The Astrophysical Journal Letters, 963:L44 (8pp), 2024 March 10 https://doi.org/10.3847/2041-8213/ad2c8c

Long-term Evolution of Relativistic Unmagnetized Collisionless Shocks

Daniel Grošelj[1,2], Lorenzo Sironi[2,3], and Anatoly Spitkovsky[4]
[1] Centre for mathematical Plasma Astrophysics, Department of Mathematics, KU Leuven, B-3001 Leuven, Belgium; daniel.groselj@kuleuven.be
[2] Department of Astronomy and Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA
[3] Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA
[4] Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA
Received 2024 January 4; revised 2024 February 20; accepted 2024 February 21; published 2024 March 7

Abstract

We study a relativistic collisionless electron–positron shock propagating into an unmagnetized ambient medium using 2D particle-in-cell simulations of unprecedented duration and size. The shock generates intermittent magnetic structures of increasingly larger size as the simulation progresses. Toward the end of our simulation, at around 26,000 plasma times, the magnetic coherence scale approaches $\lambda \sim 100$ plasma skin depths, both ahead and behind the shock front. We anticipate a continued growth of λ beyond the time span of our simulation, as long as the shock accelerates particles to increasingly higher energies. The post-shock field is concentrated in localized patches, which maintain a local magnetic energy fraction $\varepsilon_B \sim 0.1$. Particles randomly sampling the downstream fields spend most of their time in low field regions ($\varepsilon_B \ll 0.1$) but emit a large fraction of the synchrotron power in the localized patches with strong fields ($\varepsilon_B \sim 0.1$). Our results have important implications for models of gamma-ray burst afterglows.

Unified Astronomy Thesaurus concepts: High energy astrophysics (739); Shocks (2086); Non-thermal radiation sources (1119); Plasma astrophysics (1261); Gamma-ray bursts (629)

Supporting material: animation

1. Introduction

Gamma-ray bursts (GRBs) are powerful cosmic explosions driven by the collapse of a massive star or the merger of binary neutron stars (Mészáros 2002; Piran 2004; Kumar & Zhang 2015). The prompt emission stemming from the internal dissipation of the GRB jet is followed by an afterglow powered by the external relativistic collisionless shock, which expands into the surrounding medium. The afterglow unfolds from a few seconds up to decades after the explosion, as we are now witnessing for GRB 221009A, the "brightest of all time" (Burns et al. 2023; Laskar et al. 2023; Williams et al. 2023). The emission from GRB afterglows is distinctly nonthermal and broadband. It ranges from the radio to the gamma-ray band (Abdo et al. 2009; Ackermann et al. 2010, 2014; De Pasquale et al. 2010; Chandra & Frail 2012; Perley et al. 2014) and has been recently observed even at TeV energies (Abdalla et al. 2019; MAGIC Collaboration et al. 2019a, 2019b; H.E.S.S. Collaboration et al. 2021; Huang et al. 2022; LHAASO Collaboration et al. 2023).

GRB afterglow shocks propagate into a very weakly magnetized ambient medium. The ratio of the magnetic to the particle energy density of the cold ambient medium (i.e., the magnetization) is $\sigma \lesssim 10^{-5}$ for massive stellar progenitors (Crowther 2007) and could be as low as $\sigma \sim 10^{-9}$ if the shock expands into an interstellar-like medium (e.g., Sironi et al. 2013; Grošelj et al. 2022). A mechanism for magnetic field generation must be present at the shock in order to facilitate efficient particle scattering and their acceleration via the Fermi process (e.g., Bell 1978; Blandford & Ostriker 1978; Drury 1983; Blandford & Eichler 1987; Achterberg et al. 2001) and to explain the resulting broadband emission. Kinetic particle-in-cell (PIC) simulations of relativistic collisionless shocks (Spitkovsky 2008a, 2008b; Martins et al. 2009; Nishikawa et al. 2009; Haugbølle 2011; Sironi et al. 2013) have shown that the Weibel filamentation instability (Fried 1959; Weibel 1959; Medvedev & Loeb 1999; Silva et al. 2003; Achterberg et al. 2007; Bret et al. 2014; Takamoto et al. 2018; Lemoine et al. 2019a) generates the required fields, which serve as the scattering agents for the Fermi process. However, Weibel-generated fields grow on plasma microscales and decay quickly downstream from the shock (Gruzinov 2001; Chang et al. 2008; Lemoine 2015), which is at odds with the relatively high field strengths inferred for some GRB afterglows (e.g., Wijers & Galama 1999; Panaitescu & Kumar 2002), including the recent GRB 221009A (Laskar et al. 2023). Moreover, particle scattering in microscale fields presents a challenge for the acceleration to very high energies (Kirk & Reville 2010; Sironi et al. 2013; Reville & Bell 2014; Asano et al. 2020; Huang et al. 2022). The largest PIC simulations reported in the literature (Keshet et al. 2009) found that the acceleration of particles to higher energies drives the formation of magnetic fields on progressively larger scales over time, which slows down the decay of the post-shock magnetic fields. However, despite the computational progress, previous as well as present PIC simulations need to be extrapolated to orders-of-magnitude-longer timescales to make direct contact with observations, and so the long-term evolution of the shock remains a subject worth investigating.

In this Letter, we revisit the microphysics of relativistic unmagnetized collisionless shocks using PIC simulations of unprecedented duration and size. We observe the generation of intense magnetic structures in the upstream flow, reaching sizes up to ~100 plasma skin depths. The bulk of the post-shock plasma becomes magnetized, which leads to particle trapping in the magnetostatic turbulence, and prevents the fast decay of magnetic fields predicted for untrapped particles (Chang et al. 2008;

