey to my audience the sense of excitement and enjoyment that attends this research.

To set the stage for my research, I will describe in Chapter 2 the historical and observational background of cosmic-ray physics, and some key portions of the relevant theoretical approaches. The last few decades of research have been active, and I will provide an overview of the various topics being investigated by the cosmic-ray community. If I perform my task well, I will have painted a picture of a large-scale collaborative effort in which the small contributions of individuals combine to produce an understanding, continuously being refined, of cosmic rays.

There are many facets of modern cosmic-ray research, both theoretical and experimental. My own efforts have focused on the former, attempting to improve our understanding of the distant environments in which some cosmic rays are thought to originate. These efforts have centered on the use of simulations performed on powerful computers. Of course, the validity of the results obtained depends on the accuracy and applicability of the simulations that produce them. Because it is my principal resource in these investigations, I will describe the particle-in-cell simulation technique – its capabilities and its limitations – in Chapter 3.

I have used particle-in-cell simulations to model two distinct types of physical systems relevant to cosmic rays. In both systems, interacting astrophysical plasmas evolve in ways that are difficult or impossible to describe analytically, but can be explored numerically (i.e., by simulation). The first system is the interstellar medium just outside (or “upstream” of) the expanding supernova shock wave produced in the explosive death of a star, which plays a critical role in the acceleration process by which ordinary particles are thought to become cosmic rays. The second is the interior of such a shock, which is a very compact region by astrophysical standards, but nevertheless has internal structure that the particle-in-cell technique is particularly well-suited to probe.

In Chapter 4, I will describe the motivation for the interstellar-medium studies, which was the prediction of significant magnetic-field amplification in the presence of a cosmic-ray current. I will also discuss the results obtained by my research group in the early days of my involvement, which were presented in our first journal publication on the topic, Niemiec et al. (2008). Because we did not see the predicted effect, our results were somewhat at odds with those being reported by other contemporary investigators. Ours was among the first attempts at applying the particle-in-cell method to that particular problem, however, so it was not automatically understood that *our* results were the incorrect ones. Our work clearly was not finished.

After our publication of the first paper, I made some modifications to the particle-in-cell program and began running my own simulations, still aimed at studying the evolution of the magnetic field upstream of supernova shocks, and at clarifying the results of our previous study.

In my simulations, which differed from our earlier ones chiefly by the choices of parameters explored, we did see the magnetic-field amplification that had been predicted, bringing us into agreement with the broader cosmic-ray research community. This led to our second publication, [Stroman et al. \(2009\)](#), which I have adapted for inclusion in this dissertation as Chapter 5.

In parallel to my simulations for that publication, our group was exploring whether the effect might be observed in a different environment, in which the shock moved at much greater speed. We found that the behavior could resemble either the prediction or that observed in our 2008 paper, depending on the assumptions made. This work was published as [Niemić et al. \(2010\)](#).

I then shifted my focus to the second system I mentioned above, the shock itself: the comparatively thin region dividing upstream from downstream, so narrow that in practice it is often (safely) approximated as an abrupt discontinuity in the bulk plasma properties such as density, speed, magnetization, and temperature. As a shock propagates, it is in this transition layer that physical processes go about converting swept-up upstream material into shocked downstream material: the same matter but often with greatly different properties. This choice of environment was motivated by the recognition that many shocks are associated by observation and/or theory with the presence of cosmic rays: we thought it worthwhile to investigate whether cosmic rays might have any measurable effect on processes governing the shock transition layer. This work is detailed in Chapter 6, which is currently being prepared for publication as a separate paper.

This dissertation is brought to completion in Chapter 7, in which I will conclude with a summary of the contribution I have made to cosmic-ray research, and some comments on the outlook of the field. Since I have now provided this outline of what is to come, there is little left to be said in introduction of this document, so without further ado, I will proceed to the science!

CHAPTER 2. Historical and theoretical background

The heart of this dissertation is the work documented in Chapters 5 and 6, each of which has its own introduction. However, those chapters were originally composed for publication in a peer-reviewed journal, and by necessity they assume a basic familiarity on the reader's part with cosmic-ray physics. This chapter's purpose is to provide an overview of the relevant concepts so that the subsequent chapters can be understood: how and when cosmic rays were discovered; the properties we can measure and what we have learned from those measurements; why supernovae are thought to be capable of supplying enough energy to cosmic-ray production to account for the observed abundance; how diffusive shock acceleration (DSA) provides a plausible explanation for cosmic-ray production in supernova remnants; why some amplification of the magnetic field above typical values is necessary for DSA to produce cosmic rays at the energies we expect; and finally, why a particular mechanism for magnetic-field amplification has been a hot topic in the cosmic-ray theory community for the last half-decade or so.

2.1 The discovery of cosmic rays

Nearly a full century has elapsed since the discovery of cosmic rays, as of the time of this writing. Interestingly, although we are still grappling with such basic questions as “what are they made of?” and “where are they coming from?” – which is certainly not to say that we don't have ideas or partial answers, or else this would be one very brief document – a number of important discoveries in particle physics were a result of early measurements of cosmic rays. I could add a lot of length to this dissertation by attempting to do justice to the whole story, but I will limit myself to a brief summary of particularly relevant parts of the story as told by [Longair \(1992\)](#).

At the beginning of the 20th century, radioactivity was a hot topic (pun not altogether intended), having been discovered as an intrinsic property of some forms of matter only in the mid-1890s. Energetic, subatomic particles that are released in the spontaneous decay of atomic nuclei were found to interact with air molecules, which they ionized. An ionized atmosphere is electrically conductive, which in the early 1900s could be quantified using an electroscope. Thus it became evident that even when all known sources of radioactivity were removed, the atmosphere remained ionized. Much of this residual ionization could be accounted for by the presence of trace amounts of naturally occurring radioactive elements.

It was an Austrian physicist by the name of Victor Hess who demonstrated the existence of an extraterrestrial source of radiation that was ionizing the atmosphere from above. In 1912 he risked his life and/or funding to carry electroscopes aloft in manned balloon excursions, to altitudes as high as 5 km. He found that above 1 km, the ionization increased with increasing height. Two years later, German physicist Werner Kolhörster extended Hess's measurements to an altitude of 9 km. (I note with the slightest touch of envy that both of these men made their historic balloon flights before they reached the age of thirty.) Iowa native and Nobel laureate Robert Millikan, who had measured the electron charge, introduced the name "cosmic rays" to refer to the phenomenon whose discovery eventually earned Hess a Nobel prize of his own, in 1936.

2.2 Cosmic rays: what we know

Increasingly sophisticated instruments, adapted to various measurement techniques, have provided us with an intriguing picture of cosmic radiation. As I said in the introduction, cosmic rays are electrically charged subatomic particles moving nearly at the speed of light, constantly hitting the Earth's atmosphere from all directions. Their composition is generally similar to that of stars and the interstellar medium, being dominated by single protons (ionized hydrogen) with some helium and trace amounts of heavier elements such as iron, as well some electrons – and small amounts of antimatter, chiefly positrons and some antiprotons.

It would be misleading of me to continue lumping all cosmic rays together into a single

population. This is because cosmic rays have an enormous range of kinetic energies (roughly a dozen orders of magnitude), and something that is true of cosmic rays at one energy (their composition, for example) may not apply to those at a significantly different energy. For example, virtually all of the lowest-energy cosmic rays we detect most likely originate within our solar system, since to arrive from elsewhere requires them to overcome the outward-flowing solar wind continuously being produced by the Sun. Conversely, the highest-energy cosmic rays (ultra-high-energy cosmic rays, UHECRs) are thought to be extragalactic in origin. Between those extremes, spanning somewhere between six and nine decades in energy, are the cosmic rays of Galactic origin (GCRs). The low-end cutoff of the GCR energy range actually fluctuates, as their arrival rate is modulated by variations in the solar wind; the upper end is much less distinct, reflecting the limitations of cosmic-ray sources within the Galaxy.

Measuring the flux of cosmic rays over the entire range of energies has required the use of multiple techniques, for a reason that becomes clear upon inspection of Figure 2.1: the thirty or so orders of magnitude between the flux of low-energy GCRs and that of the UHECRs. At the low end, cosmic rays are so abundant at the top of the atmosphere that a high-altitude balloon or satellite can carry instruments for the direct detection and precise measurement of the energy and composition – sensitive even to different isotopes of the same element. Beyond roughly $10^{14.5}$ eV, however, cosmic rays are so rare that direct detection is forbiddingly improbable, and instead an indirect detection method must be employed: arrays of ground-based detectors measure the arrival of the secondary particles and light in the “extensive air shower” produced when a cosmic ray collides with an atmospheric molecule, and reconstruct the mass, energy, and arrival direction of the primary cosmic ray.

Also evident in Figure 2.1 is that the slope of the cosmic-ray energy spectrum (on a plot with logarithmic axes) is uniform over large ranges of energy. This behavior, generally indicative of a power-law spectrum in which the flux at a particular energy is proportional to some algebraic power of the energy: $\mathcal{F}(E) \propto E^{-\alpha}$, where α is called the spectral index or power-law index. At energies below 3×10^{15} eV, the so-called “knee” of the spectrum, $\alpha \approx 2.7$; at energies above the knee and below the 6×10^{18} eV “ankle” the spectrum steepens (“softens”) to an index of




Figure 2.1 The energy spectrum of cosmic rays, from [Swordy \(2001\)](#). For comparison, the Large Hadron Collider is designed to accelerate protons to an energy of 7×10^{12} eV.

$\alpha \approx 3.2$. Beyond the ankle, the spectrum hardens, then rolls off due to the Greisen-Zatsepin-Kuzmin limit (Sokolsky and for the HiRes Collaboration, 2010): the sea of cosmic microwave background photons that permeates the universe is opaque to cosmic-ray protons whose energy exceeds roughly 6×10^{19} eV (Greisen, 1966; Zatsepin and Kuz'min, 1966), preventing them from traveling more than a few tens of megaparsecs.

I mention the wide range in cosmic-ray energy and flux because it is one of their most interesting features: at the time of this writing, our most powerful artificial particle accelerator (the Large Hadron Collider at CERN) is designed to accelerate protons to an energy of 7×10^{12} eV, but hasn't yet operated at that level. Meanwhile, something "out there" is outdoing our best efforts by a factor of ten million or more (Bird et al., 1995). But at this point, I will restrict the scope of my discussion to the Galactic cosmic rays, with energies in the GeV-PeV range.

Because cosmic rays are charged particles in motion, their trajectories are susceptible to deflection by magnetic fields. A magnetic field whose orientation varies stochastically with position will have the effect of scattering cosmic rays in random directions, and the Galaxy has such a magnetic field. This has two prominent effects: the first is that although we can measure what direction a cosmic ray was moving when we detect it, we know next to nothing about its path from source to detector. The second effect is that the local cosmic-ray distribution becomes isotropic: the sky is equally "bright" in cosmic rays in all directions, to within a tenth of a percent (Abdo et al., 2009; Abbasi et al., 2010). If we are to learn anything about the sources of cosmic rays from local measurements, we have to be clever.

One thing we can learn is the cosmic rays' mean age. The collisions of cosmic rays with interstellar material results in the production of secondary cosmic rays, whose composition includes unstable nuclei whose half-life can be measured in the laboratory, as well as stable nuclei that are nevertheless too fragile to survive in any significant quantity in stellar furnaces. The relative abundances of these nuclei, correcting for the relativistic reduction in their decay rates, allows us to "date" the cosmic-ray population. We find in this way that the cosmic rays arriving at Earth have spent roughly 15 million years en route from their sources (Amsler et al.,

2008).

2.3 Supernovae: the cosmic-ray connection

Having a timescale for cosmic-ray residence in the Galaxy allows us to make a crude estimate of the power needed to replenish the supply against escape and other losses. Assuming a uniform energy density of 1 eV cm^{-3} filling the Galactic disk of diameter 35 kpc, thickness 1 kpc and being replaced every fifteen million years, we find that roughly $10^{41} \text{ erg s}^{-1}$ ($3 \times 10^{48} \text{ erg yr}^{-1}$) is necessary to sustain their numbers. Although the Sun supplies the lowest-energy cosmic rays, main-sequence stars end up being inadequate for energetic-particle output at this scale, which is far greater than the individual bolometric luminosity of all but the most massive stars – in fact, this requirement is demanding enough that it restricts our options for cosmic-ray source types to a manageable number.

From an energetics standpoint, stars do make generous contributions to the Galactic coffers when they die. A star whose mass exceeds that of the Sun by a factor of eight or more will explode as a core-collapse supernova, injecting 10^{51} erg into an expanding shock wave and fireball. For a few weeks, the fireball shines with an intensity that rivals the combined output of the host galaxy's 10^{11} remaining stars, as the photograph of Supernova 1994D in Figure 2.2 illustrates. In galaxies similar in size to the Milky Way, supernova explosions occur at a rate of a few per century. The energy supplied by supernovae is an order of magnitude greater than that needed for cosmic-ray production.

There is observational evidence that cosmic rays may have some connection to supernovae. Although the intense, expanding fireball fades from view over the course of several months, a strong shock wave continues expanding at a few percent of the speed of light for hundreds or even thousands of years. One effect of the shock is to compress and heat the shocked material to temperatures of order 10^8 K , making supernova remnants bright sources at X-ray wavelengths via the emission lines of heavy elements in the heated ejecta. However, the KeV-scale line emission of thermal gas is not the evidence I speak of.

Continuum emission from nonthermal particle populations is also visible in the vicinity of




Figure 2.2 Supernova 1994D, the fourth supernova detected in 1994, is the bright spot in the lower left corner of this Hubble Space Telescope image of the lenticular galaxy NGC 4526. Image credit: NASA, ESA, The Hubble Key Project Team, and The High-Z Supernova Search Team.

the shocks of supernova remnants. Figure 2.3 is an X-ray image of the shell-type remnant of the nearby Supernova 1006, located some 2.2 kpc above Earth’s southern hemisphere. This approximately spherical structure is 20 parsecs in diameter and expanding at approximately 5000 km s^{-1} (Katsuda et al., 2009). Roughly the same size on the sky as the Sun and Moon, the morphology of SN1006 is spatially resolved at all observational wavelengths. While much of the visible area of the remnant is thermal emission, the northeast and southwest limbs contain thin filaments of X-ray synchrotron emission, produced by energetic electrons in a magnetic field. For this particular remnant, the electrons may have energies as high as 10^{14} eV (Reynolds, 1998). More recently, TeV γ -rays have been detected from these parts of SN1006, another signature of particles with energy at least that high, although whether this radiation is of leptonic or hadronic origin has not been conclusively demonstrated (Acero et al., 2010).

The association of nonthermal, highly energetic particle populations with strong shocks is by no means conclusive proof that supernova remnants are a significant source of cosmic rays, but the possibility becomes even more enticing in light of the beautiful and simple theory of diffusive shock acceleration, which describes a mechanism for the production of a power-law distribution of energetic particles at astrophysical shocks.

2.4 Diffusive shock acceleration (DSA)

Early efforts at explaining the existence of very energetic particles recognized a basic property of astrophysical plasma: the density is so low (just a handful of particles per cubic centimeter) that binary interactions between particles are virtually unheard-of. This is seldom true of terrestrial matter, where collisions rapidly redistribute energy until the particle population relaxes to a Maxwell-Boltzmann distribution. But in the near-vacuum of interstellar space, any particle that happens to acquire some extra energy is unlikely to be compelled to share it through any collisional process. Interactions are instead collective: the forces acting on any one particle are the result of the electromagnetic field produced by the coordinated movements of a large number of particles.



Figure 2.3 At an age of approximately one thousand years, the remnant of supernova 1006 displays a morphology resembling a spherical shell some 20 parsecs in diameter. Of particular interest to cosmic-ray researchers are the bright rims of emission from nonthermal particle distributions, seen here in the northeast (upper-left) and southwest parts of the shock. Image credit: NASA/CXC/Rutgers/J.Hughes et al.

Fermi (1949) proposed that localized regions of enhanced magnetic field, embedded in the turbulent interstellar medium, could act as scattering centers, exchanging energy with fast particles. For isotropic turbulence, a head-on collision is slightly more probable than a collision with a receding magnetic irregularity, so on average a population of energetic particles in such an environment will gain energy. The acceleration rate is proportional to the square of the speed of the turbulent motion of the scattering centers, and this mechanism is thus called second-order Fermi acceleration.

For the nonrelativistic velocity dispersion typical of interstellar plasma, second-order Fermi acceleration is not an especially fast or effective means of producing large numbers of particles with high energy. An astrophysical shock wave introduces a different distribution of scattering centers, favorable for a more efficient particle-acceleration process. The effect can be understood even using a highly simplified model, in which a shock moves at constant speed V_{sh} into a uniform environment.

A steady, one-dimensional hydrodynamical strong shock is a planar discontinuity that divides space into upstream and downstream regions, i.e., the material before and after the shock's passage; the unshocked upstream medium has different bulk properties such as density, velocity, and temperature from the shocked material downstream. In the reference frame in which the shock is stationary, material crosses the shock from upstream to downstream at a constant rate, so from the law of conservation of mass we can easily see that the upstream density and velocity ρ_1 , V_1 are related to their downstream counterparts by

$$\rho_1 V_1 = \rho_2 V_2. \quad (2.1)$$

The ratio of upstream density to downstream density is the compression ratio r , whose value is important in determining the properties of the distribution of accelerated particles. It can be written (e.g. Shu, 1992) as

$$r = \frac{\gamma_a + 1}{\gamma_a - 1 + 2\mathcal{M}^{-2}} \rightarrow \frac{\gamma_a + 1}{\gamma_a - 1} \text{ as } \mathcal{M} \rightarrow \infty, \quad (2.2)$$

where γ_a is the adiabatic index of the material, which has the value 5/3 for a nonrelativistic monatomic ideal gas, and the Mach number $\mathcal{M} = V_s/c_s$ is the ratio of the shock speed to the

speed of hydrodynamical waves, but which plays a role of diminishing importance the larger it gets. Thus for “strong” shocks (by definition, a shock in which $\mathcal{M} \gg 1$) moving through nonrelativistic gases, the compression ratio is asymptotically close to a “universal” value, $r = 4$

Now let us allow such a shock to pass through a region containing magnetic irregularities embedded in a turbulent interstellar medium, such as the type where second-order Fermi acceleration of particles might be possible. The systematic velocity difference between all upstream scattering centers and all downstream scattering centers will be $\Delta V = V_s(1 - r^{-1})$, and if this is very large compared with the characteristic velocities of the magnetic scattering centers, then for simplicity we can approximate them as stationary relative to their respective upstream or downstream environment.

The realization that such a configuration could accelerate energetic particles came nearly simultaneously to several researchers in the late 1970s (Bell, 1978a; Blandford and Ostriker, 1978; Krymskii, 1977; Axford et al., 1977). The essence of the mechanism is that by bouncing back and forth between the upstream and downstream regions, an energetic particle effectively has only head-on collisions, giving rise to an acceleration that is proportional to $\Delta V/c$, rather than the $(\Delta V/c)^2$ behavior that arises from colliding with approaching and receding scattering centers at nearly equal rates.

Because magnetic fields only provide elastic scattering, a charged particle changes direction only, and not energy, when it interacts with stationary magnetic irregularities. But which irregularities are “stationary” is dependent on one’s frame of reference. As seen from either upstream or downstream, the far side of the shock holds a stampede of nonstationary scattering centers, which *can* impart kinetic energy to a cosmic ray. See Figure 2.4 for a simplified sketch of this process, in which the same sequence of events is pictured twice: once from the perspective of the upstream reference frame, and once from that of the downstream. Observers upstream and downstream might disagree on *when* the particle is actually gaining energy, but they agree that it’s happening.

Thus a nonrelativistic shock may be able to produce cosmic rays efficiently through the so-called first-order Fermi acceleration mechanism. This is also commonly called diffusive shock

acceleration (DSA). In addition to its efficiency, a particularly enticing aspect of DSA is that it produces a power-law spectrum of cosmic-ray momenta, such that the spectral index depends only on the shock compression ratio. Recalling that all strong shocks have approximately the same compression ratio $r = 4$, we might expect a common spectral index – and because the measured cosmic-ray energy spectrum has a power-law functional form over much of its range, the prospect of a “universal” spectrum of accelerated particles is an appealing one. I will describe the mechanism very briefly.

We begin by assuming that in the reference frame of the shock, there is a population of cosmic rays with speed v and density N , with a distribution that is approximately isotropic at the shock position. Although the flux of cosmic rays across the shock from upstream to downstream is $Nv/4$, the flux of cosmic rays escaping to downstream infinity is NV_2 , so the probability of an individual cosmic ray escaping rather than bouncing back toward the upstream region is $\mathcal{P} = 4V_2/v$ (Bell, 1978a). Meanwhile, the average cosmic-ray particle that makes a round trip across the shock and back increases its momentum by $\langle \Delta p \rangle = 4p\Delta V/3v$ (Drury, 1983). For a cosmic ray to increase its momentum by a given amount requires it to cross the shock some number of times, and the larger this number, the more opportunities it has to escape downstream. The balance between acceleration and escape results in a momentum distribution $f(p) \propto p^{-3r/(r-1)}$. When $r = 4$, this produces the spectrum $p^2 f(p) \propto p^{-2}$. This is nearly, but not exactly, the $p^{-2.7}$ behavior seen in the cosmic-ray spectrum below 3×10^{15} eV. However, our observations are being made a great distance from the possible sources, and the precise momentum dependence of propagation-related effects, such as the probability of escape from the Galaxy, is the focus of ongoing research (Strong et al., 2007).

