Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **485,** 5105–5119 (2019)
Advance Access publication 2019 January 30

doi:10.1093/mnras/stz232

Kinetic simulations of mildly relativistic shocks – I. Particle acceleration in high Mach number shocks

P. Crumley [1,2]★ D. Caprioli,[3] S. Markoff[2] and A. Spitkovsky[1]

[1]*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08554, USA*
[2]*Anton Pannekoek Institute for Astronomy, University of Amsterdam & GRAPPA, PO Box 94249, NL-1090 GE Amsterdam, the Netherlands*
[3]*Department of Astronomy & Astrophysics, University of Chicago, Chicago, IL 60637, USA*

Accepted 2019 January 16. Received 2019 January 15; in original form 2018 September 26

ABSTRACT

We use fully kinetic particle-in-cell simulations with unprecedentedly large transverse box sizes to study particle acceleration in weakly magnetized mildly relativistic shocks travelling at a velocity $\approx 0.75c$ and a Mach number of 15. We examine both subluminal (quasi-parallel) and superluminal (quasi-perpendicular) magnetic field orientations. We find that quasi-parallel shocks are mediated by a filamentary non-resonant (Bell) instability driven by returning ions, producing magnetic fluctuations on scales comparable to the ion gyroradius. In quasi-parallel shocks, both electrons and ions are accelerated into non-thermal power laws whose maximum energy grows linearly with time. The upstream heating of electrons is small, and the two species enter the shock front in rough thermal equilibrium. The shock's structure is complex; the current of returning non-thermal ions evacuates cavities in the upstream that form filaments of amplified magnetic fields once advected downstream. At late times, 10 per cent of the shock's energy goes into non-thermal protons and $\gtrsim$10 per cent into magnetic fields. We find that properly capturing the magnetic turbulence driven by the non-thermal ions is important for properly measuring the energy fraction of non-thermal electrons, ε_e. We find $\varepsilon_e \sim 5 \times 10^{-4}$ for quasi-parallel shocks with $v = 0.75c$, slightly larger than what was measured in simulations of non-relativistic shocks. In quasi-perpendicular shocks, no non-thermal power-law develops in ions or electrons. The ion acceleration efficiency in quasi-parallel shocks suggests that astrophysical objects that could host mildly relativistic quasi-parallel shocks – for example, the jets of active galactic nuclei or microquasars – may be important sources of cosmic rays and their secondaries, such as gamma-rays and neutrinos.

Key words: acceleration of particles – plasmas – radiation mechanism: non-thermal; shock waves.

1 INTRODUCTION

Collisionless shocks are capable of producing non-thermal particles in power laws extending for many orders of magnitude in energy. Krymskii (1977), Axford, Leer & Skadron (1977), Bell (1978), and Blandford & Ostriker (1978) realized that collisionless shocks would be a manifestation of a first-order Fermi/diffusive shock acceleration (DSA) process (Fermi 1949); particles can gain energy from the velocity difference between the upstream and downstream. The particles can reach large energies by being scattered back towards the shock front by magnetic turbulence, thereby crossing the shock front many times. The observed non-thermal radiation from astrophysical systems with strong shocks such as supernova remnants and gamma-ray burst afterglows is strong evidence that shocks in astrophysical systems do indeed accelerate particles.

Classic DSA theory has been extremely successful at explaining how non-thermal particles can be accelerated into a power law extending for decades in energy, but it is unable to determine why a particle becomes non-thermal in the first place (see e.g. Drury 1983; Blandford & Eichler 1987; Malkov & Drury 2001 for reviews). That is, a particle may need to be pre-energized before it is 'injected' into DSA. Because DSA theory is unable to answer how particles are injected, it cannot predict the fraction of the shock's energy that is placed into non-thermal particles. Constraining the non-thermal energy fraction of ions (ε_p), electrons (ε_e), or magnetic field (ε_B) requires either modelling multiwavelength observations or numerical simulations. While the values of ε_p, ε_e, and ε_B depend on several shock parameters (e.g. sonic and Alfvénic Mach numbers, magnetic field inclination), both observation and simulations show

★ E-mail: pcrumley@princeton.edu

a striking difference in ε_e between non-relativistic and relativistic shocks. Relativistic shocks appear to create non-thermal electrons more effectively than non-relativistic ones.

Observational evidence for a changing ε_e with shock velocity comes from a wide range of astrophysical systems. The values obtained are model dependent, but there is a global trend of increasing ε_e with shock velocity. The non-relativistic shocks of the young supernova remnant *Tycho* is an inefficient electron accelerator, with $\varepsilon_e \sim 10^{-3}\varepsilon_p$, or $\varepsilon_e \sim 10^{-4}$ if $\varepsilon_p = 0.1$ (Morlino & Caprioli 2012). A low value of $\varepsilon_e/\varepsilon_p$ in non-relativistic shocks is supported by observations of cosmic rays at Earth. For every cosmic ray electron with ∼30 GeV, several hundred protons are detected (e.g. Aguilar et al. 2015, 2014). On the other hand, the ultra-relativistic shocks that produce gamma-ray burst afterglows are well fit with the much larger $\varepsilon_e \sim 0.1$ (Wijers & Galama 1999; Panaitescu & Kumar 2000; Santana, Barniol Duran & Kumar 2014). Particle-in-cell simulations of electron-ion shocks agree with this trend. The simulations of 1D non-relativistic shocks travelling at $v = 0.1c$ in Park, Caprioli & Spitkovsky (2015) show ε_e of $\sim 10^{-4}$ and the relativistic electron–ion shocks in Sironi & Spitkovsky (2011) have an $\varepsilon_e \sim 0.1$ at $\gamma = 15$.

The production of non-thermal particles depends on the magnetic turbulence at the shock front. In collisionless shocks, binary collisions between particles are not capable of isotropizing the upstream bulk motion of the plasma to subsonic speeds. Instead, the shock must be mediated by electric and magnetic fields that are generated by collective motions in the plasma. There are different mediation mechanisms for shocks, but we focus on the two instabilities that mediate mildly relativistic high Alfvénic Mach number shocks: the Weibel instability and the Bell instability. The Weibel instability occurs when initially current-neutral counter-streaming plasma undergoes transverse filamentation into skin-depth wide oppositely directed current filaments, which enhance the transverse magnetic field and eventually isotropize the particles (Weibel 1959). The Weibel instability mediates weakly-magnetized, quasi-parallel, collisionless relativistic shocks (Medvedev & Loeb 1999; Spitkovsky 2008a; Sironi & Spitkovsky 2011). The Bell instability is caused by a current of high-energy ions that are freely streaming through a magnetized plasma. The current in this beam travels along B_0 and amplifies transverse waves. The current is balanced by a returning current carried by the background plasma. Because of the different rigidity of energetic ions and compensating electrons, regions experiencing transverse magnetic fluctuations $\delta \boldsymbol{B}$ lose current balance, and the background plasma feels a $\boldsymbol{J} \times \delta \boldsymbol{B}$ force that pushes it outwards, growing the fluctuation (Bell 2004, 2005). In its non-linear stage, the Bell instability is capable of producing magnetic fluctuations the size of the ion's gyroradius, and therefore, it may play an important role in scattering the highest energy particles (Caprioli & Spitkovsky 2014b,c).

Mildly relativistic shocks are a good place to accelerate particles to high energies because both the shock speed (β_s) and Lorentz factor (γ_s) are of order unity. A large β_s means the particle gains more energy with each crossing of the shock front; $\Delta E/E \propto \beta_s$ in non-relativistic shocks, and $\Delta E/E \sim 1$ after the first crossing in relativistic shocks (Bell 1978; Quenby & Lieu 1989; Achterberg et al. 2001). However, in ultra-relativistic shocks, the component of the magnetic field perpendicular to the shock front is magnified by a factor γ_s due to a Lorentz transformation. Therefore, nearly all magnetized ultra-relativistic shocks will be quasi-perpendicular. The quasi-perpendicular magnetic field geometry may severely suppress particle injection in magnetized relativistic shocks because it is superluminal, i.e. particles confined to background magnetic field lines would need to be travelling faster than the speed of light to return upstream (Begelman & Kirk 1990; Sironi & Spitkovsky 2009; Sironi, Spitkovsky & Arons 2013). Mildly relativistic shocks with $\gamma_s\beta_s \sim 1$ have an advantage in producing high-energy particles because the acceleration time is short, but the shock is subluminal for a larger range of upstream magnetic field inclinations than ultra-relativistic shocks.

To accelerate the highest energy particles, it is not enough to have the particle gain $\Delta E/E \sim 1$ at the shock front. The particles also must return to the shock front quickly enough, so they cross the shock many times in a short time-scale. The time it takes to accelerate a particle to energy E in a non-relativistic shock depends on the diffusion coefficient D and is $t_{acc} = c^2 D(E)/\beta_s^2$ (e.g. Drury 1983). If the turbulence has large magnetic field fluctuations at all scales, the diffusion coefficient depends linearly on energy, and the maximum energy grows linearly with time. Previous work has shown that when the magnetic turbulence is at small scales, such as the upstream fluctuations that occur in unmagnetized shocks, the maximum energy increases more slowly as $\sqrt{t}$ (e.g. Sironi et al. 2013). The difference between the two scalings may be the difference in whether an object can be an ultra-high energy cosmic ray progenitor.

This paper is the first in a series of papers that kinetically examines mildly relativistic shocks ($\gamma_s\beta_s \sim 1$). In this paper, we show the necessary ingredients to properly capture the magnetic field structure at the shock front in a 2D simulation of weakly magnetized, electron–ion shocks, and to measure ε_e. We examine mildly relativistic shocks travelling at $\gamma\beta \approx 1$ using self-consistent, fully kinetic, particle-in-cell simulations. Most previous studies of the mildly relativistic shocks have used Monte Carlo and semi-analytical methods that require assumptions about the turbulence and the fraction of non-thermal particles (Kirk et al. 2000; Keshet & Waxman 2005; Morlino, Blasi & Vietri 2007; Summerlin & Baring 2012; Ellison, Warren & Bykov 2013), or fully kinetic simulations that were too short to study the late-time evolution of the shock structure (Dieckmann, Shukla & Eliasson 2006; Bohdan et al. 2017; Dieckmann & Bret 2018). We do not vary the shock velocity in this paper, but we plan on using the results of this paper to examine how ε_e, ε_B, and ε_p change in the transition between non-relativistic and relativistic shocks in a future work.

Our paper is organized as follows. First, we describe our simulation set-up in Section 2. In Sections 3 and 4, we analyse the structure of and particle acceleration in a mildly relativistic quasi-parallel shock. In Section 5, we show how resolving the ion gyroradius is important to capture the turbulence at the shock front and measuring ε_e. In Section 6, we show that superluminal shocks are unable to accelerate particles into a power law. Finally, we conclude by summarizing our results and discussing their astrophysical relevance. Many astrophysical objects host mildly relativistic outflows capable of producing strong shocks. However, the magnetized mildly relativistic shocks studied in this paper are most relevant to the internal shocks in jets with modest Lorentz factors like the jets observed as bright sources of non-thermal radiation in microquasars and low-luminosity active galactic nuclei (AGNs).

2 METHODOLOGY

We use the parallel electromagnetic PIC code *TRISTAN-MP* (Buneman 1993; Spitkovsky 2005) to simulate mildly relativistic electron–ion shocks (Buneman 1993; Spitkovsky 2005). The computational domain is a 2D Cartesian xy grid, with length L_x and

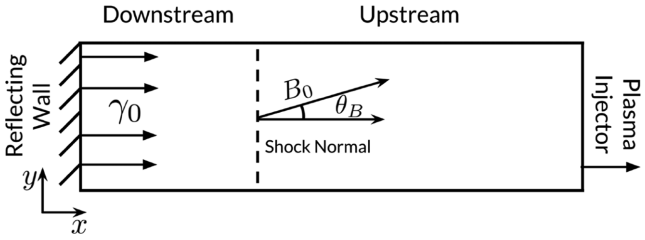

Figure 1. A schematic of our shock simulation. Plasma is injected at rest at the right boundary and a reflecting wall is pushed into the plasma with Lorentz factor γ_0, forming the shock. The injected plasma is initialized with a background magnetic field that lies in the $x-y$ plane. We use periodic boundaries in the y-direction.

width L_y. While the computational domain is 2D, we track all three spatial components of the field quantities and the particle momenta. We use periodic boundary conditions along the y-direction and we place reflecting walls at the left and the right boundary of the computational domain. A schematic of our simulation set-up is shown in Fig. 1.