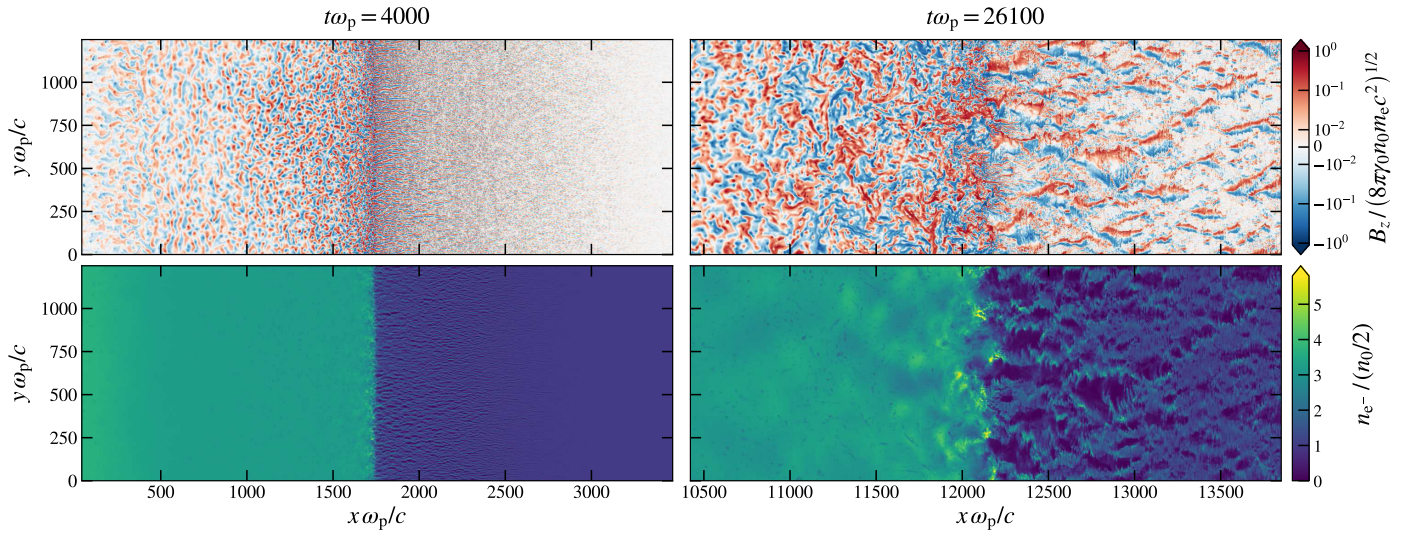

Figure 1. Structure of a collisionless relativistic pair plasma shock, propagating into an unmagnetized upstream medium with a bulk Lorentz factor $\gamma_0 = 10$, at $t\omega_p = 4000$ (left) and at $t\omega_p = 26{,}100$ (right). On the top we show the out-of-plane magnetic field B_z, and on the bottom we show the electron number density n_{e^-}. The animation lasts 40 s and shows the spatiotemporal evolution of B_z (top panel) and n_{e^-} (bottom panel) from the start ($t\omega_p = 0$) to the end ($t\omega_p = 26{,}100$) of the simulation.

(An animation of this figure is available.)

Lemoine 2015). The downstream field is concentrated in localized patches, which occupy roughly 1% of the total volume but contain about half of the magnetic energy. Our results have direct implications for particle acceleration and emission from GRB afterglow shocks.

2. Method

We perform 2D simulations of relativistic collisionless shocks using the PIC code OSIRIS 4.0 (Fonseca et al. 2002, 2013). In order to evolve the simulations for as long as possible, we employ for simplicity an electron–positron pair composition. As is routinely done (e.g., Sironi et al. 2013), the shock formation is driven by the reflection of the cold upstream flow (moving to the left) from the left simulation boundary at $x = 0$. The rest frame of the simulation coincides with the downstream frame, and the shock propagates in the positive x-direction. In order to reduce reflecting boundary artifacts, the initial upstream-flow Lorentz factor transitions linearly from unity to the far-upstream value $\gamma_0 = 10$ over a distance of 400 c/ω_p from $x = 0$, where $c/\omega_p = (\gamma_0 m_e c^2/4\pi n_0 e^2)$ is the plasma skin depth and n_0 is the simulation-frame density of the upstream electrons and positrons. At $t = 0$, the upstream plasma fills the transition layer between $x = 0$ and $x = 400\,c/\omega_p$. As the simulation progresses, the region $x > 400\,c/\omega_p$ is gradually populated with upstream plasma introduced from a moving particle injector, which moves to the right at light speed c until it reaches the end of the box.

Our fiducial simulation has a box size of $20{,}000 \times 1248\,c/\omega_p$ and is evolved up to $t\omega_p = 26{,}100$. The cell size is $\Delta x = 0.2\,c/\omega_p$, and our time step is $\Delta t = 0.5\Delta x/c$. The upstream plasma is represented with 144 particles per cell per species with cubic spline shapes. The total number of particles at the end of the simulation exceeds 0.4 trillion. Particle noise is further reduced by low-pass filtering the electric currents at each step. The fields are advanced using a modified stencil (Blinne et al. 2018) with improved numerical dispersion of electromagnetic waves to suppress the numerical Cherenkov instability (see Grošelj et al. 2022). For improved computational performance, we employ a vectorized particle push and deposit, together with a hybrid MPI/OpenMP parallelization strategy, all of which are readily available in the OSIRIS code (Fonseca et al. 2013).

3. Shock Evolution

Figure 1 shows the evolution of the shock structure in our fiducial simulation, as observed through the snapshots of the out-of-plane magnetic field (top panels) and the electron density (bottom). The particles returning from the shock into the upstream flow drive plasma-streaming instabilities, which grow and amplify magnetic fields. The upstream region populated with the self-generated fields defines a turbulent precursor, which expands with time as the returning particles propagate farther into the upstream. At early times ($t\omega_p \lesssim 4000$; left panels in Figure 1), the upstream plasma near the shock front develops a microscale filamentary structure, typical of shocks mediated by the Weibel instability (e.g., Lemoine et al. 2019a). In contrast, at later stages of the shock evolution (right panels in Figure 1), the upstream Weibel filaments undergo a secondary nonlinear instability (Honda et al. 2000; Peterson et al. 2021, 2022; Grošelj et al. 2022), which results in large-scale magnetized plasma cavities.[5]

The structure of the precursor in the plasma cavity-dominated phase is illustrated in Figure 2. Cavity formation is triggered by the development of a net electric current and charge in the returning population of electrons and positrons (defined as having $\gamma\beta_x > 0$), which stream against the incoming flow (Figure 2(e)). For example, at the time of Figure 2, the returning particles carry an excess positive charge. The asymmetry of the returning particles is then imprinted onto the current filaments of positive ($J_x > 0$) or negative ($J_x < 0$) polarity in the incoming background plasma (Figure 2(d)),

[5] Magnetized plasma cavities can be also spotted in earlier simulations of relativistic pair shocks (see Figure 1(b) from Keshet et al. 2009). However, their role in the long-term evolution of pair shocks was not fully appreciated at the time.

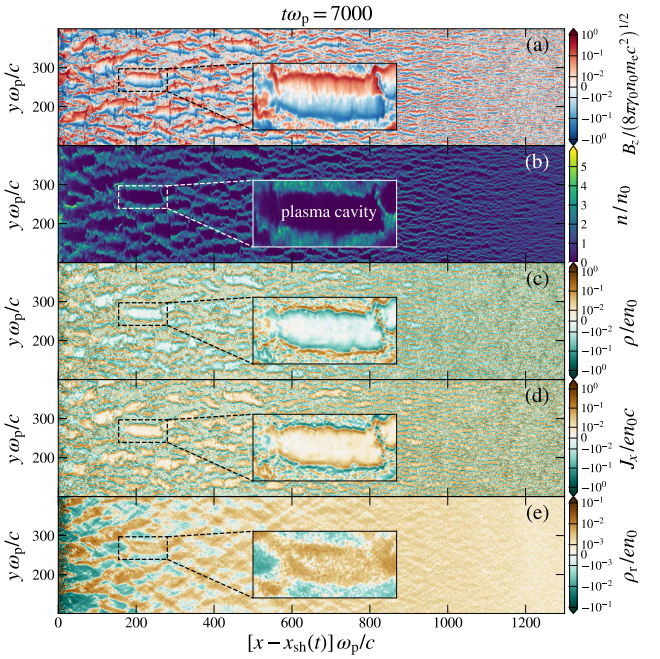

Figure 2. Structure of the shock precursor populated with magnetized plasma cavities at $t\omega_p = 7000$. Only a fraction of the box is shown in the range $100 < y\omega_p/c < 400$ and $0 < [x - x_{sh}(t)]\omega_p/c < 1300$, where $x_{sh}(t)$ is the location of the propagating shock front. From top to bottom we show (a) the z-component of the magnetic field B_z, (b) the particle number density n, (c) the total electric charge density ρ, (d) the x-component of the electric current J_x, and (e) the charge density ρ_r of the returning particles with $\gamma\beta_x > 0$. The inset plots zoom in on the structure of a single cavity.