2.5 The limits of DSA

Diffusive shock acceleration of cosmic rays is a result of their repeated crossings of a shock. How frequently an individual particle completes a round-trip voyage across the shock and back determines its average rate of energy gain. But a cosmic ray meandering through the interstellar medium, guided by the gently curving magnetic-field lines typical of the Galaxy,

could spend an arbitrarily long time between successive shock crossings. Shocks, particularly those resulting from a supernova explosion, do not remain strong for an arbitrarily long time: several hundred years or a few thousands years, at most.

The obvious question is whether supernova remnant shocks can sustain sufficiently efficient DSA during their lifetimes to account for the observed spectrum of Galactic cosmic rays. The efficiency of the process depends on how effectively the cosmic rays are confined the shock environment. The diffusion coefficient that describes their propagation (and hence their confinement) is itself a consequence of the distribution of irregularities in the magnetic field. For a given magnetic-field amplitude B , a particle of momentum p and charge q moving perpendicular to the field orientation (pitch $\mu = \cos\theta = 0$) has a gyroradius $r_g = p/qB$. The length scale on which the field changes direction determines the mean free path of the particle trajectory's guiding center between pitch-angle scattering events. When this length is small, the diffusive transport of cosmic rays can become approximately isotropic with a small diffusion coefficient, the best possible confinement for a given magnetic field strength.

Prospects looked bleak early on: [Lagage and Cesarsky \(1983\)](#) investigated various detailed models of the upstream diffusive transport, and found that even in the best-case scenario, for the typical Galactic magnetic-field amplitude of $\sim 1\text{--}3\mu\text{G}$, cosmic rays produced in a supernova remnant could not exceed $E_{max} \sim 10^{14}$ eV; for a more realistic scenario they estimated E_{max} an order of magnitude lower. For SNR shocks to remain a contender in the quest for the cosmic-ray sources, some improvement or perhaps some loophole would be necessary. For example, [Voelk and Biermann \(1988\)](#) considered the effect of the supernova progenitor on its surroundings prior to the explosion: massive stars can produce strong winds, resulting in a radially varying density profile of the circumstellar medium; they concluded that SNR shocks propagating through this wind could indeed produce cosmic rays up to at least knee energies.

Other efforts investigated the effect that a significant energy density in cosmic rays might have on the shock environment. It was demonstrated that though comparatively few in number, the cosmic rays' contribution to the pressure could modify the macroscopic shock characteristics, perhaps in a time-dependent way ([Falle and Giddings, 1987](#)), and improve the acceleration

efficiency for particles at higher energy (Drury and Völk, 1981). It was also understood that the cosmic rays in the upstream region could excite magnetohydrodynamic waves through resonant interactions (Skillington, 1975a,b,c; Bell, 1978a; Achterberg, 1983); in so doing, cosmic rays could potentially modify their own diffusive transport.

The interaction of cosmic rays with magnetohydrodynamic waves attracts additional attention as a result of supernova remnant observations that imply magnetic field amplitudes much larger than the few microgauss typical of the Galaxy, perhaps even to hundreds of microgauss (Reynolds, 2008). If a shock propagates in a direction perpendicular the local magnetic-field orientation, the magnetic field lines' separation decreases by the compression ratio r , amplifying the field by a modest factor not greater than 4, so mere shock compression is insufficient to explain the field amplitudes inferred for these sources. Lucek and Bell (2000) and Bell and Lucek (2001), using numerical simulations and analytical calculations, respectively, suggested that resonant interactions between cosmic rays and Alfvén waves co-propagating along a guiding magnetic field could amplify such waves to considerable amplitudes, possibly allowing very young, freely expanding supernova remnants to accelerate cosmic rays to energies well beyond the knee. Considering a wider range of shock velocities and thus of SNR ages, Ptuskin and Zirakashvili (2003) concluded that the maximum cosmic-ray energy decreases rapidly in time, but that supernova remnants were nevertheless likely to be a significant contributor to the Galactic cosmic-ray population.

2.6 A special case: Bell's non-resonant streaming instability

Interactions between cosmic rays and Alfvén waves are thought to be capable of amplifying the magnetic field in the upstream environment, thereby improving the confinement of cosmic rays to the shock vicinity and thus increasing the efficiency and maximum energy of diffusive shock acceleration. These interactions can be resonant, when a cosmic ray gyrating around a mean magnetic field amplifies hydromagnetic waves whose wavelength is similar to the cosmic-ray gyroradius (Wentzel, 1974; Achterberg, 1983), or they can be non-resonant, in which ion beams excite the growth of magnetic fluctuations with a wavelength much less than the ion

gyroradius (Winske and Leroy, 1984). Bell (2004) revealed that for parameters typical of young supernova remnants, a non-resonant instability could grow faster than resonant instabilities and attain $\delta B \gg B_0$; that is, the amplitude of the magnetic turbulence resulting from the interactions of cosmic rays with small seed fluctuations in the existing interstellar magnetic field might far exceed the initial amplitude of that field, by an order of magnitude or more. An accompanying MHD simulation produced localized amplification by a factor ~ 100 . I will describe this instability in greater detail in Chapter 4.

With the potential for such impressive numbers, Bell’s non-resonant instability attracted considerable attention (most notably, it prompted the investigation that comprises much of the work on which I am basing this dissertation). Pohl et al. (2005) quickly recognized that even a substantially amplified field would not persist indefinitely once it had been swept up by the shock, but that this cascading of magnetic energy could explain the presence of thin filaments of X-ray synchrotron emission seen near the shocks of some supernova remnants. But the truly pressing question was whether Bell’s instability could in fact attain that sort of amplification, or if some response of the system might cause the growth of the field to saturate at a lower amplitude. Pelletier et al. (2006) analyzed the relative growth rates and saturation of both the resonant and non-resonant instabilities, accounting for a response of the cosmic rays, and identified that saturation of the non-resonant mode could occur when the field reaches an amplitude such that the cosmic-ray gyroradius decreases until the cosmic rays are effectively resonant with the waves they have generated. In a related work, Marcowith et al. (2006) evaluated the impact of the magnetic-field amplification on the upstream diffusive cosmic-ray transport, including the transfer of energy from cosmic rays into amplified field energy, and claimed good agreement between the predicted spectrum and the observed cosmic-ray spectrum up to the knee.

Supernova shocks aren’t the only environments where cosmic rays are expected to have a significant streaming presence. The relativistic shocks postulated in gamma-ray burst jets can also generate a cosmic-ray current strong enough to amplify the magnetic field (Reville et al., 2006; Niemiec et al., 2010), where its effect may permit the production of cosmic rays with

energy beyond the ankle of the spectrum (Li and Waxman, 2006).

Giacalone and Jokipii (2007) presented an independent mechanism for magnetic-field amplification, in which a shock encountering large-scale, low-amplitude magnetic turbulence could result in substantial amplification factors. This process could operate in concert with Bell’s modes, but represents an alternative scenario that may apply even when amplification by streaming instability does not. Reville et al. (2007) placed environmental constraints on the non-resonant instability, and found that when the ionization fraction of the surrounding medium is low, such as when shocks interact with molecular clouds, the unstable modes do not operate efficiently.

Bell’s calculations had been performed using the formalism of magnetohydrodynamics. Blasi and Amato (2008) presented a kinetic derivation of the same unstable mode, and its attendant potential for considerable magnetic-field amplification. Meanwhile, Niemiec and Pohl (2008) brought forth the early results of the first effort at modeling the non-resonant instability using three-dimensional particle-in-cell simulations, in which they observed only modest amplification via a different instability. Zirakashvili et al. (2008) reproduced Bell’s results with further MHD simulations, and in Zirakashvili and Ptuskin (2008) included this amplification in simulations of DSA to discriminate between different types of supernova when determining the maximum cosmic-ray energy. An independent MHD simulation aimed at exploring the nature of the diffusive transport demonstrated that the non-resonant instability directly contributes to cosmic-ray confinement by reducing the diffusion coefficient (Reville et al., 2008).

Niemiec et al. (2008), presenting in more detail their particle-in-cell results, reported that Bell’s non-resonant instability was not observed. However, Amato and Blasi (2009), expanding on their earlier kinetic results, attributed this absence of the Bell mode to an inapplicable choice of parameters. Riquelme and Spitkovsky (2009) concurred, using particle-in-cell simulations to reproduce both Bell’s and Niemiec’s results by varying the initial level of magnetization, although the saturation of the non-resonant mode occurred at a somewhat lower amplitude than in the MHD approach (closer to 10 than to 100). The kinetic calculations of Luo and Melrose (2009) supported this somewhat lower saturation level, and nay-sayers

no longer, [Stroman et al. \(2009\)](#) (reproduced here as Chapter 5) confirmed the presence of the non-resonant mode with amplification consistent with the contemporary results just mentioned. Additional support for the non-resonant mode with amplification factor ~ 10 came from the hybrid simulations of [Gargaté et al. \(2010\)](#)

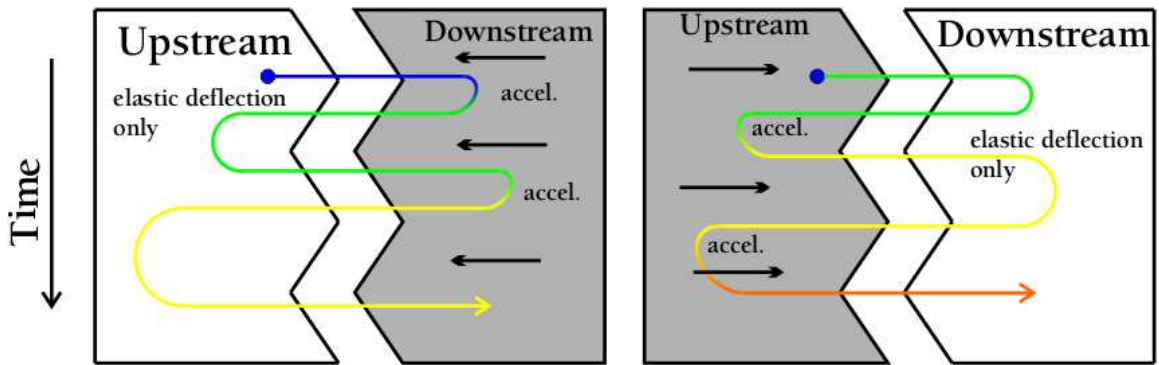


Figure 2.4 Sketch of first-order Fermi acceleration, also called diffusive shock acceleration. In the reference frame in which the upstream (downstream) medium is at rest, the perspective of the left (right) panel above, a charged particle scatters elastically off the embedded magnetostatic irregularities upstream (downstream), which changes its direction but not its energy. Across the shock, however, the particle interacts with oncoming downstream (upstream) embedded irregularities, increasing its energy.

CHAPTER 3. Particle-in-cell plasma simulation

Having laid the scientific groundwork for the research to be discussed in the coming chapters, I will shift my focus to a description of my principal tool. This chapter will address the methods involved in preparing, running, and analyzing the plasma simulations at the heart of my research.

3.1 The particle-in-cell method with TRISTAN

As the name implies, the particle-in-cell (PIC) technique represents matter as discrete particles (or “macroparticle” ensembles) occupying positions within a lattice of “cells.” Time advances stepwise, and during each timestep, the electromagnetic field is interpolated from the cell vertices to the particles’ positions, the force on the particles is calculated, their velocity is updated accordingly, and the current produced by moving charges is distributed to the cell vertices in a self-consistent manner.

There are various PIC implementations available, with one of the primary variable elements being the explicit inclusion or omission of binary interactions (particle collisions, such as those mediated by the Coulomb force between point charges). Astrophysical plasmas are typically of such a low density that collisions are exceedingly unlikely, so PIC solutions for these *collisionless* plasmas can safely assume that only collective effects play a significant role in the equations of motion for individual particles. The collisionless particle code TRISTAN ([Buneman, 1993](#)) is at the core of the simulations performed for the present research, so the remaining discussion of the PIC method will be informed largely from the perspective of that particular framework.

A PIC simulation contains several particles per cell, but interactions between particle pairs are negligible. Thus the particles accelerate only in response to the force produced by the elec-

tric and magnetic fields \mathbf{E} and \mathbf{B} stored on the grid of cells, according to the usual (relativistic) Lorentz force

$$\frac{d\mathbf{p}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}). \quad (3.1)$$

Under the influence of electromagnetic forces, a particle's momentum $\mathbf{p} = m\gamma\mathbf{v}$ changes at a rate proportional to its charge q .

Of particular significance is the dependence of acceleration on the ratio of charge to mass: as long as we maintain a constant q/m , we can use relatively few *macroparticles* to represent compact ensembles of many identical particles without sacrificing any dynamically significant effects. Unlike physical point-like particles, the distribution of charge in these ensembles has a finite spatial extent. These “clouds” of charge can overlap and approach each other arbitrarily closely without any of the energetic unpleasanties associated with the Coulomb potential barrier between point charges (Birdsall and Langdon, 1991). The shape and size of the charge distribution are free parameters in PIC, which are chosen according to the desired balance between speed (fewer calculations) and accuracy (less numerical noise). Although they have spatial extent, these clouds have no additional degrees of freedom such as rotation.

The electric and magnetic fields are updated each timestep according to Maxwell's equations:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad (3.2)$$

$$\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right) \quad (3.3)$$

where μ_0 and ϵ_0 are the permeability and permittivity of free space, respectively, and \mathbf{J} is the current density. Of course Maxwell's equations are four in number, not two, but the remaining equations ($\nabla \cdot \mathbf{B} = 0$; $\nabla \cdot \mathbf{E} = \rho_e/\epsilon_0$) will remain satisfied if satisfied initially, provided the charge continuity equation $\nabla \cdot \mathbf{J} + \partial\rho_e/\partial t = 0$ is obeyed at all times (Villasenor and Buneman, 1992).

Maxwell's differential equations are recast in a finite-difference formulation that, like the charge-cloud shapes, may involve fewer terms for high speed or more terms for high accuracy. Keeping these equations properly space- and time-centered is accomplished by recording each

of the six electromagnetic field components at a different point on the unit cell, and with the understanding that the field has that particular value at either the beginning of a timestep or its midpoint (Buneman, 1993). The value of a given field component at some point in space is obtained by interpolation, but the values being interpolated have locations unique to that field component.

The moving charge-clouds (which for simplicity will hereafter be called “particles” once more) contribute to the current density in each cell. There are various ways a particle’s translation can be, er, translated into current density, with the preferred method often depending on the particle’s shape, but there are typically many particles in each cell and it is their cumulative effect that determines \mathbf{J} in Equation 3.3. By this equation, current density and magnetic-field curl produce the incremental change in the electric field from one timestep to the next. Likewise, the electric-field curl produces the incremental change in the magnetic field via Equation 3.2.

The particles that move through these changing fields are accelerated each timestep according to Equation 3.1, with the field components being interpolated from the cell grid to the coordinates of the particles’ centers of mass. Special relativity is automatically obeyed when the Lorentz factor $\gamma = (1 - v^2/c^2)^{-1/2}$ is properly included in the definition of momentum.

This captures the essence of a relativistic, electromagnetic PIC code without binary interactions (i.e., suitable for collisionless systems): particles and fields are coupled through Newton’s laws of motion, Maxwell’s equations, and Einstein’s special theory of relativity. Such an approach is in principle capable of simulating astrophysical plasmas with a high degree of accuracy. But the computational cost of PIC simulations is quite high. A highly detailed and realistic model of even a modest physical system can easily exceed the computational resources available in the early 21st century. But there are compromises to make this powerful tool a practical one.

3.1.1 The necessity of compromise

A particle-in-cell simulation using realistic physical parameters is so expensive as to be of limited practical applicability. As long as appropriate care is taken, however, some parameters can be changed quantitatively while preserving their qualitative character, or at least without introducing relevant changes into the dynamics of the system being studied.

These compromises generally fall into at least three categories, at least in terms of what is sacrificed; the benefit in each case is a reduction in the computational resources needed. The first and most precarious compromise is a change of a constant of nature. Two examples employed in the present research and discussed below are the reduction of the proton-electron mass ratio, and the reduction of the number of spatial dimensions. Another type of compromise is adjusting the balance of different elements in the system, such as raising the temperature or increasing the density of one species of particle relative to another. A third class of “compromise” is merely the simplification or approximation of the system being studied. This is a common practice throughout physics and is not unique to simulation. For example, a real population of ions may contain protons and heavier particles including partially or fully ionized nuclei, but we may lump all the different ions together into a single species of particle. Likewise we may neglect subtle variations of a large-scale field in favor of a spatially homogeneous one.

The proton mass is approximately 1836 times that of the electron. The proton plasma frequency $\omega_{pi} = \sqrt{q^2 N_p / m_p \epsilon_0}$, inversely proportional to the square root of the mass, is thus about 43 times slower than it is for electrons of equal density. If a magnetic field is present, the proton gyrofrequency is a full 1836 times slower than the electron gyrofrequency (assuming the particles are nonrelativistic). In a magnetic field of 3 microgauss and an electron density of 1 cm^{-3} , electrons will undergo roughly two million plasma oscillations in the time it takes a proton to complete one gyration. If a PIC simulation is to include both effects, it will need several million timesteps per proton gyroperiod.

It is therefore attractive to employ an artificially low mass ratio in PIC simulations, but this can affect the relative growth rates of the various instabilities. [Bret and Dieckmann \(2010\)](#) investigated the impact of altered values of m_p/m_e on the hierarchy of unstable growth modes,

and found that some systems are more sensitive than others to changes in the mass ratio, because the unstable modes may have a different functional dependence on m_p/m_e . Even if the linear-stage hierarchy is preserved, the eventual nonlinear state can vary. Thus it is understood that simulations with a low mass ratio may mislead us by over- or underemphasizing some physical processes relative to their importance at $m_p/m_e = 1836$.

The incentive to reduce the mass ratio is therefore limited, but the cost of higher mass ratios is more than just the increased number of timesteps mentioned above. The range of physically relevant length scales also depends on the mass ratio, so higher ratios require a larger number of cells *in each spatial dimension* along with more timesteps, resulting in a dramatic increase in computational expense for a given physical problem with each increment in mass ratio. A lower limit on the scale of a PIC simulation is given by the requirement that the cell size Δ be much smaller than the length scale of the shortest physically relevant phenomenon, and that the timesteps δt are likewise short enough to resolve the fastest fluctuations. For astrophysical plasmas, the electron plasma frequency ω_{pe} and skin depth $\lambda_{se} \equiv c/\omega_{pe}$ set the time and length scales if electrons are expected to be dynamically important: as long as $\omega_{pe}\delta t \ll 1$ and $\Delta/\lambda_{se} \ll 1$, the electron plasma oscillations will be well resolved. But if ion dynamics are also of interest, it is essential that the simulation be large enough to include their skin depth and plasma oscillation period: for protons, a factor $\sqrt{m_p/m_e} \approx 43$ more cell lengths and timesteps than what is needed for electrons. For heavier ions, this factor is larger still.

This brings us to the other compromise: resolving fewer than three spatial dimensions. When the effect of interest does not require a fully three-dimensional environment to occur, a simulation in one or two spatial dimensions can provide useful insight into the scaling behavior of a system's evolution. Simulating a line or plane instead of a volume reduces the additional expense incurred by increasing the mass ratio. However, the restrictions imposed on the evolution of multi-dimensional effects again remind us that such simulations are intended merely to provide insights into what physics may be governing realistic, three-dimensional systems. As [Vladimirov et al. \(2008\)](#) point out, a PIC simulation that resolves electron dynamics, yet is large and long enough to produce even a modest acceleration of cosmic rays self-consistently

and from first principles, may require $\sim 10^{11}$ cells in each dimension and $\sim 10^{14}$ timesteps – as those authors so aptly state, “obviously far beyond any conceivable computing capabilities.” Nevertheless, simulations of reduced dimensionality are quite tractable on today’s computing architectures and allow for a superior resolution of separate spatial scales via a larger (more realistic) mass ratio.

In addition to using altered values of physical constants, it is common practice to adjust the parameters of the system being studied. This may be for the sake of exaggerating a certain physical phenomenon in order to observe it more clearly: for example, the growth rate of an instability may be proportional to the density ratio between two plasma populations, whose realistic value is in the neighborhood of 10^{-8} . Simulating this environment with a density ratio of e.g. 10^{-2} will reduce the growth time of this hypothetical instability by a factor of a million, while still satisfying the general criterion that one plasma be very dilute relative to the other. Making this type of adjustment may bring about undesirable effects by enhancing other physical processes, but this may also be mitigated somewhat by adjusting other parameters as needed until a suitable balance is found.

3.1.2 PIC versus alternative approaches

We have just discussed the principal shortcoming of the particle-in-cell technique: namely, that the use of realistic parameters limits the problem size so dramatically, given the capabilities of computing resources for the foreseeable future, as to be of limited practical application. If compromises are necessary to use PIC simulations, one may ask, what advantage do they offer that make them attractive as an alternative to other options?

After all, there *are* other options for simulating plasma, and each has strengths and limitations. The magnetohydrodynamical approach (MHD) approximates the plasma system as one or more fluids. Each fluid element has bulk properties such as density, pressure, or temperature, and MHD simulations are a powerful tool for probing systems spanning a large range of length scales, provided the fluid element represents a volume large enough for the fluid approximation to apply. But this approximation requires an assumption about the particle

distribution function, and physical processes that affect or are sensitive to the details of the distribution function are unlikely to be represented with accuracy.

Some plasma simulations employ a fluid representation of the electrons, or cold plasma electrons and ions together, while treating the ions and/or cosmic rays with a particle representation. If the detailed distribution function of one or more particle species is not expected to be important, this hybrid approach can reduce the computational expense relative to a full PIC code, while treating some particle populations with greater accuracy than MHD allows. However, interactions between particle and fluid elements may be difficult to treat self-consistently.

