

The synchrotron maser emission from relativistic magnetized shocks: dependence on the pre-shock temperature

Aliya-Nur Babul[★] and Lorenzo Sironi[★]

Department of Astronomy, Columbia University, 550 W 120th St, New York, NY 10027, USA

Accepted 2020 August 21. Received 2020 August 20; in original form 2020 June 5

ABSTRACT

Electromagnetic precursor waves generated by the synchrotron maser instability at relativistic magnetized shocks have been recently invoked to explain the coherent radio emission of fast radio bursts. By means of 2D particle-in-cell simulations, we explore the properties of the precursor waves in relativistic electron–positron perpendicular shocks as a function of the pre-shock magnetization $\sigma \gtrsim 1$ (i.e. the ratio of incoming Poynting flux to particle energy flux) and thermal spread $\Delta\gamma \equiv kT/mc^2 = 10^{-5} - 10^{-1}$. We measure the fraction f_{ξ} of total incoming energy that is converted into precursor waves, as computed in the post-shock frame. At fixed magnetization, we find that f_{ξ} is nearly independent of temperature as long as $\Delta\gamma \lesssim 10^{-1.5}$ (with only a modest decrease of a factor of 3 from $\Delta\gamma = 10^{-5}$ to $\Delta\gamma = 10^{-1.5}$), but it drops by nearly two orders of magnitude for $\Delta\gamma \gtrsim 10^{-1}$. At fixed temperature, the scaling with magnetization $f_{\xi} \sim 10^{-3} \sigma^{-1}$ is consistent with our earlier 1D results. For our reference $\sigma = 1$, the power spectrum of precursor waves is relatively broad (fractional width $\sim 1 - 3$) for cold temperatures, whereas it shows pronounced line-like features with fractional width ~ 0.2 for $10^{-3} \lesssim \Delta\gamma \lesssim 10^{-1.5}$. For $\sigma \gtrsim 1$, the precursor waves are beamed within an angle $\simeq \sigma^{-1/2}$ from the shock normal (as measured in the post-shock frame), as required so they can outrun the shock. Our results can provide physically grounded inputs for FRB emission models based on maser emission from relativistic shocks.

Key words: magnetic fields – masers – radiation mechanisms: non-thermal – shock waves – stars: neutron.

1 INTRODUCTION

Relativistic shocks are invoked as candidate sources of non-thermal particles in pulsar wind nebulae, gamma-ray bursts, and active galactic nuclei jets, and as possible accelerators of ultrahigh-energy cosmic rays (e.g. Blandford & Eichler 1987). However, relativistic shocks are generally quasi-perpendicular, i.e. with pre-shock field orthogonal to the shock direction of propagation, a configuration that – if the magnetic field is sufficiently strong – inhibits efficient particle acceleration (e.g. Begelman & Kirk 1990; Sironi, Spitkovsky & Arons 2013). While poor particle accelerators, relativistic magnetized perpendicular shocks can channel an appreciable fraction of the incoming flow energy into semicoherent electromagnetic waves propagating back into the upstream medium (hereafter, ‘precursor waves’ moving ahead of the shock). The waves are attributed to the synchrotron maser instability (Alsop & Arons 1988; Hoshino & Arons 1991). The instability is sourced by a population inversion that naturally occurs at the shock front, where a coherent ring-like distribution is constantly produced as the incoming particles gyrate in the shock-compressed field.

Recently, the discovery of fast radio bursts (FRBs; for recent reviews, see Cordes & Chatterjee 2019; Petroff, Hessels & Lorimer 2019; Platts et al. 2019) has revived the interest in the synchrotron maser. FRBs are bright (~ 1 Jy) pulses of millisecond duration

detected in the \sim GHz band. Their extremely high brightness temperature, $T_{\text{b}} \sim 10^{37}$ K, requires a coherent emission mechanism (e.g. Katz 2016). The synchrotron maser at relativistic shocks has been invoked as one of the candidate emission mechanisms within the so-called ‘magnetar scenario’ (Lyubarsky 2014; Murase, Kashiyama & Mészáros 2016; Beloborodov 2017; Waxman 2017; Beloborodov 2019; Metzger, Margalit & Sironi 2019; Margalit et al. 2020a; Margalit, Metzger & Sironi 2020b), where magnetars are invoked as FRB progenitors, a hypothesis recently confirmed by the detection of FRBs from a Galactic magnetar (Bochenek et al. 2020; Scholz & CHIME/FRB Collaboration 2020). In response to motions of the magnetar crust, the above-lying magnetosphere is violently twisted and a strongly magnetized pulse is formed, which propagates away through the magnetar wind. In the shock maser scenario, the FRB is generated at ultrarelativistic shocks resulting from the collision of the magnetized pulse with the steady wind that is produced by the magnetar spin-down luminosity or by the cumulative effect of earlier flares.

The fundamental properties of the precursor waves generated by the synchrotron maser – i.e. their efficiency, power spectrum, angular distribution, and polarization – can be quantified with self-consistent particle-in-cell (PIC) simulations. PIC simulations of relativistic magnetized shocks focusing on the synchrotron maser emission have been performed both for electron–positron plasmas (Langdon, Arons & Max 1988; Gallant et al. 1992; Sironi & Spitkovsky 2009; Iwamoto et al. 2017, 2018; Plotnikov, Grassi & Grech 2018; Plotnikov & Sironi 2019) and electron–proton or electron–positron–proton plasmas (Hoshino et al. 1992; Amato & Arons 2006;

[★] E-mail: avb2131@columbia.edu (A-NB); lsironi@astro.columbia.edu (LS)

Lyubarsky 2006; Hoshino 2008; Stockem et al. 2012; Iwamoto et al. 2019).

These works are primarily focused on low magnetizations $\sigma \lesssim 1$, where σ is the ratio of upstream Poynting flux to kinetic energy flux. On the other hand, FRBs are expected to originate from extreme environments where the energy content of the plasma is dominated by magnetic fields, as in magnetar winds. In Plotnikov & Sironi (2019), we performed 1D PIC simulations of electron–positron shocks and investigated how the properties of the synchrotron maser depend on the flow magnetization, in the $\sigma \gtrsim 1$ regime most relevant for FRB sources. We found that the shock converts a fraction $f_\xi \approx 2 \times 10^{-3} \sigma^{-1}$ of the total incoming energy into the precursor waves, as measured in the post-shock (downstream) frame. At $\sigma \gtrsim 1$, we showed that the shock structure displays two solitons separated by a cavity, and the peak frequency of the spectrum corresponds to an eigenmode of the cavity. We also found that the efficiency and spectrum of the precursor waves do not depend on the bulk Lorentz factor of the pre-shock flow.

The results in Plotnikov & Sironi (2019) were obtained assuming that the pre-shock flow has small thermal spread, $\Delta\gamma \equiv kT/mc^2 = 10^{-4}$. In fact, with the exception of the study by Amato & Arons (2006) – who focused on non-thermal lepton acceleration in pair-proton plasmas, rather than on the properties of the precursor waves – all prior studies were conducted in the limit of negligible upstream temperatures. In this work, we discuss the dependence of the precursor waves generated by the synchrotron maser on the upstream temperature, by means of 2D PIC simulations of relativistic magnetized electron–positron shocks. We focus on magnetically dominated plasmas ($\sigma = 1$ and 3) and explore thermal spreads in the range $\Delta\gamma = 10^{-5} - 10^{-1}$. All our simulations are evolved for sufficiently long ($\gtrsim 4000 \omega_p^{-1}$) so that the properties of the precursor waves, such as their Poynting flux and power spectrum, attain a steady state. At fixed magnetization, we find that the efficiency f_ξ is nearly independent of temperature as long as $\Delta\gamma \lesssim 10^{-1.5}$ (with only a modest decrease of a factor of three from $\Delta\gamma = 10^{-5}$ to $\Delta\gamma = 10^{-1.5}$), but it drops by nearly two orders of magnitude for $\Delta\gamma \gtrsim 10^{-1}$ (the drop occurs at the same $\Delta\gamma$ regardless of the flow Lorentz factor). For our reference $\sigma = 1$, the power spectrum of precursor waves is relatively broad (fractional width $\sim 1-3$) for cold temperatures, whereas it shows narrow line-like features with fractional width ~ 0.2 for $10^{-3} \lesssim \Delta\gamma \lesssim 10^{-1.5}$. For $\sigma \gtrsim 1$, the precursor waves are beamed within an angle $\simeq \sigma^{-1/2}$ from the shock normal (as measured in the post-shock frame), as required so they can outrun the shock.