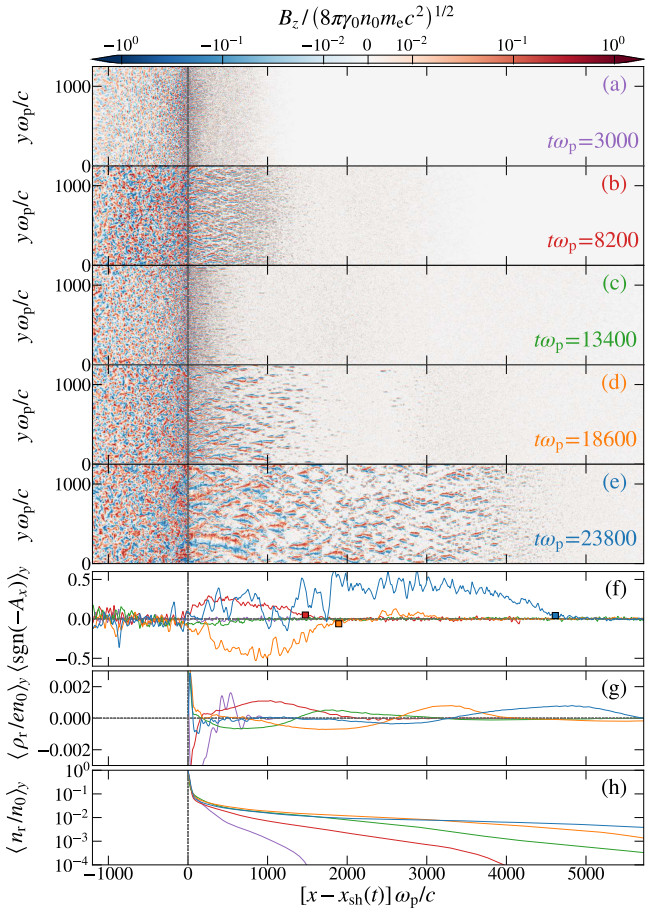

Figure 3. Evolution of the magnetic field and the emergence of symmetry breaking. Panels (a)–(e) show snapshots of B_z at 5 times in the simulation. In panel (f) we show the "order parameter" $Q(x) = \langle \mathrm{sgn}(-A_x)\rangle_y$. Values of $Q(x)$ significantly above or below 0 indicate the presence of cavities of a given magnetic polarity (see Section 3 for details). Panel (g) shows the transversely averaged upstream charge density of the returning particles with $\gamma\beta_x > 0$. In (h) we show the transversely averaged number density of the returning particles. Squares in (f) show the location of cavity birth (see Appendix).

which tries to neutralize the returning-particle charge and current. The current filaments with the same polarity as the returning-particle current ($J_x > 0$ in Figure 2(d)) become subject to the cavitation instability and expand in the transverse (y-) direction; the filaments of the opposite polarity ($J_x < 0$ in Figure 2(d)) do not expand. In the rest frame of the simulation, the electric current inside the cavities is sustained by a low-density population of incoming particles of the opposite charge sign to the returning-particle charge density (Figures 2(c)–(e)). The current inside the cavities is left essentially unscreened, which enables the generation of intense magnetic fields (Grošelj et al. 2022). The newly born plasma cavities then grow and merge as they propagate toward the shock (Figure 2(b)), which allows the magnetic field to cascade to progressively larger scales (Figure 2(a)). At the end of our simulation, the largest individual cavities reach transverse sizes of about 200 plasma skin depths (Figure 1, right panels).

The current and charge of the returning particles, which drive cavity formation, exhibit oscillations in sign over time and in space. This gives rise to bursts (or cycles) of cavity generation with alternating polarity, as demonstrated in Figure 3. The returning-particle charge fluctuations originate near the shock with a time-varying sign and spread into the upstream (Figure 3(g)). The amplitude of the charge fluctuations amounts to a significant fraction of en_r, where n_r is the returning-particle number density (Figure 3(h)). In response to the charge oscillations in the stream of returning particles, the polarity of the magnetized plasma cavities changes over time (Figures 3(a)–(e)). To quantify the symmetry breaking in the induced magnetic field structure, we introduce an "order parameter" $Q(x)$, defined as $Q(x) = \langle \mathrm{sgn}(-A_x)\rangle_y$, where A_x is the longitudinal component of the magnetic vector potential (Figure 3(f)). We solve for A_x in the Coulomb gauge so that $-A_x$ can be related via Ampere's law to the current as $\nabla^2(-A_x) = (4\pi/c)J_x$. Essentially, $Q(x)$ measures the relative difference in the volume occupied by magnetic structures with $A_x < 0$ versus those with $A_x > 0$, and as such it provides a convenient diagnostic for tracking the changes in polarity over time and in space. In our fiducial run, the first cavity cycle develops a positive polarity (see Figure 3(b) and the corresponding profile of $Q(x)$). It starts around $t\omega_p \approx 4500$ and lasts until $t\omega_p \approx 11{,}000$ or so. The second cycle displays a negative polarity (Figure 3(d)). It starts around $t\omega_p \approx 13{,}500$, continues up to $t\omega_p \approx 23{,}000$, and transitions into the third cycle, which features cavities of positive polarity (Figure 3(e)). The third cavity cycle is still ongoing at the end of the run (Figure 1, right panels).

The origin of the returning-particle charge and current oscillations is connected to the mechanism of particle reflection at Weibel-mediated shocks, which was recently studied in Parsons et al. (2023). The authors found that returning particles of a given species (electrons or positrons) need to "find" current channels with a matching sign ($J_x < 0$ for electrons and $J_x > 0$ for positrons when the shock travels in the positive x-direction, as in our run) in order to efficiently move away from the shock

into the upstream. This is supported by Figures 2(d)–(e), which show that the local sign of the returning-particle charge density ρ_r correlates with the local sign of the current J_x, within a few hundred c/ω_p from the shock. The injection of particles back into the upstream through the current channels provides a mechanism for spontaneous symmetry breaking. Namely, if the distribution of the background plasma current develops a seed asymmetry, one of the returning species finds current channels with a matching sign more easily and is more successful at moving away from the shock, and so the returning population as a whole develops a net current and charge (Figure 2(e)). The returning particles of the species that sustains the excess charge then propagate farther into the upstream and drive the formation of cavities, which amplifies the original seed asymmetry of the background plasma current, reinforcing the positive feedback loop. However, the continuous promotion of returning particles with a given charge sign generates a growing electrostatic potential between the shock and far upstream. The electrostatic field pulls the species that provides the excess charge toward the shock and feeds back negatively on the asymmetry. Instead of continuous generation of cavities with a fixed polarity, we therefore expect oscillations with a time-varying polarity, which are indeed observed in our simulations. Note that when the composition includes ions, the difference in inertia between electrons and ions acts as a natural mechanism of symmetry breaking (Lemoine & Pelletier 2011; Grošelj et al. 2022). Thus, the exact nature of cavity cycles in electron–ion shocks could differ from that in pair shocks.[6]

4. Magnetic Field Statistics

In Figure 4 we characterize the evolution of the magnetic field behind (left column) and ahead of (right column) the shock. Owing to cavity cycles (see Section 3), the evolution of magnetic field properties is not strictly monotonic. On the other hand, well-defined trends in the evolution can still be identified when the full time span of the simulation is considered. In particular, the magnetic field transverse coherence scale[7] λ in front of the shock (Figure 4(d)) grows by an order of magnitude over the duration of the simulation and approaches $\lambda \sim 100\, c/\omega_p$ at late time. The corresponding late-time magnetic spectrum (Figure 4(f)) is broadband; the magnetic energy uniformly fills the transverse wavenumber range $0.02 \lesssim k_y c/\omega_p \lesssim 0.5$. The size of structures in front of the shock correlates with the distance Δx_{cav} between the shock and the upstream location where the cavities arriving at the shock are born. In Figure 4(d) we compare measurements of Δx_{cav} (gray dots) with the evolution of λ and find empirically an approximately linear scaling $\lambda \approx 0.012 \Delta x_{cav}$.[8] We discuss the possible implications of this result in Section 6.