Our simulation set-up is very similar to previous PIC simulations of collisionless shocks, e.g. Spitkovsky (2008a), Sironi & Spitkovsky (2011), and Sironi et al. (2013), with one major difference: the simulation's rest frame. Our simulations are performed in the upstream rest frame, while previous simulations were performed in the downstream rest frame. We find that having a stationary upstream helps reduce numerical heating that occurs when simulating a moving cold plasma beam. At the right boundary, we inject electrons and ions drawn from stationary Maxwell–Jüttner distributions with thermal spread $\Delta\gamma \equiv m_i c^2/k_B T$, where k_B is the Boltzmann constant, T is the temperature, and m_i is the ion mass. The two species are in thermal equilibrium with $T_i = T_e$. The left wall travels to the right with a Lorentz factor $\gamma_0 = 1.5$, $v_0 \approx 0.75c$, driving a shock. We truncate the computational domain left of the wall, and measure all distances relative to the left wall.[1] The right boundary expands and we do not let any shock reflected particles reach the right boundary. To compare our results to downstream simulations, we Lorentz boost some quantities like the downstream spectra to the left-wall rest frame.

TRISTAN-MP uses a normalized unit system, like most particle-in-cell codes. Time is measured in units of the upstream inverse plasma frequency, $\omega_{pe}^{-1} = \sqrt{m_e/(4\pi n q^2)}$, where m_e is the electron mass, n is the upstream number density of electrons, and q is the charge of electrons. Distances are measured in units of the electron skin depth in the upstream plasma, c/ω_{pe}. We report lengths in units of the ion skin depth, $c/\omega_{upi} = c\omega_{pe}^{-1}\sqrt{m_i/m_e}$, and we measure times in terms of the upstream ion plasma frequency (ω_{pi}^{-1}). We use a reduced mass ratio, m_i/m_e, set to 64 in the fiducial run, easing computational expense of reaching late times in terms of $\omega_{pi}t$. We checked for convergence in mass ratio and show our results in Appendix A. We resolve the upstream electron skin depth with a different number of cells depending on the simulation. We use the criterion that we must resolve both the upstream Debye length and the downstream electron skin depth with at least one cell. The Debye length is $\lambda_D = c\omega_{pe}^{-1}\sqrt{\frac{m_i}{m_e}\Delta\gamma}$. In the fiducial run, this criterion is satisfied when resolving the upstream electron skin depth with four cells, with a time-step of $\Delta t = 0.1125\omega_{pe}^{-1}$. The low resolution is required to capture the relevant ion length and time-scales. We use four particles per cell (two ions and two electrons) and filter particles' contributions to the current 32 times per time-step to reduce noise, as is routinely done (e.g. Spitkovsky 2005). We checked convergence in the number of particles-per-cell (ppc) and spatial resolution, running smaller simulations with ppc $= 64$ and $c/\omega_{pe} = 8$ (see Appendix C).

We initialize an upstream magnetic field B_0 with inclination θ_B to the shock normal. A completely parallel shock has $\theta_B = 0$ and a completely perpendicular shock has $\theta_B = 90^\circ$. We use two different magnetic field orientations in this paper, one is quasi-parallel ($\theta_B \approx 10^\circ$), and the other is quasi-perpendicular ($\theta_B \approx 55.^\circ$). In the downstream rest frame the angle, θ_B', is 15° and 65°, respectively ($\gamma_0 \tan\theta_B = \tan\theta_B'$). We choose a magnetic field orientation such that perpendicular component of the initialized magnetic field lies entirely in the x–y plane, i.e. $B_{z,0} = 0$.

The strength of the magnetic field is described by σ_0, the ratio of the energy density of the upstream magnetic field to the kinetic energy of the upstream flow as measured in the downstream frame: $\sigma_0 = v_A'^2/(c^2[\gamma_0 - 1])$ where v_A' is the upstream Alfvén velocity as measured in the downstream frame. Re-writing everything in upstream quantities:

$$\sigma_0 = \frac{B_0^2(\gamma_0^2 \sin^2\theta_B + \cos^2\theta_B)}{4\pi\gamma_0 n_0 m_i c^2(\gamma_0 - 1)(1 + m_e/m_i)}, \quad (1)$$

In this paper, we fix $\sigma_0 = 0.007$. The magnetization fixes the Alfénic Mach number, which is the shock speed divided by upstream Alfvénic velocity. Using conservation of particle number $\beta_s = r\beta_0/(r-1)$, where r is the compression ratio of the shock, so the Alfvénic Mach number is

$$M_A \approx \frac{r}{r-1}\frac{\beta_0}{\sqrt{\sigma_0(\gamma_0 - 1)}}. \quad (2)$$

Since our shock is strong, the compression ratio of the shock, $r \sim 4\gamma_0^2$, in the upstream frame and for $\beta_0 = 0.75$, $\sigma_0 = 0.007$, we find $\beta_s \approx 0.83$ and $M_A \approx 15$.

Plasma beta, β_p, is the ratio between the thermal pressure and the magnetic pressure in the upstream plasma:

$$\beta_p = \frac{4c^2\Delta\gamma}{v_A^2\left(1 + \frac{m_e}{m_i}\right)}, \quad (3)$$

where v_A is the Alfvénic velocity in the upstream frame. Increasing the plasma beta for a fixed magnetization will decrease the sonic Mach number of the shock. We assume equilibrium between the magnetic pressure and thermal pressure of the upstream plasma, that is $\beta_p = 1$. For $\gamma_0 = 1.5$, $M_A = 15$, $m_i/m_e = 64$, and $\theta_B' = 15^\circ$, $\beta_p = 1$ corresponds to a temperature $\Delta\gamma = 1.27 \times 10^{-3}$ and sonic Mach number $M_s \approx 15$.

We measure three quantities that capture how much of the upstream kinetic energy goes into magnetic fields, non-thermal ions, and non-thermal electrons: ε_B, ε_p, and ε_e, respectively. To calculate ε_B, we Lorentz boost the total magnetic field $\boldsymbol{B}$ into a frame moving at γ_0 in the $+x$-direction and use

$$\varepsilon_B = \frac{|\boldsymbol{B}'|^2(1 + m_e/m_i)^{-1}}{4\pi\gamma_0 n_0 m_i c^2(\gamma_0 - 1)}. \quad (4)$$

Here, ε_B and σ_0 are the only field quantities we report in the downstream rest frame. We measure ε_p and ε_e by first Lorentz boosting the spectra into the downstream rest frame and then measuring the fraction of energy in the non-thermal particles. This ratio is converted to ε_p and ε_e by multiplying by the average energy

[1] Measuring the distances relative to the moving left wall causes the upstream to appear to be moving to the left with time even though its velocity is zero in the simulations.

Table 1. Summary of simulations. See Section 2 for more information.

B_0 inclination	B_0 orientation	m_i/m_e	L_y Number of cells	c/ω_{pe} Number of cells
Quasi-para 10.13°	In plane	64	16000	4
	In plane	64	800	4
	In plane	16	3500	8
	In plane	160	4000	4
	Out of plane	64	3220	4
Quasi-perp 55°	In plane	64	1600	4
	Out of plane	64	1600	4

Note: All simulations had $\gamma_0 = 1.5$, $\sigma = 0.007$, and $\beta_p = 1$ ($M_A \approx M_s \approx 15$).

carried by the respective particle downstream and dividing by the average kinetic energy of an incoming ion. Or more explicitly:

$$\varepsilon_{p,e} = \frac{\int_{E_{inj}}^{\infty} E' \frac{dn'}{dE'} dE'}{\int_0^{\infty} E' \frac{dn'}{dE'} dE'} \times \frac{m_{i,e}(\langle \gamma'_{i,e} \rangle - 1)}{m_i(\gamma_0 - 1)}. \quad (5)$$

The transverse size of our fiducial run is 16 000 cells, equivalent to 4000 electron skin depths or 500 ion skin depths for the reduced mass ratio 64, significantly wider than previous PIC shock simulations (e.g. Sironi & Spitkovsky 2009, 2011; Niemiec et al. 2012; Sironi et al. 2013; Park et al. 2015). For example, while Bohdan et al. (2017) also studied electron acceleration in high Mach number, mildly relativistic shocks, our simulations are ∼20 times wider in terms of ion skin depths and were run three times longer in terms of ion gyrofrequency (even when accounting for the different rest frames). The sizes of our simulations are more comparable to hybrid-MHD (kinetic ions-fluid electrons) simulations. We require a large size to resolve the gyroradii of the non-thermal ions. We summarize the simulations run in this work in Table 1.

3 STRUCTURE OF A BELL-MEDIATED SHOCK

In this section, we discuss the structure of a quasi-parallel, electron–ion, magnetized, mildly relativistic shock. With our large simulations, we are able to capture the large-scale structure of magnetic turbulence upstream. For the first time in a PIC simulation, we use a large enough box to resolve the gyroradius of the highest energy ion as the shock transitions between the early times – when the Weibel instability mediates the shock – to late times – when Bell instability is the dominant instability.

The evolution of density and B_z with time is shown in Fig. 2 for a small slice of our simulation. The shock forms because of the Weibel instability, which starts the Fermi process. Once they have escaped upstream, particles are scattered back to the shock front due to small angle scattering, accelerating protons and electrons (Spitkovsky 2008b; Sironi & Spitkovsky 2011). Sufficiently high-energy particles are not efficiently scattered by Weibel instability, and are able to escape far upstream. There is a net positive current $\boldsymbol{J}$ carried by non-thermal ions that non-resonantly grows transverse $\boldsymbol{B}$ waves, and the shock transitions from being Weibel-mediated to Bell-mediated.

The growth rate of the fastest growing wavelength of the Bell instability depends only on the current carried by non-thermal ions upstream, which are travelling at roughly the shock speed β_s. In the upstream frame, the growth time of the Bell waves and the wavenumber of the fastest growing mode is given by equations (16)

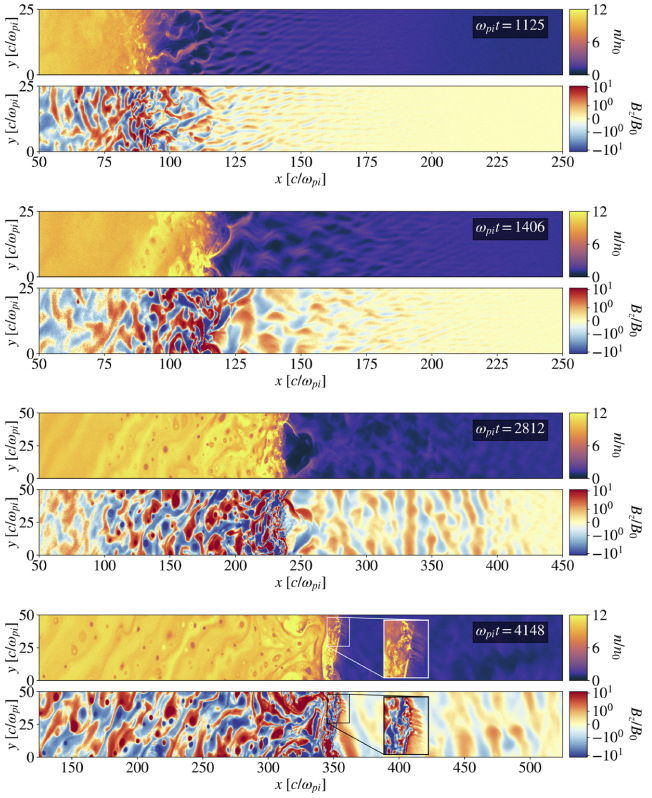

Figure 2. The transition from a Weibel-mediated shock to Bell-mediated shock. A small section of the simulation around the shock is followed with time. Initially the upstream filaments are on the order of a few ion skin depths with similar-sized magnetic perturbations. As the shock becomes Bell-mediated, cavities are evacuated upstream and advected downstream. The presence of non-thermal ions at the shock front increases the compression ratio by ∼10 per cent to $\sim 4.6\gamma_0^2$, in good agreement with hybrid simulations of quasi-parallel shocks (Caprioli & Spitkovsky 2014a). The magnetic field amplification is significantly larger at the walls of the cavities, while the centre of the cavities has a small magnetic field. The size of the cavities grows with time. At $\omega_{pi}t = 4148$, there are no cavities close to the shock front in this section of the simulation (one is forming around $x \sim 475c/\omega_{pi}$), and the magnetic amplification is large enough such that the shock is locally quasi-perpendicular with a small Alfvénic Mach number. Whistler waves can be observed at the shock (see inset). Whistler waves may be an effective preheating source for injection of electrons (Riquelme & Spitkovsky 2011).