The paper is organized as follows. In Section 2, we present the numerical method and the simulation set-up. We then discuss the main results of our investigation, as regard to shock structure (Section 3), precursor efficiency (Section 4) and beaming and power spectrum (Section 5). We summarize our findings in Section 6 and discuss their astrophysical implications.

2 SIMULATION SET-UP

We use the 3D electromagnetic PIC code TRISTAN-MP (Spitkovsky 2005) to perform simulations of relativistic magnetized shocks in pair plasmas. We perform simulations in 2D spatial domains, but all three components of particle velocities, electric currents, and electromagnetic fields are retained.

The simulations are performed in the post-shock frame. The upstream flow, consisting of electrons and positrons, drifts in the $-\hat{x}$ direction with speed $-\beta_0 \hat{x}$, where $\beta_0 = (1 - 1/\gamma_0^2)^{1/2}$. For the simulations presented in this paper, we employ $\gamma_0 = 10$, but we have

verified that larger values of γ_0 (up to $\gamma_0 = 80$) do not change our conclusions (see also Plotnikov & Sironi 2019, for an investigation of the dependence on γ_0 with 1D simulations). The incoming flow reflects off a wall at $x = 0$. The shock is formed by the interaction of the incoming and reflected flows and propagates along $+\hat{x}$. The upstream temperature is cast in terms of the thermal spread $\Delta\gamma \equiv kT/mc^2$, which we vary from $\Delta\gamma = 10^{-5}$ up to $\Delta\gamma = 10^{-1}$. Here, m is the electron mass and c is the speed of light.

The pre-shock plasma carries a frozen-in magnetic field $\mathbf{B}_0 = B_0 \hat{z}$ orthogonal to the xy plane of our simulations, and the associated motional electric field $\mathbf{E}_0 = -\beta_0 B_0 \hat{y} \equiv -E_0 \hat{y}$. Our field configuration parallels the one employed by Iwamoto et al. (2017), and corresponds to a perpendicular shock with out-of-plane upstream field. The magnetic field strength is parametrized via the magnetization, i.e. the ratio of upstream Poynting flux to kinetic energy flux:

$$\sigma = \frac{B_0^2}{4\pi\gamma_0 N_0 m c^2} = \left(\frac{\omega_c}{\omega_p}\right)^2 = \left(\frac{c/\omega_p}{r_L}\right)^2, \quad (1)$$

where N_0 is the number density of the upstream plasma (including both species), $\omega_c = eB_0/\gamma_0 mc$ is the Larmor frequency, $\omega_p = (4\pi N_0 e^2/\gamma_0 m)^{1/2}$ is the plasma frequency, c/ω_p is the plasma skin depth, and $r_L = \gamma_0 mc^2/eB_0$ is the Larmor radius of particles with Lorentz factor γ_0 . Here, e is the electron charge. We explore two values of magnetization, $\sigma = 1$ and 3. We use the plasma skin depth c/ω_p as our unit of length and the inverse plasma frequency ω_p^{-1} as our unit of time.

We employ periodic boundary conditions in the y -direction. The incoming plasma is injected through a ‘moving injector,’ which moves along $+\hat{x}$ at the speed of light. The simulation box is expanded in the x -direction as the injector approaches the right boundary of the computational domain. This permits us to save memory and computing time, while following the evolution of all the upstream regions that are causally connected with the shock (for details see e.g. Spitkovsky 2005; Sironi & Spitkovsky 2009). Over time, the distance between the shock and the injector increases, and the incoming flow might suffer from the so-called numerical Cerenkov instability (e.g. Dieckmann et al. 2006). By employing a fourth-order spatial stencil for Maxwell’s equations (Greenwood et al. 2004), we find no evidence of numerical Cerenkov instability within the time span covered by our simulations.

The leftmost edge of the downstream region, which is a conducting boundary for electromagnetic fields and a reflecting wall for particles (hereafter, the ‘wall’), is initially located at $x = 0$. The focus of this work is on upstream-propagating waves generated by the shock, rather than on the properties of the shocked plasma. In order to save memory and computing time, we choose to periodically jump the wall towards the shock, such that the average speed of the wall is ~ 5 per cent lower than the shock speed and the wall always stays safely behind the shock (by at least a few tens of c/ω_p). After every jump, we enforce conducting boundary conditions for the electromagnetic fields at the new position of the wall, and we discard particles to the left of the wall. By performing a few tests without the ‘jumping wall,’ i.e. retaining the whole downstream region, we have verified that this strategy does not artificially affect any property of the precursor waves.¹

¹It may be argued that, by removing a significant portion of the downstream region, one may inhibit the downstream plasma from, e.g. relaxing to isotropy (which in turn would affect the shock speed). However, our choice of a perpendicular out-of-plane field is by itself suppressing relaxation to full

We now describe the numerical parameters used in our work. We employ a high spatial resolution, with $c/\omega_p = 100$ cells. This ensures that the high frequency/wavenumber part of the power spectrum of precursor waves (i.e. $kc/\omega_p \gg 1$) is properly captured (Iwamoto et al. 2017). A few tests with a lower spatial resolution of $c/\omega_p = 50$ cells have shown good agreement with our production runs, so a resolution of $c/\omega_p = 50$ cells might also be sufficient. The transverse size of the box is 1440 cells, corresponding to $\sim 14 c/\omega_p$. This is sufficient to capture genuine 2D effects in the properties of the shock and of the precursor waves (e.g. filamentation of the upstream density, see Section 3). Experiments with even larger boxes give essentially the same results.

The numerical speed of light is 0.45 cells time-step $^{-1}$. We evolve our simulations for a few thousands of ω_p^{-1} , which is sufficient to study the steady-state properties of precursor emission. Our longest simulations have been run for $\sim 6000 \omega_p^{-1}$, corresponding to more than 1.3 million time-steps.

Our 2D simulations are typically initialized with $N_0 = 16$ particles per cell (including both species) for $\sigma = 1$ and 4 particles per cell for $\sigma = 3$. For $\sigma = 1$, we have also performed simulations with $N_0 = 4$ for all the temperature values we investigated, finding excellent agreement with our reference $N_0 = 16$ cases (in fact, in Figs 7 and 8 we use $N_0 = 4$ simulations). The simulations with the hottest upstream plasma, $\Delta\gamma = 10^{-1}$, employ 32 particles per cell (for both $\sigma = 1$ and 3), since the precursor emission is very weak, and so harder to properly capture (see Section 4). For this temperature, we have checked that simulations with even larger $N_0 = 128$ give the same results. In order to further reduce numerical noise in the simulations, we filter the electric current deposited to the grid by the particles, effectively mimicking the role of a larger number of particles per cell (Spitkovsky 2005; Belyaev 2015). We typically apply $N_{sm} = 20$ passes of a binomial three-point low-pass digital filter at every time-step. In Appendix A, we show the effects of N_0 and N_{sm} on the precursor wave spectrum.

In addition to 2D simulation, which constitute the bulk of this paper, we have also performed a suite of 1D simulations with $\sigma = 1$ and the same range of temperatures as in 2D. In 1D simulations, we typically employ $N_0 = 40$ particles per cell, a spatial resolution of $c/\omega_p = 112$ cells and a numerical speed of light of 0.5 cells time-step $^{-1}$. For the hottest temperature $\Delta\gamma = 10^{-1}$, the number of particles per cell is increased to $N_0 = 400$.

3 SHOCK STRUCTURE

In Figs 1, 2, and 3, we present the 2D shock structure from our simulations with $\sigma = 1$ and three values of upstream thermal spread, respectively, $\Delta\gamma = 10^{-5}$, 10^{-2} , and 10^{-1} . From top to bottom, in each figure we present the transverse magnetic field B_z/B_0 , the transverse electric field E_y/E_0 , the particle number density in units of the upstream density N_0 , and the particle number density averaged along the y -direction.