Magnetic structures generated over the precursor are compressed at the shock and transmitted into the downstream. Structures of growing size ahead of the shock thus imply that more energy is transferred to longer wavelength modes in the downstream, which results in a slower overall decay of the magnetic energy behind the shock (Figures 4(a) and (i)). The post-shock fields at wavenumbers $k_y \gtrsim \omega_p/c$ decay within a few hundred skin depths from the shock, leaving behind the field at $k_y \lesssim 0.1\, \omega_p/c$ (Figure 4(e)), where the decay is slow because the particles are marginally magnetized. The Larmor radius $R_L(\gamma)$ of a particle with typical Lorentz factor $\gamma \sim \gamma_0$ is $R_L(\gamma_0) \simeq (2\varepsilon_B)^{-1/2}\, c/\omega_p$, where $\varepsilon_B = B_z^2/8\pi\gamma_0 n_0 m_e c^2$ is the magnetic energy fraction. The mean of ε_B, measured in the same downstream slab as the magnetic spectrum, is $\langle \varepsilon_B \rangle \gtrsim 10^{-3}$ (Figure 4(a)), except at early times ($t\omega_p \lesssim 4000$). For $\langle \varepsilon_B \rangle \sim 10^{-3}$, particles with $\gamma \sim \gamma_0$ become marginally magnetized at a wavenumber $k_y \sim \pi/R_L(\gamma_0) \sim 0.1\, \omega_p/c$, which is near the peak of the magnetic spectrum (Figure 4(e)). This confirms that the bulk of the post-shock plasma becomes magnetized in the long-time regime. A similar finding was reported by Keshet et al. (2009) in a simulation evolved up to $t\omega_p \approx 12{,}600$.

The decaying magnetic field behind the shock evolves into a patchy configuration composed of long-lived structures with strong fields (reaching values up to $\varepsilon_B \sim 0.1$) and a background with weaker fields ($\varepsilon_B \ll 0.1$). This is reflected in the measured

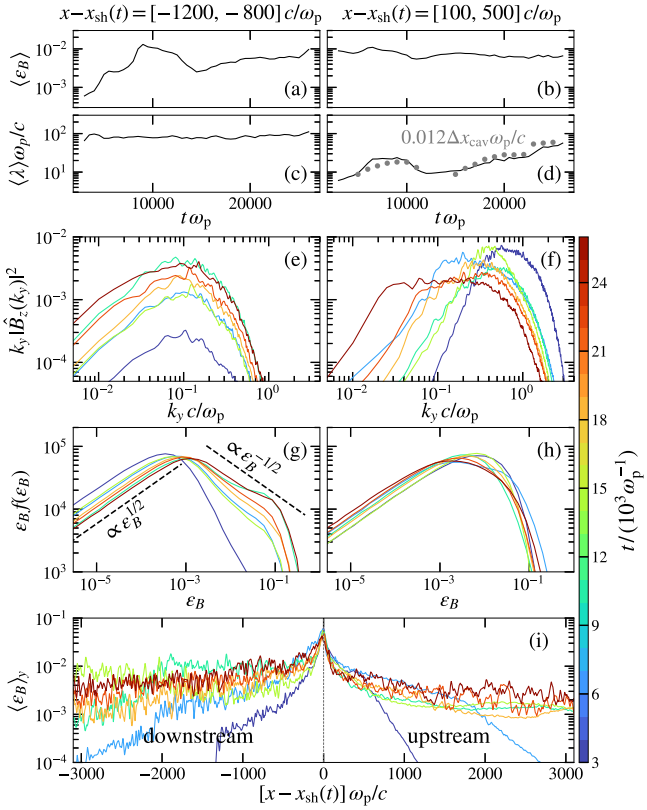

Figure 4. Evolution of the magnetic field behind (left) and ahead of (right) the shock. Panels (a)–(d) show the average magnetic energy fraction $\varepsilon_B = B_z^2/8\pi\gamma_0 n_0 m_e c^2$ and the field transverse coherence scale λ as a function of time in a fixed downstream ($x - x_{sh} = [-1200, -800]\, c/\omega_p$) and upstream ($x - x_{sh} = [100, 500]\, c/\omega_p$) slab. Panels (e)–(f) show the magnetic spectrum as a function of the transverse wavenumber k_y in the downstream (e) and upstream (f) slab (colors represent time). The spectrum is compensated by k_y to highlight the energy content per scale. In (g)–(h) we show the distribution of ε_B in the downstream (g) and upstream (h) slab. Finally, in (i) we show the transversely averaged profile of ε_B. Gray dots in panel (d) measure the distance $\Delta x_{cav}\omega_p/c$ (rescaled by a numeric prefactor ≈0.012) between the shock and the upstream location of cavity birth (see Appendix for details).

[6] We performed limited-duration electron–ion shock simulations and still found reversals of cavity polarity over time. However, the comparatively short duration of these simulations prevents us from making reliable extrapolations to longer times.

[7] We define $\lambda = \pi \int k_y^{-1} P_B(k_y) dk_y / \int P_B(k_y) dk_y$, where k_y is the transverse wavenumber and $P_B(k_y) = |\hat{B}_z(k_y)|^2$ is the 1D magnetic energy spectrum as a function of k_y.

[8] For details on the Δx_{cav} measurement see Appendix.

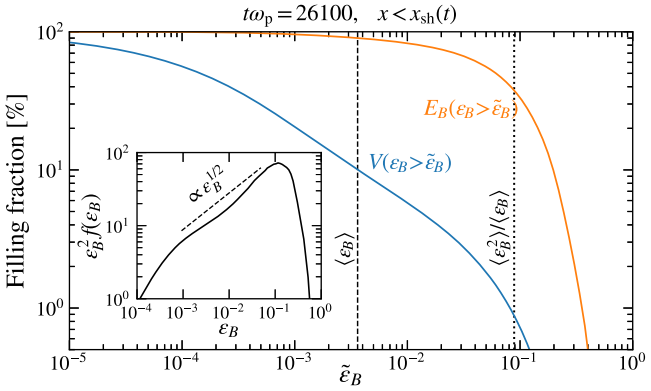

Figure 5. Fraction of the volume (blue) and of the magnetic energy (orange) contained in downstream ($x < x_{sh}$) regions with ε_B larger than a given threshold $\tilde{\varepsilon}_B$. The value $\varepsilon_{B*} = \langle \varepsilon_B^2 \rangle / \langle \varepsilon_B \rangle$ (dotted vertical line) is much greater than the mean $\langle \varepsilon_B \rangle$ (dashed line); ε_{B*} represents the typical local energy density of intense magnetic fluctuations, which make a significant contribution to the synchrotron emission (see Section 4 for details). The inset shows the magnetic energy content (arbitrary units) per logarithmic interval in ε_B.