and (17) of Reville, Kirk & Duffy (2006):

$$\omega_{pi} t_{Bell} \approx \frac{2}{\xi_{cr}\beta_{cr}}; \quad \frac{k_{max}c}{\omega_{pi}} = \frac{1}{2}\xi_{cr}\frac{v_{cr}}{v_{\mathcal{A}}}. \quad (6)$$

ξ_{cr} is the number fraction of non-thermal ions in the upstream, and v_{cr}, β_{cr} are the drift velocity of the current carrying ions. Following Caprioli & Spitkovsky (2014b), we measure the drift velocity of the returning ions and find it is equal to the shock speed, $\beta_s \approx 0.83$. We also measure the number density of returning ions and find $\xi_{cr} \approx 0.029$. For our $\mathcal{M}_{\mathcal{A}} \sim 15$ shock, we estimate the fastest growing mode to have a wave number of $\sim 0.22\omega_{pi}/c$, which is in good agreement with the peak of the Fourier power spectrum plotted in Fig. 3. The growth time of the fastest growing mode is $\sim 80/\omega_{pi}$. Transverse waves start to appear around $\omega_{pi}t \sim 560$, and the shock completely transitions to being Bell-mediated around $\omega_{pi}t \sim 1500$ or ∼20 times the growth time of the Bell instability. The Bell modes are visible as a striped B_z and B_y pattern upstream, transverse to the shock normal (see Fig. 2). We measured the polarity of the

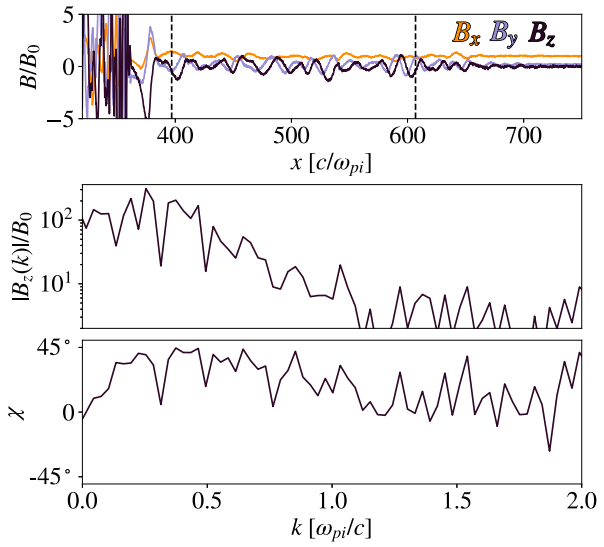

Figure 3. The polarization of the upstream waves at time $\omega_{\rm pi}t = 4148$. The top panel shows a 1D slice ($y = 250c/\omega_{\rm pi}$) of the magnetic field. Vertical lines mark the region where the Fourier transform of the fields is calculated. The middle panel shows the Fourier transform of B_z, and the bottom panel is a graph of Stokes polarization angle χ (measured following Park et al. 2015). Non-resonant Bell waves are right-handed circularly polarized ($\chi = 45^\circ$), as are the waves in the upstream of the simulation. The power in the wave peaks at $kc \approx 0.25\omega_{\rm pi}$, in good agreement with the fastest growing mode of the non-resonant Bell instability.

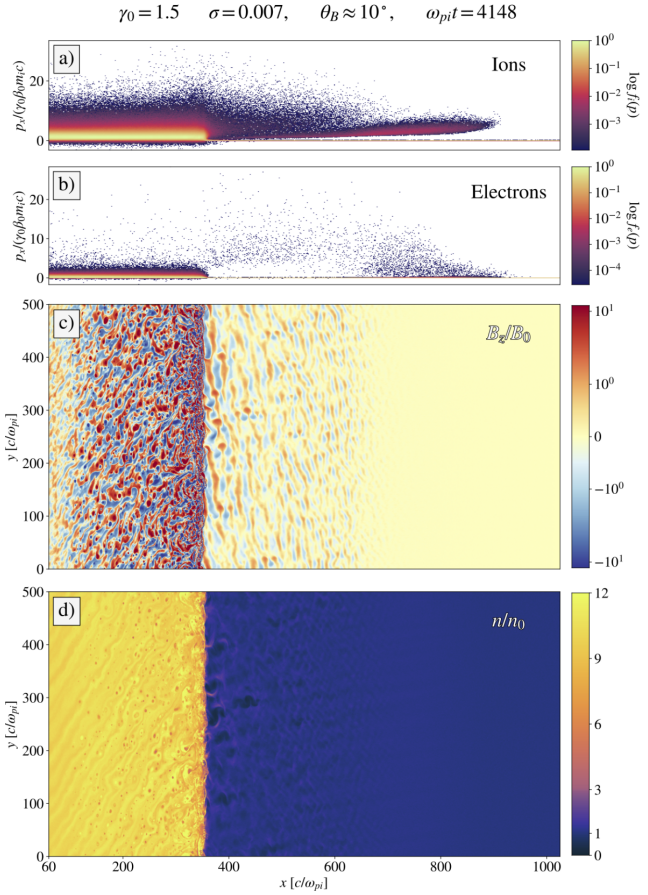

Figure 4. The fiducial run at time $\omega_{\rm pi}t = 4148$. Distances are measured in ion skin depths. The top two graphs are $x - p_x$ phase diagrams for the ions and electrons, respectively. The bottom two graphs show B_z and the density domain. The upstream location where Bell instability starts to grow ($x \sim 650c/\omega_{\rm pi}$) is visible in all four subplots. In the ion phase diagram, it is where the ions transition from a narrow beam in phase space to a more diffuse, scattered structure. In the electron phase diagram, Bell is marked by the dearth of reflected electrons with energies below $\sim\gamma_0^2 m_i c^2$ between $350c/\omega_{\rm pi}$ and $650c/\omega_{\rm pi}$. These two features are aligned with the transverse, circularly polarized magnetic field amplification seen in the B_z plot, and the corresponding evacuated cavities in the density.

upstream waves and found they are right hand circularly polarized, as expected for non-resonant ion waves (see Fig. 3).

The ion current $\boldsymbol{J}$ also drives a filamentary mode. Our simulation shows a similar structure to the one found in hybrid simulation of non-relativistic quasi-parallel shocks (see Caprioli & Spitkovsky 2013). To keep the plasma quasi-neutral, there is a returning negative current in the background plasma that is aligned with B_0. Therefore, the background plasma feels a $-\boldsymbol{J} \times \delta\boldsymbol{B}$ force, which pushes the background plasma away from the regions of strongest current and focuses the current carrying ions, growing the instability further. The size of these upstream cavities grows with time, and they contain a large magnetic field at their walls. The transverse size is limited to be less than the size of gyroradius of the current-carrying ions.

We show the entire simulation domain of our fiducial, magnetized quasi-parallel shock at the final time-step $\omega_{\rm pi}t = 4148$ in Fig. 4. The upper two panels show the ion and electron $x - p_x$ phase diagrams, and lower two panels show B_z and the density. The upstream location where the returning ions start to non-resonantly drive circularly polarized transverse magnetic field fluctuations, $x \sim 650c/\omega_{\rm pi}$, is also visible in structural changes in the top two panels of the phase diagrams. The returning ions transition from being a relatively narrow beam of particles to having a more diffuse structure as a result of the efficient scattering by the magnetic field. In the electron phase diagram the transition is marked by a minimum momentum required to participate in the particle acceleration. The minimum momentum is visible as a lack of non-thermal electrons with Lorentz factors $\lesssim\gamma_0^2 m_{\rm i}/m_{\rm e}$ between $350c/\omega_{\rm pi}$ and $650c/\omega_{\rm pi}$. All of the non-thermal electrons have energies comparable to the non-thermal ions. The upstream magnetic field amplification is accompanied by the evacuation of small cavities in the density far upstream that grow to larger worm-like cavities closer to the shock front. The amplified magnetic fields are largest at the walls of these cavities. Downstream, the cavities become long filaments with $\varepsilon_B \gtrsim 0.2$ and small hole of low densities and large B_z, $\varepsilon_B \gtrsim 1$.

When swept downstream, the cavity forms two underdense holes that contain large magnetic fields of opposite B_z polarity, and are surrounded by a ring of current. A close-up of a cavity being advected downstream can be seen in Fig. 5. In 3D, the two holes would likely be connected by a flux tube. Some holes merge downstream, and in 3D, the disruption of possible flux tubes may be an important site of magnetic dissipation. Whether the flux tubes do indeed form in 3D simulation, and the importance of any magnetic dissipation during their lifetimes is left to future work.

At the latest times, whistler waves are observed in the ion foot of the shock. Whistler waves are small-scale, dispersive waves that can occur in quasi-perpendicular shocks with low Alfvénic Mach numbers. Whistler waves can grow in non-relativistic quasi-perpendicular shocks when $M_{\mathcal{A}} \lesssim \sqrt{m_i/m_e}$ (e.g. Krasnoselskikh et al. 2002; Riquelme & Spitkovsky 2011). This inequality is satisfied in our simulation when $\delta B_\perp/B_0 \gtrsim 2$, which is indeed the case in the regions where whistler waves appear in our simulation

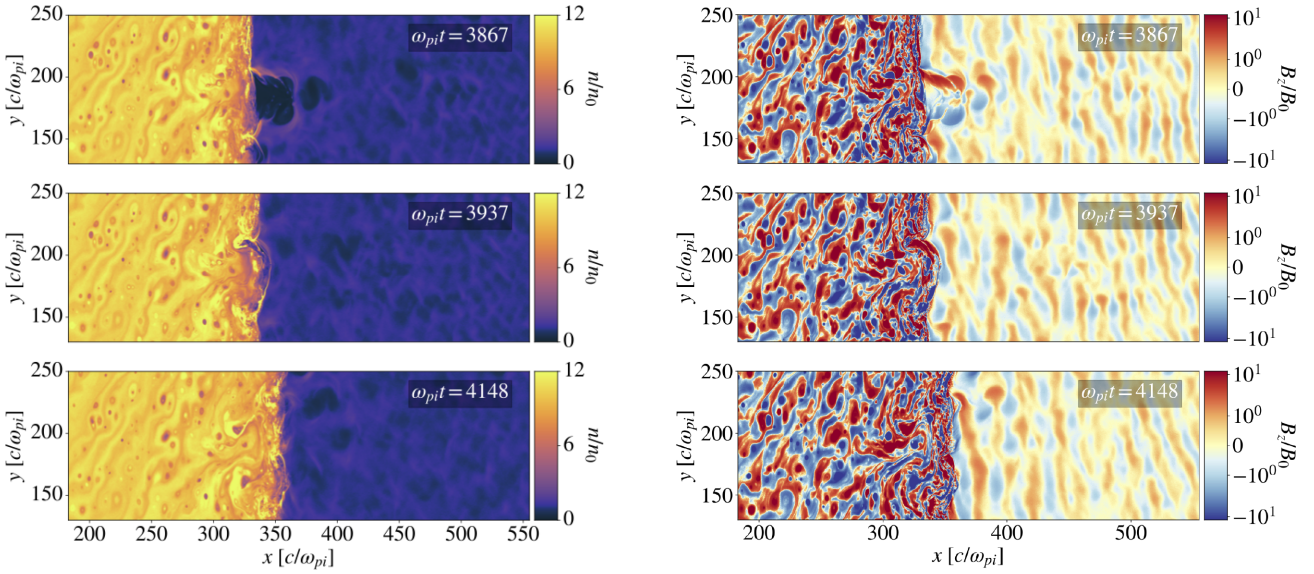

Figure 5. Advection of an evacuated upstream cavity into downstream. The left panels show the density and the right show the out-of-plane magnetic field B_z. Once advected downstream, the upstream cavities forms two holes with opposite B_z sign filled with a large amount of magnetic field.