The shock front is at $x = x_{sh}$. The upstream flow is on the positive side ($x - x_{sh} > 0$) and the downstream plasma is on the negative side ($x - x_{sh} < 0$). The existence of a well-developed shock is confirmed by the jump in number density and in B_z at the front location. The shock front itself exhibits a soliton-like structure, as revealed by the density spikes in panels (d) at $x - x_{sh} \sim 0$ (see e.g. Alsop & Arons 1988). The density spike in the soliton is higher for colder plasmas

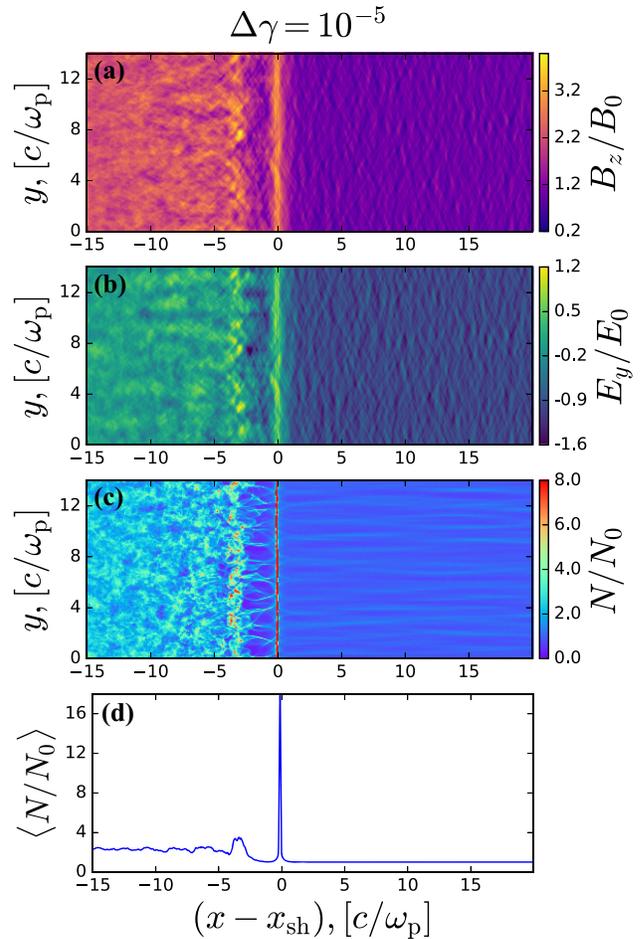


Figure 1. Shock structure from the 2D PIC simulation with $\sigma = 1$ and $\Delta\gamma = 10^{-5}$ at $\omega_p t = 2000$, when the precursor emission has reached a steady state. We focus on the vicinity of the shock. We present (a) the transverse magnetic field B_z/B_0 ; (b) the transverse electric field E_y/E_0 ; (c) the particle number density, in units of the upstream density N_0 ; (d) the particle number density averaged along the y -direction.

[compare panels (d) among the three figures], as derived analytically by Chiueh & Lai (1991). In the soliton, the incoming particles gyrate around the compressed magnetic field and form a semicoherent ring in momentum space. As shown in Figs 4, 5, and 6, where we plot, for different values of $\Delta\gamma$, the $\gamma\beta_x - \gamma\beta_y$ momentum space of particles populating the density spike, the thickness of the ring depends on the pre-shock temperature. A cold well-defined ring appears for low temperatures ($\Delta\gamma = 10^{-5}$ in Fig. 4), whereas the centre of the ring is nearly filled with particles for hot flows ($\Delta\gamma = 10^{-1}$ in Fig. 6). The radius of the ring is $\sim \gamma_0\beta_0 \sim 10$, corresponding to the bulk four-velocity of incoming particles.

The synchrotron maser instability, and the resulting precursor waves, is believed to be sourced by the population inversion in the ring (Alsop & Arons 1988; Hoshino & Arons 1991).² Such a population inversion tends to disappear for hot flows, as shown in Fig. 6. We then expect that the synchrotron maser emission will become inefficient in hot plasmas (see also Amato & Arons 2006). As a result of the synchrotron maser instability, a train of semicoherent

isotropy (the particles only isotropize in the xy plane perpendicular to the field), regardless of whether we employ the jumping wall or not.

²We remark that the continuous flow of plasma through the shock ensures that the population inversion is steadily maintained.

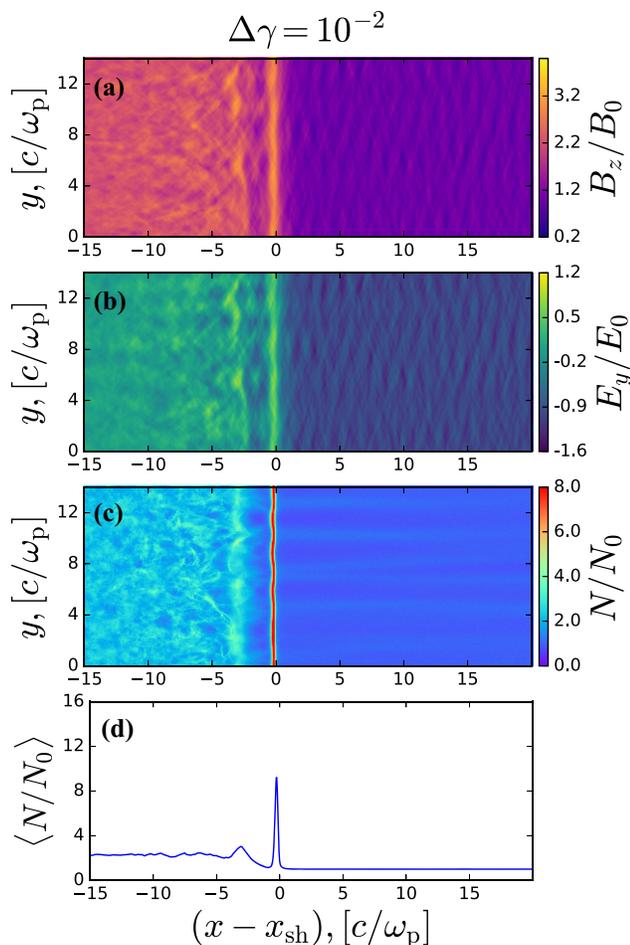


Figure 2. Same as in Fig. 1, but for a 2D simulation with $\sigma = 1$ and upstream plasma temperature of $\Delta\gamma = 10^{-2}$ at time $\omega_p t = 2000$.

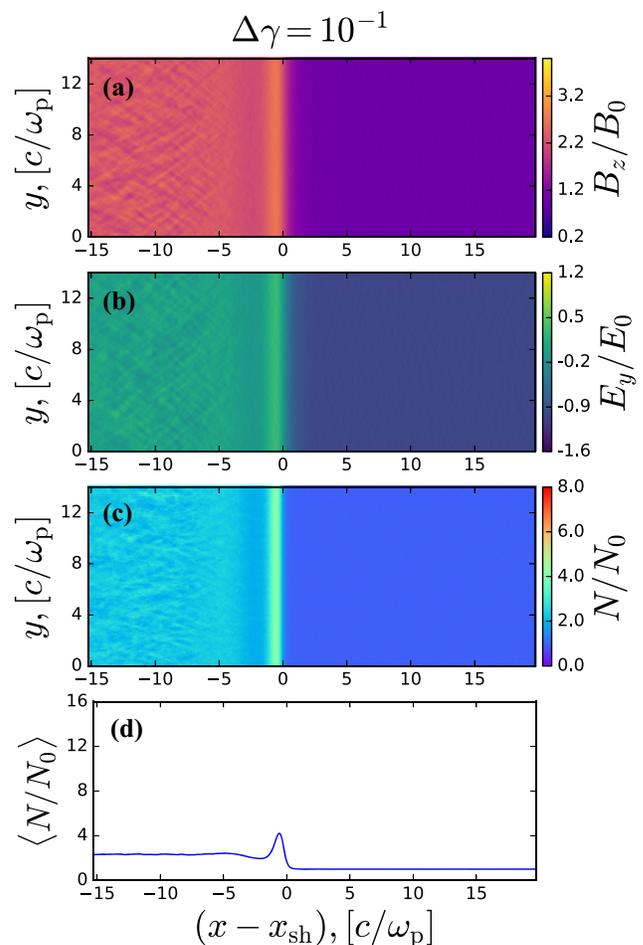


Figure 3. Same as in Fig. 1, but for a 2D simulation with $\sigma = 1$ and upstream plasma temperature of $\Delta\gamma = 10^{-1}$ at time $\omega_p t = 1500$.

large-amplitude electromagnetic precursor waves is emitted towards the upstream, as shown in the B_z/B_0 and E_y/B_0 panels, for cold ($\Delta\gamma = 10^{-5}$ in Fig. 1) and moderate ($\Delta\gamma = 10^{-2}$ in Fig. 2) temperatures. As expected, no evidence of precursor waves is seen in hot plasmas ($\Delta\gamma = 10^{-1}$ in Fig. 3).