As an alternative to a fluid description of plasma, a fully kinetic description is one that does take explicit account of the particle distributions in both position and momentum. A very smooth (low noise) and detail-sensitive approach involves solving the Vlasov equation for small phase-space volume elements. This fine-grained approach can offer very high accuracy, and can probe physical processes on a very small scale. But like PIC, it is expensive, and for many systems the extra accuracy is not needed.

Although it is not particularly well suited to self-consistent, first-principles studies, the Monte Carlo style of simulation is frequently employed for studies of particle acceleration at astrophysical shocks. Physical processes that can be characterized in terms of probabilistic outcomes are modeled with a random-number generator. For example, if it is expected from theory that some fraction of a population of particles will be accelerated by a particular amount within a given interval of time, then the acceleration of each particle will be selected randomly at the individual level to produce that result – without no attempt made at justifying why this particle, but not that particle, and so on. This random-selection process does have a place within PIC simulations if there is an element of chance (such as particles with a finite lifetime against decay), or when choosing how to represent an abstract, smooth distribution function (for example, a spatially uniform gas with a Maxwell-Boltzmann momentum distribution) with a finite number of particles. But while the other simulations mentioned above proceed in a deterministic fashion (there are no degrees of freedom once the initial conditions have been established), the outcome of a Monte Carlo simulation is effectively the result of numerous

rolls of virtual dice.

There are numerous variations on each approach; to elaborate on any of them beyond what is here would require expanding the scope of the present discussion. But the important point is where the PIC technique has the advantage: self-consistent, first-principles accuracy when this can be attained with a finite particle number, for a growing range of problem sizes.

3.2 Parallel simulation with THISMPI

The code `THISMPI` has its origins in Buneman’s “Tri-dimensional Stanford” code `TRISTAN` (Buneman, 1993). Based on the prescription for parallelization in Cai et al. (2003), a high-performance adaptation using Message Passing Interface (MPI) called `TRISMPI` was prepared by Jacek Niemiec for the earliest simulations discussed here (see Chapter 4 and Niemiec et al. (2008)). This code included the charge-conserving current-deposition algorithm of Umeda et al. (2003).

My principal contribution to the simulation effort was the adaptation of `TRISMPI` to the truly 2.5-dimensional code `THISMPI`: the “Two-and-a-Half-dImensional Stanford code with MPI.” The fractional dimension count refers to the practice of allowing three-dimensional velocity and field vectors but confining all spatial coordinates to the $x - y$ plane; it is also expressed as “2D3V.” These awkward designations serve to distinguish the approach from pure 2-D constructions in which currents and fields are also confined to the plane. We had been approximating 2D3V simulations with `TRISMPI` by reducing the neglected dimension to a thickness of three cells, the minimum allowed by our algorithms.

Eliminating the third spatial dimension reduced by 1/6 the memory and disk space required for the particles’ instantaneous position and velocity coordinates. It also reduced by 2/3 the space required for the electric and magnetic field arrays (on account of no longer requiring three cells in the neglected dimension). Since there was now no z -dependence for any field quantity, any terms containing derivatives in z dropped out of Maxwell’s equations and increased the speed of the associated field-update subroutines. The current-deposition routine became slightly faster as well, since the z -component of their motion could be translated directly to a

current instead of being distributed into multiple layers of cells.

In addition to flattening the simulation domain, I improved the accuracy of the Maxwell field-update routines by implementing the fourth-order algorithm of [Greenwood et al. \(2004\)](#). I also improved the efficiency of post-processing. A typical simulation would run on hundreds of processors, with each producing a set (possibly numbering in the hundreds) of diagnostic output files at prescribed times. After the run was complete, it was necessary to run a separate program just to collect the data from each processor and combine these into a single file for each timestep represented. I incorporated this collection and reduction into the main program, so that the files it produced were ready for analysis with no intermediate processing.

I also increased the number of diagnostic quantities reported. The densities and current densities for each particle species, as well as the electric and magnetic fields, and finally for a random sample of particles, x , v_x and $\sqrt{v_y^2 + v_z^2}$ were being reported previously. I updated the per-particle readout to include the y coordinate and v_y and v_z separately. Certain other quantities, such as the drift velocity or temperature, were calculated from the random sample of particles. To improve the accuracy especially of the temperature¹ calculation, I introduced a routine to calculate it based on all particles. The simulation domain is divided into chunks of `ksize` \times `ksize` cells, where `ksize` is an input parameter. The average velocity of particles (of a given species) within a chunk is calculated, and then a three-dimensional Lorentz transformation subtracts the systematic component from their velocities. Within this frame, the mean kinetic energy per particle is computed and output to the diagnostic file. Smaller values of `ksize` allow for a more precise transformation into the local flow frame when motion is turbulent, but correspondingly decrease the number of particles being used to calculate both the mean velocity and mean random kinetic energy. Larger values ensure an adequate number of particles for a good average, but increase the likelihood of averaging over a region in which systematic velocity differences are present; this would inflate the reported temperature because some bulk motion would not be eliminated in the Lorentz transformation. I employed a `ksize` of 10 when using roughly 20 computational particles per cell.

¹More precisely, the mean random kinetic energy per particle.

CHAPTER 4. Amplification of magnetic fields by Bell’s non-resonant streaming instability

I have multiple objectives in this chapter. First, I will provide a proper introduction to Bell’s non-resonant instability (Bell, 2004), whose impact on cosmic-ray research I described in Chapter 2.6 without going into any detail about the actual mechanism. Second, I will summarize the early stages of our investigation into whether the prediction of this instability would be supported by particle-in-cell (PIC) simulations. This work was published as Niemiec et al. (2008), and made use of parameters that enabled a different unstable mode to dominate the evolution of the system. We then made different choices for our parameters, and that work is described in Chapter 5, which was published as Stroman et al. (2009).

4.1 Bell’s non-resonant streaming instability

An unstable physical system responds to a perturbative displacement from equilibrium by increasing the perturbation. A tangible example is a ball balanced precariously on a hilltop: a slight nudge in any direction will result in the ball’s subsequent acceleration down the hill. When the ball reaches the bottom of the hill, further acceleration will cease. The instability is said to have saturated. In the astrophysical context of interest, namely, streaming instabilities, the unstable physical system consists of the electromagnetic field and the distribution functions of the various particle species present: the ions and electrons of the interstellar medium, and cosmic rays. The collective streaming motion of the cosmic rays relative to the interstellar medium is like the force of gravity in the ball-on-hill analogy: if the electromagnetic field is “nudged” in a certain way, by the introduction of small-amplitude wave-like perturbations, the streaming particles interact with the waves and amplify them.

Although there are multiple mathematical approaches to describing an instability, it is common practice to identify the dispersion relation linking the wave vector (\mathbf{k} , measured in radians per unit length; the wavelength is $\lambda = 2\pi/k$) of a perturbation to its angular frequency ω . If the perturbation has an amplitude of ξ , a wave vector \mathbf{k} , and a frequency ω , it can be written (neglecting an arbitrary phase offset) as

$$\xi(\mathbf{r}, t) = \xi \exp(i\mathbf{k} \cdot \mathbf{r} - i\omega t). \quad (4.1)$$

The frequency is in general allowed to be a complex number $\omega = \omega_r + i\gamma$, though for ordinary traveling waves and standing waves, the imaginary component is negligible. For some systems, there are solutions of the dispersion relation for which the imaginary part of ω is comparable to or even exceeds the real part. Introducing the explicit complex notation into Equation 4.1 reveals the effect of the imaginary frequency γ : exponential growth or decay with the passage of time. For the amplification of magnetic perturbations, we are interested in systems whose dispersion relations have complex solutions in which γ is non-negligible and of the appropriate sign for growth.

In the rest frame of a steady shock, the distribution function of existing cosmic rays relaxes to a stationary, spatially varying distribution upstream, in which diffusive transport away from the shock is balanced by advection with the flow of the incoming upstream medium. In the simplest one-dimensional case, where the upstream and downstream diffusion coefficients $D_{1,2}$ are constant (though unequal), the continuity equation for the upstream cosmic-ray distribution function $f_1(x)$ is

$$V_1 \frac{\partial f_1}{\partial x} + \frac{\partial}{\partial x} \left(D_1 \frac{\partial f_1}{\partial x} \right) = 0; \quad (4.2)$$

that is, for every position x , bulk advection of cosmic rays (at speed V_1 , for we remain in the reference frame where the shock is stationary) is balanced by diffusion against a density gradient. This equation has the general solution

$$f_1 = A_1 + B_1 e^{\frac{V_1}{D_1} x} \quad (4.3)$$

upstream, with A_1, B_1 being constants, and the downstream solution has the same functional form. We require the distribution to be continuous at the shock (taken to be at position $x = 0$)

so $f_1(x=0) = f_2(x=0)$. Far ahead of the shock we expect no cosmic rays: $f_1(x \rightarrow -\infty) \rightarrow 0$ fixes $A_1 = 0$, and far behind the shock we expect the cosmic-ray distribution to remain finite, fixing $B_2 = 0$. To summarize, in a one-dimensional shock environment with uniform upstream and downstream media, the cosmic-ray distribution is constant downstream, while upstream it is exponentially attenuated with increasing distance from the shock.

The upstream *cosmic-ray precursor* has two properties that are important to the role it may play in exciting streaming instabilities that lead to magnetic-field amplification. One of these properties is the variation of the “typical” cosmic-ray energy with distance from the shock. If the diffusion coefficient has an energy dependence, it will affect the upstream attenuation length $D_1(E)/V_1$. When the scattering of cosmic rays is accomplished by magnetic fields, higher-energy particles have a larger mean free path than lower-energy particles. They diffuse farther from the shock, on average, meaning that at increasing distances upstream, one expects to find a cosmic-ray distribution that increasingly favors high-energy particles relative to the distribution near the shock. With their large Lorentz factors, these precursor cosmic rays are “stiff” enough that any instability they do excite in the upstream plasma will have an appreciable amount of time to grow before having any back-reaction on the cosmic rays that drive it.

The second significant property of the precursor is that it carries an electric current. We saw in Chapter 2 that diffusive shock acceleration produces a power-law distribution in particle momentum, irrespective of the mass composition of accelerated particles. The particles with higher momentum can penetrate farther upstream on account of the energy-dependent diffusion coefficient. At a given relativistic momentum, the electron-to-ion density ratio may be of order 10^{-2} (Bell, 1978b), so the outer regions of the cosmic-ray precursor are positively charged. Seen from the upstream rest frame, this ion-dominated streaming population of cosmic rays is neither charge- nor current-neutral, but the high conductivity of interstellar plasma prevents the development of any large-scale electric fields. Electrons, possibly representing both the thermal and the accelerated electron populations, respond rapidly to local charge imbalances and maintain the expected neutrality.

Bell (2004) predicted that the positively-charged precursor, moving relative to the upstream medium and approximately parallel to a uniform interstellar magnetic field $B_0 = B_{\parallel}$, could produce a strong electric current. Incorporating the strength of the cosmic-ray current j_{\parallel} into the parameter $\zeta \equiv |B_{\parallel} j_{\parallel}| r_{g1} / \rho V_s^2$ (where r_{g1} is the lower bound to the range of cosmic-ray gyroradii), Bell found the dispersion relation

$$\omega^2 - v_A^2 k^2 \pm \zeta V_s^2 k / r_{g1} = 0$$

(where $v_A = B_{\parallel} / \sqrt{\mu_0 \rho}$ is the speed of Alfvén waves) for wave vectors satisfying $1 < kr_{g1} < \zeta V_s^2 / v_A^2$ and oriented parallel to the magnetic field, itself parallel to the cosmic-ray current. In this regime, the imaginary component of the frequency dwarfs the real component, so magnetic fluctuations of appropriate wavelengths experience almost pure growth. This growth rate γ reaches its highest value $\gamma_{max} = k_{max} v_A$ when the wave vector $k = k_{max} = \zeta V_s^2 / 2v_A^2 r_{g1}$.

Bell’s analysis and subsequent simulation used a fluid description of the plasma-CR system. Amato and Blasi (2009), taking instead a kinetic approach, considered different models for the return current and found that two different but plausible models for the return current (having a beam of CR electrons drifting with the CR ions, or else distributing the current among all the plasma electrons, resulting in a slight drift relative to the plasma ions) produced the same dispersion relation. Also considering the instability from a kinetic standpoint, Luo and Melrose (2009) confirmed that it is the return-current electrons, not the cosmic rays themselves, that drive the initial linear instability, and also calculated that this mechanism could produce significant amplification of the magnetic field before reaching saturation. Magnetohydrodynamical simulations demonstrated that this instability could produce a turbulent magnetic field whose RMS amplitude exceeded B_0 by more than an order of magnitude (Bell, 2004; Zirakashvili et al., 2008). If correctly understood, this mechanism could explain the strong fields inferred for some SNR from X-ray observations (Vink and Laming, 2003; Völk et al., 2005); moreover, a sufficiently amplified magnetic field might increase the maximum energy for particle acceleration in supernova shocks from the limit calculated by Lagage and Cesarsky (1983), perhaps extending the source spectrum to the observed “knee” at a few PeV.

4.2 First results of PIC simulations

The magnetohydrodynamical simulations of Bell’s non-resonant instability, in which a cosmic-ray current reportedly drove the amplification of the magnetic field to levels much greater than its initial value, approximated that current as being strictly constant. This is a perfectly reasonable assumption during the linear stage of unstable growth, but as the turbulent magnetic field becomes strongly amplified, it is important to treat the cosmic rays in a self-consistent manner. In order to observe the back-reaction of the amplified field and turbulent interstellar plasma on the cosmic rays producing the instability, we made an early attempt at reproducing Bell’s instability with a particle-in-cell (PIC) simulation (Niemić et al., 2008).

These simulations were in either three dimensions or effectively two-and-a-half dimensions (using a three-dimensional simulation domain with the minimum possible number of cells in the neglected dimension). In order to manage the computational expense of three dimensions, the simulation parameters were chosen with an emphasis on efficiency, with the intent of increasing the wavenumber k_{max} and growth rate γ_{max} of the instability, given by Bell (2004) but expressed in terms of simulation variables as

$$k_{max} = \frac{1}{2} \frac{N_{CR}}{N_i} \frac{V_s}{v_A} \lambda_{si}^{-1}$$

and $\gamma_{max} = v_A k_{max}$. The Alfvén speed $v_A = |B_{\parallel}| / \sqrt{\mu_0 (m_i N_i + m_e N_e)}$; varying this value changes k_{max} but not γ_{max} . The ion skin depth λ_{si} is a factor $\sqrt{m_i/m_e}$ larger than the electron skin depth λ_{se} , which is one of our basic simulation parameters (set in this case to seven cells in length).

To reduce the number of cells and timesteps needed to observe the growth of Bell’s instability, we set N_{CR}/N_i as high as 1/3, set the shock velocity V_s to $0.3c$ (still nonrelativistic, but an order of magnitude greater than typical supernova shock velocities), and decreased m_i/m_e to values as low as 10. For the 3-D simulation, the parameters resulted in $\gamma_{max} = 0.0137c/\lambda_{se}$. The Alfvén speed v_A had to be small enough to allow k_{max} to remain reasonably large; the choice of $v_A = 0.0157c$ resulted in the wavelength $\lambda_{max} \equiv 2\pi/k_{max} = 50\Delta$, which could fit a reasonable number of times in each of three directions in our simulation box.

The predicted growth of circularly polarized magnetic fluctuations, with a flow-aligned wave vector of magnitude k_{max} and an e -folding time of γ_{max}^{-1} , was not observed in the simulations we performed for [Niemic et al. \(2008\)](#). Instead, we saw an oblique filamentary mode whose growth was a factor 5 slower. The filaments evolved into turbulence, the magnetic field was amplified to a level $\delta B \sim B_0$, and this turbulent field facilitated the transfer of bulk momentum from the cosmic rays to the background plasma. This resulted in the convergence of drift velocities (see [Figure 4.1](#)), at which point the cosmic-ray current relative to the plasma ceased to be sufficient to drive further amplification of the field.




Figure 4.1 Drift velocity of all particle species in primary 3-D simulation of [Niemic et al. \(2008\)](#). The convergence of velocities correlates with the saturation of magnetic-field amplification.

Upon a closer examination of our choice of parameters, we concluded that we had not satisfied one of the assumptions of [Bell \(2004\)](#). The dispersion relation had been calculated in the asymptotic limit $\gamma/\Omega_i \rightarrow 0$: the growth rate must be small in comparison to the ion gyrofrequency. However, our parameters yielded $\gamma/\Omega_i \sim 3$. In other words, the plasma was not sufficiently magnetized. [Riquelme and Spitkovsky \(2009\)](#) investigated the threshold conditions and confirmed that when $V_s(N_{CR}/N_i) \geq 2v_A$ (equivalent to $\gamma_{max} \geq \Omega_i$), the filamentation mode we observed competes with or even suppresses the parallel-wave mode. In light of the

importance of satisfying $\gamma_{max} \ll \Omega_i$, we prepared a new series of simulations. Chapter 5 is our peer-reviewed report on that work ([Stroman et al., 2009](#)).

CHAPTER 5. Kinetic simulations of turbulent magnetic-field growth by streaming cosmic rays

Adapted from a paper published in *The Astrophysical Journal* (Stroman et al., 2009)

Thomas Stroman, Martin Pohl, Jacek Niemiec

5.1 Abstract

Efficient acceleration of cosmic rays (via the mechanism of diffusive shock acceleration) requires turbulent, amplified magnetic fields in the shock's upstream region. We present results of multidimensional particle-in-cell simulations aimed at observing the magnetic field amplification that is expected to arise from the cosmic-ray current ahead of the shock, and the impact on the properties of the upstream interstellar medium. We find that the initial structure and peak strength of the amplified field are somewhat sensitive to the choice of parameters, but that the field growth saturates in a similar manner in all cases: the back-reaction on the cosmic rays leads to modification of their rest-frame distribution and also a net transfer of momentum to the interstellar medium, substantially weakening their relative drift while also implying the development of a modified shock. The upstream medium becomes turbulent, with significant spatial fluctuations in density and velocity, the latter in particular leading to moderate upstream heating; such fluctuations will also have a strong influence on the shock structure.

5.2 Introduction

The forward shocks of young shell-type supernova remnants may be efficient acceleration sites for Galactic cosmic rays. The theory of diffusive shock acceleration (DSA) provides a promising mechanism by which a small fraction of particles can attain nonthermal energies in

such an environment, provided they are confined to the shock vicinity; the review by [Reynolds \(2008\)](#) and its references have treated the subject in detail. This confinement may be supplied to some degree by the interstellar magnetic field, but the estimated upper limit in energy falls short of what is observed in cosmic rays of Galactic origin. The implied need for a much stronger, turbulent upstream field, together with observational evidence for strong magnetic fields immediately behind the forward shock, has attracted attention to the question of upstream magnetic-field amplification.

It is well known that ion beams can generate magnetic turbulence via resonant and non-resonant interactions ([Winske and Leroy, 1984](#)). [Bell \(2004\)](#) (hereafter B04) suggested that non-resonant magnetic field amplification may also be caused by the cosmic-ray current expected in the cosmic-ray precursor of a quasi-parallel shock. Although substantial field amplification has occurred in magnetohydrodynamical (MHD) simulations that assume a constant cosmic-ray current (B04; also [Zirakashvili et al., 2008](#)), the approximations of MHD may not be appropriate for modeling the non-linear evolution and eventual saturation of this instability; we must use kinetic methods to simulate these later stages with accuracy, including any back-reaction on the cosmic rays.

Our earlier work ([Niemić et al., 2008](#), hereafter N08) used 2-D and 3-D particle-in-cell (PIC) simulations to model the growth and saturation of a current-driven instability. We relaxed the condition assumed in the calculations of B04 that the rate of unstable growth must be much less than the ion gyrofrequency, and we found that an oblique filamentary mode dominated the initial magnetic field growth. [Riquelme and Spitkovsky \(2009\)](#) confirmed the findings of N08, and in particular investigated the parameters for the transition between the filamentation mode seen in N08 and the parallel-wave mode seen in the MHD simulations, and verified that the parallel mode appears only in the regime with $\omega \ll \Omega_i$. However, despite the difference in initial unstable modes, the non-linear characteristics of the system observed in both works (N08; [Riquelme and Spitkovsky, 2009](#)) were similar. In particular, both investigations observed only modest amplification of the magnetic field, and saturation by the same mechanism. The analytical calculations of [Luo and Melrose \(2009\)](#), with a kinetic treatment of

the current-driven instability, predict a saturation level consistent with the kinetic simulations.

In this paper, we report more recent work, in which we simulate the current-driven instability and its saturation in a parameter regime in which the non-relativistic ion gyrofrequency far exceeds the (complex) frequency of the current-driven instability. Such parameters are particularly relevant to the environment of non-relativistic strong shocks whose cosmic-ray precursors are sparse in comparison with the magnetized ambient interstellar medium (ISM); an understanding of the limits of field amplification in such an environment is therefore essential to the formulation of a self-consistent theory of cosmic-ray acceleration in the Galaxy. We find that the initial evolution of the magnetic field in this parameter regime is in agreement with the predictions of B04 and consistent with the transition parameter, ω/Ω_i , identified by [Riquelme and Spitkovsky \(2009\)](#). We observe that the field growth is limited to modest values by the same saturation mechanism as before (N08; [Riquelme and Spitkovsky, 2009](#)). In addition to the magnetic field, we examine the evolution of the upstream medium-precursor system in terms of bulk and local properties that are relevant for modeling the shock acceleration process. Notable findings include the rise of strong turbulent motion relative to the bulk upstream rest frame, and of motional electric fields associated with the magnetic fields that permeate the turbulent flow. The temperature of the upstream medium, measured in the local rest frame as determined with high spatial resolution, attains significant values during the evolutionary phase characterized by the turbulent motion, such that the ratio of magnetic to thermal energy density may in fact decrease as a result of this precursor-ISM interaction. Such modifications to the upstream medium may lead to substantial effects on the propagation of the subshock, and on the properties of the downstream medium: both the pre-existing upstream turbulence after advection and shock compression, and further magnetic-field generation via turbulent dynamo processes, may be quite sensitive to turbulence in the upstream medium ([Giacalone and Jokipii, 2007](#); [Inoue et al., 2009](#)).