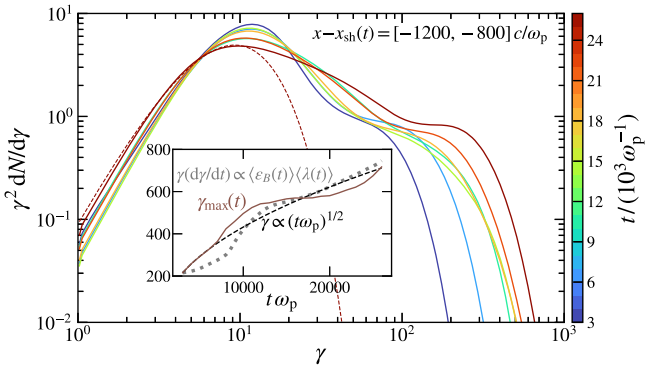

Figure 6. Evolution of the particle energy spectrum in a fixed slab behind the shock (colors represent time). The dashed red curve shows a Maxwellian fit to the low-energy part of the spectrum at the end of the simulation. The inset shows the evolution of the maximum Lorentz factor $\gamma_{max}(t)$ (brown curve) and compares the measurement with theoretical expectations for particle scattering in a magnetic field with fixed $\langle \varepsilon_B \rangle \langle \lambda \omega_p / c \rangle \approx 0.4$ (black dashed curve) vs. scattering in a field where $\langle \varepsilon_B \rangle \langle \lambda \omega_p / c \rangle$ is time evolving (gray dotted curve; see Section 5 for details). The maximum Lorentz factor is determined by the value of γ where $\gamma dN/d\gamma$ drops by 10^5 below the peak.

probability distribution of ε_B behind the shock (Figure 4(g)). The distribution (with logarithmic bins for ε_B) rises as $\varepsilon_B f(\varepsilon_B) \propto \varepsilon_B^{1/2}$ on the low-energy side toward the peak around $\varepsilon_B \sim 10^{-3}$. The 1/2 slope is consistent with a normal Gaussian distribution for the fluctuating field B_z.[9] Due to intermittent magnetic structures, the distribution does not cut off exponentially above the peak but instead transitions into a hard power-law tail with a slope around $-1/2$ (i.e., $\varepsilon_B f(\varepsilon_B) \propto \varepsilon_B^{-1/2}$). The tail extends up to $\varepsilon_B \sim 0.1$ at late times.

The relatively hard power-law tail above the peak of the probability distribution of ε_B in the downstream has an immediate implication for models of GRB afterglows. Namely, a significant fraction of the synchrotron power $P_{syn} \propto \gamma^2 \varepsilon_B$ from particles randomly sampling the downstream field will be emitted at the localized structures with $\varepsilon_B \sim 0.1$. This is demonstrated in Figure 5, which shows the fraction of magnetic energy and the volume fraction contained in downstream regions with ε_B larger than a given threshold $\tilde{\varepsilon}_B$. Roughly ~50% of the total magnetic energy, and hence of the total synchrotron power, is contributed by structures with $\varepsilon_B \gtrsim 0.1$, which occupy about ~1% of the downstream volume. That a significant fraction of the emission is contributed by the high field regions is further supported by the downstream magnetic energy distribution (inset of Figure 5, showing $\varepsilon_B^2 f(\varepsilon_B)$), which peaks near $\varepsilon_B \sim 0.1$.

Our results motivate the design of two-zone synchrotron emission models (e.g., Kumar et al. 2012; Khangulyan et al. 2021), where particles spend most of their time in low field regions but emit an order unity fraction of their synchrotron power intermittently at the locations with strong fields.[10] In the scenario implied by our simulations, the average synchrotron cooling rate (for a given γ) is determined by the mean value $\langle \varepsilon_B \rangle \ll 0.1$ of the magnetic energy fraction in the emission region (see Figure 5). In contrast, the photon energy $E_{syn} \sim \gamma^2 \hbar e B_* / m_e c$ at the peak of the emitted spectral energy distribution is set by the high field strength B_*, corresponding to the peak of $\varepsilon_B^2 f(\varepsilon_B)$ at $\varepsilon_{B*} \sim 0.1$ (inset of Figure 5), since it is the value at the peak of $\varepsilon_B^2 f(\varepsilon_B)$ that makes the single largest contribution to the emission. A convenient analytic estimate of ε_{B*} can be given based on the synchrotron-power-weighted mean of ε_B, which equals $\langle \varepsilon_B^2 \rangle / \langle \varepsilon_B \rangle$ (for fixed γ). In our simulation, indeed $\varepsilon_{B*} \sim \langle \varepsilon_B^2 \rangle / \langle \varepsilon_B \rangle \sim 0.1$. Using our analytic estimate, we can express $E_{syn} \sim \kappa^{1/2} \gamma^2 \hbar e B_{rms} / m_e c$, where $\kappa = \langle \varepsilon_B^2 \rangle / \langle \varepsilon_B \rangle^2$ is the kurtosis of the magnetic fluctuations and $B_{rms} \propto \langle \varepsilon_B \rangle^{1/2}$ is the root mean square field strength. The downstream field at the end of our simulation has $\kappa \approx 25$ on average, which is much greater than the value $\kappa = 3$ for normal Gaussian distributions. As a result, E_{syn} is raised by a factor of ~3 compared to the emission in a nonintermittent field with $\kappa = 3$. The factor could be increased further for more intermittent distributions with even larger κ. Note that the factor $\sim \kappa^{1/2}$ applies to all the emitting particles, including the highest-energy ones. If the highest-energy particles are accelerated at maximum efficiency (i.e., close to the Bohm regime with $\lambda \sim R_L(\gamma)$), the intermittency of the post-shock magnetic field enables the production of synchrotron photons beyond the classical burn-off limit (de Jager et al. 1996; Kumar et al. 2012; Khangulyan et al. 2021).

5. Particle Acceleration

Finally, we study particle acceleration at the relativistic shock. Figure 6 shows the evolution of the particle energy spectrum ($\gamma^2 dN/d\gamma$) in a fixed slab behind the shock at $x - x_{sh} = [-1200, -800]\, c/\omega_p$. Particles are accelerated to higher energies as the simulation progresses, consistent with earlier works (e.g., Spitkovsky 2008a). At select times that correlate with bursts of plasma cavity generation (around $t\omega_p \sim 10^4$ and toward the end of the simulation; see Figure 3), the particle maximum Lorentz factor in the slab grows faster than the canonical $\gamma_{max}(t) \propto (t\omega_p)^{1/2}$ scaling (Sironi et al. 2013; Plotnikov et al. 2018), expected for Weibel-mediated relativistic shocks (see inset of Figure 6). However, in between the intervals of fast acceleration lies a quiescent phase with a very slow growth of $\gamma_{max}(t)$ so that the maximum Lorentz factor still traces out a $\sim (t\omega_p)^{1/2}$ envelope on average. The

[9] If $f(b) \propto \exp(-ab^2)$, where $b = B_z/(8\pi\gamma_0 n_0 m_e c^2)^{1/2}$, $b^2 = \varepsilon_B$, and a is a constant, then $\varepsilon_B f(\varepsilon_B) \propto \varepsilon_B^{1/2} \exp(-a\varepsilon_B)$.