When the precursor emission is efficient, electromagnetic waves are seen not only in the upstream region ($x > x_{\text{sh}}$), but also right behind the leading soliton, in the density cavity at $-2c/\omega_p \lesssim x - x_{\text{sh}} \lesssim 0$. As discussed in Plotnikov & Sironi (2019), this cavity is a peculiarity of $\sigma \gtrsim 1$ shocks. It plays an important role in setting the properties of precursor waves since the peak frequency of the wave spectrum is observed to correspond to an eigenmode of the cavity, i.e. the precursor waves might be resonantly amplified by the density cavity. The hot case in Fig. 3 does not display such a density cavity, and in fact its precursor emission is strongly suppressed (see Section 4). For $\sigma \gtrsim 1$, the precursor waves appear to be generated by an oscillating current localized near the downstream side of the cavity (at $x \sim x_{\text{sh}} - 2c/\omega_p$). This is generally not accounted for within the standard description of the synchrotron maser instability (Alsop & Arons 1988; Hoshino 2001). A characterization of the cavity, and its role in setting the oscillating current that ultimately drives the precursor waves, is left for future work.

The wave vector \mathbf{k} of the precursor waves is nearly aligned with the shock direction of propagation. The fluctuating magnetic field is along z (i.e. along the same direction as the upstream field \mathbf{B}_0), and

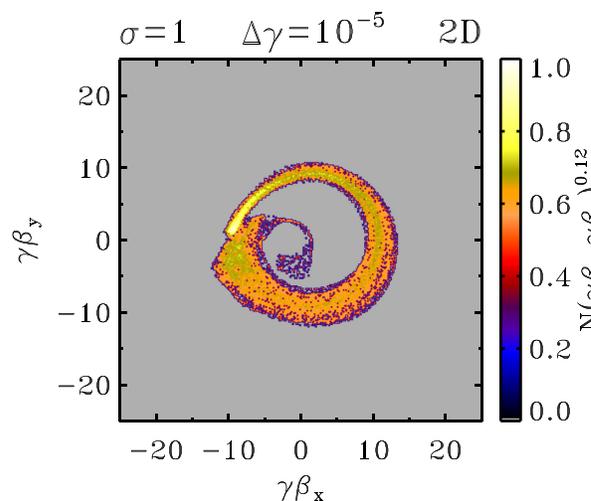


Figure 4. Momentum space $\gamma\beta_x - \gamma\beta_y$ of positrons, from the 2D simulation with $\Delta\gamma = 10^{-5}$. The particles are selected at time $\omega_p t = 2000$ to be located near the soliton-like structure at the shock front, in the range . The histogram is normalized such that $N(\gamma\beta_x, \gamma\beta_y) = 1$ in the pixel with the highest value, and the colour scale is stretched with 0.12 power to emphasize weak phase space structures.

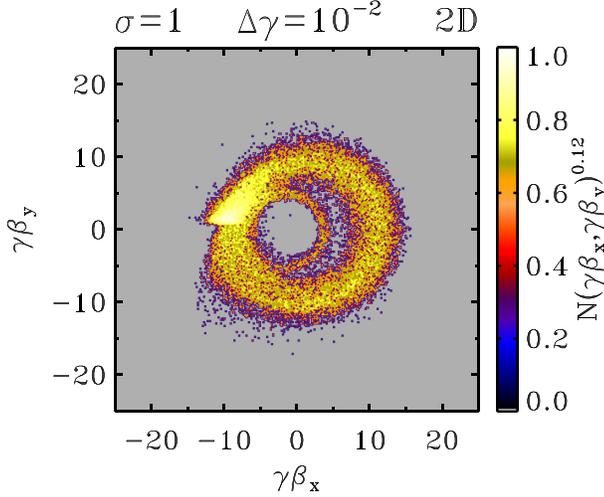


Figure 5. Same as in Fig. 4, but for the 2D simulation with $\Delta\gamma = 10^{-2}$ at time $\omega_p t = 2000$.

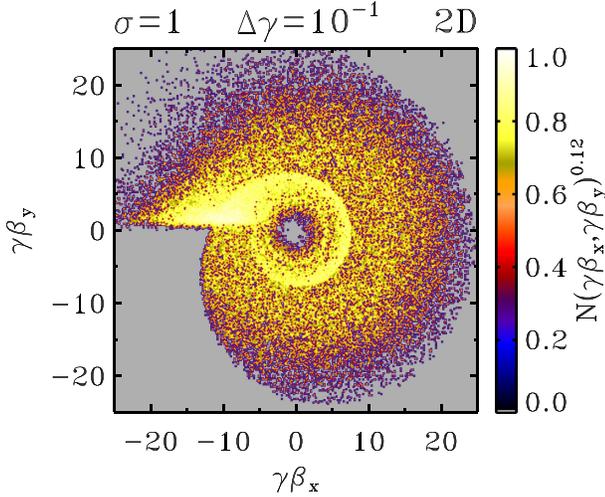


Figure 6. Same as in Fig. 4, but for the 2D simulation with $\Delta\gamma = 10^{-1}$ at time $\omega_p t = 1500$.

the fluctuating electric field is perpendicular to both \mathbf{k} and \mathbf{B}_0 . The wave is then linearly polarized and identified with the extraordinary mode (X-mode). We remark that 2D simulations with out-of-plane fields do not allow for the excitation of the ordinary mode (O-mode). Iwamoto et al. (2018) performed 2D simulations of weakly magnetized shocks ($\sigma \lesssim 1$) with in-plane upstream fields and showed that O modes are stronger than X modes at low magnetizations ($\sigma \lesssim 10^{-2}$), but they become weaker as σ increases. A detailed 3D investigation of the dependence of the precursor emission (in terms of both O and X modes) on the pre-shock temperature will be presented in a forthcoming study.

Figs 1 and 2 anticipate that the precursor wave spectrum in cold plasmas peaks at higher frequencies/wavenumbers than in warmer plasmas (see Section 5 below). This is suggested by comparing panels (a) between Figs 1 and 2. Smaller scale structures in B_z are seen for colder plasmas, i.e. the power contained in high-frequency fluctuations is larger in colder plasmas.

Whenever the precursor emission is efficient (i.e. in Figs 1 and 2, but not 5), filamentary structures are observed in the density (panels c), elongated along the shock direction of propagation. Following Iwamoto et al. (2017) and Plotnikov et al. (2018), we attribute these structures to the self-focusing and filamentation of the high-amplitude electromagnetic wave when it propagates through the upstream plasma, as studied analytically in electron–proton plasmas by Max, Arons & Langdon (1974) and Drake et al. (1974). The density filaments in the warm case (Fig. 2) appear less sharp than in the cold case (Fig. 1), suggesting that the precursor waves are not interacting as strongly with the upstream flow. The dominant wavelength is also larger for the warm case than for the cold case, in agreement with the analytical results of Drake et al. (1974) (which, however, were obtained for an electron–proton plasma).

4 PRECURSOR EFFICIENCY

In this section, we quantify how the wave efficiency depends on the flow temperature and magnetization. We measure the wave intensity in a region between $5 c/\omega_p$ and $30 c/\omega_p$ ahead of the shock front.³ This region is far enough from the shock not to be affected by the front structure itself, and it contains a large number of precursor wavelengths so that we can obtain a solid measure of the precursor average properties. The wave intensity is then calculated as the spatial average

$$\langle \delta B_z^2 \rangle = \langle (B_z - B_0)^2 \rangle, \quad (2)$$

and we define the normalized wave energy as

$$\xi_B = \frac{\langle \delta B_z^2 \rangle}{B_0^2}. \quad (3)$$

We have verified that in all our simulations $\langle \delta B_z^2 \rangle / B_0^2 \simeq \langle \delta E_y \delta B_z \rangle / (E_0 B_0)$, i.e. the parameter ξ_B also quantifies the ratio of wave Poynting flux to incoming Poynting flux. Here, $\delta E_y = E_y + \beta_0 B_0 = E_y + E_0$.

In Fig. 7, we show for different temperatures the time evolution of the normalized wave energy, for 2D simulations with $\sigma = 1$ (see legend). Following a transient, all the curves reach a steady state at $\omega_p t \gtrsim 1000$. The precursor efficiency is nearly independent of temperature as long as $\Delta\gamma \lesssim 10^{-1.5}$ (with only a modest decrease of a factor of three from $\Delta\gamma = 10^{-5}$ to $\Delta\gamma = 10^{-1.5}$), but between $\Delta\gamma = 10^{-1.5}$ and $\Delta\gamma = 10^{-1}$ it drops by nearly two orders of magnitude (red curve). We have confirmed this result with 3D simulations (Sironi et al., in preparation). By running a 2D simulation with $\gamma_0 = 40$ and $\Delta\gamma = 10^{-1}$, we have checked that the efficiency drop occurs at the same $\Delta\gamma$ regardless of the flow Lorentz factor.