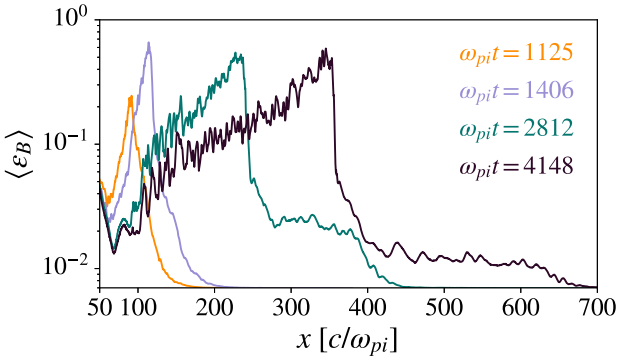

Figure 6. The evolution of the y-averaged ε_B with time in the fiducial run. The times correspond to the times shown in Fig. 2. After the shock becomes mediated by the Bell instability ($\omega_{\rm pi}t \sim 1500$), the magnetic field amplification at the shock front increases. In addition, the amplified magnetic field increases further upstream and downstream with time. *Note* that the increase of ε_B at $x \lesssim 60\,\omega_{\rm pi}$ is an artefact from the left boundary that initializes the shock.

(see the inset of bottom panels of Fig. 2). Whistler waves are an efficient way to transfer energy from ions to electrons (Riquelme & Spitkovsky 2011). For realistic mass ratios, the condition for the shock to be unstable to whistlers is more easily achieved, and the whistlers may play a role in helping the electrons reach a sizable fraction of the ion temperature. However, we do not see good evidence that this increase in whistler waves increases electron injection into DSA. As increasing the mass ratio does not increase $\varepsilon_{\rm e}$ (see Appendix A).

The change in magnetic field amplification with time is best seen in Fig. 6. After the shock becomes Bell-mediated, at $t \sim 1500\omega_{\rm pi}$, there is significantly larger magnetic field amplification both upstream and downstream compared to earlier times when the shock was Weibel-mediated. The extent of the amplification both upstream and downstream grows with time. The average value of magnetic field amplification saturates at $\sim 2B_0$ upstream with a peak value of $\gtrsim 5B_0$ inside of the filaments close to the shock front. We note that ε_B upstream decreases by ~33 per cent from $\omega_{\rm pi} = 2812$ and $\omega_{\rm pi}t = 4148$. This decrease may be because the initial beam of ions that triggers the Bell instability is more anisotropic than the later returning ions that are scattered more efficiently. The change from a beam to a diffuse structure in phase space is clearly visible in the ion phase diagram in Fig. 4. At later times, the importance of the organized beam will decrease and the returning particles may be more diffuse.

4 PARTICLE ACCELERATION IN A BELL-MEDIATED SHOCK

In the previous section, we showed that the transition from a Weibel-mediated shock to Bell-mediated shock changes the upstream magnetic field turbulence and shock geometry. In this section, we examine the implications of the Bell turbulence on particle acceleration.

The downstream spectral evolution with time is shown in Fig. 7. Both the ions and electrons show a Maxwellian distribution with significant non-thermal populations. In both the ion and electron spectra, the maximum energy grows with time. The electrons equilibrate with a temperature ≈0.23 times the ion temperature. The power law of electron distribution is consistent with $f(p) \propto p^{-4.2}$ ($\mathrm{d}n/\mathrm{d}E \propto E^{-2.2}$). A $p^{-4.2}$ spectrum is predicted for DSA in ultra-relativistic shocks with $\gamma_0 \gg 1$, while p^{-4} is predicted for non-relativistic shocks (e.g. Achterberg et al. 2001). In the ions, the downstream spectral index appears softer than the electrons, but the spectral index is not well constrained. A proper measurement of the ion spectrum would require a simulation that was run significantly longer with a larger extent of the non-thermal ion power law. In our simulation, $\varepsilon_{\rm p}$, the fraction the shock's energy that is put into non-thermal ions is ~0.1. The downstream electron temperature is 30 per cent of the ion temperature and $\varepsilon_{\rm e} \sim 5 \times 10^{-4}$. To measure $\varepsilon_{\rm e}$ and $\varepsilon_{\rm p}$, we integrate the non-thermal spectrum starting at a momentum $p_{\rm inj} = 2\gamma_0\beta_0 m_i c$ as measured in the downstream frame.

The upstream spectral evolution with time is shown in Fig. 8. Both the ions and electrons show two-humped spectra. The lower

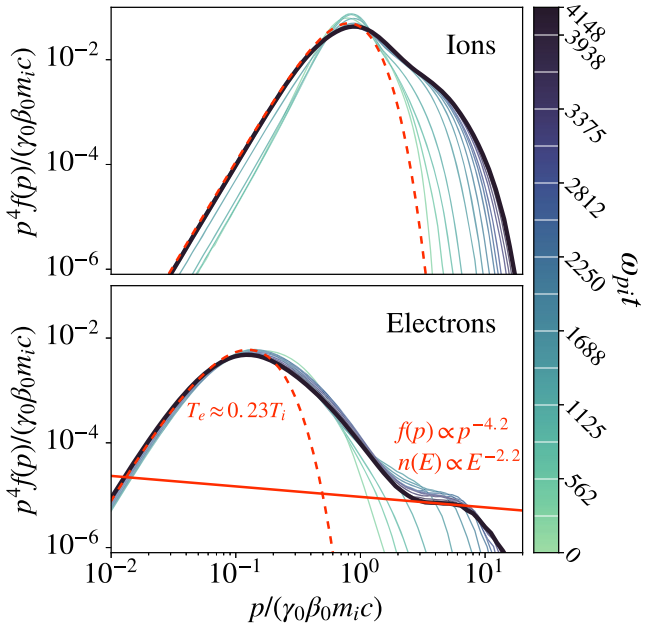

Figure 7. Time evolution of the downstream spectrum for the fiducial run. The spectra are boosted into the downstream rest frame and extracted from an area $10-80c/\omega_{\rm pi}$ downstream of the shock. When the shock had not yet travelled far enough to the right, we take the leftmost section of the downstream, between max $(0, x_{\rm s} - 200)$ and max $(5, x_{\rm s} - 100)$. The white lines in the colour bar mark the times when spectra are plotted, including the final tick mark. The ions have a temperature $\sim 0.15 m_i c^2$, and the electrons equilibrate at a temperature $\sim$23 per cent of the ions.

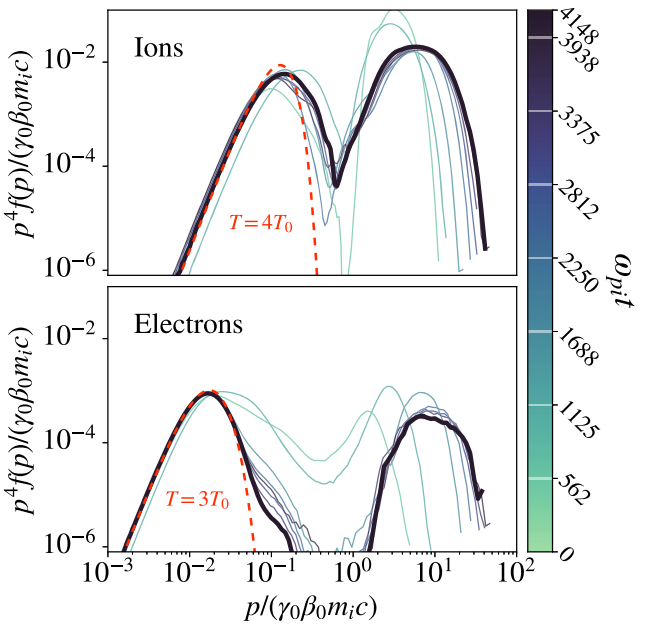

Figure 8. Time evolution of the upstream spectrum for the fiducial run. The spectra shown in the simulation (i.e. upstream) rest frame and extracted from an area $10-20c/\omega_{pi}$ upstream of the shock. Both the ions and electrons show two-humped spectra. The left hump is the incoming beam of particles, while the right hump is the returning particles. In the spectra we show Maxwellian distributions fit by eye, with temperatures in terms of T_0, the injected temperature of the plasma. When the shock is Weibel-mediated, the electrons are pre-heated significantly, which can be seen because the left hump is not fit well by a Maxwellian at early times. After Bell instability kicks in at around $t \sim 1500\omega_{pi}$, both the ions and electrons are preheated only by a factor of around a few.

energy bump corresponds to the incoming beam of particles, while the higher energy bump is the returning particles. The ions rapidly converge to having a minor amount of pre-heating, as lower bump is well fit by a Maxwellian. In contrast, the electrons show strong heating at times $\omega_{\rm pi}t = 562$ and 1125, which we attribute to the Weibel filaments that transfer a large fraction of the ions' kinetic energy to the electrons upstream, causing a significant amount of preheating of the electrons, increasing the electron gyroradius (e.g. Kumar, Eichler & Gedalin 2015). The initial Weibel filaments are effective at injecting electrons not only due to this pre-heating but also because the perturbations cause the shock to be highly random in magnetic field obliquity, providing locally quasi-perpendicular field orientation that is conducive to electron injection via shock drift acceleration (SDA; Ball & Melrose 2001; Park et al. 2015). At later times, after the shock switches to being Bell-mediated, the upstream magnetic fluctuations grow to larger size scales, and the energy transfer between reflected ions and incoming electrons decreases tremendously. This can be seen by the fact that the lower energy electron hump in Fig. 8 becomes increasingly better fit by a simple Maxwellian.

The two species enter the shock front in approximate thermal equilibrium. Since the momenta of the electrons entering the shock are far less than that of the ions, the scale separation between incoming electrons and ions is preserved. There is a minimum energy electrons must attain to escape far upstream, $E_{\rm min} \sim \gamma_0^2 m_i c^2$, visible in the hole in the electron phase diagram (see Fig. 4), and in the time evolution of the upstream spectra, see Fig. 8. All of the reflected electrons have energies comparable to the reflected ions. As we argue in Section 5, the Bell instability creates regions of the shock that are locally superluminal. The size of these regions is comparable to the ion gyroradius. Electrons need to attain energies comparable to the ion gyroradius to escape these regions. We leave a detailed study of the orbits of the non-thermal electrons to future work. However, before electrons escape far upstream, they undergo several cycles of SDA. Once they reach the regions of quasi-parallel fields, they can escape upstream. This hints that the acceleration of electrons in our simulation is likely done through a combination of SDA and then DSA as seen in the simulations of Park et al. (2015).

For both ions and electrons, the highest energy particle in the box for each species has roughly the same kinetic energy. The two species are expected to have the same maximum energy if the rate of energy gain depends only on the rigidity of the particle and not the charge, like in DSA. The maximum energy grows linearly with time and does not saturate, $E_{\rm max} \propto t$ (see Fig. 9). The rate at which the maximum energy particle gains energy is consistent with the Bohm limit in standard DSA theory ($E \propto t$, $D \propto E$, where D is the diffusion coefficient). The linear scaling is much faster than $E_{\rm max} \propto \sqrt{t}$ scaling found from small-angle scattering in relativistic Weibel-mediated shocks (Plotnikov, Pelletier & Lemoine 2011; Sironi et al. 2013), and makes a large difference the potential maximum energy particle produced by an astrophysical system.

5 IMPORTANCE OF THE FILAMENTARY MODE

We find that to properly recover the correct shock structure, and measure $\varepsilon_{\rm e}$, the simulation domain must be wide enough so that the transverse waves cannot fill the entire box. This criterion is satisfied when the size of the box is much larger than the gyroradius of the

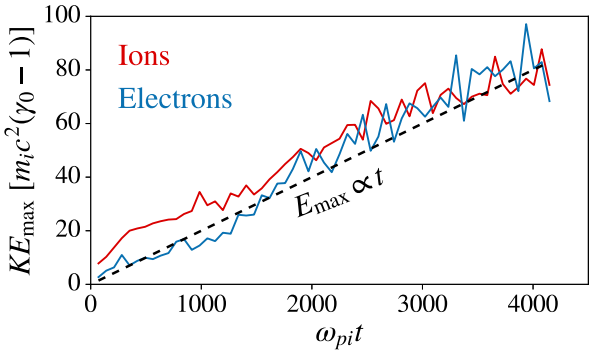

Figure 9. The energy of the highest energy ion and electron in the simulation is plotted as a function of time. After the shock becomes Bell-mediated (around $\omega_{pi}t \sim 1500$), the maximum energy of both the electrons and ions is comparable to each other and increases linearly in time with no sign of saturation, consistent with the Bohm scaling.

ions carrying the majority of the current upstream. When using a transverse size that is too small, the Bell modes can grow until they fill the transverse size. When this happens, ε_e is suppressed. The characteristic angle of the amplified magnetic field, $\theta_B \equiv \tan^{-1}|B_\perp/B_\parallel|$, is $\gtrsim 45^\circ$. For $\beta_s = 0.83$, the angle at which the shock becomes superluminal is $\theta_{crit} = \cos^{-1}\beta_s \approx 34^\circ$, where θ is measured in the upstream frame (Begelman & Kirk 1990; Sironi & Spitkovsky 2011). The existence of locally quasi-parallel regions of the shock front allow electrons to escape back upstream.