We begin by describing in §5.3.1 the simple physical scenario that captures all the ingredients necessary for the streaming instability we want to simulate, with §5.3.2 summarizing our computational representation of that environment, and the specific parameterizations being

outlined in §5.3.3. We present our simulation results in §5.4 along with discussion of their significance, and finally summarize in §5.5.

5.3 Simulation setup

5.3.1 Model

To study the growth and saturation of magnetic turbulence in response to a cosmic-ray current, we simulate the evolution of a system representative of the upstream environment of a non-relativistic strong parallel shock. We assume that the interstellar medium is a homogeneous collisionless plasma composed entirely of ions and electrons (number density N_i and N_e , respectively) of opposite but equal charge $\pm e$, permeated with a uniform magnetic field $B_{\parallel 0}$. A population of cosmic ray ions (density $N_{CR} = N_e - N_i$) drifts at the shock velocity v_{sh} along the magnetic field relative to the plasma and carries a current density $j_{CR} = eN_{CR}v_{sh}$; the slight electron excess in the plasma maintains charge neutrality, and the electrons drift relative to the ions with $v_d = v_{sh}N_{CR}/N_e$ to provide a return current density $j_{ret} = -eN_e v_d = -j_{CR}$. The ISM ion and electron populations are each distributed according to Maxwell-Boltzmann statistics in their respective rest frames with the same temperature, given by the electron thermal velocity $v_{e,th}$.

Two assumptions in our setup, that only electrons provide the return current and that the streaming cosmic rays consist only of ions, arise from the large mass ratio m_i/m_e between physical ions and electrons. The electrons will respond to any local charge or current imbalance on a much shorter timescale than the ions will. Moreover, at high energies, ions outnumber electrons (on account of their greater rigidity) by a factor that depends on the spectral index of their momentum distribution but is generally large enough that cosmic-ray electrons play a negligible role in the precursor physics. Cosmic rays upstream undergo pitch-angle scattering off magnetic inhomogeneities and approach a quasi-isotropic distribution, forming a steady cosmic-ray precursor in the shock rest frame. Ions with lower energies will be advected back toward the shock sooner than ions with higher energies, so at some distance upstream one expects to find a population of high-energy cosmic-ray ions drifting with the shock, up to some

cutoff energy and/or distance (Reville et al., 2009) associated with particles that escape the system.

For our simulation, we simplify the cosmic-ray distribution function. We assume that the isotropic, scattered cosmic rays of the precursor are energetically much more significant than any escaping cosmic rays that happen to be passing through the region of our simulation, so we neglect the latter altogether. Our box size is assumed to be small in comparison both to the scale length of the cosmic-ray precursor, such that we may neglect spatial variations in their distribution, and to the Larmor radii of any cosmic rays present, such that the details of their energy distribution are similarly negligible, and we replace the already-narrow range of energies with a single energy (in the CR population rest frame), with Lorentz factor γ_{CR} . The quasi-linear calculations in B04 do not depend explicitly on cosmic-ray energy, and tests with various values of γ_{CR} have not produced significantly different outcomes, so it is unlikely that our choice of energy distribution obscures behavior that a more realistic model would illuminate.

In considering the current carried by cosmic rays through a cold ambient plasma parallel to a uniform magnetic field, the calculations in B04 predict an instability in which the dominant mode, circularly polarized magnetic fluctuations with wavenumber $k_{\parallel max}$, grows at a rate $\Im\omega = \gamma_{max} = v_A k_{\parallel max}$, $\Re\omega \approx 0$, where

$$k_{\parallel max} = \frac{eN_{CR}v_{sh}B_{\parallel 0}}{2N_i m_i v_A^2} \approx \frac{\mu_0 J_{CR}}{2B_{\parallel 0}} \quad (5.1)$$

and $v_A = \left[B_{\parallel 0}^2 / \mu_0 (N_e m_e + N_i m_i) \right]^{1/2}$ is the Alfvén velocity of the plasma. As mentioned above, this result is subject to the restriction that the fluctuation satisfy $\omega \ll \Omega_i$; that is, that the growth rate of the instability must be much less than the nonrelativistic ion gyrofrequency.

5.3.2 Implementation

Our simulations employ a modified version of the particle-in-cell code TRISTAN (Buneman, 1993). The code is updated to work in a high-performance computing environment and uses the charge-conserving current deposition algorithm of Umeda et al. (2003). We have incorporated a fourth-order-accurate algorithm for updating the electromagnetic fields based on Maxwell's

equations (Greenwood et al., 2004), and adapted the code to a truly 2.5-dimensional simulation.¹

The shortest relevant physical time- and length-scales are given by the electron plasma period $\omega_{pe}^{-1} = [m_e \epsilon_0 / e^2 N_e]^{1/2}$ and electron skin depth $\lambda_{se} = c / \omega_{pe}$, respectively. To ensure that physical small-scale electrostatic effects are distinct from those arising from the artificial discretization of space and time, for our simulations we choose parameters such that $\lambda_{se} \geq 4\Delta$, where Δ is the length of one cell, and $\omega_{pe}^{-1} \geq 8.9\delta t$, where δt is the timestep size. The wavelength $\lambda_{max} \equiv 2\pi / k_{\parallel max}$ of the predicted magnetic-field fluctuations is of order $10^2 \lambda_{se}$ for all simulations (see Table 5.1). We perform our simulations in a rectangular computational grid, with both the homogeneous magnetic field and the direction of cosmic-ray drift aligned with the longer dimension. Periodic boundary conditions in both directions allow particles and fields to “wrap around” to the far side of the grid, but this also imposes an upper limit to the size of structures that can be represented accurately. Having observed a migration of the dominant structures toward increasing size scales in N08, for one simulation (Run A) we chose a grid slightly larger than $20\lambda_{max}$ in the drift direction, and $5\lambda_{max}$ in the transverse direction. We ran additional simulations on smaller grids to explore the influence of variations in our choice of parameters on the initial amplification of the magnetic field and the turbulent properties of the field and plasma.

To minimize statistical noise associated with the particles, we employ a combination of large per-cell particle counts and low-pass filtering of the electric currents arising from their motion. We initialize the simulations with either 20 or 25 plasma ions per cell, and a density ratio of $N_i / N_{CR} = 50$ plasma ions per cosmic-ray ion. However, in most runs we split each cosmic ray into ten simulated particles that preserve the charge-to-mass ratio but provide better statistics, and the contribution of each particle to currents and densities is weighted accordingly.

Electrons, less massive than the ions by a factor of $m_i / m_e = 50$ in all tests, are physically one population but simulated as two populations. In a simulation for which $N_i = 25$, each

¹As reported in N08, we have previously approximated two dimensions by simulating a volume with a thickness of only three grid cells.




Figure 5.1 Evolution of the magnetic field components in Run B displaying the growth and saturation of the instability predicted in B04. The magnetic field amplitude is in units of the initial homogeneous field, and time is in dimensionless units set by the calculated inverse growth rate γ_{max}^{-1} from that instability.

cell contains 25 “full” electrons as well as five simulated particles, which add up to half of a split electron. The two simulated electron populations are initialized according to the same distribution function, and behavioral comparisons of the full electrons and the split electrons in the simulations reveal no differences between the two populations, to within statistical limits.

The initial temperature of the plasma is set artificially high, with $v_{e,th} = 0.01c$, in order to mitigate Buneman-type electrostatic effects arising from the drift between the stationary population of ions and the slowly drifting electrons. Our simulations in N08 demonstrated that cooler plasmas were heated through such effects on a much shorter timescale than any turbulent magnetic-field amplification, but the resulting anisotropy of the ion distribution function in particular persisted on such timescales. A higher initial temperature leads to better preservation of isotropy against the intra-plasma drift. Additionally, using a density ratio N_i/N_{CR} of 50 instead of 3 (the value in N08) significantly reduces v_d .

The cosmic rays are initialized to be mono-energetic and isotropic in the shock rest frame. This frame drifts in the $-x$ -direction, antiparallel to the homogeneous magnetic field, at a speed $v_{sh} = 0.4c$ or $0.3c$. We set the magnetic field strength such that the Alfvén speed $v_A = 0.01c$, or equivalently that the plasma frequency exceeds the Larmor frequency by a factor of $\omega_{pe}/\Omega_e \approx 14.14$.

5.3.3 Run-specific parameters

In Table 5.1 we summarize the initial conditions and some observations from each of the simulations we discuss. The default parameters are described below, followed by a list of each secondary run’s purpose and any deviations from those parameters.

Run A is our large-grid run for this paper, and is the basis for many subsequent “test” simulations. The cosmic rays are isotropic and mono-energetic with Lorentz factor $\gamma_{CR} = 50$ in the CR rest frame, which drifts at $v_{sh} = 0.4c$ in the $-x$ -direction in the simulation frame. There are initially 25 plasma ions per cell, and each cosmic-ray ion is split into ten equal parts, for a net of five simulated particles per cell. For these parameters, $\lambda_{max} = 444\Delta$. In order to capture the non-linear stage when the dominant structure size exceeds this, we simulate on a

grid of $9000\Delta \times 2400\Delta$, corresponding to $20.26\lambda_{max} \times 5.4\lambda_{max}$.

Run A' uses the same physical parameters as Run A, but on a smaller grid, fitting $6.75\lambda_{max} \times 3.6\lambda_{max}$.

Run A uses a large grid to ameliorate the limitations imposed by periodic boundary conditions, which is particularly relevant during the later evolution when structures in the magnetic field begin to grow quite large. However, for the purpose of analyzing the properties of the upstream medium during the growth and saturation, we simulate a smaller region with a greater detail level. Run B is our test with increased spatial resolution, in which the electron skin length has been increased to $\lambda_{se} = 6\Delta$, and the particle density reduced to $N_i = 20$, while holding the other parameters fixed. The grid is $6.75\lambda_{max} \times 2.7\lambda_{max}$.

Runs C through F are tests with varying cosmic-ray Lorentz factors. All use the intermediate value for electron skin length of $\lambda_{se} = 5\Delta$.

Run C has a reduced shock speed $v_{sh} = 0.3c$. Its cosmic rays also have a reduced Lorentz factor, $\gamma_{CR} = 10$. It is unique among our runs in that it does not split the cosmic rays. The grid for Run C admits $5.4\lambda_{max} \times 2.03\lambda_{max}$. Run C' is identical to Run C but splits the cosmic rays.

Runs D, E, and F use the usual shock speed $v_{sh} = 0.4c$ and cosmic-ray Lorentz factor $\gamma_{CR} = 10, 25, \text{ and } 200$, respectively. Their grid measures $5.76\lambda_{max} \times 2.16\lambda_{max}$.

5.4 Results and discussion

The spatially-averaged amplitude of the turbulent components of the magnetic field, parallel and perpendicular to the initial homogeneous field $B_{\parallel 0}$, are plotted as a function of time in Figure 5.1 for Run B, representing well the evolution in all runs. The unit of choice for time is the inverse predicted growth rate γ_{max}^{-1} . When $t\gamma_{max} \approx 5$, the field enters a period of exponential growth. As seen in Table 5.1, the peak growth rate observed is often 10%–20% less than γ_{max} ; this small discrepancy may be primarily a thermal effect. As discussed in N08, the finite temperature of the plasma tends to reduce the maximum growth rate from its value in the cold-plasma limit.




Figure 5.2 The perpendicular magnetic field in units of $B_{\parallel 0}$ (top) and density of plasma ions (bottom) at $t_{\gamma_{max}} = 7.0$ for Run A, when $\langle \delta B/B_{\parallel 0} \rangle \approx 0.50$. The parallel-mode structure is clearly visible in the magnetic field. Localized density fluctuations about the mean (25 ions per cell) are smaller, with an amplitude of $\langle \delta N_i/N_i \rangle \approx 0.08$.




Figure 5.3 As in Figure 5.2, the perpendicular magnetic field and density of plasma at $t\gamma_{max} = 14.0$ in Run A. The plasma density fluctuations are strong and well correlated with the magnetic field structures during the non-linear stage of magnetic field growth. The structures move rapidly on oblique trajectories at this time in the simulation (see Figure 5.4).

Initially, the growth of the magnetic field is stationary, in the form of non-propagating waves with wave-vectors oriented parallel to the cosmic-ray drift and similar in magnitude to $k_{\parallel max}$ (Figure 5.2). The cosmic-ray Larmor radius $r_{CRg} \gg \lambda_{max}$, so this is not the resonant interaction of cosmic rays with self-generated Alfvén waves, but is instead the non-resonant generation of an approximately purely growing wave mode. However, density fluctuations begin to appear while $\delta B \ll B_{\parallel 0}$. The root-mean-square magnitude of the normalized density fluctuations $\langle \delta N_i / N_i \rangle$ during the linear growth stage is smaller than $\langle \delta B / B_{\parallel 0} \rangle$, typically at a level of 20 to 50 % between $t_{\gamma max} = 5$ and $t_{\gamma max} = 10$. As the turbulent magnetic field becomes comparable in strength to the homogeneous field, the plasma begins to move and the cosmic rays slow down, reducing the current and saturating the magnetic-field amplification at a level of 10–20 $B_{\parallel 0}$.

Any density and velocity fluctuations that arise may contribute to further deviation from a purely growing parallel mode, and may change the environment enough that further field amplification is not favored. We observe that when $B_{\perp} \sim B_{\parallel 0}$, the growth rate is only gently decreasing and the field still has approximately a parallel-mode structure, but when $B_{\perp} \sim 10B_{\parallel 0}$, the spatial organization of the field no longer resembles a parallel mode. As seen in Figure 5.3, the dominant length scale of the fluctuations grows and significantly exceeds λ_{max} . Substantial variation appears along the direction transverse to the drift, reaching the maximum length scale allowed by the periodic boundary conditions of even our largest computational grid. The field growth comes to a stop, saturating generally between 10 and 20 times the amplitude of the initial homogeneous field.

The reduction of the field growth rate accompanies the onset of turbulent plasma motion. Parcels of ambient plasma whose size is comparable to λ_{max} move and collide, giving rise to considerable density fluctuations. Figure 5.4 (left) indicates that the plasma parcels can attain a considerable range of speeds in the directions transverse to the cosmic-ray drift. The strong perpendicular field and turbulent plasma allow bulk momentum transfer from the cosmic rays to the ambient plasma, so superimposed on the turbulent motions, the plasma begins to drift more and more rapidly in the direction of the cosmic rays, while the cosmic rays

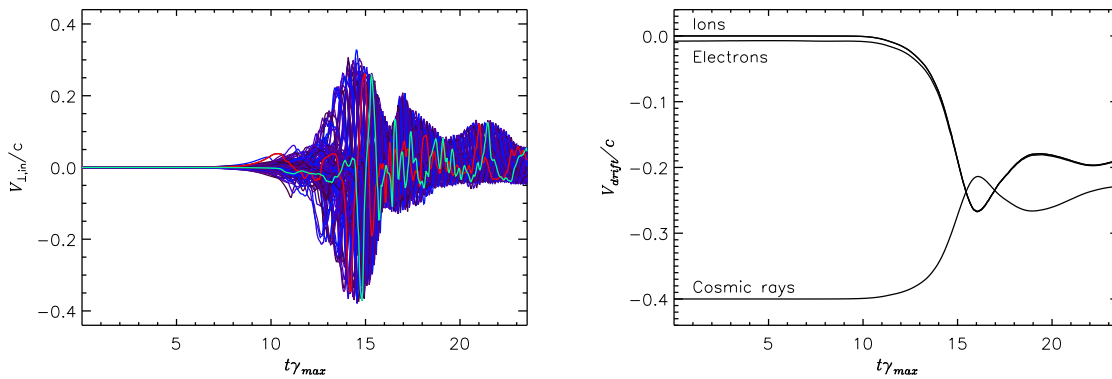


Figure 5.4 **Left:** From Run B, the transverse in-plane velocity of the plasma in regions of size $(0.45\lambda_{max})^2$. Two regions have been selected and superimposed in lighter colors for clarity. **Right:** In the same run, momentum transfer from the cosmic rays to the background plasma removes their relative drift, limiting the amplification of the magnetic field. The “overshoot” at $t\gamma_{max} \sim 16$ appears in the small-grid runs but is not present in Run A, with the larger grid, so it is likely a consequence of the periodic boundary conditions

simultaneously slow down (Figure 5.4, right), consistent with the quasi-linear predictions of [Winske and Leroy \(1984\)](#) for non-resonant modes. In our simulations, the speeds converge such that the relative drift is roughly an order of magnitude smaller than its initial value. [Riquelme and Spitkovsky \(2009\)](#) suggest that convergence to a relative velocity approximating the Alfvén speed is sufficient to inhibit further magnetic-field amplification and thus that saturation may take place. Although defining the Alfvén speed for a turbulent magnetic field is complicated, we find that by considering only the projection of the magnetic field onto the drift direction along with the local plasma density averaged over regions of a few λ_{se}^2 in area, we do observe a “mean” Alfvén speed parallel to the cosmic-ray drift that differs from the post-saturation velocity separation only by a factor of order unity.

In addition to slowing down in bulk, the instantaneous rest-frame distribution of cosmic rays is modified, as shown in Figure 5.5. Some anisotropy is introduced, which may have consequences for the modeling both of radiation and particle acceleration. The anisotropy appears in all simulations in a qualitatively similar manner, though the rate at which it develops

(as well as the extent to which the CR bulk speed changes) decreases with increasing γ_{CR} . The fastest development of rest-frame anisotropy occurs for cosmic rays whose trajectory is transverse to the drift direction in the simulation. During the linear stage of magnetic-field growth, the stationary parallel-wave mode with wavelength λ_{max} will be favorable for resonant scattering of cosmic rays whose pitch angle μ is given by $\lambda_{max}\Omega_{CR}/2\pi = v_{\parallel} = \mu c \left(1 - \gamma_{CR}^{-2}\right)^{1/2}$. For the parameters in our simulations and in agreement with the work of [Luo and Melrose \(2009\)](#), μ is of order 10^{-2} , corresponding to those cosmic rays in which the aforementioned anisotropy is first seen. Additional scattering and acceleration of cosmic rays may occur in the electric fields associated with magnetic fields being transported in the turbulent flow during the nonlinear stage. Second-order Fermi acceleration is possible in such a turbulent environment. The electric fields we observe during saturation are typically an order of magnitude smaller than the magnetic fields. A Lorentz transformation into the *global* plasma rest frame reduces the electric fields only from about $t\gamma_{max} = 15$ onward, in agreement with the behavior of the plasma bulk speed. If we transform the EM fields into the *local* plasma rest frame, thereby accounting for its turbulent motion, we see a reduction of the electric fields beginning at $t\gamma_{max} \approx 10$. After about $t\gamma_{max} = 15$, less than 20% of the original electric field remains, demonstrating that the dominant fraction of the electric field in our simulation arises from the transport of magnetic fields. This indicates that pure elastic scattering may be a poor assumption, if one considers the diffusive transport of high-energy cosmic rays in the upstream region of young SNR. Runs C–F with varying cosmic-ray Lorentz factors illustrated principally that “stiffer” cosmic rays took slightly longer to begin slowing down and were marginally more resistant to the onset of rest-frame anisotropy, but there was no qualitative difference in the linear growth of the magnetic field nor in the saturation mechanism.

The interactions between the shock itself and any spatial variations not only in plasma density but also in velocity are known to introduce vorticity in the downstream region that may amplify the magnetic field ([Giacalone and Jokipii, 2007](#); [Inoue et al., 2009](#)). They may also have a non-negligible impact on the injection process or the scattering properties of the downstream medium, and consequently the resulting distribution of energetic particles

(Malkov and O’C Drury, 2001). If the injection mechanism depends primarily on highly localized processes, as opposed to being a product of the average conditions over large spatiotemporal scales, then the presence of appropriately distributed inhomogeneities might reduce the availability of readily accelerated particles in one region, with no guarantee of a corresponding increase in a neighboring region. In particular, if injection is ultimately more efficient at a parallel shock than a perpendicular shock, the development of a turbulent magnetic field in the formerly parallel region via this streaming instability would reduce the overall quasi-parallel fraction of the shock’s surface if comparable turbulence is not established in the initially quasi-perpendicular regions. Beyond injection, the cosmic-ray transport properties upstream and downstream contribute significantly to the acceleration efficiency and the maximum energy cosmic rays can attain. Marcowith et al. (2006) investigated the sensitivity of the cosmic-ray diffusion coefficient to the character of the turbulence, and found that, in particular, anisotropic turbulence may lead to somewhat higher maximum particle energy.

In addition to possible implications for the spectrum of accelerated particles, the character of the upstream magnetic turbulence also eventually contributes to the emission properties of the downstream medium. The presence of anisotropy in the turbulence after being shock-compressed may lead to a larger degree of polarization in the radio synchrotron emission than might otherwise be expected (Stroman and Pohl, 2009).