The steady-state values of the normalized wave energy are shown as a function of temperature in the top panel of Fig. 8. There, we present results from 1D and 2D simulations with $\sigma = 1$ (the blue and black points, respectively) and from 2D simulations with $\sigma = 3$ (the red points). All the values are extracted from a time range when the precursor has achieved a steady state, more specifically at $\omega_p t \gtrsim 2000$ for $\sigma = 1$ and $\omega_p t \gtrsim 3000$ for $\sigma = 3$. As shown in the top panel of Fig. 8, 1D results for $\sigma = 1$ generally follow the same trend as our reference 2D simulations (compare the blue and black lines), with two notable differences. At very high temperatures ($\Delta\gamma = 10^{-1}$), the drop in wave energy is much more dramatic in 2D than in 1D, by roughly one order of magnitude. At cold temperatures

³Our results do not appreciably change if the region is extended further upstream.

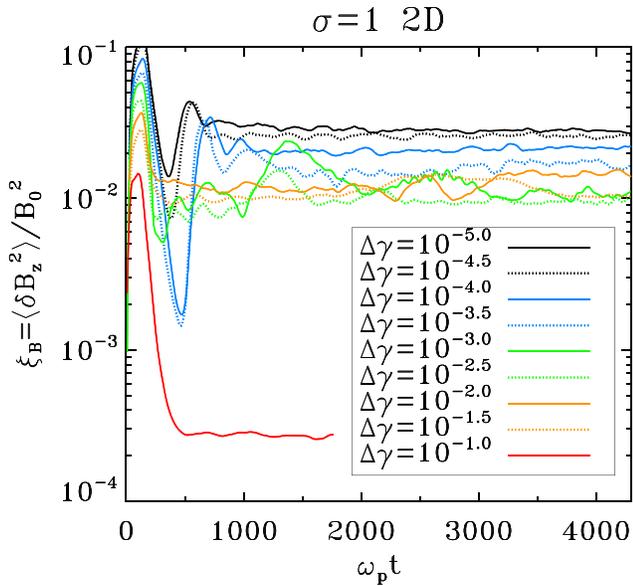


Figure 7. Time evolution of the normalized precursor wave energy, $\xi_B = \langle \delta B_z^2 \rangle / B_0^2$, for different values of the upstream thermal spread $\Delta\gamma = kT/mc^2$, as indicated in the legend. The precursor wave energy was extracted from a $25 c/\omega_p$ -wide slab located at $5 c/\omega_p < x - x_{\text{sh}} < 30 c/\omega_p$. The time evolution of the normalized Poynting flux associated with the precursor waves, $\langle \delta E_y \delta B_z \rangle / E_0 B_0$, is indistinguishable. All the curves in this figure refer to simulations with $N_0 = 4$, apart from the red line ($\Delta\gamma = 10^{-1}$) that uses $N_0 = 32$.

($\Delta\gamma \lesssim 10^{-5.5}$), the precursor waves are stronger in 1D than in 2D. We attribute this difference in efficiency at low temperatures to the longitudinal heating induced in 2D by the filamentation mode, which is absent in 1D (see also Appendix B, for the effect of longitudinal dispersion on the precursor strength). Note that the 1D simulations presented by Plotnikov & Sironi (2019) employed $\Delta\gamma = 10^{-4}$, for which 1D and 2D results only differ by ~ 50 per cent.

As regard to the dependence on magnetization, the top panel of Fig. 8 shows that the normalized wave energy is roughly the same between $\sigma = 1$ and $\sigma = 3$: for cold and moderate temperatures ($\Delta\gamma \lesssim 10^{-1.5}$), ξ_B is a few percent, while it drops by nearly two orders of magnitude for $\Delta\gamma \gtrsim 10^{-1}$. In the range of temperatures where the precursor is efficient, the minimum of the normalized wave energy is attained for $\Delta\gamma \sim 10^{-2.5}$ in $\sigma = 1$ and for $\Delta\gamma \sim 10^{-3.5}$ in $\sigma = 3$. We shall use these values to roughly distinguish between cold cases and warm cases, which, as shown in Section 5 for $\sigma = 1$, display different spectral properties.

From the normalized wave energy in the top panel of Fig. 8, we can extract the so-called strength parameter:

$$a = \frac{e \delta E_y}{mc\omega}. \quad (4)$$

At $\sigma \gtrsim 1$, where the typical wave frequency in cold plasmas is $\omega \sim 3\sigma^{1/2}\omega_p$ (Plotnikov & Sironi 2019), we find that $a \sim 0.3 (\xi_B/10^{-2})^{1/2} (\gamma_0/10)$. The strength parameter could also be measured directly from the transverse motion (along y) of the upstream particles in the field of the wave, since the oscillations in $\gamma\beta_y$ are directly related to the wave strength parameter (e.g. Lyubarsky 2006; Iwamoto et al. 2017). Notice that we expect the particle transverse oscillations to become relativistic at $\gamma_0 \gtrsim 30$. However, we have performed 2D simulations with $\sigma = 1$ and $\Delta\gamma = 10^{-4}$ for pre-shock Lorentz factors up to $\gamma_0 = 80$ and we do not find that this changes

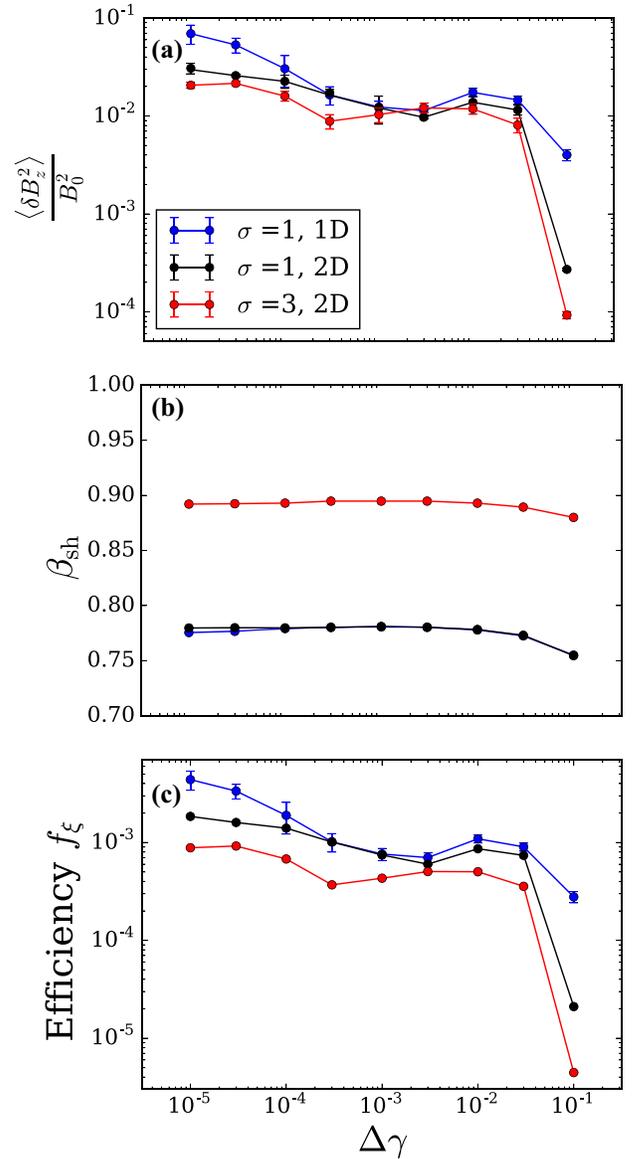


Figure 8. Dependence of the steady-state precursor properties on magnetization, temperature, and dimensionality of the simulation domain. In each panel, the red points refer to 2D simulations with $\sigma = 1$, the blue points to 1D simulations with $\sigma = 1$, and the red points to 2D simulations with $\sigma = 3$ (see legend). 2D simulations employ $N_0 = 4$, whereas 1D simulations use $N_0 = 40$, apart from the hottest temperature ($\Delta\gamma = 10^{-1}$) where 2D runs have $N_0 = 32$ and 1D runs use $N_0 = 400$. The data and error bars in each panel are the mean and standard deviation in the time range $\omega_p t \geq 2000$ for $\sigma = 1$ and $\omega_p t \geq 3000$ for $\sigma = 3$. Top panel: normalized precursor wave energy $\xi_B = \langle \delta B_z^2 \rangle / B_0^2$. In all cases, the values of the time-averaged normalized Poynting flux $\langle \delta E_y \delta B_z \rangle / E_0 B_0$ are indistinguishable. Middle panel: shock velocity in the downstream frame, in units of the speed of light. Bottom panel: precursor efficiency f_ξ measured in the downstream frame, as defined in equation (7).

the precursor dynamics or efficiency, in agreement with our earlier 1D results (see fig. 8 in Margalit et al. 2020b).