[10] The high field regions in our simulation amount to a small fraction of the downstream volume but contain roughly half of the total magnetic energy, which differs from the model by Kumar et al. (2012) where the fraction of energy in high field regions is small. The scenario implied by our simulations is closer to the model proposed by Khangulyan et al. (2021).

rapid changes in the instantaneous rate of particle acceleration can be largely attributed to the evolution of the magnetic field properties, which set the particle scattering frequency as $\nu_{\rm scatt} \sim \lambda c/R_{\rm L}^2 \sim \varepsilon_B(\lambda\omega_{\rm p}/c)(\gamma/\gamma_0)^{-2}\omega_{\rm p}$ (e.g., Plotnikov et al. 2011). An integration of $d\gamma/dt = 0.25\,\nu_{\rm scatt}\gamma$ (using 0.25 as an ad hoc prefactor) with a time-dependent $\langle\varepsilon_B\rangle\langle\lambda\omega_{\rm p}/c\rangle$, measured in the same downstream slab as the particle spectrum (see Figure 4), gives a result broadly consistent with the measured $\gamma_{\rm max}(t)$ (gray dotted curve in the inset of Figure 6). A reasonable overall agreement is also obtained for a fixed $\langle\varepsilon_B\rangle\langle\lambda\omega_{\rm p}/c\rangle \approx 0.4$ (dashed black curve).

A novel feature observed in the particle spectrum at late times is the formation of a distinct suprathermal component, which connects the Maxwellian peak (near $\gamma \sim \gamma_0 = 10$) to the high-energy power-law tail with an index close to -2 at late time (i.e., $dN/d\gamma \propto \gamma^{-2}$ at $100 \lesssim \gamma \lesssim 200$). The fraction of kinetic energy contained in the suprathermal component grows with time. Toward the end of the run, the suprathermal component contains about 30% of the total kinetic energy, while the high-energy tail contains stably around 10%. The suprathermal component grows at the expense of the thermal (i.e., Maxwellian) part of the distribution, which loses energy over time. It is interesting to note that the start of the high-energy tail shifts with time to higher γ, in favor of the growing suprathermal component. In this regard, let us mention that standard treatments of the first-order Fermi process at relativistic shocks (e.g., Achterberg et al. 2001) assume a discontinuous flow profile across the shock transition. However, in practice the shock develops a broad transition region over which the incoming flow gradually slows down (Lemoine et al. 2019a). The enhanced scattering due to large-scale structures and the strongly turbulent shock transition layer observed in our simulation imply that moderate energy particles ($\gamma \sim$ a few γ_0) returning into the upstream may deflect toward the downstream before being able to probe the full velocity difference between the far upstream and downstream. We speculate that these particles represent the dominant contribution to the suprathermal component, but we leave a more detailed investigation of this aspect for the future.

6. Discussion and Conclusions

We studied the evolution of a relativistic pair plasma shock propagating into an unmagnetized medium using 2D PIC simulations of unprecedented duration and size. Toward the end of our fiducial run, at $t\omega_{\rm p} \approx 26{,}100$, the shock generates magnetic structures reaching sizes of roughly $\sim$100 plasma skin depths on both sides of the shock. The size of magnetic structures generated in our simulation may be sufficient to circumvent at least some of the outstanding issues related to the modeling of GRB afterglows. For instance, it has been argued (Huang et al. 2022) that magnetic fields with a coherence scale of about $\sim$100 ion skin depths or larger are required to explain the X-ray and gamma-ray afterglow of GRB 190829A (H.E.S.S. Collaboration et al. 2021).[11] It should be noted that further evolution of the shock structure is expected beyond the time span of our simulation, as discussed below. We also find that the bulk of the post-shock plasma becomes magnetized, which slows down the magnetic field decay, as compared to early times of the shock evolution (Section 4). The probability distribution of the downstream magnetic energy fraction ε_B is intermittent; it features a hard power-law tail that extends up to $\varepsilon_B \sim 0.1$. Particles randomly sampling the downstream field spend most of their time in low field regions ($\varepsilon_B \ll 0.1$) but emit roughly half of their synchrotron power at localized magnetic structures with $\varepsilon_B \sim 0.1$.

The shock exhibits rich spatiotemporal dynamics, characterized by bursts of magnetized plasma cavity generation and oscillations in the sign of their polarity (Section 3). The cavitation instability enables the formation of large-scale magnetic structures over the turbulent shock precursor, giving rise to a broadband magnetic wavenumber spectrum at the end of the simulation (Figure 4(f)). The spectrum extends from the large energy-containing scale, characterized by the transverse coherence scale λ, down to plasma microscales $\sim c/\omega_{\rm p}$. Empirically, we find that λ measured immediately ahead of the shock is proportional to $\Delta x_{\rm cav}$, where $\Delta x_{\rm cav}$ is the distance between the shock and the upstream location of cavity birth (Figure 4(d)). Beyond the duration of our simulation, an upper limit on $\lambda(t)$ may be obtained by assuming that the distance $\Delta x_{\rm cav}$ scales linearly with the size of the precursor $\ell_{\rm prec}$, which is set by the scattering length of the highest-energy particles. The upstream scattering length $\ell_{\rm scatt} \propto \gamma^2$ (Lemoine et al. 2019b). This implies $\lambda \propto t$, assuming the shock keeps accelerating particles to higher energy as $\gamma_{\rm max}(t) \propto t^{1/2}$ (see Section 5). A linear growth rate for λ presents a potentially favorable scenario for particle acceleration; it suggests that the ratio $\lambda/R_{\rm L}(\gamma_{\rm max}) \sim \varepsilon_B^{1/2}(\lambda\omega_{\rm p}/c)(\gamma_{\rm max}/\gamma_0)^{-1}$ grows over time and converges toward the Bohm limit ($\lambda \sim R_{\rm L}(\gamma_{\rm max})$). While these results are encouraging, we caution the reader that extrapolations beyond the time span of our simulation are uncertain, and the scaling $\lambda \propto t$ is only an optimistic estimate.

The size of magnetic structures increases from early to late time also behind the shock front (Figure 1). However, quantitatively, the increase of the field coherence scale λ in the downstream is not large (see also Keshet et al. 2009). At different times and/or at different locations behind the shock, we measure typical values of the order of $\lambda \sim 100\,c/\omega_{\rm p}$ (Figure 4(c)). This scale corresponds within factors of a few to the wavenumber $k_y \sim 0.1\,\omega_{\rm p}/c$ at which the magnetic spectrum peaks (Figure 4(e)), which also happens to be the scale at which the bulk of the post-shock plasma becomes marginally magnetized. Presumably, the demonstration of a large increase of λ in the downstream requires longer duration simulations with even wider boxes. In particular, we expect the coherence scale at a fixed distance behind the shock to grow when the size of magnetic structures arriving at the shock exceeds our "typical" downstream value of $\sim 100\,c/\omega_{\rm p}$. Then, as the growing magnetic structures are transmitted into the downstream, the coherence scale will increase on both sides of the shock.