The turbulent motion in the upstream medium may be significant for self-consistent models of cosmic-ray acceleration. In our simulations, plasma parcels attain velocities that exceed the initial Alfvén speed and collide, producing “fronts” whose effects may include further localized field amplification or, perhaps quite importantly, compression and heating of the upstream medium. We observe that heating of the ambient plasma occurs cospatially with the interfaces between rapidly moving plasma parcels and the slower regions into which they are pressing. In Run B, for instance, the average kinetic energy per ion is of order $10^{-6}m_i c^2$ at $t\gamma_{max} = 3.8$, before the linear growth has begun, and is slightly higher but still of that order at $t\gamma_{max} = 7.6$, during linear growth. But at $t\gamma_{max} = 11.5$, the beginning of saturation, the mean kinetic energy per ion has increased to $\sim 5 \times 10^{-4}m_i c^2$, and at $t\gamma_{max} = 15.3$, it has

climbed to $10^{-2}m_i c^2$. The common and convenient assumption that the entire amplification and saturation process proceeds in a low-temperature regime must be carefully considered if conditions allow for turbulence to develop rapidly in the precursor, but as [Vladimirov et al. \(2008\)](#) observe, the substantial heating expected from the dissipation of even a small fraction of upstream turbulence may significantly increase the rate at which particles are injected into the acceleration process.

The reduction of the cosmic-ray current (by the converging drift velocities of their population and the ambient plasma) serves to underscore the need to treat these particles kinetically. The back-reaction from the turbulent plasma and amplified magnetic field is non-negligible, with the most prominent consequence being saturation of the field growth at a lower amplitude than might be expected from a constant cosmic-ray current ([Ohira et al., 2009](#)).

Finally, it is worth noting that although the initial dominant mode of magnetic-field amplification agrees with the calculations of B04, the later evolution and saturation mechanism have much in common with the simulations performed in N08, in which the choice of simulation parameters exposed an oblique filamentary mode as the dominant instability. It seems it doesn't matter what form the initial linear instability takes.

An important consideration in interpreting our findings is that the ability of our simulations to accurately portray the evolution of the magnetic field and the plasma properties at very late times, beyond the saturation, is limited. In addition to the limitation imposed by the finite size of the simulation domain, the assumptions of environmental homogeneity become increasingly unsound. Changes in the plasma flow speed imply density adjustments in order to conserve mass flux ([Ohira et al., 2009](#)). Also, as the shock approaches, spatial gradients in the cosmic-ray precursor and ambient plasma properties play an increasingly important role. Furthermore, the time in which upstream amplification can occur before the system is overtaken by the shock is only $\sim 40\gamma_{max}^{-1}$ for our parameters, assuming efficient Bohm-type diffusion so that the shock sweep-up timescale is $c r_{CRg}/3v_{sh}^2$ (where r_{CRg} is the cosmic-ray gyroradius) (N08). Thus, field growth is inevitably limited even in the absence of saturation.

It bears mentioning that to a first approximation, the effects of a changing environment will

be limited. As mentioned in §5.3.1, the parameters of the streaming instability are agnostic to the details of the cosmic-ray distribution. Thus, although an approaching precursor implies a continual softening of the local cosmic-ray energy spectrum due to the arrival of particles over an increasingly inclusive energy range, it is only the increased overall flux that matters. Until surplus high-energy electrons are included, the net cosmic-ray current will increase, whose effect according to Equation 5.1 will be to amplify the parallel mode described therein on shorter time- and length-scales, potentially well separated from the large scales attained by the turbulence generated earlier. Even this effect is difficult to predict with confidence: not only will cosmic rays that arrive during saturation encounter a turbulent upstream medium, but may themselves have been influenced by prior unstable evolution nearer to the shock. In any case, since the precursor length scale is proportional to the diffusion coefficient, and hence to the energy for Bohm diffusion, cosmic ray particles of significantly lower energy (and higher flux) are brought in only shortly prior to capture by the shock.

Although the non-resonant streaming instability is concerned only with the current they carry (and, as demonstrated by Pelletier et al. (2006) and Luo and Melrose (2009), is largely produced by the return current in the plasma), this or other instabilities thus encountered may introduce anisotropy and inhomogeneity into the local cosmic-ray distribution, with potentially non-trivial consequences. As increased computational capacity and improved models become available, we can expect simulations representing a significant extent of the precursor to yield valuable insight into these interrelated and nonlinear processes occurring there.

5.5 Summary and conclusions

We have simulated the turbulent amplification of the interstellar magnetic field upstream of a nonrelativistic shock using a kinetic 2.5-dimensional particle-in-cell code.

We observe:

- That the non-resonant streaming instability (B04) is seen initially for our choice of parameters, but with fluctuations appearing in the plasma density;

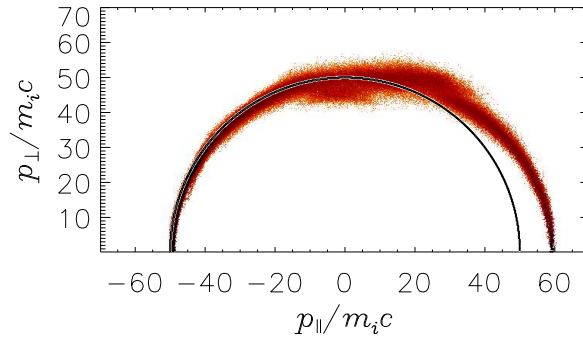


Figure 5.5 The phase-space distribution of cosmic rays in the instantaneous rest frame of the population at time $t\gamma_{max} \approx 15.7$ in Run B, compared with the initial distribution (thin semicircle). In addition to revealing the anisotropy in the forms of stretching and broadening that began during the linear magnetic-field growth, the instantaneous rest frame differs from the initial frame by approximately $0.18c$.

- That the plasma quickly evolves and saps the bulk momentum from the drifting cosmic rays;
- The reduction of relative drift effectively removes the cosmic-ray current and saturates the magnetic-field amplification to $\sim 20B_{||0}$;
- This saturation mechanism is independent of the initial linear instability, occurring as a result of the back-reaction on the cosmic rays;
- Strong transverse plasma motions arise in conjunction with the turbulent magnetic field exceeding the homogeneous background field;
- Interactions between differently-moving plasma parcels give rise to significant fluctuations in density and temperature of the upstream interstellar medium;
- The plasma begins to drift, suggestive of a cosmic ray-modified shock, until its speed approximately matches that of the cosmic-ray population.
- The phase-space distribution of cosmic rays undergoes significant changes: the development of anisotropy within the population rest frame beginning during the linear growth of

the magnetic field, and decrease in bulk speed associated temporally with the saturation of the magnetic field;

- Evolution past this point is beyond the reach of our current simulation method, as it requires additional input to reflect the changing external environment.

The growth of non-resonant streaming instabilities in the cosmic-ray precursor of supernova shocks may contribute to significant turbulent amplification of the magnetic field upstream of the shock. However, the amplitudes achieved via any one process are likely to be somewhat less than suggested in the artificial case of constant cosmic-ray current. Simulations that incorporate incrementally more of the environmental changes expected in a real precursor scenario can be expected as computational capacity increases, although it may take innovations beyond the foreseeable future to make a realistic particle-in-cell approach tractable for such a simulation ([Vladimirov et al., 2008](#)).

This research was supported in part by the National Science Foundation both through TeraGrid resources provided by NCSA ([Catlett et al., 2007](#)) and under Grant No. PHY05-51164. The work of JN is supported by MNiSW research project N N203 393034, and The Foundation for Polish Science through the HOMING program, which is supported by a grant from Iceland, Liechtenstein, and Norway through the EEA Financial Mechanism.

Table 5.1 Parameters and selected results of simulations

| Run | Grid (Δ^2) | N_i (Δ^{-2}) | λ_{se} (Δ) | v_{sh}/c | λ_{max} (λ_{se}) | γ_{CR} | $B_{\perp}/B_{\parallel 0}$ | γ/γ_{max} | γ_{max}/Ω_i |
|-----|---------------------|-------------------------|-----------------------------|------------|------------------------------------|---------------|-----------------------------|-----------------------|-------------------------|
| A | 9000×2400 | 25 | 4 | 0.4 | 111 | 50 | 13.3 | 0.96 | 0.4 |
| A' | 3000×1600 | 25 | 4 | 0.4 | 111 | 50 | 17.6 | 0.86 | 0.4 |
| B | 4500×1800 | 20 | 6 | 0.4 | 111 | 50 | 19.7 | 0.89 | 0.4 |
| C | 4000×1500 | 25 | 5 | 0.3 | 148 | 10 | 11.6 | 0.78 | 0.3 |
| C' | 4000×1500 | 25 | 5 | 0.3 | 148 | 10 | 9.9 | 0.93 | 0.3 |
| D | 3200×1200 | 25 | 5 | 0.4 | 111 | 10 | 12.7 | 0.86 | 0.4 |
| E | 3200×1200 | 25 | 5 | 0.4 | 111 | 25 | 18.4 | 0.90 | 0.4 |
| F | 3200×1200 | 25 | 5 | 0.4 | 111 | 200 | 20.1 | 0.91 | 0.4 |

Run A' reproduces the parameters of Run A on a spatial scale comparable to run B. Run C' differs from Run C only in that the cosmic-ray particles are not split into multiple simulated particles. N_i is the number of ions per cell (size Δ^2), $\lambda_{se} = c/\omega_{pe}$ is the electron skin depth, v_{sh} is the shock speed at which the cosmic-ray population drifts; γ_{CR} is the cosmic-ray Lorentz factor in the population rest frame. γ_{max}/Ω_i is the ratio of the predicted growth rate of the most unstable mode (with wavelength λ_{max}) to the nonrelativistic ion gyrofrequency. The peak measured perpendicular magnetic field and growth rate are given by B_{\perp} and γ in units of their respective reference quantities. Parameters common to all runs include an ion-to-electron mass ratio $m_i/m_e = 50$, a density ratio of ions to cosmic rays $N_i/N_{CR} = 50$, and an initial homogeneous magnetic field antiparallel to the cosmic-ray drift such that the Alfvén speed is $v_A = 0.01c$.

CHAPTER 6. Do cosmic rays affect instabilities in the transition layer of nonrelativistic collisionless shocks?

Adapted from a paper to be submitted to *The Astrophysical Journal*

Thomas Stroman, Martin Pohl, Antoine Bret, Jacek Niemiec

6.1 Abstract

There is an observational correlation between astrophysical shocks and non-thermal particle distributions extending to high energies. To investigate the possible feedback of these particles on the shock at the microscopic level, we perform particle-in-cell (PIC) simulations of uniform, interpenetrating plasmas representing the shock transition layer, both with and without an additional population of cosmic rays. We vary the relative density of the counterstreaming plasmas, the strength of a homogeneous, guiding magnetic field, and the energy density in cosmic rays. We compare the early development of the unstable spectrum for selected configurations without cosmic rays to the growth rates predicted from linear theory, for assurance that the system is well represented by the PIC technique. Within the parameter space explored, we do not detect an unambiguous signature of any cosmic ray-induced effects on the microscopic instabilities that govern the formation of a shock. We demonstrate that an overly coarse distribution of energetic particles can artificially alter the statistical noise that produces the perturbative seeds of instabilities, and that such effects can be mitigated by increasing the density of computational particles.

6.2 Introduction

Shocks form in a wide variety of astrophysical environments, from planetary bow shocks in the heliosphere to colliding clusters of galaxies. The presence of nonthermal particle populations in some of these environments is inferred from radio, X-ray, and γ -ray observations (Reynolds, 2008), and a number of theoretical mechanisms link strong shocks (such as those found in young supernova remnants) to particle acceleration processes (e.g. Drury, 1983; Webb et al., 1983). An understanding of the nonlinear coupling between the shock, the energetic particles, and the spectrum of any excited waves is essential to the proper interpretation of the observational signatures of shocked systems, as well as a correct understanding of the origin of cosmic rays.

In most cases the density of the shocked medium is sufficiently low for collisions between particles to be infrequent; such collisionless shocks are mediated through collective electromagnetic effects. The thickness of the shock transition region and the nature of the microphysics occurring therein may play a key role in the injection efficiency of thermal particles into the acceleration processes thought to accelerate cosmic rays to high energy (Bell, 1978a,b). In addition, efficient particle acceleration associated with some shock environments raises the question of the extent to which a significant cosmic-ray contribution to the local energy density can influence the structure of the shock front itself. This may be of particular importance in starburst galaxies, where massive stars and frequent supernovae increase the availability of sites for particle acceleration. The emission of TeV γ -rays from some nearby starburst galaxies suggests an abundance of cosmic rays that significantly exceeds the values observed locally (Acciari et al., 2009).

A number of mechanisms exist whereby cosmic rays may play a role in shaping the shock and its environment. A significant cosmic-ray contribution to the pressure upstream of a shock is understood to result in substantial modification to the shock environment (Eichler, 1984); this effect is expected to produce an effective shock compression ratio that is perceived as being larger by cosmic rays of higher energy, possibly hardening the local cosmic-ray source spectrum (Malkov and O’C Drury, 2001). Under some conditions, the current carried by cosmic-ray ions

diffusing ahead of the shock may excite instabilities in the upstream interstellar medium that lead to large-amplitude magnetic turbulence (Bell, 2004) or other substantial alterations to the upstream environment, such as the heating of thermal electrons (e.g., Rakowski et al., 2008). However, these effects occur on spatial scales much larger than the thickness of the subshock, and may operate quite independently of the processes occurring therein.

It is to these subshock processes that we turn our attention in this paper. In order to quantify or constrain the effect of a “spectator” population of cosmic rays (that is, any cosmic rays present that have already been accelerated elsewhere) on the formation and evolution of nonrelativistic collisionless subshocks, we design simulations to explore the initial linear and subsequent nonlinear growth of instabilities into turbulence in the shock transition region. We restrict our focus to the interpenetration layer, in which two counterstreaming plasma shells overlap in space. This is fertile ground for the growth of well-known and well-studied instabilities (Silva et al., 2003; Frederiksen et al., 2004; Bret, 2009), and the two-plasma system will respond accordingly. We repeat our simulations with an added component, a plasma consisting of highly relativistic particles, so we may compare the systems at both the early and late stages of their evolution and identify whether the presence of the cosmic rays has any effect.

Hydrodynamical models of shocks are incapable of accurately resolving length scales much smaller than a particle mean free path and must therefore approximate shocks as discontinuities. A self-consistent kinetic approach is necessary for a more complete exploration of the relevant small-scale physical processes. Here we use particle-in-cell (PIC) simulations in two spatial dimensions to model the interaction of counterstreaming plasma flows in the presence of a third, hot plasma of energetic particles. The PIC technique is a kinetic (particle-based) approach to modeling the self-consistent evolution of an arbitrary distribution of charged particles and electromagnetic fields, and thus is well-suited to the nonlinear development of unstable plasma systems. In order to isolate the subshock instabilities from larger-scale spatial effects such as those arising from systematic charge separation, all distribution functions are initially homogeneous with respect to position. We consider the simplified case in which two electron-

ion plasmas, not necessarily of equal density, move non-relativistically through each other. This movement may also be parallel to a uniform magnetic field. We first observe the evolution of the drift velocities, field amplitudes, particle distributions, and wave spectra in the case when cosmic rays are negligible. We then include cosmic rays whose energy density is an order of magnitude below the kinetic energy in the bulk flow of the plasmas. We also explore the case of an artificially high energy density in cosmic rays, exceeding the value of equipartition with the bulk flows, to gain insight into effects that may be too small to arise in the case intended to represent a more realistic environment but that may nevertheless be significant in regions where energetic particles are unusually abundant. However, our results suggest that even a high abundance of cosmic rays is not sufficient on its own to produce a significant deviation from the usual evolution of the instabilities shaping a shock.

6.3 Objectives and approach

To constrain the extent to which cosmic rays might influence the physics of the shock transition layer, we perform a series of simulations of the intra-subshock environment. Using the benchmarks of system behavior outlined in Section 6.3.3, we characterize the response of the counterstreaming-plasma system to the cascade of instabilities arising from interactions among the various components in the absence of cosmic rays. We then repeat the simulations with cosmic rays present, so as to facilitate the side-by-side comparison of the various benchmarks. The parameter choices for the systems that we explore in this manner are described in detail in Section 6.3.1, and our simulation model and its implementation are described in Section 6.3.2. Then, a test of selected parameter configurations against the predictions of analytical studies of related systems is described in Section 6.4.

6.3.1 Parameter-space configurations

The physical system under consideration is modeled as one homogeneous electron-ion plasma moving relative to another. These plasmas may in general have different densities. In addition, a uniform magnetic field may be aligned with the flow direction. Finally, a pop-

ulation of cosmic rays may be present, at rest in bulk in the center-of-momentum frame of the two plasmas. This is the reference frame chosen for the simulation, so the two plasmas flow in opposite directions; we adopt the nomenclature of “stream” and “counterstream” to distinguish between them. By our convention, the stream’s velocity is in the $-x$ -direction, antiparallel to the guiding magnetic field when this is present.

Our primary simulations are sensitive to variations in three parameters: the density ratio between the stream and the counterstream, the strength of the guiding magnetic field, and the energy density in the cosmic rays.

The inter-plasma density ratio $w \equiv n_s/n_{cs}$ takes three values and their respective designations: the “symmetric” case $w = 1$, the “intermediate” case $w = 0.3$, and the “dilute” case $w = 0.1$. The velocity of the stream is fixed at $\mathbf{v}_s = -0.2c\hat{\mathbf{x}}$, while the velocity of the counterstream obeys the relation $n_{cs}\mathbf{v}_{cs} + n_s\mathbf{v}_s = 0$; in the dilute-stream case, the relative flow speed is therefore $0.22c$. Likewise, the density of the counterstream is fixed, so the stream density alone varies; the total electron density n_e and thus the electron plasma frequency $\omega_{pe} = \sqrt{e^2 n_e / \epsilon_0 m_e}$ (where ϵ_0 is the vacuum permittivity) is largest in the symmetric case and reduced by a factor $\sqrt{1.1/2}$ in the dilute case.

The magnetic field $B_{0,x}$ may be absent, or present at either of two amplitudes given by the ratio of the electron cyclotron frequency $\Omega_e = eB_{0,x}/m_e$ to the electron plasma frequency, $b \equiv \Omega_e/\omega_{pe}$. The “absent” magnetic field refers to $b = B_{0,x} = 0$. The values designated “weak” and “strong” correspond to $b = 0.01$ and $b = 0.1$, respectively; in a plasma of electron density of e.g. $n_e \sim 1 \text{ cm}^{-3}$, a magnetic field of $\sim 300 \mu\text{G}$ is necessary for $b = 0.1$. Since the speed of Alfvén hydromagnetic waves $v_A \equiv bc/\sqrt{1 + m_i/m_e}$ is at most $0.014c$ for our choice of $m_i/m_e = 50$, all of the plasma collisions we consider have Alfvénic Mach numbers significantly larger than unity (up to ∞ in the case when $b = 0$). Note that because the electron plasma frequency is not independent of the density ratio w in our simulations, neither is the absolute magnetic field amplitude corresponding to a particular value of b .

All simulations include cosmic-ray particles consisting of electrons and ions, initialized according to an isotropic distribution function and a single speed (Lorentz factor $\Gamma_{CR} =$

50) whose statistical weight $w_{CR} \equiv n_{CR}/n_s$ is adjusted to three levels: “negligible” when $w_{CR}\Gamma_{CR} = 10^{-8}$, “present” when $w_{CR}\Gamma_{CR} = 10^{-3}$, and “abundant” when $w_{CR}\Gamma_{CR} = 10$. Since w_{CR} is defined in terms of the stream density n_s , the absolute energy density in cosmic rays also varies with the density ratio w , being ten times larger in the symmetric case than the dilute case for each value of w_{CR} . Neglecting the contribution of electrons, the bulk kinetic energy in the stream plus counterstream is of order $n_{cs}m_i v_s^2 w(1+w)/2$, while the cosmic-ray energy density is of order $n_{CR}\Gamma_{CR}m_i c^2 = 25w_{CR}\Gamma_{CR}w n_{cs}m_i v_s^2$, where we have used the relation $v_s/c = 0.2$. Thus the ratio of cosmic ray energy density to bulk kinetic energy density is $50w_{CR}\Gamma_{CR}/(1+w)$: considerably larger than unity for the “abundant” case and a few percent in the “present” case.

6.3.2 Simulation setup

Our simulations employ a modified version of the relativistic electromagnetic particle-in-cell code TRISTAN (Buneman, 1993), updated for parallel use with MPI, operating in two and a half dimensions (2D3V) with periodic boundary conditions. The charge-conserving current deposition routine of Umeda et al. (2003) and the field update algorithm with fourth-order accuracy from Greenwood et al. (2004) are the most prominent additions, as well as digital filtering of electric currents to suppress small-scale noise via an iterative smoothing algorithm.

The primary set of simulations, 27 in total, were conducted on a spatial grid in the $x - y$ plane of size $280\lambda_{se} \times 180\lambda_{se}$ (periodic boundary conditions in x and y , with elongation in the flow direction x), where $\lambda_{se} \equiv c/\omega_{pe} = 10\Delta$ is the electron skin depth, set to ten grid cells of length Δ . The electron plasma frequency ω_{pe} is determined from the sum of the stream and counterstream electron densities only and thus depends on w but not on w_{CR} . Supplementary high-resolution simulations in which $\lambda_{se} = 30\Delta$ were performed on a grid of more cells but representing a smaller physical region, $128\lambda_{se} \times 96\lambda_{se}$. The timestep δt was chosen such that $\omega_{pe}^{-1} \approx 22\delta t$ for the $\lambda_{se} = 10\Delta$ simulations, or $\omega_{pe}^{-1} \approx 66\delta t$ for the high-resolution $\lambda_{se} = 30\Delta$ simulations.