The second panel in Fig. 8 shows the shock velocity in units of the speed of light, as measured in the downstream frame. Its dependence on dimensionality and temperature (in the regime $\Delta\gamma \lesssim 10^{-1}$ we have explored) is minimal. The marginal reduction at large $\Delta\gamma$ is expected based on MHD jump conditions. The fact that for $\sigma = 1$, 1D shocks

in very cold plasmas are slightly slower than 2D shocks mirrors the different efficiency in precursor emission: in 1D the precursor takes away a larger fraction of the plasma energy, thereby slowing down the shock. The dependence on magnetization follows the expectation of MHD jump conditions. Assuming the adiabatic index of a relativistic gas with two degrees of freedom ($\gamma_{\text{ad}} = 3/2$; in fact, our particles only isotropize in the xy plane orthogonal to the magnetic field), the dimensionless four-velocity of ultrarelativistic magnetized (i.e. $\gamma_0 \gg 1$ and $\sigma \gg 1$) shocks is expected to be $\gamma_{\text{sh}}\beta_{\text{sh}} = (5\sigma/4 + 7/20)^{1/2}$ (e.g. Pétri & Lyubarsky 2007; Plotnikov et al. 2018). At $\sigma \gg 1$, this scales as $\gamma_{\text{sh}}\beta_{\text{sh}} \propto \sigma^{1/2}$. As we discuss in Section 5, this has important implications for the beaming of the precursor emission.

The data in the first and second panels are used to compute the wave efficiency f_{ξ} shown in the third panel. This is defined as the fraction of incoming total energy (electromagnetic and kinetic) that is converted into precursor wave energy. In the downstream frame of the simulations, the energy that has flown into the shock per unit area up to time t is

$$E_{\text{in}} = \gamma_0(1 + \sigma)N_0mc^2(\beta_0 + \beta_{\text{sh}})ct, \quad (5)$$

where the flux factor $(\beta_0 + \beta_{\text{sh}})$ accounts for the fact that the shock is moving towards the upstream. The energy converted into precursor waves per unit shock area is

$$E_{\text{out}} = \frac{\langle \delta B_z^2 \rangle}{4\pi}(1 - \beta_{\text{sh}})ct, \quad (6)$$

where we have assumed that the whole region between the shock and the leading edge of the precursor (moving at c) is occupied by precursor waves with uniform energy density. The efficiency is then

$$f_{\xi} = \frac{E_{\text{out}}}{E_{\text{in}}} = \xi_B \left(\frac{\sigma}{1 + \sigma} \right) \left(\frac{1 - \beta_{\text{sh}}}{\beta_0 + \beta_{\text{sh}}} \right). \quad (7)$$

Notice that in the $\sigma \gg 1$ limit this scales as $f_{\xi} \propto \xi_B \sigma^{-1}$. In the shock maser scenario for FRBs, this quantifies the fraction of blast wave energy that is converted into precursor wave energy (i.e. the candidate FRB). The ratio of precursor power to blast wave power can be obtained by accounting for the duration of the precursor emission, which in the limit $\sigma \gg 1$ is a factor of $\sim \sigma^{-1}$ shorter than the blast wave ejection duration. It follows that in the limit of high magnetizations the ratio of emitted precursor power to blast wave power is $\sim \xi_B$.

The bottom panel of Fig. 8 shows that the precursor efficiency in 2D simulations scales as $f_{\xi} \sim 10^{-3}\sigma^{-1}$ as long as $\Delta\gamma \lesssim 10^{-1.5}$, whereas it drops abruptly to $f_{\xi} \lesssim 10^{-5}$ for $\Delta\gamma \gtrsim 10^{-1}$.

5 PRECURSOR BEAMING AND POWER SPECTRUM

To characterize the spectral properties of precursor waves, we have constructed the 2D wavenumber spectrum $P_{2D}(k_x, k_y) \propto |\delta \vec{B}_z(k_x, k_y)|^2$, by taking the Fourier transform $\delta \vec{B}_z(k_x, k_y)$ of the fluctuating magnetic field $\delta B_z(x, y)$. Our spectra are computed in the post-shock frame, by extracting $\delta B_z(x, y)$ in the same region ahead of the shock ($5c/\omega_p < x - x_{\text{sh}} < 30c/\omega_p$) where we compute the precursor efficiency. We will show both the 2D power spectrum $P_{2D}(k_x, k_y)$ and the k_y -integrated 1D power spectrum $P(k_x) = \int P_{2D}(k_x, k_y) dk_y$.⁴ The power spectra are normalized such that $\int P_{2D}(k_x, k_y) dk_x dk_y = \int P(k_x) dk_x = \xi_B$.

⁴In Plotnikov & Sironi (2019), f_{ξ} showed that for 1D simulations, wavenumber spectra (k_x -spectra) and frequency spectra (ω -spectra) nearly overlap, when accounting for the dispersion relation of X modes.

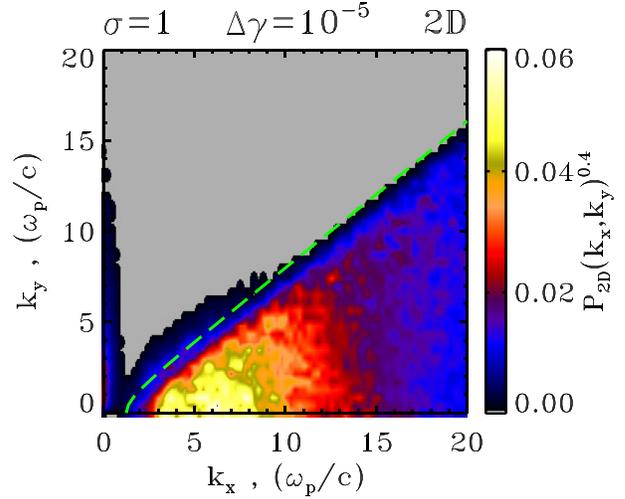


Figure 9. 2D wavenumber power spectrum $P_{2D}(k_x, k_y)$ for the 2D simulation with $\sigma = 1$ and $\Delta\gamma = 10^{-5}$, averaged in the time range $2000 \leq \omega_p t \leq 2500$. The spectrum is extracted from the same region ahead of the shock where ξ_B is computed, i.e. $5c/\omega_p < x - x_{\text{sh}} < 30c/\omega_p$, and is normalized such that at each time its integral equals ξ_B . The dashed green line indicates the theoretical upper limit in equation (10) for waves with group speed larger than the shock speed, taking the measured $\beta_{\text{sh}} = 0.78$.

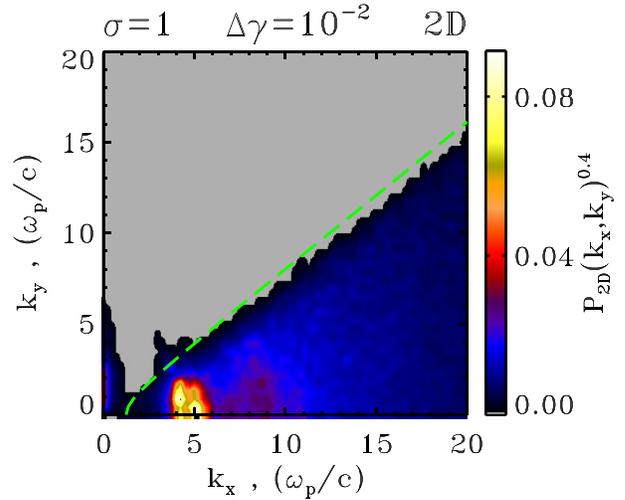


Figure 10. As in Fig. 9, for the 2D simulation with $\sigma = 1$ and $\Delta\gamma = 10^{-2}$, averaged in the time range $2000 \leq \omega_p t \leq 2500$. The dashed green line corresponds to equation (10) taking the measured $\beta_{\text{sh}} = 0.78$.