Even though the presented PIC simulation is the largest and longest of its kind, extrapolations are still required in order to make direct contact with GRB afterglows. The apparent duration of our simulation in the observer frame (Mészáros 2006) is $T_{\rm obs} \simeq T(1+z)/2\gamma_0 \approx 0.02(\gamma_0/10)^{-1}(1+z)\ \ (n_{\rm u}/1\ {\rm cm}^{-3})^{-1/2}\ (m_{\rm i}/m_{\rm e})^{1/2}$ s, where T is the duration in the (downstream) simulation frame, $n_{\rm u}$ is the upstream-frame number density of the ambient medium, z is the redshift, and $m_{\rm i}/m_{\rm e}$ is the ion-electron mass ratio ($m_{\rm i}/m_{\rm e} = 1$ in our run, but for

[11] For application to GRB shocks propagating into an electron–proton ambient medium, time and length scales should be measured in units of the ion (i.e., proton) inverse plasma frequency $\omega_{\rm pi}^{-1} = (m_{\rm i}/m_{\rm e})^{1/2}\omega_{\rm p}^{-1}$ and skin depth $c/\omega_{\rm pi} = (m_{\rm i}/m_{\rm e})^{1/2}c/\omega_{\rm p}$, where $m_{\rm i}/m_{\rm e}$ is the ion-electron mass ratio. In our simulation, $m_{\rm i}/m_{\rm e} = 1$ for computational convenience.

electron–proton shocks $m_i/m_e \approx 1836$). The transverse size of the box is $L_\perp \approx 7 \times 10^8 (n_u/1\ \mathrm{cm}^{-3})^{-1/2} (m_i/m_e)^{1/2}$ cm. Thus, our present simulation probes the shock physics over time and length scales intermediate between the plasma microscales and the global macroscales of the blast wave. In our fiducial run, spatial regions featuring cavities of a given polarity extend up to thousands of skin depths in the longitudinal direction and across the full transverse dimension of the simulation box. In a macroscopically large domain, these long-range correlations will be ultimately limited by causality. In that case, we expect the shock to form cavities of different polarity in causally disconnected regions transverse to the shock normal. Our present simulation box is too narrow to cancel out the system-wide correlations. Strictly speaking, the exact times and locations of cavity generation observed in our fiducial run represent only a particular realization of our numerical experiment. Nevertheless, the general trend of large-scale field generation is still clearly evident when considering the full time span of the simulation.

Previous PIC simulations found that as much as ∼90% of the post-shock kinetic energy is contained in thermal (i.e., Maxwellian) particles (e.g., Spitkovsky 2008a, 2008b; Martins et al. 2009). Radiative signatures of thermal electrons from GRB afterglows have been considered by a number of authors (e.g., Eichler & Waxman 2005; Giannios & Spitkovsky 2009; Ressler & Laskar 2017; Laskar et al. 2019; Warren et al. 2022). However, conclusive observational evidence for the presence of a prominent thermal component is presently lacking. In contrast to previous PIC simulations, we find that the fraction of energy contained in the thermal component drops over time. At the end of our fiducial run, the fraction of energy in the Maxwellian component drops from the "canonical" ∼90% to ∼60% (Section 5), which reduces the tension with the widely used assumption of purely nonthermal electrons in afterglow models (e.g., Sari et al. 1998; Granot & Sari 2002).

In order to evolve the shock simulations for a maximum possible duration, we focused on a pair plasma composition. Previous works showed that the physics of electron–ion relativistic unmagnetized shocks is similar to pair shocks, owing to efficient preheating of the incoming electrons (e.g., Spitkovsky 2008b; Martins et al. 2009; Sironi et al. 2013). Therefore, we expect qualitatively similar results for electron–ion shocks. A critical parameter affecting the shock evolution is the upstream magnetization of the ambient medium σ (see Sironi et al. 2013; Reville & Bell 2014; Plotnikov et al. 2018). The physical picture described here requires that the shock operates close to the $\sigma = 0$ limit. Presently, it is not well understood how small σ needs to be for the $\sigma = 0$ limit to be relevant. It has been shown that different levels of ambient magnetization can dramatically alter the behavior of collisionless shocks (e.g., Bret et al. 2017; Bret & Narayan 2018; Grošelj et al. 2022; Haggerty et al. 2022). Thus, it may be reasonably expected that only a fraction of GRBs take place in environments where a vanishingly small magnetization can be assumed, which might help explain the relatively large scatter of best-fit modeling parameters with respect to different afterglow observations.

Acknowledgments

We thank A. Vanthieghem, M. Lemoine, B. Reville, and A. Beloborodov for helpful discussions related to this work. D.G. is supported by the Research Foundation—Flanders (FWO) Senior Postdoctoral Fellowship 12B1424N. L.S. and D.G. were also supported by NASA ATP grant 80NSSC20K0565. This research was facilitated by Multimessenger Plasma Physics Center (MPPC), NSF grants PHY-2206607 and PHY-2206609. The work was supported by a grant from the Simons Foundation (MP-SCMPS-0000147, to L.S. and A.S.). L.S. acknowledges support from DoE Early Career Award DE-SC0023015. Simulations were performed on NASA Pleiades (GID: s2796, s1967, s2560). We acknowledge the OSIRIS Consortium, consisting of UCLA and IST (Lisbon, Portugal) for the use of OSIRIS and for providing access to the OSIRIS framework.

Software: OSIRIS (Fonseca et al. 2002, 2013).

Appendix
Measurement of Δx_{cav}

We employ the order parameter $Q(x) = \langle \mathrm{sgn}(-A_x) \rangle_y$ (see Figure 3) to determine the typical distance Δx_{cav} between the upstream location where the cavities are born and the shock. The distance Δx_{cav} determines the size of the cavities arriving at the shock (see Section 4). We implemented an algorithm to measure Δx_{cav} as follows. To reduce noise, we first smooth $Q(x)$ over a longitudinal scale of 100 c/ω_p, and we set $Q(x) = 0$ where $\langle \varepsilon_B \rangle_y < 10^{-6}$, since A_x cannot be reliably calculated when the magnetic fluctuations are vanishingly small. As an ad hoc condition for statistical significance we introduce $\max|Q(x > x_{\mathrm{sh}})| > \epsilon$ with $\epsilon = 0.1$. At $t\omega_p \lesssim 4000$ and during a quiescent phase around $t\omega_p \sim 12{,}000$ cavities are largely absent and $\max|Q(x > x_{\mathrm{sh}})| < \epsilon$, so that Δx_{cav} is not measured. If the condition for statistical significance is satisfied, we proceed with the calculation. Spatially separated regions of the precursor can simultaneously generate cavities of different polarity (e.g., see Figure 3(d) and the corresponding profile of $Q(x)$), but it is the cavities that make the largest contribution to the symmetry breaking, which dominate the scale of magnetic structures ultimately arriving at the shock. For this reason, we first determine the polarity of the dominant cavities as $P = \mathrm{sgn}\langle Q(x > x_{\mathrm{sh}}) \rangle_x$. We then measure Δx_{cav} as the largest distance between x_{sh} and any upstream location with $PQ(x) > \epsilon/2$. We visually inspected the magnetic field snapshots and confirmed that our method gives reasonable estimates for the upstream location $x_{\mathrm{cav}} = x_{\mathrm{sh}} + \Delta x_{\mathrm{cav}}$ of cavity birth. For reference, we show the measured locations with squares in Figure 3(f) (measurements are absent for $t\omega_p = 3000$, 13,400 because cavities are not present at a statistically significant level at those times).