The size needed to properly resolve Bell cosmic ray current driven instability depends on the size of the evacuated cavities upstream. The cavities grow until they are advected towards the shock. The size of the cavities when they impact the shock front depends on the ratio of the advection time to the growth rate of the cavities ($\Gamma \propto |B_\perp|\xi_{acc}^{1/2}$, where ξ_{acc} is the pressure of the accelerated ions, Reville & Bell 2012; Caprioli & Spitkovsky 2013). The maximum size is ultimately limited by the gyroradii of the ions carrying the current. The gyroradius of a typical shock reflected ion ($E \sim \gamma_0^2 m_i c^2$) is $r_{gyr} \sim 30\, c/\omega_{pi}$. Caprioli & Spitkovsky (2013) found that in hybrid simulations of non-relativistic shocks, the size of the filament can grow to be comparable to the gyroradius of the highest energy ion.

To show the importance of the transverse size of simulation for the measurement of ε_e, we compare our fiducial run that is much wider than the typical ion gyroradius to a simulation with a width slightly smaller than the ion gyroradius in Fig. 10. At early times, the smaller, $L_y = 25\, c/\omega_{pi}$, run agrees with the fiducial, $L_y = 500\, c/\omega_{pi}$, run in terms of magnetic structure (panel a) and particle acceleration (panel c). However, as we can see in panel b, failing to resolve the ion gyroradii significantly changes the shock structure, where the fiducial run shows filaments with width $\sim r_{gyr}$, but the smaller run does not have filaments at the shock front. The filaments in the fiducial run cause the shock front to have regions of subluminal and superluminal magnetic field orientations at the same time. Without filaments, as in the smaller run, alternating strips of subluminal and superluminal fields fill the transverse direction upstream. This difference in structure changes ε_e (see the panel c of Fig. 10). The box must be sufficiently wide to capture the filaments to properly

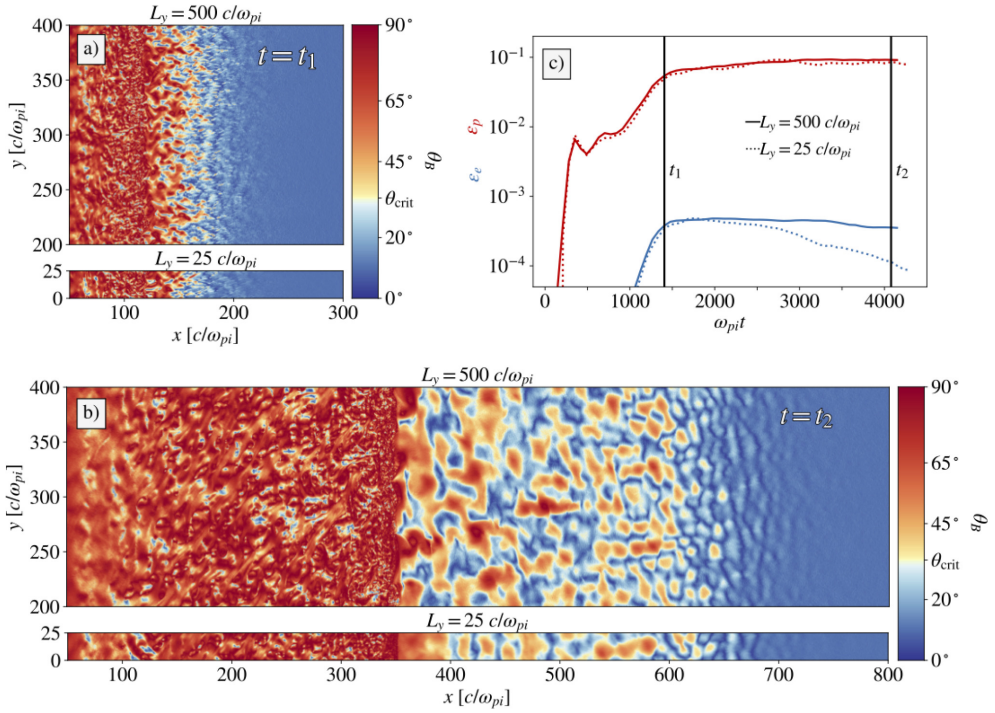

Figure 10. Magnetic field structure in the fiducial simulation (width $L_y = 500\, c/\omega_{pi}$) is compared to a simulation with a smaller transverse size ($L_y = 25\, c/\omega_{pi}$). We show the angle of the magnetic field with the shock normal, $\theta_B \equiv \tan^{-1}|B_\perp/B_\parallel|$ for two simulations at two times: $t_1 = 1400/\omega_{pi}$, when the shock begins to become Bell-mediated, and $t_2 = 4078/\omega_{pi}$. The angle at which the shock transitions between subluminal and superluminal, θ_{crit}, is marked on the colourbar. The larger simulation is shown as the upper panel in (a) and (b), and is plotted with the solid lines in (c). The smaller $L_y = 25\, c/\omega_{pi}$ simulation is shown in the lower panels of (a) and (b) in the 2D plots and dotted lines in (c). At earlier times, the two simulations agree. However, at late times, the shock front has patches of superluminal and subluminal magnetic field orientations in the fiducial run, while in the run with a smaller L_y, the shock is superluminal across the entire transverse direction. The effect of the additional magnetic turbulence on particle acceleration is shown in panel (c). When the transverse waves fill the box in the $L_y = 25\, c/\omega_{pi}$ run, ε_e decreases.

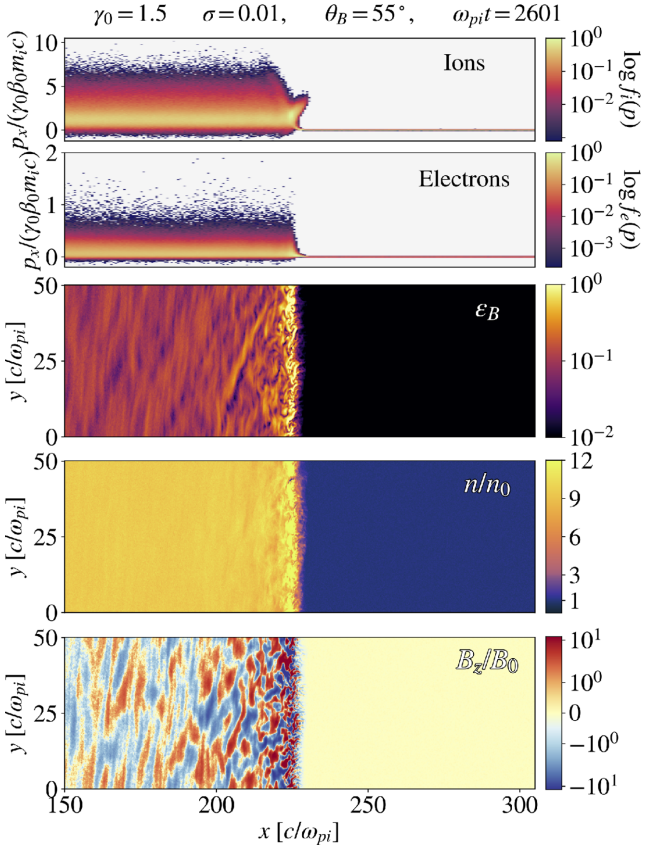

Figure 11. Structure of a superluminal mildly relativistic shock with $\theta_B = 55°$. Compare this figure to the subluminal case with $\theta_B \approx 10°$ in Fig. 4. There are no particles escaping upstream and no upstream turbulence, as expected for superluminal magnetized shocks.

measure the energy in non-thermal electrons in mildly relativistic shocks.

Electron injection into shock acceleration decreases if the transverse size of the waves can grow to the entire box. Decreasing the box size decreases the normalization of the electron power law, but not the extent of the power law in ions nor electrons. We confirmed that the decrease in ε_e is not just a matter of statistics by running small runs with larger ppc (up to 64). We also compared our fiducial, $L_y = 500\, c/\omega_{pi}$, run to slightly smaller run that still captured the gyroradius, $L_y = 200\, c/\omega_{pi}$. We saw no difference in ε_e in those runs, suggesting it is sufficient to capture a few patches of a locally quasi-parallel shocks to properly measure ε_e.

6 QUASI-PERPENDICULAR MILDLY RELATIVISTIC SHOCKS

Here, we change the magnetic field orientation to examine a quasi-perpendicular mildly relativistic shock with $\theta_B = 55°$. For $\gamma_0 = 1.5$, the angle at which the shock becomes superluminal is $\theta_{crit} \approx 34°$. Therefore with $\theta_B = 55°$, particles are unable to escape upstream and drive turbulence far from the shock. The inability of particles to escape far upstream prevents them from being injected into DSA and forming a non-thermal power law.

In Fig. 11, we show an example of a superluminal, quasi-perpendicular shock. Unlike the quasi-parallel shock, there are no reflected ions capable of driving turbulence upstream. The lack of returning ions results in a smaller compression ratio, and

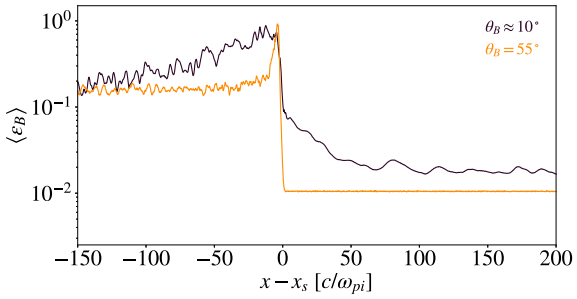

Figure 12. A comparison of the magnetic field amplification in mildly relativistic quasi-parallel and quasi-perpendicular shock at the final time-step of each respective simulation. Returning ions amplify the magnetic field upstream in the quasi-parallel shock. The additional turbulence produced by the returning ions also increases the downstream magnetic field beyond shock compression of the unamplified field in the quasi-parallel shock.

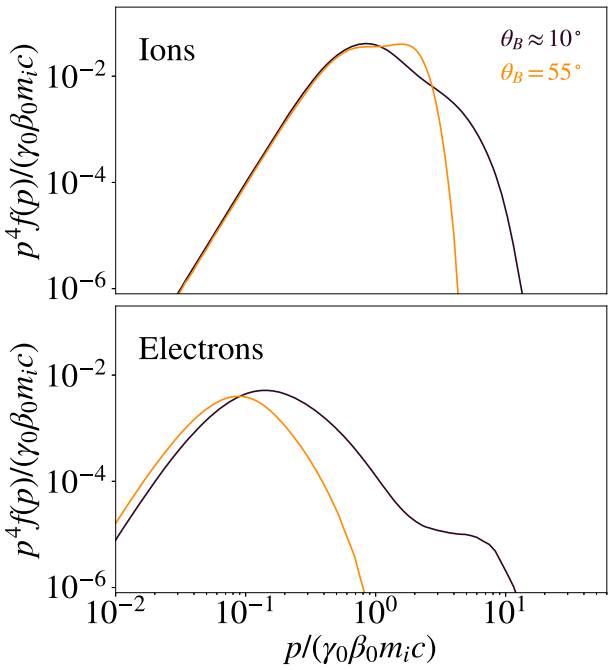

Figure 13. A comparison of the downstream spectra of a superluminal shock with $\theta_B = 55°$ and the subluminal shock with $\theta_B \approx 10°$ at time $\omega_{pi}t \approx 2600$ measured 10–80c/ω_{pi} downstream from the shock. In superluminal shocks, neither electrons nor ions show non-thermal power laws, but the ion spectrum does show a bump from shock-drift acceleration.

in weaker magnetic fields downstream, in good agreement with hybrid simulations of quasi-parallel shocks (Caprioli & Spitkovsky 2014b). We show the averaged 1D profiles of the magnetic field amplification in Fig. 12.

In the quasi-perpendicular shock, a prominent suprathermal bump from SDA is visible in the downstream ion spectra shown in Fig. 13. However, the maximum energy in both species does not grow with time or develop into a power law. The lack of returning upstream particles means that there are no particles to seed magnetic turbulence upstream. Therefore, there is no pre-heating upstream: $T_p = T_e = T_0$. Downstream, $T_p \approx 6T_e$. The ion temperature is nearly the same in the subluminal and superluminal shocks, but T_e of the superluminal shock is about half of T_e in the quasi-parallel case. The additional turbulence provided by the shock-reflected particles

in quasi-parallel shocks drives the downstream species closer to equilibrium.