In Figs 9–11, we present the 2D power spectrum for three representative simulations, having different values of magnetization and pre-shock temperature. In Fig. 9, we show results for $\sigma = 1$ and $\Delta\gamma = 10^{-5}$ (the cold case described in Figs 1 and 4), in Fig. 10 for $\sigma = 1$ and $\Delta\gamma = 10^{-2}$ (the warm case in Figs 2 and 5), and in Fig. 11 for $\sigma = 3$ and $\Delta\gamma = 10^{-5}$ (a cold case with higher magnetization).

In each of the plots, the power at $k_x \sim 0$ and $k_y \gtrsim \omega_p/c$ is attributed to wave filamentation associated with the density filaments observed in panel (c) of Figs 1 and 2. Most of the spectral power, however, resides at higher k_x , within the region delimited by the green dashed line (defined below). By comparing the figures, one sees that for cold plasmas (Figs 9 and 11) the power is distributed over a wide range of longitudinal wavenumbers ($k_x \sim 5 - 10 \omega_p/c$), whereas for warm

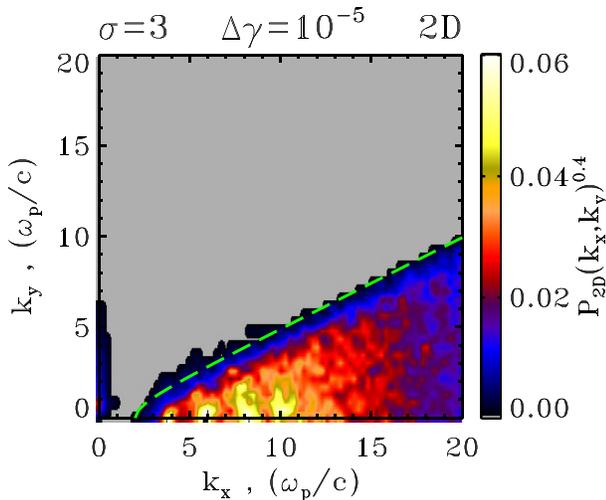


Figure 11. As in Fig. 9, for the 2D simulation with $\sigma = 3$ and $\Delta\gamma = 10^{-5}$, averaged in the time range $4000 \leq \omega_p t \leq 4500$. The dashed green line corresponds to equation (10) taking the measured $\beta_{\text{sh}} = 0.89$.

plasmas (Fig. 10) the spectrum is sharply peaked at $k_x \sim 5 \omega_p/c$. This will be further discussed below, where we show the k_y -integrated spectrum $P(k_x)$.

By comparing Figs 9 and 11, which differ in magnetization, we find that the precursor emission is beamed within a narrower angle $\theta = \arctan(k_y/k_x)$ for larger σ . As we now discuss, this follows from the requirement that the waves be able to escape ahead of the shock, which moves faster for higher magnetizations (see middle panel of Fig. 8).

The dispersion relation of the extraordinary mode (X mode) in cold plasmas in the frame where the background plasma is at rest reads (see e.g. Hoshino & Arons 1991)

$$\frac{k'^2 c^2}{\omega'^2} = 1 - \frac{\omega_p'^2}{\omega'^2 - \sigma \omega_p'^2}, \quad (8)$$

where double primed quantities are measured in the upstream rest frame. In the limit $\gamma_0^2 \gg \sigma$, the dispersion relation in the downstream frame becomes

$$k^2 c^2 \simeq \omega^2 - \omega_p^2, \quad (9)$$

which is identical to the dispersion relation of a simple electromagnetic wave propagating in an unmagnetized plasma.

The motion of the shock front imposes a cut-off below which the wave cannot escape into the upstream medium. This cut-off is obtained by equating the projection of the group velocity $v_g = d\omega/dk$ on to the shock normal with the shock speed, i.e. $v_g \cos \theta \geq \beta_{\text{sh}} c$, which leads to

$$k_x \geq \gamma_{\text{sh}} \beta_{\text{sh}} \sqrt{k_y^2 + \frac{\omega_p^2}{c^2}}. \quad (10)$$

This inequality identifies the values of (k_x, k_y) for which the wave can successfully outrun the shock (Iwamoto et al. 2017).

This has two main consequences. First, the range of k_x allowed for wave propagation has an absolute lower limit, $k_{x,\text{cut}} = \gamma_{\text{sh}} \beta_{\text{sh}} \omega_p/c$. In the limit of high magnetizations, this scales as $k_{x,\text{cut}} \simeq \sigma^{1/2} \omega_p/c$, in agreement with a comparison of Figs 9 and 11. Secondly, for each given k_x , the allowed range of k_y is constrained by equation (10), which is indicated by the dashed green lines in Figs 9–11. At $k_x, k_y \gg \omega_p/c$, this corresponds to precursor emission being confined within

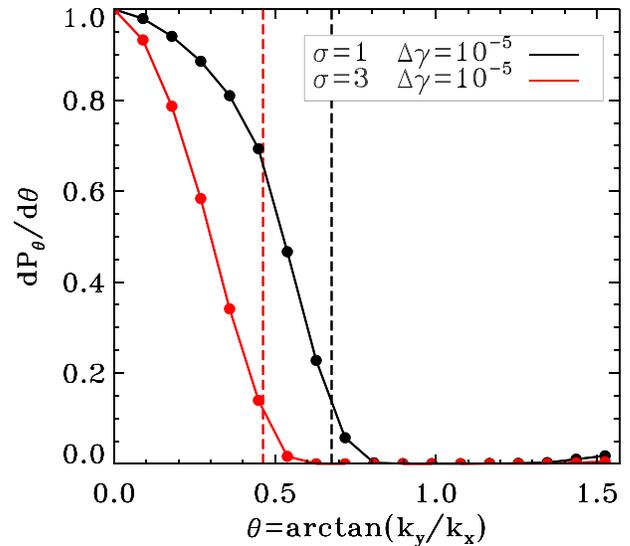


Figure 12. Angular distribution of the precursor power $dP_{\theta}/d\theta$, where $\theta = \arctan(k_y/k_x)$, obtained from the 2D power spectrum $P_{2D}(k_x, k_y)$ by integrating along lines of constant θ . We show results for the same simulations shown in Figs 9 and 11 (see legend): black for $\sigma = 1$ and $\Delta\gamma = 10^{-5}$ (Fig. 9), red for $\sigma = 3$ and $\Delta\gamma = 10^{-5}$ (Fig. 11). The curves are normalized to their respective maximum values. The vertical dashed lines (same colour coding as the solid lines) represent the boundary $\theta_{\text{crit}} = \arctan(1/\gamma_{\text{sh}}\beta_{\text{sh}})$ defined by equation (10), such that only waves with $\theta \leq \theta_{\text{crit}}$ can outrun the shock.

an angle $\theta_{\text{crit}} = \arctan(1/\gamma_{\text{sh}}\beta_{\text{sh}})$ from the shock normal. For $\sigma \gg 1$, this scales as $\theta_{\text{crit}} \simeq \sigma^{-1/2}$. So, precursor waves from more strongly magnetized shocks will be directed closer to the shock normal.

This is demonstrated in Fig. 12, by computing the angular distribution of precursor power $dP_{\theta}/d\theta$, obtained by integrating the 2D wavenumber spectrum P_{2D} along lines of constant θ . The two curves correspond to the two cases in Figs 9 and 11 (see legend), and are normalized to their respective peak values. The vertical dashed lines (same colour coding as the solid lines) correspond to $\theta_{\text{crit}} = \arctan(1/\gamma_{\text{sh}}\beta_{\text{sh}})$. The plot shows that the precursor emission at $\theta > \theta_{\text{crit}}$ is indeed negligible (the power at $\theta = \pi/2$ is contributed by the non-propagating filamentation mode, as described above). As expected, precursor waves in flows with higher magnetizations are more strongly beamed since $\theta_{\text{crit}} \simeq \sigma^{-1/2}$. Moreover, the precursor is even more beamed than the constraint in equation (10) would prescribe: the width at half maximum of $dP_{\theta}/d\theta$ is $\sim 0.7\theta_{\text{crit}}$.

Further comparison of wavenumber spectra among cases with fixed $\sigma = 1$ and varying temperatures is presented in Fig. 13, where we show the k_y -integrated wavenumber spectrum $P(k_x)$; more precisely, we show $k_x P(k_x)$ to emphasize where most of the power resides). We focus our discussion on the range $k_x < 20 \omega_p/c$ that is robust against variations of numerical parameters, see Appendix A.