Six separate particle populations from three plasmas are modeled: the “stream” moving in

the $-x$ direction, the “counterstream” moving in the $+x$ direction, and the energetic particles representing cosmic rays. Within each plasma, ions and electrons of charge $\pm e$ and mass ratio $m_i/m_e = 50$ have equal charge density and a common drift velocity so that the entire setup has no net current and no charge imbalance. Each cell in a primary (high-resolution supplementary) simulation is initialized with a total of 90 (120) computational particles: 20 stream ions, 20 counterstream ions, and 5 (20) cosmic-ray ions; and an electron for each ion. The physical density of each plasma is manipulated through the assignment of the appropriate statistical weights, w and w_{CR} , to the various particle species.

The stream and counterstream, viewed from their respective rest frames at $v_s = -0.2c\hat{x}$ and $v_{cs} = w \times 0.2c\hat{x}$, are described by a Maxwell-Boltzmann distribution in which the electrons’ most probable speed is given by $v_{th,e} = 0.01c$ and the ions are in equilibrium with the electrons. The cosmic rays, whose rest frame is the simulation frame $v_{CR} = 0$, are isotropic and each is initialized with Lorentz factor 50, regardless of whether it is an ion or an electron.

6.3.3 Behavioral benchmarks

To provide a basis for comparison among different cosmic-ray densities w_{CR} for each combination of plasma density ratio w and magnetic-field amplitude b , we select the following attributes of the system for study: the drift velocity of each particle population, the instantaneous root-mean-square amplitude of the parallel and perpendicular components of the electric and magnetic field over the entire simulation domain, the effective temperature of each stream or counterstream particle species, and the spectrum of excited wave modes in the magnetic field.

As there is no initial bulk motion in the perpendicular directions, only the parallel component V_x of drift velocity is considered. The electric field amplitudes are presented in units of the scaling electric field $E_\omega \equiv \omega_{pe} c m_e / e$ (equivalent to the field at which the electric energy density is half the electrons’ rest-mass energy density, $\epsilon_0 E^2 / 2 = N_e m_e c^2 / 2$); the magnetic field multiplied by c is expressible in the same units.

The particle distributions may not remain strictly Maxwellian throughout the duration of

the simulation. As a surrogate for temperature, therefore, the mean random kinetic energy of the electrons and ions of the stream and counterstream is calculated by determining the systematic velocity component within a $10\Delta \times 10\Delta$ region and eliminating this local bulk motion via an appropriate Lorentz transformation; the mean post-transformation Lorentz factor γ' corresponds only to the random motion.

Finally, we will explore the effect of cosmic rays on the time evolution of the spectral decomposition of the perpendicular (out-of-plane) magnetic field B_z into its spatial Fourier components, both parallel (wave number $k_x = k_{\parallel}$) and perpendicular to the drift (wave number $k_y = k_{\perp}$).

6.4 Comparison with analytical beam-plasma predictions

As a test that our simulation results were consistent with theory, we applied the methods of [Bret \(2009\)](#) to selected stream-counterstream configurations without cosmic rays. Whether magnetized or not, beam-plasma systems (in which a fast, dilute “beam” plays a role comparable to that of our stream, with the dense “plasma” representing our counterstream) are susceptible to a host of both electrostatic and electromagnetic instabilities. For flow-aligned wave vectors, electrostatic modes such as two-stream or Buneman are likely to grow. In the direction normal to the flow, the filamentation instability (sometimes referred to as “Weibel”) is usually excited as well. Finally, modes with wave vectors oriented obliquely are likewise unstable, so that the unstable spectrum is eventually at least two-dimensional.

The full spectrum has been first evaluated solving the exact dispersion equation in the cold approximation, accounting thus for a guiding magnetic field as well as finite-mass ions. It turns out that for the present configuration, ions play a very limited role (in the linear phase) and are not responsible for any unstable modes which wouldn’t be excited if they were infinitely massive. Unlike settings exhibiting Bell’s modes, for example ([Bell, 2004](#)), where a single proton beam is considered without electrons moving *at the same speed*, we are here dealing with a plasma-shells collision, where protons and electrons are comoving. As a result, the effect of finite-mass ions is simply a first-order correction to the electronic spectrum. The

dispersion equation displays the very same branches, and the growth rate is altered by a quantity proportional to $\mathcal{O}(m_e/m_i)$.

For flow-aligned wave vectors, the most unstable wave vector $k_{\parallel,m}$ and maximum growth rate γ_m read (with evaluation corresponding to the configuration displayed in Figure 6.4)

$$\begin{aligned} k_{\parallel,m} \lambda_{se} \frac{\Delta v}{c} &= \frac{\sqrt{3} \sqrt{1 + m_e/m_i}}{\sqrt{2} \Gamma_s^{3/2}} = 1.19, \\ \frac{\gamma_m}{\omega_{pe}} &= \frac{\sqrt{1 + m_e/m_i}}{\sqrt{2} \Gamma_s^{3/2}} = 0.68, \quad \text{symmetric case,} \end{aligned} \quad (6.1)$$

and,

$$\begin{aligned} k_{\parallel,m} \lambda_{se} \frac{\Delta v}{c} &= \sqrt{1 + w} = 1.04, \quad w \ll 1 \\ \frac{\gamma_m}{\omega_{pe}} &= \frac{\sqrt{3(1 + w)} w^{1/3} (1 + m_e/m_i)^{1/3}}{2^{4/3} \Gamma_s} = 0.31, \quad \text{dilute case,} \end{aligned} \quad (6.2)$$

where $\Gamma_s = (1 - v_s^2/c^2)^{-1/2}$ is the bulk Lorentz factor of the stream, which moves in the simulation frame with speed $v_s = 0.417c/(1 + w)$. As seen in Figure 6.4, panels (a) and (f), oblique modes dominate for the mildly relativistic conditions of the simulation. The full-spectrum maximum growth rate is thus slightly larger than the numerical values calculated above for modes propagating along the flow.

For wave vectors normal to the flow, the growth rate reaches its maximum for $k_{\perp} = \infty$, with

$$\frac{\gamma_m}{\omega_{pe}} = 2 \frac{v_s}{c} \sqrt{\frac{1 + m_e/m_i}{\Gamma_s}} = 0.41, \quad \text{symmetric case,} \quad (6.3)$$

and,

$$\frac{\gamma_m}{\omega_{pe}} = \frac{v_s}{c} \sqrt{\frac{w(1 + w)(1 + m_e/m_i)}{\Gamma_s}} = 0.12, \quad \text{dilute case.} \quad (6.4)$$

These filamentation data have been calculated neglecting the magnetic field, which is small ($\Omega_e = 0.01\omega_{pe}$) in the present setup. Electrostatic unstable modes propagating along the flow are rigorously insensitive to the flow-aligned magnetic field. As long as $m_e/m_i \ll 1$, the 2-D linear spectra computed with or without finite-mass protons are indistinguishable. The hot spectra, accounting for the $v_{th,e} = 0.01c$ thermal spread for the electrons, have then been calculated considering infinitely heavy protons and using the fluid approximation described in [Bret and Deutsch \(2006\)](#).

The predicted growth rates for the two-dimensional k_{\parallel}, k_{\perp} wave-vector space are plotted in Figure 6.4, panels (b) and (g) for the dilute and symmetric streams, respectively, when the stream velocity is $0.417c$ (corresponding to a bulk Lorentz factor $\Gamma_{rel} = 1.1$) *relative to* the counterstream and $m_i/m_e = 100$. The configurations described in Section 6.3.1 provide a relative drift of $0.4c$ in the symmetric case, but only $0.22c$ when the stream is dilute. For an optimal match to the calculation, we therefore repeated the early stage of our simulation with the parameters $m_i/m_e = 100$ and $\Gamma_{rel} = 1.1$, on a grid of $3840\Delta \times 3840\Delta$, and $\lambda_{se} = 30\Delta$.

To make a comparison with the analytical predictions, we extract the two-dimensional \mathbf{k} spectrum at times separated by only a short interval intended to capture the earliest linear growth of the instabilities, and compute the average growth rate, which is plotted in Figure 6.4 panels (c)–(e) and (h)–(j) for the dilute and symmetric cases, respectively. The agreement is satisfactory, both qualitatively and quantitatively, and the dominant modes are correctly rendered. As expected, the cold theoretical spectrum saturates at high k_{\perp} while the hot version displays a local extremum for an oblique wave vector, as the kinetic pressure prevents the pinching of high- k_{\perp} small filaments (Silva et al., 2002).

Moreover, for high k_{\perp} , the PIC spectrum describes wavelengths only a few cells long small enough to be affected by the smoothing algorithm. This explains why these modes' growth is slower than expected. An electron skin depth of several hundred cells would almost certainly provide sufficient separation from the filtering length for the plots to agree better with the theoretical ones in this region, but such a simulation could not be large enough, or run long enough, to observe the later evolution without prohibitive computational expense.

Modes with $k_{\perp} = 0$ are purely electrostatic, and produce no magnetic field. This is why their growth rate is much better rendered when measuring the E spectrum rather than the B spectrum. Conversely, modes with $k_{\parallel} = 0$ are mostly electromagnetic, which explains why their growth is only evident on the B spectrum.

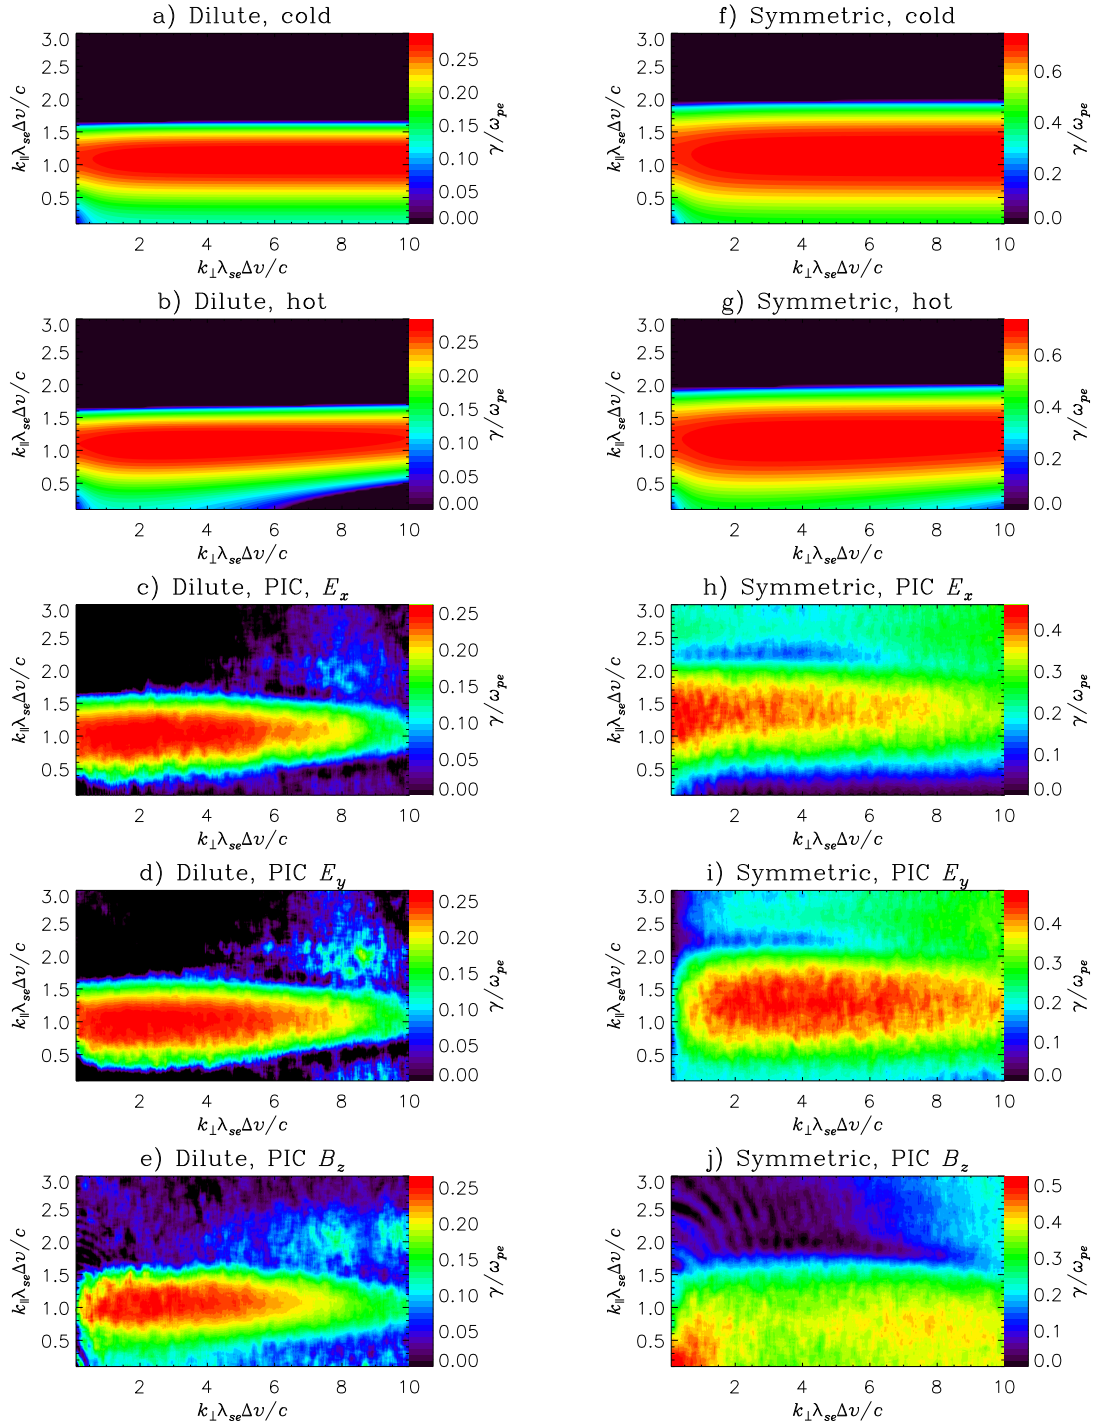


Figure 6.1 The measured or calculated growth rates (in units of electron

Figure 6.1 (Continued)

plasma frequency) of unstable modes with wave vector k_{\parallel}, k_{\perp} (scaled by $\lambda_{se}\Delta v/c$) for the dilute (left column) and symmetric (right column) stream configurations. Within each column, the uppermost plot (**a,f**) is the instantaneous growth rate calculated in the zero-temperature limit. The next plot (**b,g**), a finite thermal spread of $v_{th,e} = 0.01c$ has been included in the calculation. The remaining three rows are the average growth rate measured in the early stages of high-resolution PIC simulations, for the growth of perturbations in E_x (**c,h**), E_y (**d,i**), and B_z (**e,j**). The growth-rate measurement is performed by calculating the fractional increase in the instantaneous 2-D Fourier power spectrum for the selected field component between times t_0 and t_1 . For the dilute configuration, $\omega_{pe}t_0 \approx 15.1$ and $\omega_{pe}t_1 \approx 18.9$; for the symmetric configuration, $\omega_{pe}t_0 \approx 0.2$ and $\omega_{pe}t_1 \approx 7.6$. By virtue of being an average over a finite interval, the growth rate obtained in this way is not a pure, instantaneous quantity like the top two rows of plots. The changing conditions in the stream and counterstream alter the instantaneous growth rate during the measurement interval, resulting in the minor differences that appear between the upper and lower parts of each column.

6.5 Results

The qualitative behavior of the counterstreaming-plasma system in the absence of cosmic rays exhibits a dependence on the density ratio w particularly in the earliest stages of evolution, and at later times the impact of the magnetic field b becomes prominent. Although the details differ from one simulation to the next, the general behavior is similar to that seen in the three-dimensional simulations of [Frederiksen et al. \(2004\)](#), in which the collision of electron-ion plasmas is characterized by the formation of current channels first in the electrons, and later in the ions, which proceed to merge and grow. In our experiments, this growth of the channels and associated structure in the magnetic field reaches a size comparable to the simulation domain, at which point the imposed periodic boundary conditions prevent further growth. However, the time necessary for this to occur – hundreds or thousands of ω_{pe}^{-1} – may exceed the residence time of a given particle in the subshock.

The benchmark behavior for our simulations of negligible cosmic-ray energy density is plotted in [Figure 6.5](#) (magnetic field absent), [Figure 6.3](#) (weak magnetic field), and [Figure 6.4](#) (strong magnetic field); the three columns of each plot correspond to the symmetric, intermediate, and dilute density ratios w . The role of the density ratio is prominent in the action of the two-stream instability on the drift velocity of the counterstreaming electron populations

(Medvedev and Loeb, 1999). For all considered values of magnetic field b , there is an abrupt deceleration of the electrons, both the stream and counterstream, when $20 < \omega_{pe}t < 30$. In the $w = 1$ symmetric case, this initial deceleration strips the electrons of nearly 90% of their relative drift, but in the dilute case, less than half the drift is removed. The electric and magnetic fields are amplified substantially at this point, but the reduced electron drift suppresses further immediate amplification. The ions are slower to respond, and their evolution differs qualitatively from that of the electrons: they form spatially alternating long-lived channels of current: as in lanes of vehicular traffic, ions moving one direction become spatially separated from those moving the opposite direction, and this greatly lengthens the time for the ion drift velocities to converge, though considerable heating of the ions takes place prior to any significant systematic deceleration.

The convergence of ion drift velocities reveals the primary effect of the magnetic field: the ion speeds remain well separated throughout the simulation lifetime in the absent or weak magnetic-field cases, but the strong magnetic field brings them together within roughly $10^4 \omega_{pe}^{-1}$ for all density ratios. It may be that at least in this two-dimensional simulation, the transverse motion necessary for separating the ions into current channels is slightly inhibited by the guiding magnetic field, increasing the extent to which the counterstreaming ion populations are forced to interact. However, it is worth noting once more that the late-term behavior of the simulations is suspect on account of the structure size becoming comparable to the domain boundaries, an artificial upper limit.

6.5.1 Behavior including cosmic rays

For the configurations we considered, the presence of cosmic rays does not appear to result in any significant deviations from the behavior observed in their absence. When their energy density is given a value intended to represent conditions typical of the Galactic disk, no differences from the negligible-cosmic-ray configuration are observed. When cosmic rays are given an exaggerated abundance, subtle effects do appear in the simulation, but upon inspection they are disregarded for one of two reasons: they can be dismissed as numerical effects arising

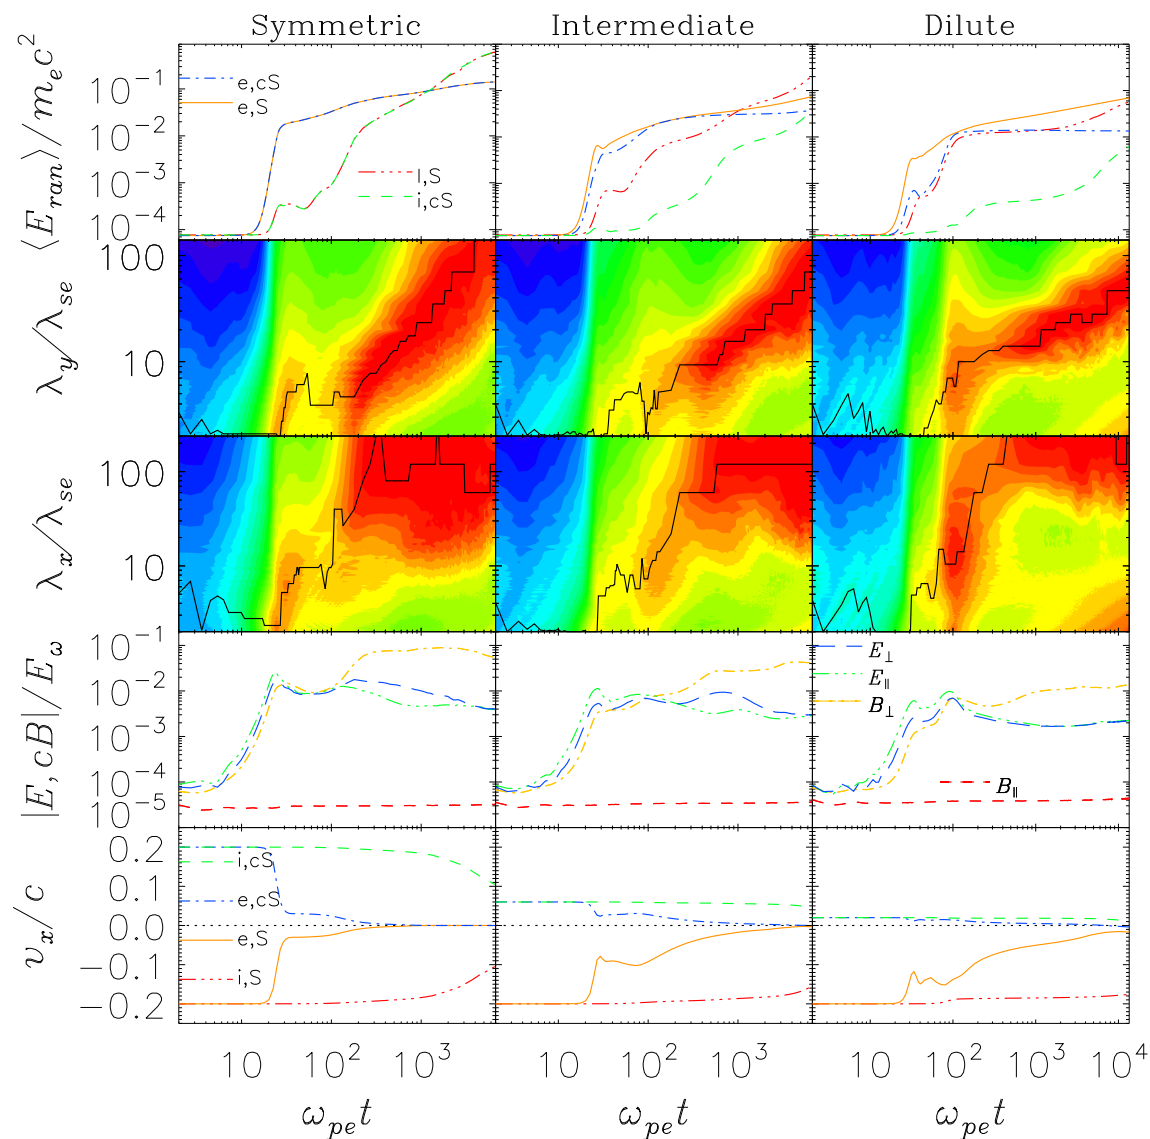


Figure 6.2 Selected results for the “negligible-absent” family of simulations, in which cosmic rays are negligible and an initial homogeneous magnetic field is absent. From left to right, the density ratio of the stream to the counterstream is 1, 0.3, and 0.1, corresponding to the “symmetric,” “intermediate,” and “dilute” cases, respectively. The bottom row shows the drift velocity of the stream electrons (solid line), stream ions (dot-dot-dot-dashed line), counterstream electrons (dot-dashed line), and counterstream ions (dashed line). A dotted line is included at zero velocity, representing the cosmic-ray ions and electrons. The second-to-bottom row displays the root-mean-square amplitudes of the electric and magnetic field parallel and

Figure 6.2 (Continued)

perpendicular to the drift axis, where the magnetic field has been multiplied by c and all components are expressed in units of $E_\omega \equiv \omega_{pe} c m_e / e$. The second and third rows from the top illustrate the time evolution of the Fourier spectra of B_z both parallel and perpendicular to the drift, with wavelengths measured in units of the electron skin depth λ_{se} . A solid black line traces the evolution of the dominant values of λ_x and λ_y . The top row plots the mean random kinetic energy (i.e. thermal energy) per particle, $E_{ran} = (\gamma' - 1) m c^2$ where γ' is the particle Lorentz factor in the local plasma rest frame (measured over a square-shaped region of area $100\Delta^2$). The lines in the top row represent the same populations as in the bottom row.

from the finite number of cosmic-ray particles employed in our simulation, or else they result in minor quantitative changes in the evolution that are of greatest prominence only when the effect of the periodic boundaries is already non-negligible. When the cosmic-ray weight is comparable to that of the plasma particles, their speed (approximately c) maximizes the amplitude of current-density fluctuations resulting from the statistically expected local departures from homogeneity.