As we show in Appendix B, when $\theta_B \sim 10^\circ$, the shock structure and particle acceleration do not depend on the orientation of the magnetic field with respect to the simulation plane. However in quasi-perpendicular shocks, the turbulence right at the shock front and downstream depends strongly on the orientation of the field with respect to the simulation plane. In the out-of-plane configuration, ions gyrate around B_z in the $x-y$ plane, reducing the effective adiabatic index, causing downstream magnetic fluctuations, and decreasing the importance of SDA in the ions. But for both in-plane and out-of-plane superluminal magnetic field configurations, particles do not return upstream and a non-thermal power law does not develop.

7 CONCLUSIONS

This paper is the first in a series that examines how particle acceleration in shocks changes when the shock transitions from non-relativistic to relativistic. We used large computational domains and a fully kinetic method. We captured the filamentary behaviour of the Bell instability. In doing so, we have followed the shock from formation to the shock's later stages, resolving all kinetic ion and electron scales. In this paper, we have shown that properly capturing the non-resonant instability is crucial for measuring the fraction of the shock's energy that goes into non-thermal electrons, ε_e.

There is observational and theoretical evidence that both relativistic and non-relativistic quasi-parallel shocks are efficient ion accelerators with a $\varepsilon_p \sim 0.1$. We find that the fraction of the shock's energy in non-thermal electrons, ε_e, to be $\sim 5 \times 10^{-4}$ for a mildly relativistic, magnetized shock with $\theta_B = 10^\circ$, $\gamma_0 = 1.5$, $\beta_0 \approx 0.75$, and $M_{\mathcal{A}} \approx 15$. Our measurement of ε_e can be directly compared to other PIC simulations in the literature of shocks with $\sigma \sim 0.01$ or Alfvénic Mach numbers $\sim$10. Our reported ε_e is consistent with than the 1D measurement of $\varepsilon_e \sim$ a few $\times 10^{-4}$ for non-relativistic shocks with $\beta_0 = 0.1$ by Park et al. (2015), and less than the $\varepsilon_e \sim 0.1$ measured in relativistic shocks by Sironi & Spitkovsky (2011). Our simulations are also consistent with the trend of increasing ε_e as the shock becomes more relativistic, but a set of directly comparable simulations spanning the non-relativistic to relativistic regime is the goal of a future work.

Our simulations had a moderately large magnetization $\sigma_0 = 0.007$, $M_{\mathcal{A}} \approx M_s \approx 15$: a set-up consistent with possible internal shocks in the jets of microquasars or quasars. We show that magnetized, subluminal, mildly relativistic shocks are capable of producing non-thermal ions and electrons in hard power-laws whose maximum energy grows linearly with time.

The evolution of a quasi-parallel weakly magnetized shock can be summarized as follows: the shock forms via the Weibel instability that starts the Fermi process and reflects particles into the upstream. The Weibel instability efficiently pre-heats incoming electrons to an energy comparable to the ions. Therefore, both electrons and ions are efficiently injected in a Weibel-mediated shock. However, the Weibel instability produces small-scale field fluctuations that only efficiently scatter particles back towards the shock if the gyroradius of the particle is comparable to the size of the fluctuations. A typical filament with a transverse size of 10 ion skin depths will not large-angle-scatter particles who have energies much larger than the shock reflected ions (whose gyroradii are $\sim$30 ion skin depths). High-energy particles are only deflected by a small angles, and can easily stream far away from the shock due to the quasi-parallel field geometry. The escaping high-energy ions drive a current in the upstream plasma, which grows transverse waves via the Bell cosmic ray filamentary instability. After about $\sim$20 growth times, the Bell instability becomes the dominant instability that mediates the shock. Once mediated by the Bell filamentary instability, the shock and non-thermal particles change in the following ways:

(i) For $\mathcal{M}_{\mathcal{A}} = 15$, there is large magnetic field amplification both upstream and downstream. The y-averaged field amplification upstream is as $\sim 2B_0$, and close to the shock the most amplified field can be $\gtrsim 5B_0$. The largest upstream field amplification occurs at the walls of worm-like cavities evacuated by energetic ions. A $\boldsymbol{J} \times \delta \boldsymbol{B}$ force acts on a current carried by the background plasma, which is balancing the current of high-energy ions that escape far upstream. The force evacuates these worm-like cavities. The transverse size and extent of these structures grows with time. Downstream, ε_B, is both larger and extends further away from the shock than when the shock is Weibel-mediated (see Figs 6 and A4).

(ii) The shock structure is highly turbulent and magnetic field amplification downstream happens in two structures: magnetized filaments with a characteristic $\varepsilon_B \gtrsim 0.1$ and compact underdense holes with a large out-of-plane magnetic field and an $\varepsilon_B \sim 1$, confirming the turbulent structure found by Caprioli & Spitkovsky (2013) using hybrid simulations. When an upstream cavity is advected downstream, it forms two underdense holes of opposite B_z sign. In 3D they would likely form a connected twisted flux tube. The holes last for a long time and even merge in 2D. They may be an important place for magnetic dissipation, but the dynamics needs to be checked in 3D.

(iii) The filamentary nature of the Bell instability plays a vital role in electron injection in mildly relativistic shocks. At early times, the Weibel instability efficiently reflects both ions and electrons at the shock front, as in relativistic shocks (e.g. Sironi et al. 2013). After the shock becomes Bell-mediated, $\omega_{pi} t \sim 1500$ in our fiducial run, electron reflection and injection into DSA becomes less effective. If the simulation box is not wide enough to capture the gyroradius of the non-thermal ions, ε_e will be suppressed significantly in shocks with $\gamma\beta \sim 1$.

(iv) The maximum energy of both the ions and electrons in the quasi-parallel simulation increases linearly after the shock becomes Bell-mediated. The maximum ion and electron energy is the same. The linear increase in energy is consistent with the Bohm scaling, $E_{max} \propto t$, and is much faster than the scaling of $E_{max} \propto t^{1/2}$ in Weibel-mediated, high-Mach number, ultra-relativistic shocks (Sironi et al. 2013). Because E_{max} grows much faster in mildly relativistic shocks, sources with mildly relativistic shocks like FRI type AGN may be more likely progenitors of ultra-high energy cosmic rays than the ultra-relativistic unmagnetized shocks in GRBs.

(v) Electrons are not pre-heated to a large fraction of the reflected ion energy in trans-relativistic shocks. In relativistic and mildly relativistic Weibel-mediated shocks, the reflected ions are capable of transferring a sizeable portion of their energy to the electrons upstream, so that the rigidity of the two species is similar at the shock front. We find that once the shock becomes Bell-mediated, ions and electrons are pre-heated by similar amounts; at one ion gyroradius in front of the shock, the two temperatures are $T_e \sim 3T_0$ and $T_p \sim 4T_0$, where T_0 is the initialized temperature of the plasma. Therefore, the two species enter the shock region in approximate thermal equilibrium, with thermal momenta far less than the bulk momentum of an ion. Thus, the large difference between ion and electron rigidity at the shock front is preserved. Downstream and right at the shock front, the ions do transfer their energy

to the electrons, and the downstream electrons equilibrate with a temperature ∼30 per cent of the ions' downstream temperature.

(vi) To be reflected upstream, the electrons must have an energy comparable to a shock-reflected ion, meaning electrons must gain a lot of energy close to the shock front, possibly through an SDA process as seen in 1D simulations of Bell-mediated electron–ion shocks (Park et al. 2015), but we did not study the orbits of the non-thermal electrons. We leave that to a future work.

(vii) Magnetized, quasi-parallel shocks are efficient ion accelerators. Approximately, 10 per cent of the shock's energy goes into non-thermal ions. The power-law index of ions appears to be softer than $f(p) \propto p^{-4}$ ($\mathrm{d}n/\mathrm{d}E \propto E^{-2}$), but the spectral index is not well constrained by our simulation. Constraining the ion spectrum would require much longer simulations.

(viii) Quasi-parallel shocks are capable of injecting electrons into a hard power law with a spectral index consistent with $f(p) \propto p^{-4.2}$ ($\mathrm{d}n/\mathrm{d}E \propto E^{-2.2}$). The power law starts at $\sim 2\gamma_0\beta_0 m_i c$. For $\mathcal{M}_A = 15, \theta_B = 10^\circ, \gamma_0 = 1.5, \beta_0 \approx 0.75$, the amount of the shock's energy that is put into the non-thermal electron power law is small, $\varepsilon_e \sim 5 \times 10^{-4}$. This ε_e corresponds to a number ratio of non-thermal electrons to total electrons of 5×10^{-5}. However, the value of ε_e may depend on the magnetic inclination upstream, σ_0, and the speed of the shock. We will address these trends in a future work.

Finally, we examined particle acceleration in superluminal mildly relativistic shocks, with $\theta_B = 55^\circ$. Since particles were unable to escape upstream, neither species was capable of entering DSA. Instead, a fraction of ions gained energy via SDA while crossing the shock front, resulting in a second suprathermal bump in the downstream ion spectrum. In mildly relativistic superluminal shocks, the maximum energy of both species did not increase with time. Throughout this paper, we have used a reduced mass ratio of $m_i/m_e = 64$. We find that m_i/m_e must be fairly large, or else Bell instability will be suppressed. We also find that 64 is sufficiently large to capture the shock structure and ε_e, as shown in Appendix A.

Our results on acceleration in mildly relativistic shocks have the following implications for simulations of electron–ion shocks: a small-reduced mass ratio, numerical heating, and too short of simulation duration all work in concert to suppress the Bell instability. If Bell does become the main instability, the simulation must have a large transverse size to capture the upstream filaments, and these filaments play a big role in injecting electrons in mildly relativistic shocks. These conclusions are very general and likely apply equally well to non-relativistic and relativistic particle-in-cell simulations of quasi-parallel shocks.

For non-relativistic shocks travelling at $v \ll c$, the filamentary mode of the Bell instability may not be as important for acceleration of electrons because even the amplified magnetic field is subluminal. This may explain the close agreement between ε_e measured in this work and the 1D simulations of $v = 0.05c$ shocks in Park et al. (2015). A set of directly comparable, 2D simulation of shocks ranging from non-relativistic to relativistic is left to future work.

Since our shocks were mildly relativistic, one needs to apply caution when extrapolating our conclusions regarding Bell-mediated shocks to relativistic or non-relativistic collisionless electron–ion quasi-parallel shocks. The growth rates of the microphysical instabilities may depend on physical parameters like magnetization and shock velocities. If σ_0 is large, the shock is likely mediated by the resonant ion instability, and when $\sigma_0 = 0$ or is very small, there is not enough magnetic field on large scales to initiate Bell instability. If the field orientation is quasi-perpendicular with little upstream turbulence, there may not be enough ion injection to cause the Bell instability. We will address how electron acceleration efficiency and magnetic amplification depends on v/c and σ_0, constraining the models that assume equipartition between ε_B and ε_e in future work.

In the future, it will be important to determine if 2D PIC simulations can properly capture the physics of a Bell-mediated shock. The reduced dimensionality of the simulations may allow for structures that do not form in 3D. A 3D simulation will be expensive because to be fully 3D, the z-direction should be much larger than that of the typical ion gyroradius and run for a long time. However, we can use the 3D results of prior hybrid simulations to inform how PIC simulations might change if run in 3D. Caprioli & Spitkovsky (2014a) found that in 3D the returning cosmic rays still evacuated cavities upstream, and the magnetic field amplification and ion injection and diffusion was largely the same. Therefore, our general picture of a transition from a Weibel- to a Bell-mediated quasi-parallel shock with an $\varepsilon_p \sim 10$ per cent, $E_{max} \propto t$ should hold. The characteristic values for the magnetic field amplification upstream and the less-amplified magnetic filaments downstream would likely remain in a 3D PIC simulation as well. However, our simulations have features not capable of being captured in hybrid simulations – the electron preheating of a factor of $\sim 3T_0$ and $\varepsilon_e \sim 5 \times 10^{-4}$. A fully 3D PIC simulation would need to be run to answer how our results may change.