ORCID iDs

Daniel Grošelj https://orcid.org/0000-0002-5408-3046
Lorenzo Sironi https://orcid.org/0000-0002-1227-2754
Anatoly Spitkovsky https://orcid.org/0000-0001-9179-9054

References

Abdalla, H., Adam, R., Aharonian, F., et al. 2019, Natur, 575, 464
Abdo, A. A., Ackermann, M., Arimoto, M., et al. 2009, Sci, 323, 1688
Achterberg, A., Gallant, Y. A., Kirk, J. G., & Guthmann, A. W. 2001, MNRAS, 328, 393
Achterberg, A., Wiersma, J., & Norman, C. A. 2007, A&A, 475, 19
Ackermann, M., Ajello, M., Asano, K., et al. 2014, Sci, 343, 42
Ackermann, M., Asano, K., Atwood, W. B., et al. 2010, ApJ, 716, 1178
Asano, K., Murase, K., & Toma, K. 2020, ApJ, 905, 105
Bell, A. R. 1978, MNRAS, 182, 147
Blandford, R., & Eichler, D. 1987, PhR, 154, 1
Blandford, R. D., & Ostriker, J. P. 1978, ApJL, 221, L29

Blinne, A., Schinkel, D., Kuschel, S., et al. 2018, CoPhC, 224, 273
Bret, A., & Narayan, R. 2018, JPlPh, 84, 905840604
Bret, A., Pe'Er, A., Sironi, L., Sądowski, A., & Narayan, R. 2017, JPlPh, 83, 715830201
Bret, A., Stockem, A., Narayan, R., & Silva, L. O. 2014, PhPl, 21, 072301
Burns, E., Svinkin, D., Fenimore, E., et al. 2023, ApJL, 946, L31
Chandra, P., & Frail, D. A. 2012, ApJ, 746, 156
Chang, P., Spitkovsky, A., & Arons, J. 2008, ApJ, 674, 378
Crowther, P. A. 2007, ARA&A, 45, 177
de Jager, O. C., Harding, A. K., Michelson, P. F., et al. 1996, ApJ, 457, 253
De Pasquale, M., Schady, P., Kuin, N. P. M., et al. 2010, ApJL, 709, L146
Drury, L. O. 1983, RPPh, 46, 973
Eichler, D., & Waxman, E. 2005, ApJ, 627, 861
Fonseca, R. A., Silva, L. O., Tsung, F. S., et al. 2002, LNCS, 2331, 342
Fonseca, R. A., Vieira, J., Fiuza, F., et al. 2013, PPCF, 55, 124011
Fried, B. D. 1959, PhFl, 2, 337
Giannios, D., & Spitkovsky, A. 2009, MNRAS, 400, 330
Granot, J., & Sari, R. 2002, ApJ, 568, 820
Grošelj, D., Sironi, L., & Beloborodov, A. M. 2022, ApJ, 933, 74
Gruzinov, A. 2001, ApJL, 563, L15
Haggerty, C. C., Bret, A., & Caprioli, D. 2022, MNRAS, 509, 2084
Haugbølle, T. 2011, ApJL, 739, L42
H.E.S.S. Collaboration, Abdalla, H., Aharonian, F., et al. 2021, Sci, 372, 1081
Honda, M., Meyer-ter-Vehn, J., & Pukhov, A. 2000, PhPl, 7, 1302
Huang, Y., Hu, S., Chen, S., et al. 2022, GCN, 32677, 1
Huang, Z.-Q., Kirk, J. G., Giacinti, G., & Reville, B. 2022, ApJ, 925, 182
Keshet, U., Katz, B., Spitkovsky, A., & Waxman, E. 2009, ApJL, 693, L127
Khangulyan, D., Aharonian, F., Romoli, C., & Taylor, A. 2021, ApJ, 914, 76
Kirk, J. G., & Reville, B. 2010, ApJL, 710, L16
Kumar, P., Hernández, R. A., Bošnjak, Ž., & Barniol Duran, R. 2012, MNRAS, 427, L40
Kumar, P., & Zhang, B. 2015, PhR, 561, 1
Laskar, T., Alexander, K. D., Margutti, R., et al. 2023, ApJL, 946, L23
Laskar, T., van Eerten, H., Schady, P., et al. 2019, ApJ, 884, 121
Lemoine, M. 2015, JPlPh, 81, 455810101
Lemoine, M., Gremillet, L., Pelletier, G., & Vanthieghem, A. 2019a, PhRvL, 123, 035101
Lemoine, M., & Pelletier, G. 2011, MNRAS, 417, 1148
Lemoine, M., Pelletier, G., Vanthieghem, A., & Gremillet, L. 2019b, PhRvE, 100, 033210
LHAASO Collaboration, Cao, Z., Aharonian, F., et al. 2023, Sci, 380, 1390
MAGIC Collaboration, Acciari, V. A., Ansoldi, S., et al. 2019a, Natur, 575, 459
MAGIC Collaboration, Acciari, V. A., Ansoldi, S., et al. 2019b, Natur, 575, 455
Martins, S. F., Fonseca, R. A., Silva, L. O., & Mori, W. B. 2009, ApJL, 695, L189
Medvedev, M. V., & Loeb, A. 1999, ApJ, 526, 697
Mészáros, P. 2002, ARA&A, 40, 137
Mészáros, P. 2006, RPPh, 69, 2259
Nishikawa, K. I., Niemiec, J., Hardee, P. E., et al. 2009, ApJL, 698, L10
Panaitescu, A., & Kumar, P. 2002, ApJ, 571, 779
Parsons, J., Spitkovsky, A., & Vanthieghem, A. 2023, arXiv:2310.12950
Perley, D. A., Cenko, S. B., Corsi, A., et al. 2014, ApJ, 781, 37
Peterson, J. R., Glenzer, S., & Fiuza, F. 2021, PhRvL, 126, 215101
Peterson, J. R., Glenzer, S., & Fiuza, F. 2022, ApJL, 924, L12
Piran, T. 2004, RvMP, 76, 1143
Plotnikov, I., Grassi, A., & Grech, M. 2018, MNRAS, 477, 5238
Plotnikov, I., Pelletier, G., & Lemoine, M. 2011, A&A, 532, A68
Ressler, S. M., & Laskar, T. 2017, ApJ, 845, 150
Reville, B., & Bell, A. R. 2014, MNRAS, 439, 2050
Sari, R., Piran, T., & Narayan, R. 1998, ApJL, 497, L17
Silva, L. O., Fonseca, R. A., Tonge, J. W., et al. 2003, ApJL, 596, L121
Sironi, L., Spitkovsky, A., & Arons, J. 2013, ApJ, 771, 54
Spitkovsky, A. 2008a, ApJL, 682, L5
Spitkovsky, A. 2008b, ApJL, 673, L39
Takamoto, M., Matsumoto, Y., & Kato, T. N. 2018, ApJL, 860, L1
Warren, D. C., Dainotti, M., Barkov, M. V., et al. 2022, ApJ, 924, 40
Weibel, E. S. 1959, PhRvL, 2, 83
Wijers, R. A. M. J., & Galama, T. J. 1999, ApJ, 523, 177
Williams, M. A., Kennea, J. A., Dichiara, S., et al. 2023, ApJL, 946, L24