In order to verify that the observed differences resulted from our choice of representation and not from some underlying physics, we repeated an “abundant” simulation with a tenfold increase in computational particles representing the same physical cosmic-ray density. Figure 6.5.1 illustrates via the electromagnetic field amplitudes that the statistical noise levels arising at the earliest times saturate at a level $\sqrt{10}$ lower when cosmic rays are represented by 50 particles per cell instead of 5, bringing both the noise level and the detailed time-evolution into better agreement with the “negligible” case. A further significant increase in the computational particle count is too expensive for side-by-side comparison with the simulations in Figure 6.5.1. Using a smaller computational grid, a simulation with 500 cosmic-ray particles per cell (leaving the other plasmas at their original 20 per cell) illustrates the continuation of the trend observed with 50 per cell. Nevertheless, the remaining difference in non-noise behavior is already nearly imperceptible at just fifty computational particles per cell. This effect is paralleled in the other aspects of the system’s evolution in which abundant cosmic rays appeared to result in minor differences, such as drift velocities and wave spectra.

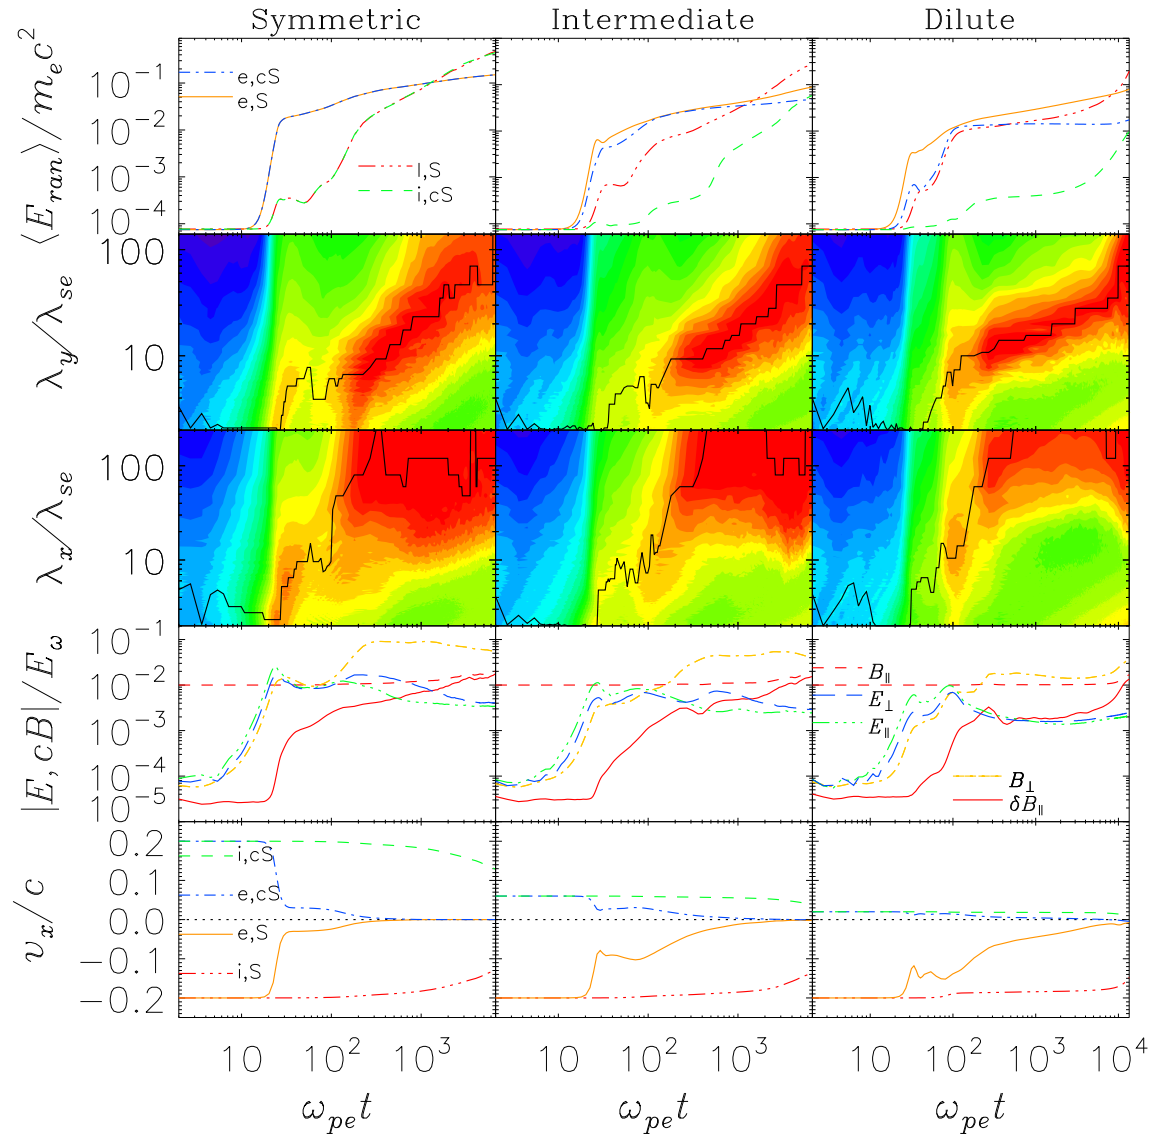


Figure 6.3 Selected results for the “negligible-weak” family of simulations. The initial magnetic field is set such that the electron cyclotron frequency is $\Omega_e \equiv eB/m_e = 0.01\omega_{pe}$. See the caption to Figure 6.5 for a detailed description of each plot.

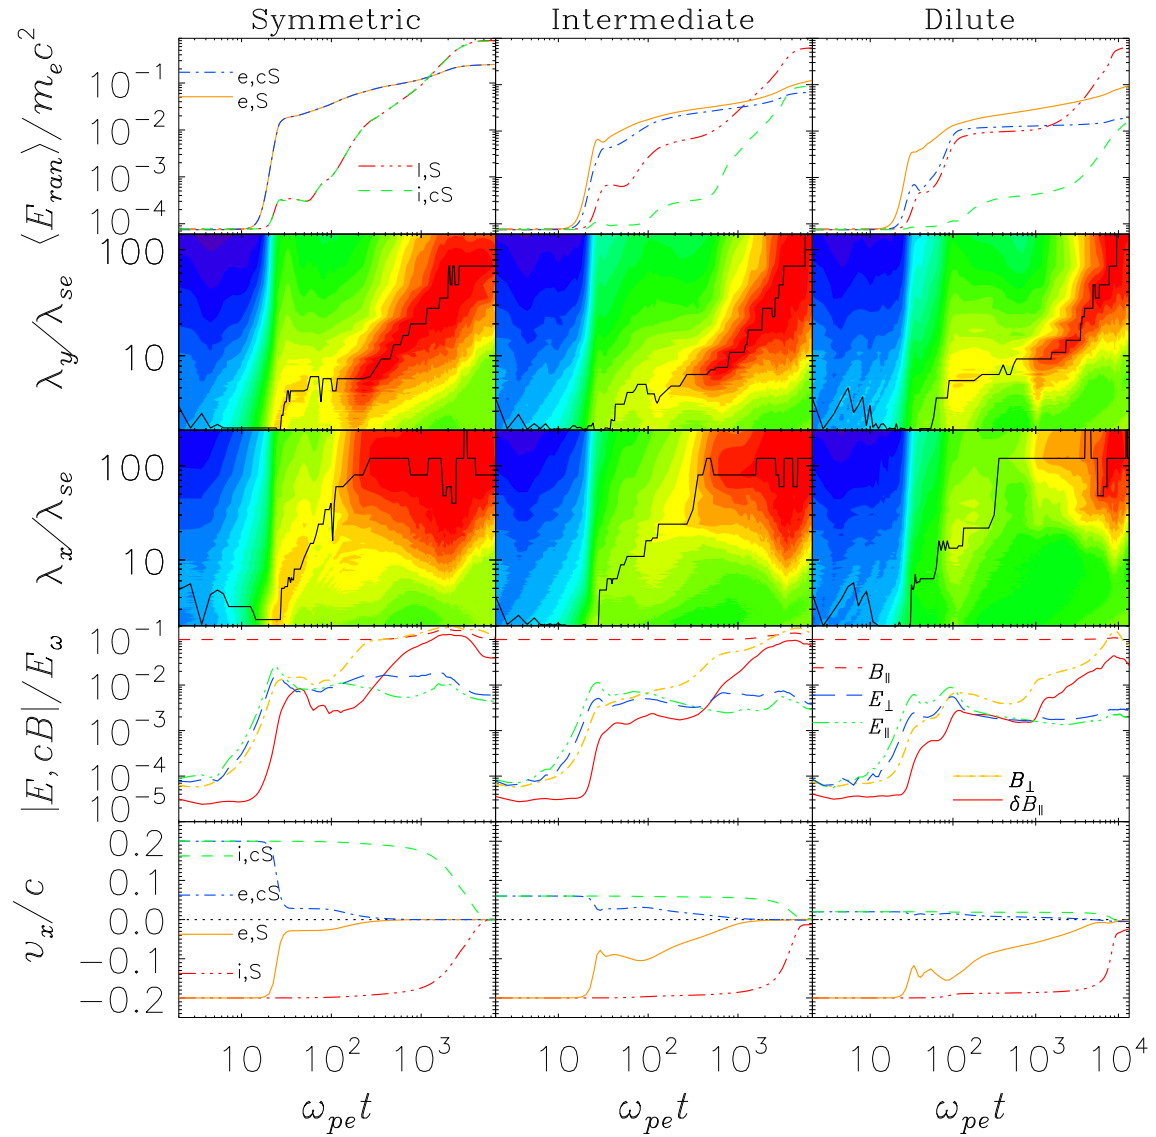


Figure 6.4 Selected results for the “negligible-strong” family of simulations. The initial magnetic field is set such that the electron cyclotron frequency is $\Omega_e = 0.1\omega_{pe}$. See the caption to Figure 6.5 for a detailed description of each plot.

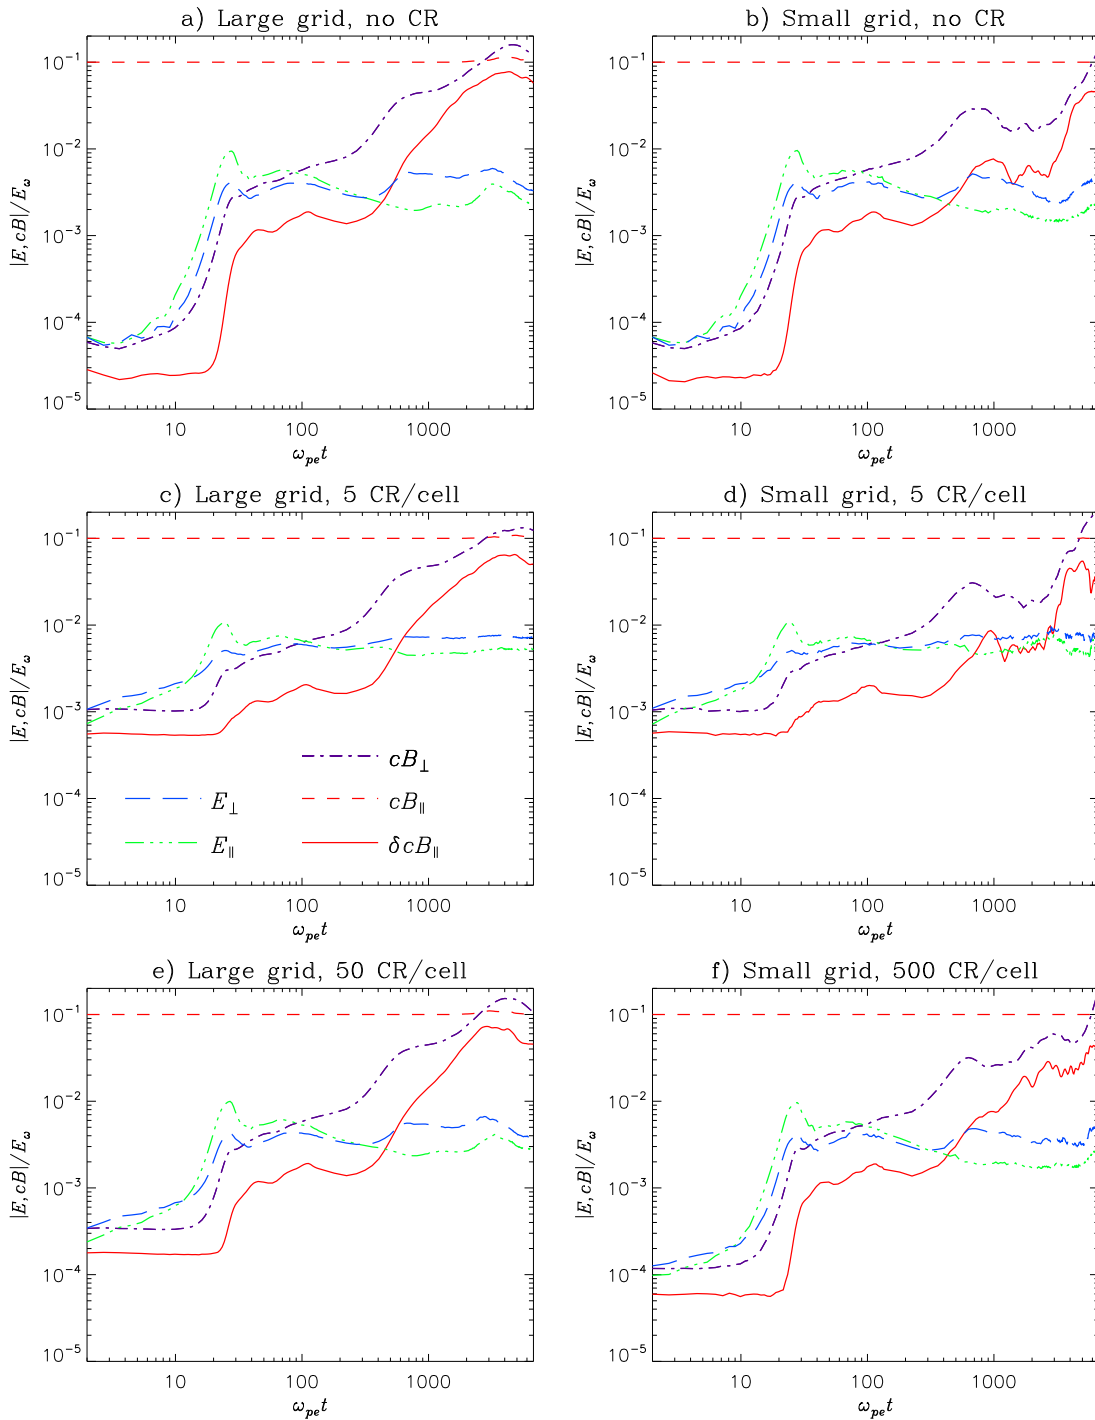


Figure 6.5 The evolution of the electric and magnetic field components in

Figure 6.5 (Continued)

the configuration $b = 0.1$, $w = 0.3$, as a function of the cosmic rays' computational representation. Panel (a) is the case when cosmic rays are negligible. The “abundant” ($w_{CR}\Gamma_{CR} = 10$) cosmic-ray distribution is represented by only five computational particles per cell in panel (c). In panel (e), the number of particles has been increased tenfold, and the initial field amplitudes decrease by approximately $\sqrt{10}$, suggesting that statistical fluctuations in the cosmic-ray contribution to the local current density result in a high noise level in the field components. The column on the right reproduces the conditions of the left column on a grid whose dimensions have been reduced by 90% in each direction, allowing a full-length simulation with 500 cosmic-ray particles per cell. Panels (b) and (d) are comparable to panels (a) and (c), respectively, until $\omega_{pe}t \sim 1000$; the difference beyond this is a result of the smaller grid's periodic boundaries inhibiting the continued spatial growth of structures earlier. In panel (f), the initial noise level is reduced by a factor of $\sqrt{100}$ in all field components, corroborating our earlier supposition that the higher amplitudes are the effect of statistical noise.

6.6 Discussion and conclusions

Motivated by an interest in possible effects of cosmic rays on the physics governing the development of collisionless astrophysical shocks, we have performed several 2.5-dimensional particle-in-cell (PIC) simulations of counterstreaming plasmas with various density ratios and magnetic-field strengths, both with and without a background population of energetic cosmic rays. Before cosmic rays are added to the picture, the system resembles the subject of numerous beam-plasma or interpenetrating-plasmas studies, where the initial behavior of the system can be understood in terms of known instabilities. Most prominently, the counterstreaming electron populations are the first to interact, via symmetric or asymmetric two-stream instabilities, as seen in, e.g., [Medvedev and Loeb \(1999\)](#) and [Frederiksen et al. \(2004\)](#). In particular, the three-dimensional simulations of unmagnetized electron-ion plasma collisions by [Frederiksen et al. \(2004\)](#) demonstrate the formation of merging and growing current filaments in first the electrons and subsequently the ions, and the nuanced relationships among the various populations.

One limitation of our approach is the reduced electron-ion mass ratio, $m_e/m_i = 1/50$. [Bret and Dieckmann \(2010\)](#) have explored the effect of the mass ratio on the hierarchy of unstable modes in beam-plasma systems. By stating that a mass ratio different from $1/1836$ should not change the nature of the most unstable mode, they were able to articulate a criterion

defining the largest acceptable mass ratio. In the present case, the unstable linear spectrum depends very weakly on the mass ratio, when no cosmic rays are introduced. As stated in Section 6.4, finite-mass ions do not add any extra unstable branches to the dispersion equation. As a result, the Bret-Dieckmann criterion is necessarily fulfilled in our configuration without cosmic rays. Turning now to the simulations including cosmic rays, we found no significant cosmic-ray effects with our current mass ratio within the cosmic-ray population. It is thus likely that a real mass ratio would result in an even weaker effect, leaving our conclusions unchanged.

Even at just fifty times the electron mass, however, the behavior of the ions is markedly different. For the most part, the electrons produce and respond to turbulent small-scale electromagnetic fields that serve to mix their distribution functions; the electrons act in concert as a single population while the ions of the stream are still easily distinguished from the ions of the counterstream on account of the filaments. Only when the filamentary structures have enlarged until relatively few repetitions are contained within the simulation plane do the ions begin to converge, and in some of our simulations that remains speculative on account of their finite duration. While it would be possible to extend the simulations ostensibly to observe the eventual convergence, this would be of limited value without simultaneously increasing the dimensions of the simulation domain to mitigate the effect of the periodic boundary conditions. An appreciable increase in size, however, may require us to abandon the simplicity of our present configuration by taking into account large-scale spatial variations, perhaps involving a clear distinction between upstream and downstream regions and an unstable charge-separation layer resulting from differences in the electrons' and ions' evolution.