In conclusion, we have used fully kinetic particle-in-cell simulations to show that quasi-parallel, magnetized, mildly relativistic electron–ion shocks travelling at velocity $v \sim 0.7c$ are capable of self-consistently generating the necessary magnetic turbulence upstream to accelerate particles to high energies quickly. In addition, these shocks put a large fraction of their energy into hadrons. These characteristics suggest that astrophysical objects that host such shocks may be an important source of high-energy hadronic radiation, cosmogenic neutrinos, and possibly a source of the ultra-high energy cosmic rays. Therefore, objects like low-luminosity AGN like BL Lacs and FRI, as well as micro-quasars, which may have mildly relativistic shocks inside of their jets, are important potential multimessenger sources for the upcoming high-energy gamma-ray telescopes such as the Cherenkov Telescope Array and also existing and forthcoming high-energy neutrino detectors like IceCube and KM3NeT.

ACKNOWLEDGEMENTS

We thank the anonymous referee for their comments. We thank Jaehong Park and Lorenzo Sironi for valuable discussions and their assistance with *TRISTAN-MP*. PC acknowledges financial support from the WARP programme of the Netherlands Organisation for Scientific Research (NWO). SM is supported by an NWO VICI grant (no. 639.043.513). PC thanks Rahul Kumar and Illya Plotnikov for valuable conversations. The simulations in this paper were performed using computational resources at the TIGRESS high performance computer centre at Princeton University, the University of Chicago Research Computing Center, and the Texas Advanced Computing Center (TACC) at The University of Texas at Austin (TG-AST100035). This work was supported by NSF grant AST-1517638. AS is supported by the Simons Foundation (grant 267233).

REFERENCES

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

Aguilar M. et al., 2014, Phys. Rev. Lett., 113, 221102

Aguilar M. et al., 2015, Phys. Rev. Lett., 114, 171103
Axford W. I., Leer E., Skadron G., 1977, Int. Cosmic Ray Conf., 11, 132
Ball L., Melrose D. B., 2001, PASA, 18, 361
Begelman M. C., Kirk J. G., 1990, ApJ, 353, 66
Bell A. R., 1978, MNRAS, 182, 147
Bell A. R., 2004, MNRAS, 353, 550
Bell A. R., 2005, MNRAS, 358, 181
Blandford R., Eichler D., 1987, Phys. Rep., 154, 1
Blandford R. D., Ostriker J. P., 1978, ApJ, 221, L29
Bohdan A., Niemiec J., Kobzar O., Pohl M., 2017, ApJ, 847, 71
Buneman O., 1993, Computer Space Plasma Physics: Simulation Techniques and Software. Terra Scientific Publishing Company, Tokyo
Caprioli D., Spitkovsky A., 2013, ApJ, 765, L20
Caprioli D., Spitkovsky A., 2014a, ApJ, 783, 91
Caprioli D., Spitkovsky A., 2014b, ApJ, 794, 46
Caprioli D., Spitkovsky A., 2014c, ApJ, 794, 47
Dieckmann M. E., Bret A., 2018, MNRAS, 473, 198
Dieckmann M. E., Shukla P. K., Eliasson B., 2006, New J. Phys., 8, 225
Drury L. O., 1983, Rep. Prog. Phys., 46, 973
Ellison D. C., Warren D. C., Bykov A. M., 2013, ApJ, 776, 46
Fermi E., 1949, Phys. Rev., 75, 1169
Keshet U., Waxman E., 2005, Phys. Rev. Lett., 94, 111102
Kirk J. G., Guthmann A. W., Gallant Y. A., Achterberg A., 2000, ApJ, 542, 235
Krasnoselskikh V. V., Lembège B., Savoini P., Lobzin V. V., 2002, Phys. Plasmas, 9, 1192
Krymskii G. F., 1977, Akademiia Nauk SSSR Doklady, 234, 1306
Kumar R., Eichler D., Gedalin M., 2015, ApJ, 806, 165
Malkov M. A., Drury L. O., 2001, Rep. Prog. Phys., 64, 429
Matsumoto Y., Amano T., Hoshino M., 2012, ApJ, 755, 109
Medvedev M. V., Loeb A., 1999, ApJ, 526, 697
Morlino G., Caprioli D., 2012, A&A, 538, A81
Morlino G., Blasi P., Vietri M., 2007, ApJ, 658, 1069
Niemiec J., Pohl M., Bret A., Wieland V., 2012, ApJ, 759, 73
Panaitescu A., Kumar P., 2000, ApJ, 543, 66
Park J., Caprioli D., Spitkovsky A., 2015, Phys. Rev. Lett., 114, 085003
Plotnikov I., Pelletier G., Lemoine M., 2011, A&A, 532, A68
Quenby J. J., Lieu R., 1989, Nature, 342, 654
Reville B., Bell A. R., 2012, MNRAS, 419, 2433
Reville B., Kirk J. G., Duffy P., 2006, Plasma Phys. Controlled Fusion, 48, 1741
Riquelme M. A., Spitkovsky A., 2011, ApJ, 733, 63
Santana R., Barniol Duran R., Kumar P., 2014, ApJ, 785, 29
Sironi L., Spitkovsky A., 2009, ApJ, 698, 1523
Sironi L., Spitkovsky A., 2011, ApJ, 726, 75
Sironi L., Spitkovsky A., Arons J., 2013, ApJ, 771, 54
Spitkovsky A., 2005, in Bulik T., Rudak B., Madejski G., eds, AIP Conf. Ser. Vol. 801, Astrophysical Sources of High Energy Particles and Radiation. Am. Inst. Phys., New York, p. 345
Spitkovsky A., 2008a, ApJ, 673, L39
Spitkovsky A., 2008b, ApJ, 682, L5
Summerlin E. J., Baring M. G., 2012, ApJ, 745, 63
Weibel E. S., 1959, Phys. Rev. Lett., 2, 83
Wijers R. A. M. J., Galama T. J., 1999, ApJ, 523, 177

APPENDIX A: SCALING WITH MASS RATIO

In the main text of paper, we have employed a reduced mass ratio $m_i/m_e = 64 \ll 1836$. Here, we examine the consequences of using a smaller mass ratio on the shock structure and particle acceleration. We find that using too small of a mass ratio can suppress the cosmic ray filamentation instability with the following implications – there is too little magnetic field amplification upstream and downstream, the maximum particle energy grows too slowly, the energy transfer between electrons and ions is too efficient, and artificially small m_i/m_e boosts ε_e. While our fiducial mass ratio of $m_i/m_e = 64$ agrees quantitatively with a higher mass ratio $m_i/m_e = 160$ simulation in terms of the magnetic field amplification and maximum particle energy, it only qualitatively agrees in the other respects. Using $m_i/m_e = 64$ slightly underpredicts the importance of the cosmic ray filamentation instability, with smaller, weaker filaments in the upstream, more electron heating, and an ε_e that is slightly bigger in the $m_i/m_e = 64$ simulation than $m_i/m_e = 160$. In the larger mass ratio simulation, whistler waves may be more important, as they will cross the $\mathcal{M}_A \lesssim \sqrt{m_i/m_e}$ barrier at smaller values of magnetic field amplification.

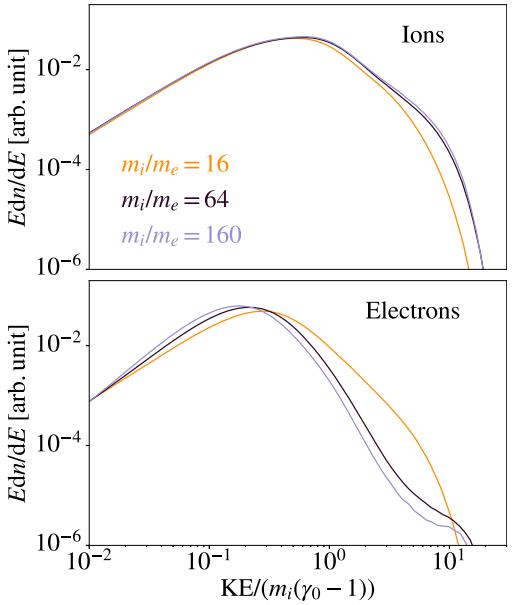

Figure A1. The two panels show the downstream particle spectra in the downstream rest frame extracted 10–80 c/ω_{pi} downstream of the shock at $\omega_{pi}t \approx 2000$, shortly after the time the shock becomes Bell-mediated in the fiducial, $m_i/m_e = 64$, simulation. As the mass ratio increases to more realistic values, the energy transfer to electrons from ions decreases, with electrons having a lower temperature downstream. For $m_i/m_e = 16$, the cosmic ray filamentary instability is significantly suppressed and ε_e is far too large at $\omega_{pi}t \sim 2000$.

To see the effect of our choice of reduced mass ratio, 64, we compare our fiducial run to a simulation with a smaller mass ratio of 16 (the mass ratio used in Sironi & Spitkovsky 2011), and a larger mass ratio of 160 in Figs A1–A4. The higher mass ratio run had the same set-up as the fiducial run except it used a smaller computational box with a width of $\approx 79\, c/\omega_{pi}$. The $m_i/m_e = 16$ simulation had the same set-up except it used a higher resolution of $c/\omega_{pe} = 8$ cells and a box width of $\approx 109\, c/\omega_{pi}$. The higher resolution was required to ensure that the smaller Debye length was resolved ($\lambda_D \propto \sqrt{m_i/m_e}$).

Figs A1 and A2 show the effect of the reduced mass ratio on the downstream and upstream spectral properties, respectively. It is clear that using a mass ratio of 16 is insufficient to correctly capture the spectral properties of a Bell-mediated, electron-ion shock at time $\omega_{pi}t \sim 2000$ in a $\gamma_0 = 1.5$ shock. However, Sironi & Spitkovsky (2011) do see circularly polarized waves in their simulations of relativistic magnetized shocks with $m_i/m_e = 16$, $\sigma = 0.1$, and $\gamma_0 = 15$. The difference may be an issue of relativistic shocks evolving faster, or the higher magnetization ($\sigma = 0.1$) used in those simulations.

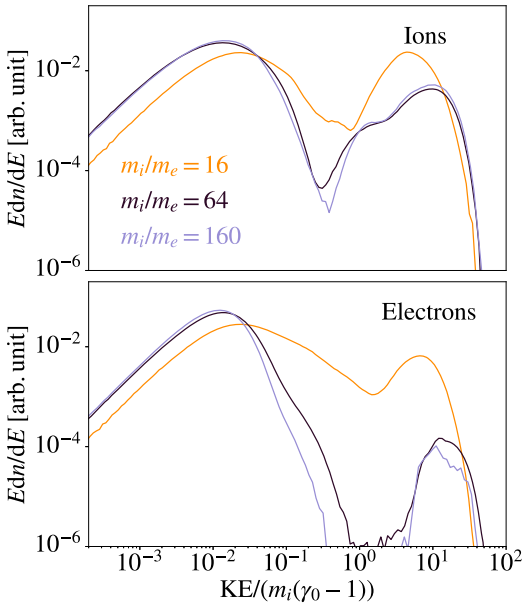

Figure A2. Three simulations that illustrate the importance of a large mass ratio. The two panels show the particle spectra in the upstream rest frame extracted 10–20 $c/\omega_{\rm pi}$ upstream of the shock at $\omega_{\rm pi}t \approx 2000$. With a small of mass ratio, electrons are heated significantly more upstream before entering the shock.

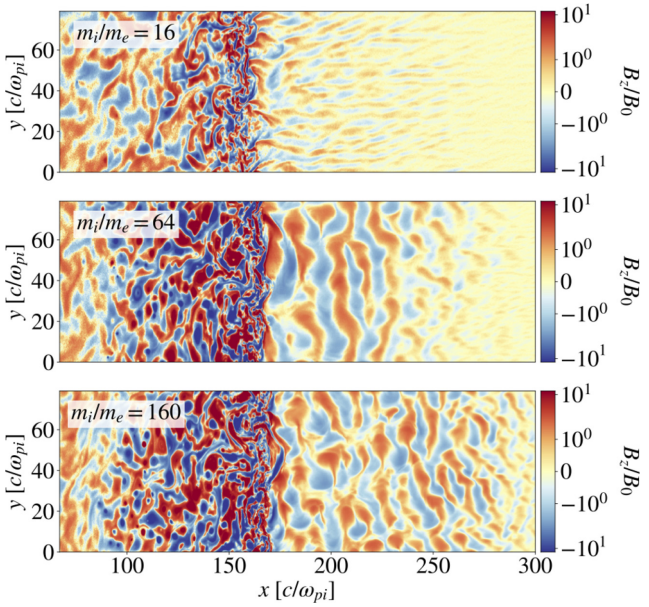

Figure A3. The effect of varying the mass ratio on the filamentary structures in the upstream. The three panels show B_z near the shock front at time $\omega_{\rm pi}t \approx 2000$. The middle panel shows a representative slice of the fiducial simulation. Increasing $m_{\rm i}/m_{\rm e}$ increases the importance of the cosmic ray filamentation instability.