When we compare the system evolution in the presence of cosmic rays with that in their absence, we find that for the physical configurations studied, cosmic rays do not introduce a statistically significant departure from the unperturbed results described in Section 6.5. This may be a consequence of the comparatively large mean free path and characteristic time scale for evolution of the cosmic-ray distribution: even with the amplification of electric and magnetic fields within the transition layer, cosmic rays of modest energy apparently do not couple to the

dynamics of thermal electrons and ions in any appreciable way. We surmise that at least for unmagnetized or parallel shocks, the impact of cosmic rays – even when their energy density is unusually large – on the instabilities mediating the shock transition is negligible.

Cosmic rays may still indirectly affect various properties of the shock, by modifying the upstream environment from its quiescent characteristics (Stroman *et al.*, 2009): a shock will of necessity propagate differently through a turbulent, heated upstream medium – perhaps with a greatly amplified magnetic field – from the comparatively clean case of a uniform, cold interstellar medium in a gently fluctuating Galactic magnetic field. Irrespective of the cosmic-ray abundance, the rapid development of turbulence in the shock transition layer and the associated heating of the electrons in particular may provide an enlarged pool of candidates for injection into the standard diffusive shock acceleration mechanism.

The combinations of parameters we explored did not yield any appreciable effects that could be attributed to the presence of cosmic rays. While we chose parameters intended to be relevant to nonrelativistic astrophysical shocks in environments where the presence of cosmic rays is suspected, it is nevertheless possible that in other, more exotic environments, cosmic rays may yet play some unforeseen role in the shock microphysics. Future simulations in three dimensions and assigning additional degrees of freedom to the magnetic field and the cosmic-ray population may uncover effects that eluded the present analysis. In any case, the dearth of differences between even the grossly exaggerated cosmic-ray energy density and the system in which it was negligible provide a sense of reassurance that the physics of perhaps a majority of astrophysical shock-forming instabilities can in principle be understood without invoking some direct microscopic interference by a spectator population of cosmic rays.

This research was supported in part by the National Science Foundation both through TeraGrid resources provided by NCSA (Catlett *et al.*, 2007) and under Grant No. PHY05-51164. The work of AB is supported by projects FIS 2006-05389 of the Spanish Ministerio de Educación y Ciencia and PAI08-0182-3162 of the Consejería de Educación y Ciencia de la Junta de Comunidades de Castilla-La Mancha. The work of JN is supported by MNiSW research project N N203 393034, and The Foundation for Polish Science through the HOMING

program, which is supported by a grant from Iceland, Liechtenstein, and Norway through the EEA Financial Mechanism.

CHAPTER 7. Conclusion

How are cosmic rays produced? With the work described in this dissertation, we're making another small step toward the answer. Much of the progress in this field has been in the form of numerous incremental advances, punctuated by the occasional discovery of something particularly invigorating.

Let us briefly review what was accomplished here. The context was the origin of cosmic rays. The observed cosmic-ray flux depends on energy via the functional form of a broken power law. The spectral steepening near 3×10^{15} eV known as the “knee” is thought to represent the beginning of a transition from cosmic rays of Galactic origin to those of extragalactic origin, possibly produced by exotic particle decay or in active galactic nuclei. The theory of diffusive shock acceleration (DSA) provides a mechanism for cosmic-ray production at supernova remnant shocks, of which several are known to exist in the Galaxy, and where the energy budget is favorable for the replenishment of cosmic rays against escape and energy loss.

However, it seems that without some sort of magnetic-field amplification, the maximum energy of cosmic rays produced via DSA falls at least an order of magnitude short of the knee (Lagage and Cesarsky, 1983). Much of the work described within this document (Niemić et al., 2008; Stroman et al., 2009, / Chapter 5) was motivated by recent predictions of significant amplification of the magnetic field in the upstream cosmic-ray precursor (Bell, 2004). Although our initial simulations were dominated by a different plasma instability from the one supposedly responsible for the dramatic amplification, we found that an appropriate adjustment of our parameters brought about the predicted amplification mechanism.

We observed that the magnetic field began to grow at approximately the predicted rate and length scale, but that small density fluctuations appeared when plasma parcels began to move

and collide with each other, producing a turbulent, heated, upstream medium in which motional electric fields could not be eliminated by any Lorentz transformation. The characteristic length scale of the turbulent structures steadily increased as well, until limitations of the periodic boundary conditions prevented further development. If these are indeed the consequences of magnetic-field amplification by Bell's mechanism, they may be relevant in future self-consistent models of particle acceleration.

Some Monte Carlo efforts have already begun to include the nonresonant amplified magnetic field, and explore competition between amplification (at particular, possibly-increasing length scales) and possibly anisotropic cascading of turbulent energy from larger to smaller length scales (Vladimirov et al., 2009). But as the detailed energy- and wavevector-dependence of the cascading is still a matter of approximation, there is still clearly a role for the type of investigation carried out here, albeit on a larger scale.

Other recent work has transported Bell's instability from the SNR-shock environment for consideration elsewhere, particularly where cold-plasma approximations may not be valid or the initial magnetic field may be different from the canonical interstellar value. For instance, Zweibel and Everett (2010) found that even in the intergalactic medium, Bell's instability may operate, driven by the flux of cosmic rays escaping from galaxies, although the characteristic scales bear little resemblance to those of the instability in supernova environments. It may be interesting to see if the intergalactic medium evolves similarly to the interstellar medium upstream of supernova shocks in the presence of a strong cosmic-ray flux, or if effects of other nearby galaxies must be taken into account, for the typical intergalactic spacing is proportionally much closer than young supernova remnants within a galaxy.

The next few years of theoretical efforts are likely to provide us with a much better understanding of the role played by Bell's instability in cosmic-ray acceleration, especially since its ability to generate strong turbulence some distance upstream of the shock has a significant effect on what sort of medium the shock encounters when it finally arrives. The detailed evolution of the upstream plasma and magnetic field between the Bell-active region and the shock itself may introduce a discernable signature in observable downstream quantities, but determining

what exactly to look for must be the objective of the next generation of simulations.

The same tools that allowed us to explore magnetic-field amplification by Bell’s mechanism can also be used, with minimal adaptation, to probe the inner workings of a nonrelativistic collisionless shock. We were certainly not the first to consider the general setup of counter-streaming plasmas or beam-plasma systems ([Stockem et al. \(2008\)](#) and [Bret \(2009\)](#) are both recent studies that have some overlap with our parameters, and contain numerous further references). But in light of the growing awareness of energetic particles being spatially correlated with shocks, both through theoretical mechanisms for local acceleration and through observations of nonthermal emission spectra from shocked regions, we thought it prudent to ascertain whether the microscopic physics governing the transition from upstream to downstream might be affected by the presence of those energetic particles.

The results of our many simulations to this end seem to provide a reassurance that if our understanding of collisionless shocks does have any gaping holes in it, it is unlikely that the role played by cosmic rays is hiding in such a hole. Even in galaxies with a cosmic-ray abundance significantly higher than our own, it appears that local shock formation processes may safely be generalized to those exotic environments.

The increasing availability of computational power makes it likely that a number of the scenarios considered here, both of magnetic-field amplification and shock formation, will be revisited in greater detail in the future. Particle-in-cell simulation is a versatile tool, capable of producing useful results (when properly interpreted) for a wide variety of problems relevant to astrophysics. The amount of guesswork involved in patching PIC results into the products of other simulation techniques (themselves improving in accuracy, for computational improvement benefits all) will continue its downward trend.

Meanwhile, ground-based and space-based observatories continue to push the limits of sensitivity and resolution when observing possible sources of cosmic rays. It seems an altogether fair statement that as we approach the centennial of Victor Hess’s balloon flights, prospects are good and optimism is high that we are at last drawing near to identifying the long sought-after extraterrestrial particle accelerators.

BIBLIOGRAPHY

- Abbasi, R., Abdou, Y., Abu-Zayyad, T., Adams, J., Aguilar, J. A., Ahlers, M., Andeen, K., Auffenberg, J., Bai, X., Baker, M., and et al.: 2010, *ApJ* **718**, L194
- Abdo, A. A., Allen, B. T., Aune, T., Berley, D., Casanova, S., Chen, C., Dingus, B. L., Ellsworth, R. W., Fleysher, L., Fleysher, R., Gonzalez, M. M., Goodman, J. A., Hoffman, C. M., Hopper, B., Hüntemeyer, P. H., Kolterman, B. E., Lansdell, C. P., Linnemann, J. T., McEnery, J. E., Mincer, A. I., Nemethy, P., Noyes, D., Pretz, J., Ryan, J. M., Parkinson, P. M. S., Shoup, A., Sinnis, G., Smith, A. J., Sullivan, G. W., Vasileiou, V., Walker, G. P., Williams, D. A., and Yodh, G. B.: 2009, *ApJ* **698**, 2121
- Acciari, V. A. et al.: 2009, *Nature* **462**, 770
- Acerro, F., Aharonian, F., Akhperjanian, A. G., Anton, G., Barres de Almeida, U., Bazer-Bachi, A. R., Becherini, Y., Behera, B., Beilicke, M., Bernlöhr, K., Bochow, A., Boisson, C., Bolmont, J., Borrel, V., Brucker, J., Brun, F., Brun, P., Bühler, R., Bulik, T., Büsching, I., Boutelier, T., Chadwick, P. M., Charbonnier, A., Chaves, R. C. G., Cheesebrough, A., Conrad, J., Chounet, L., Clapson, A. C., Coignet, G., Dalton, M., Daniel, M. K., Davids, I. D., Degrange, B., Deil, C., Dickinson, H. J., Djannati-Ataï, A., Domainko, W., O’C. Drury, L., Dubois, F., Dubus, G., Dyks, J., Dyrda, M., Egberts, K., Eger, P., Espigat, P., Fallon, L., Farnier, C., Fegan, S., Feinstein, F., Fiasson, A., Förster, A., Fontaine, G., Füßling, M., Gabici, S., Gallant, Y. A., Gérard, L., Gerbig, D., Giebels, B., Glicenstein, J. F., Glück, B., Goret, P., Göring, D., Hauser, D., Hauser, M., Heinz, S., Heinzlmann, G., Henri, G., Hermann, G., Hinton, J. A., Hoffmann, A., Hofmann, W., Hofverberg, P., Holleran, M., Hoppe, S., Horns, D., Jacholkowska, A., de Jager, O. C., Jahn, C., Jung, I., Katarzyński, K.,

- Katz, U., Kaufmann, S., Kerschhaggl, M., Khangulyan, D., Khélifi, B., Keogh, D., Klochkov, D., Kluźniak, W., Kneiske, T., Komin, N., Kosack, K., Kossakowski, R., Lamanna, G., Lemoine-Goumard, M., Lenain, J., Lohse, T., Marandon, V., Marcowith, A., Masbou, J., Maurin, D., McComb, T. J. L., Medina, M. C., Méhault, J., Moderski, R., Moulin, E., Naumann-Godo, M., de Naurois, M., Nedbal, D., Nekrassov, D., Nicholas, B., Niemiec, J., Nolan, S. J., Ohm, S., Olive, J., de Oña Wilhelmi, E., Orford, K. J., Ostrowski, M., Panter, M., Paz Arribas, M., Pedalletti, G., Pelletier, G., Petrucci, P., Pita, S., Pühlhofer, G., Punch, M., Quirrenbach, A., Raubenheimer, B. C., Raue, M., Rayner, S. M., Reimer, O., Renaud, M., de Los Reyes, R., Rieger, F., Ripken, J., Rob, L., Rosier-Lees, S., Rowell, G., Rudak, B., Rulten, C. B., Ruppel, J., Ryde, F., Sahakian, V., Santangelo, A., Schlickeiser, R., Schöck, F. M., Schönwald, A., Schwanke, U., Schwarzburg, S., Schwemmer, S., Shalchi, A., Sushch, I., Sikora, M., Skilton, J. L., Sol, H., Stawarz, L., Steenkamp, R., Stegmann, C., Stinzing, F., Superina, G., Szostek, A., Tam, P. H., Tavernet, J., Terrier, R., Tibolla, O., Tluczykont, M., van Eldik, C., Vasileiadis, G., Venter, C., Venter, L., Vialle, J. P., Vincent, P., Vink, J., Vivier, M., Völk, H. J., Volpe, F., Vorobiov, S., Wagner, S. J., Ward, M., Zdziarski, A. A., and Zech, A.: 2010, *A&A* **516**, A62+
- Achterberg, A.: 1983, *A&A* **119**, 274
- Amato, E. and Blasi, P.: 2009, *MNRAS* **392**, 1591
- Amsler, C. et al.: 2008, *Phys. Lett.* **B667**, 1
- Axford, W. I., Leer, E., and Skadron, G.: 1977, in *International Cosmic Ray Conference*, Vol. 11 of *International Cosmic Ray Conference*, pp 132–137
- Bell, A. R.: 1978a, *MNRAS* **182**, 147
- Bell, A. R.: 1978b, *MNRAS* **182**, 443
- Bell, A. R.: 2004, *MNRAS* **353**, 550
- Bell, A. R. and Lucek, S. G.: 2001, *MNRAS* **321**, 433

- Bird, D. J., Corbato, S. C., Dai, H. Y., Elbert, J. W., Green, K. D., Huang, M. A., Kieda, D. B., Ko, S., Larsen, C. G., Loh, E. C., Luo, M. Z., Salamon, M. H., Smith, J. D., Sokolsky, P., Sommers, P., Tang, J. K. K., and Thomas, S. B.: 1995, *ApJ* **441**, 144
- Birdsall, C. K. and Langdon: 1991, *Plasma Physics via Computer Simulation (Series on Plasma Physics)*, Taylor & Francis
- Blandford, R. D. and Ostriker, J. P.: 1978, *ApJ* **221**, L29
- Blasi, P. and Amato, E.: 2008, in *International Cosmic Ray Conference*, Vol. 2 of *International Cosmic Ray Conference*, pp 235–238
- Bret, A.: 2009, *ApJ* **699**, 990
- Bret, A. and Deutsch, C.: 2006, *Physics of Plasmas* **13(4)**, 042106
- Bret, A. and Dieckmann, M. E.: 2010, *Physics of Plasmas* **17(3)**, 032109
- Buneman, O.: 1993, *Computer Space Plasma Physics: Simulation Techniques and Software*, ed. H. Matsumoto & Y. Omura, pp 67–84, Tokyo: Terra
- Cai, D., Li, Y., Nishikawa, K., and et al.: 2003, in J. Büchner, C. Dum, & M. Scholer (ed.), *Space Plasma Simulation*, Vol. 615 of *Lecture Notes in Physics*, Berlin Springer Verlag, pp 25–53
- Catlett, C. et al.: 2007, *TeraGrid: Analysis of Organization, System Architecture, and Middleware Enabling New Types of Applications*, IOS Press
- Drury, L. O.: 1983, *Reports on Progress in Physics* **46**, 973
- Drury, L. O. and Völk, J. H.: 1981, *ApJ* **248**, 344
- Eichler, D.: 1984, *ApJ* **277**, 429
- Falle, S. A. E. G. and Giddings, J. R.: 1987, *MNRAS* **225**, 399
- Fermi, E.: 1949, *Physical Review* **75**, 1169

- Frederiksen, J. T., Hededal, C. B., Haugbølle, T., and Nordlund, Å.: 2004, *ApJ* **608**, L13
- Gargaté, L., Fonseca, R. A., Niemiec, J., Pohl, M., Bingham, R., and Silva, L. O.: 2010, *ApJ* **711**, L127
- Giacalone, J. and Jokipii, J. R.: 2007, *ApJ* **663**, L41
- Greenwood, A. D., Cartwright, K. L., Luginsland, J. W., and Baca, E. A.: 2004, *Journal of Computational Physics* **201**, 665
- Greisen, K.: 1966, *Physical Review Letters* **16**, 748
- Inoue, T., Yamazaki, R., and Inutsuka, S.-i.: 2009, *ApJ* **695**, 825
- Katsuda, S., Petre, R., Long, K. S., Reynolds, S. P., Winkler, P. F., Mori, K., and Tsunemi, H.: 2009, *ApJ* **692**, L105
- Krymskii, G. F.: 1977, *Akademiia Nauk SSSR Doklady* **234**, 1306
- Lagage, P. O. and Cesarsky, C. J.: 1983, *A&A* **125**, 249
- Li, Z. and Waxman, E.: 2006, *ApJ* **651**, 328
- Longair, M. S.: 1992, *High energy astrophysics. Vol.1: Particles, photons and their detection*
- Lucek, S. G. and Bell, A. R.: 2000, *MNRAS* **314**, 65
- Luo, Q. and Melrose, D.: 2009, *MNRAS* **397**, 1402
- Malkov, M. A. and O’C Drury, L.: 2001, *Reports on Progress in Physics* **64**, 429
- Marcowith, A., Lemoine, M., and Pelletier, G.: 2006, *A&A* **453**, 193
- Medvedev, M. V. and Loeb, A.: 1999, *ApJ* **526**, 697
- Niemiec, J. and Pohl, M.: 2008, in *International Cosmic Ray Conference*, Vol. 2 of *International Cosmic Ray Conference*, pp 279–282
- Niemiec, J., Pohl, M., Bret, A., and Stroman, T.: 2010, *ApJ* **709**, 1148

- Niemiec, J., Pohl, M., Stroman, T., and Nishikawa, K.: 2008, *ApJ* **684**, 1174
- Ohira, Y., Reville, B., Kirk, J. G., and Takahara, F.: 2009, *ApJ* **698**, 445
- Pelletier, G., Lemoine, M., and Marcowith, A.: 2006, *A&A* **453**, 181
- Pohl, M., Yan, H., and Lazarian, A.: 2005, *ApJ* **626**, L101
- Ptuskin, V. S. and Zirakashvili, V. N.: 2003, *A&A* **403**, 1
- Rakowski, C. E., Laming, J. M., and Ghavamian, P.: 2008, *ApJ* **684**, 348
- Reville, B., Kirk, J. G., and Duffy, P.: 2006, *Plasma Physics and Controlled Fusion* **48**, 1741
- Reville, B., Kirk, J. G., and Duffy, P.: 2009, *ApJ* **694**, 951
- Reville, B., Kirk, J. G., Duffy, P., and O'Sullivan, S.: 2007, *A&A* **475**, 435
- Reville, B., O'Sullivan, S., Duffy, P., and Kirk, J. G.: 2008, *MNRAS* **386**, 509
- Reynolds, S. P.: 1998, *ApJ* **493**, 375
- Reynolds, S. P.: 2008, *ARA&A* **46**, 89
- Riquelme, M. A. and Spitkovsky, A.: 2009, *ApJ* **694**, 626
- Shu, F. H.: 1992, *The physics of astrophysics. Volume II: Gas dynamics.*
- Silva, L. O., Fonseca, R. A., Tonge, J. W., Dawson, J. M., Mori, W. B., and Medvedev, M. V.:
2003, *ApJ* **596**, L121
- Silva, L. O., Fonseca, R. A., Tonge, J. W., Mori, W. B., and Dawson, J. M.: 2002, *Physics of
Plasmas* **9**, 2458
- Skilling, J.: 1975a, *MNRAS* **172**, 557
- Skilling, J.: 1975b, *MNRAS* **173**, 245
- Skilling, J.: 1975c, *MNRAS* **173**, 255

- Sokolsky, P. and for the HiRes Collaboration: 2010, *ArXiv e-prints*
- Stockem, A., Lazar, M., Shukla, P. K., and Smolyakov, A.: 2008, *Journal of Plasma Physics* **75**, 529
- Stroman, T., Pohl, M., and Niemiec, J.: 2009, *ApJ* **706**, 38
- Stroman, W. and Pohl, M.: 2009, *ApJ* **696**, 1864
- Strong, A. W., Moskalenko, I. V., and Ptuskin, V. S.: 2007, *Annual Review of Nuclear and Particle Science* **57**, 285
- Swordy, S. P.: 2001, *Space Sci. Rev.* **99**, 85
- Umeda, T., Omura, Y., Tominaga, T., and Matsumoto, H.: 2003, *Computer Physics Communications* **156**, 73
- Villasenor, J. and Buneman, O.: 1992, *Computer Physics Communications* **69**, 306
- Vink, J. and Laming, J. M.: 2003, *ApJ* **584**, 758
- Vladimirov, A. E., Bykov, A. M., and Ellison, D. C.: 2008, *ApJ* **688**, 1084
- Vladimirov, A. E., Bykov, A. M., and Ellison, D. C.: 2009, *ApJ* **703**, L29
- Voelk, H. J. and Biermann, P. L.: 1988, *ApJ* **333**, L65
- Völk, H. J., Berezhko, E. G., and Ksenofontov, L. T.: 2005, *A&A* **433**, 229
- Webb, G. M., Axford, W. I., and Terasawa, T.: 1983, *ApJ* **270**, 537
- Wentzel, D. G.: 1974, *ARA&A* **12**, 71
- Winske, D. and Leroy, M. M.: 1984, *J. Geophys. Res.* **89**, 2673
- Zatsepin, G. T. and Kuz'min, V. A.: 1966, *Soviet Journal of Experimental and Theoretical Physics Letters* **4**, 78
- Zirakashvili, V. N. and Ptuskin, V. S.: 2008, *ApJ* **678**, 939

Zirakashvili, V. N., Ptuskin, V. S., and Völk, H. J.: 2008, *ApJ* **678**, 255

Zweibel, E. G. and Everett, J. E.: 2010, *ApJ* **709**, 1412