The $m_{\rm i}/m_{\rm e} = 16$ run predicts an $\varepsilon_{\rm e}$ that is >10 times larger than the runs with larger mass ratios, whereas the $\varepsilon_{\rm e}$ in the $m_{\rm i}/m_{\rm e} = 64$ and 160 runs agree within a factor of 2. There is no sign of an increase in energy transfer in the high-mass run due to whistler waves. As the mass ratio increases, the upstream energy exchange between electrons and ions is less efficient. In all of the runs, the downstream thermal electrons achieve a significant fraction of the ion energy.

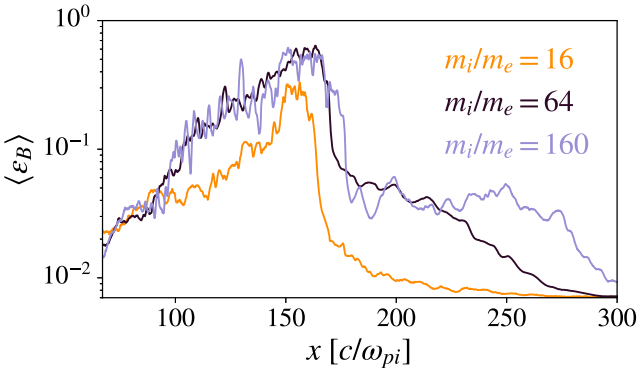

Figure A4. The effect of varying the mass ratio on the y-averaged ε_B at time $\omega_{\rm pi}t \approx 2000$. When more realistic mass ratios are used, the cosmic ray filamentation instability becomes more important, increasing the magnetic field amplification both upstream and downstream of the shock front. Note the good quantitative agreement between $m_{\rm i}/m_e = 64$ and 160 in terms of the magnitude of the magnetic amplification upstream and downstream.

The downstream ion spectrum in the $m_{\rm i}/m_{\rm e} = 16$ run does not extend to as high of energy as the larger mass ratio runs, so $\varepsilon_{\rm p}$ is not well-constrained, but it appears similar to the larger mass ratio runs. Note that the non-thermal tail extends to larger energies in both the electrons and ions for the larger mass ratio runs as compared to the $m_{\rm i}/m_{\rm e} = 16$. The high-energy particles are scattered back to the shock more efficiently because the larger mass ratio runs show more upstream magnetic field amplification and structures at larger scales (see Figs A3 and A4).

In terms of the overall magnetic field amplification, the $m_{\rm i}/m_{\rm e} = 160$ and the fiducial run agree very well both upstream and downstream, the only difference is that the extent of the amplification is larger upstream for $m_{\rm i}/m_{\rm e} = 160$. The difference in extent is because the larger mass ratio run transitions from Weibel-mediated to Bell earlier than the fiducial run (in terms of $\omega_{\rm pi}^{-1}$). The smaller mass ratio of $m_{\rm i}/m_{\rm e} = 16$ has not yet transitioned to being Bell-mediated.

APPENDIX B: DEPENDENCE ON MAGNETIC FIELD ORIENTATION WITH RESPECT TO SIMULATION PLANE

In this appendix, we show that for both the quasi-parallel and superluminal shocks considered in this paper, the orientation of the magnetic field with respect to the simulation plane does not significantly affect our results.

In the main text of the paper, we used an initial magnetic field B_0 that was in the $x-y$ plane. There is a growing body of evidence that 2D simulations initialized with $B_\perp$ along z as opposed to y may be more efficient at accelerating electrons in high Mach number, perpendicular shocks (e.g. Matsumoto, Amano & Hoshino 2012; Sironi et al. 2013; Bohdan et al. 2017). In the quasi-parallel shocks, the out-of-plane configuration initially reflects more electrons, resulting in an $\varepsilon_{\rm e}$ that is ~2 times larger than the 5×10^{-4} reported in the main paper. However, once the turbulence upstream saturates, the amplified magnetic field is the dominant field at the shock front. Then, the two B_0 configurations have approximately the same $\varepsilon_{\rm e}$.

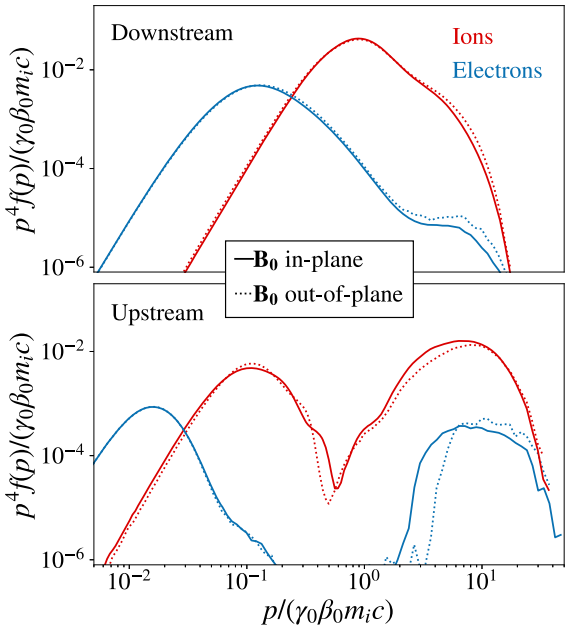

Figure B1. Comparison of upstream and downstream spectra with different initial magnetic field orientations with respect to the simulation plane at time $\omega_{\rm pi}t = 4148$. Solid lines correspond to the fiducial run with an in-plane B_0 field (i.e. $B_{0,\perp}$ along y). Dotted lines correspond to a smaller run with out-of-plane orientation ($B_{0,\perp}$ along z) and a width of 100 $c/\omega_{\rm pi}$. The upper panel is the particle spectra in the simulation rest frame extracted 20–30 $c/\omega_{\rm pi}$ upstream of the shock, while the lower panel is the spectra in the downstream rest frame extracted 10–80 $c/\omega_{\rm pi}$ downstream of the shock. While the out-of-plane inclination initially reflects more electrons, when turbulence B field upstream grows larger, the difference between the two inclinations electron acceleration efficiency is small.

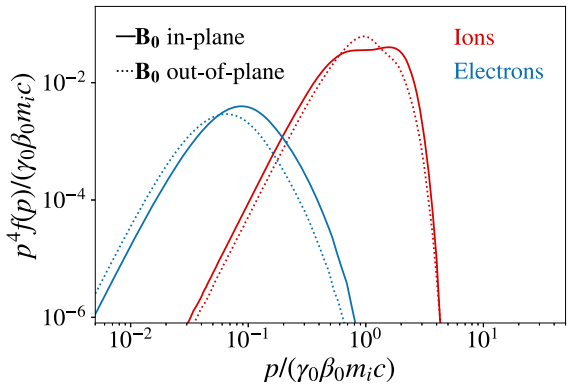

Figure B2. Comparison of downstream superluminal spectra with different initial magnetic field orientations with respect to the simulation plane at time $\omega_{\rm pi}t \approx 2600$. Solid lines correspond to superluminal shocks with $\theta_B = 55^\circ$ and in-plane B_0 field (i.e. $B_{0,\perp}$ along y). Dotted lines correspond a superluminal shock with out-of-plane orientation ($B_{0,\perp}$ along z). The spectra are extracted 10–80 $c/\omega_{\rm pi}$ downstream of the shock and are calculated in the downstream rest frame. The upstream is not shown as the spectra are simply Maxwellians with $T = T_0$. Neither configurations show non-thermal acceleration. The in-plane configuration has hotter downstream ions and electrons and a larger SDA bump in the ion spectrum due to its larger compression ratio (see Fig. B3).

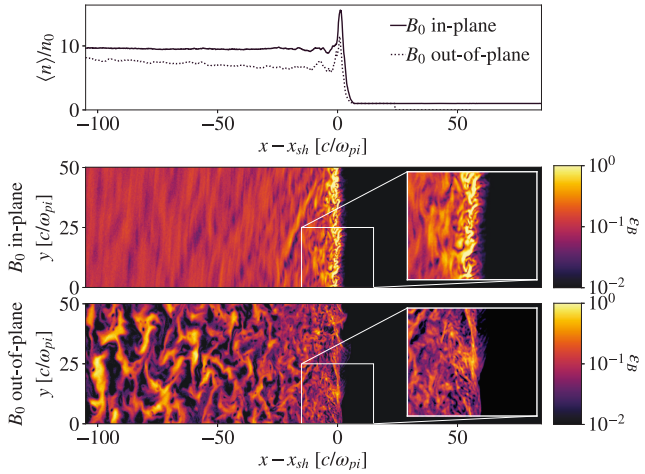

Figure B3. Comparison of field quantities close to the shock between superluminal in-plane and out-of-plane magnetic field configuration. As can be seen in the top panel, the two configurations have a different compression ratios, and therefore different shock speeds. The lower panels show ε_B for the in-plane and out-of-plane configuration. The out-of-plane shock takes longer to thermalize the plasma, showing less compression and a more turbulent downstream. The out-of-plane shock front has structures on the scale of the ion gyroradius. The in-plane shock thermalizes the plasma very quickly and has a higher compression ratio.

For the superluminal shocks, the field configuration is more important. The in-plane shock more quickly thermalizes the plasma at the shock front, and therefore has a larger compression ratio. The larger compression ratio results in the ion spectrum having a larger SDA bump, and hotter electrons (see e.g. Caprioli & Spitkovsky 2014a). The out-of-plane configuration does not effectively scatter ions in the z direction decreasing the adiabatic index of the shock. The smaller adiabatic index leads to a smaller compression ratio. However, ions do gyrate efficiently in the $x-y$ plane at the shock front, causing more turbulence at the shock front and downstream. In the out-of-plane configuration, the turbulence causes the electrons to be heated far downstream. However, neither configuration is capable of producing upstream magnetic turbulence nor accelerating particles to high energies.

APPENDIX C: CONVERGENCE IN RESOLUTION AND NUMBER OF PARTICLES

In this section, we briefly show that our choice of resolution and particles per cell is sufficient to measure $\varepsilon_{\rm e}$ and $\varepsilon_{\rm p}$. We compare our fiducial run with $c/\omega_{\rm pe} = 4$ cells, ppc $= 4$ to a run with twice the resolution $c/\omega_{\rm pe} = 8$ cells, and a different run with 64 particles per cell. The downstream spectra are shown in Fig. C1. $\varepsilon_{\rm p}$ and $\varepsilon_{\rm e}$ agree with a factor of 2 for all the runs. The shock structure, magnetic field turbulence, and magnetic amplification was very similar in these runs.

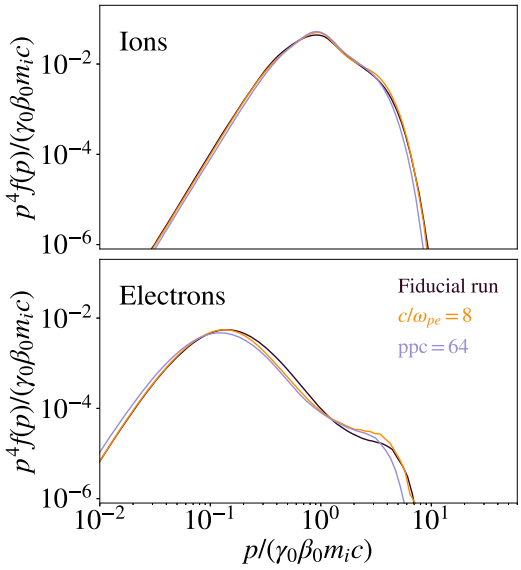

Figure C1. The figure shows the particle spectra in the downstream rest frame extracted 10–50 $c/\omega_{\rm pi}$ downstream of the shock at $\omega_{\rm pi}t \approx 1500$, shortly after the time the shock becomes Bell-mediated in the fiducial, $c/\omega_{\rm pe} = 4$, ppc = 4, simulation. We compare our fiducial run to a simulation with twice as many cells per skin depth and another simulation with 16 times as many particles per cell. $\varepsilon_{\rm e}$ and $\varepsilon_{\rm p}$ in all the runs agree within a factor of $\sim$2.

This paper has been typeset from a TEX/LATEX file prepared by the author.