Once the precursor efficiency settles to a steady state (see Fig. 7), the spectral shape is also nearly time independent. Regardless of the pre-shock temperature, the spectra share some common features: (i) the range of longitudinal wavenumbers has a sharp cut-off at $k_{x,\text{cut}} \simeq 2 \omega_p/c$, which descends from the constraint in equation (10) for $k_y = 0$; (ii) the spectral shape at high wavenumbers ($k_x \gtrsim 5 \omega_p/c$) resembles a power law $P(k_x) \propto k_x^{-2}$.

Despite these similarities, sharp differences exist between cases with different pre-shock temperatures. For cold temperatures, the spectrum peaks at $k_x \sim 5 - 10 \omega_p/c$ and is relatively broad, with fractional width $\sim 1-3$. This is common to all cases with $\Delta\gamma \lesssim$

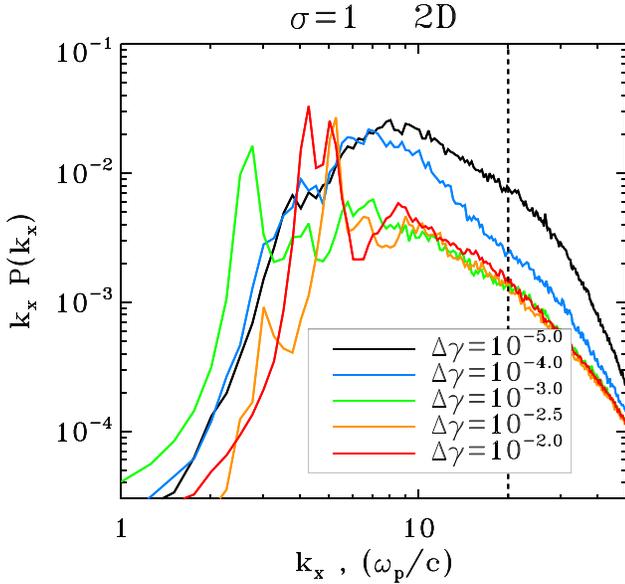


Figure 13. 1D wavenumber power spectrum $P(k_x) = \int P_{2D}(k_x, k_y) dk_y$ for 2D simulations with $\sigma = 1$ and varying $\Delta\gamma$ (see legend). The spectrum is extracted from the same region ahead of the shock where ξ_B is computed, i.e. $5c/\omega_p < x - x_{\text{sh}} < 30c/\omega_p$. At each time, the spectrum is normalized such that its integral equals ξ_B . Each curve is obtained by averaging the spectra in the time interval $2000 \leq \omega_p t \leq 2500$. The vertical dashed line indicates the boundary $k_{\text{max}} = 20\omega_p/c$ beyond which our spectra may be affected by numerical artefacts (see Appendix A). The Nyquist wavenumber for our simulations is $k_N \simeq 314\omega_p/c$.

$10^{-3.5}$ (we show $\Delta\gamma = 10^{-5}$ in black and $\Delta\gamma = 10^{-4}$ in blue). In this temperature range, the spectral power at high wavenumbers ($k_x \gtrsim 10\omega_p/c$) gets increasingly suppressed for larger thermal spreads, as discussed by Amato & Arons (2006). For warm temperatures ($10^{-3} \lesssim \Delta\gamma \lesssim 10^{-1.5}$), the spectrum shows pronounced line-like features with fractional width ~ 0.2 . The line-like features are located at the low-wavenumber end of the spectrum, at $k_x \sim 3 - 5\omega_p/c$. We show spectra for $\Delta\gamma = 10^{-3}$ (green), $\Delta\gamma = 10^{-2.5}$ (yellow) and $\Delta\gamma = 10^{-2}$ (red). The spectrum for $\Delta\gamma = 10^{-1.5}$ (not shown) is very similar to the $\Delta\gamma = 10^{-2}$ case, whereas we remind that the precursor efficiency is strongly suppressed for even hotter temperatures ($\Delta\gamma \gtrsim 10^{-1}$).

6 SUMMARY AND DISCUSSION

In this work, we have investigated by means of 2D PIC simulations the physics of the precursor waves emitted by perpendicular relativistic electron–positron shocks with out-of-plane upstream fields. We have focused on the high magnetization regime $\sigma \gtrsim 1$ appropriate for magnetar winds, motivated by the shock-powered synchrotron maser scenario proposed for FRBs (Lyubarsky 2014; Murase et al. 2016; Beloborodov 2017; Waxman 2017; Metzger et al. 2019; Beloborodov 2019; Margalit et al. 2020 a,b). We have explored the efficiency and spectrum of the precursor waves as a function of the pre-shock thermal spread $\Delta\gamma = kT/mc^2$ in the range $\Delta\gamma = 10^{-5} - 10^{-1}$. All our simulations have been run for a sufficiently long time ($\gtrsim 4000\omega_p^{-1}$) that the precursor emission has achieved a steady state. Our main results are as follows:

(i) By measuring the fraction f_ξ of total (i.e. electromagnetic and kinetic) incoming energy that is converted into precursor waves, as

computed in the post-shock frame, we can quantify the efficiency of precursor emission. At fixed temperature, the scaling with magnetization $f_\xi \sim 10^{-3}\sigma^{-1}$ at $\sigma \gtrsim 1$ is consistent with our earlier 1D results (Plotnikov & Sironi 2019).

(ii) At fixed magnetization, the precursor efficiency is nearly independent of temperature as long as $\Delta\gamma \lesssim 10^{-1.5}$ (with only a modest decrease of a factor of three from $\Delta\gamma = 10^{-5}$ to $\Delta\gamma = 10^{-1.5}$), but between $\Delta\gamma = 10^{-1.5}$ and $\Delta\gamma = 10^{-1}$ it drops by nearly two orders of magnitude (the drop occurs at the same $\Delta\gamma$ regardless of the flow Lorentz factor). We have confirmed this result with dedicated 3D simulations (Sironi et al., in preparation). So, shocks propagating in hot plasmas with $\Delta\gamma \gtrsim 10^{-1}$ are unlikely to power FRBs.

(iii) For $\sigma \gtrsim 1$, the precursor waves are beamed within a cone of half-opening angle $\theta_{\text{crit}} = \arctan(1/\gamma_{\text{sh}}\beta_{\text{sh}}) \simeq \sigma^{-1/2}$ around the shock normal (as measured in the post-shock frame). This stems from the fact that only the waves whose group velocity projected along the shock normal is larger than the shock speed can outrun the shock. More precisely, the width at half maximum of the angular distribution of precursor power is $\sim 0.7\theta_{\text{crit}}$.

(iv) For $\sigma = 1$, we have compared the power spectrum $P(k_x)$ of precursor waves (integrated over the transverse wavenumber k_y) among different $\Delta\gamma$. For cold temperatures, the spectrum peaks at $k_x \sim 5 - 10\omega_p/c$ and is relatively broad, with fractional width $\sim 1 - 3$. In contrast, for warm temperatures ($10^{-3} \lesssim \Delta\gamma \lesssim 10^{-1.5}$) it shows pronounced line-like features with fractional width ~ 0.2 . The line-like features are located at the low-wavenumber end of the spectrum, at $k_x \sim 3 - 5\omega_p/c$. For both cold and warm flows, the high-wavenumber part at $k_x > 10\omega_p/c$ can be roughly modelled as a power law $P(k_x) \propto k_x^{-2}$.

Our simulations employ 2D computational domains initialized with out-of-plane magnetic fields. This configuration only allows for the excitation of the X mode (and not of the O mode), so in our case the precursor waves are 100 per cent linearly polarized, with fluctuating magnetic field along the same direction as the upstream mean field. A 3D investigation of the dependence of the precursor emission (in terms of both O and X modes) on the pre-shock temperature will be presented elsewhere.

The shocks investigated in this work propagate in an electron–positron plasma. The physics of electron–proton shocks will qualitatively differ, as discussed analytically by Lyubarsky (2006) and investigated with PIC simulations by Hoshino (2008), Sironi & Spitkovsky (2011), and Iwamoto et al. (2019). In magnetized electron–proton flows, the electron synchrotron maser instability generates a train of electromagnetic precursor waves propagating into the upstream. In response to the waves, the guiding-centre velocity of the incoming electrons decreases, since they experience relativistic transverse oscillations in the strong field of the wave. Due to their high mass, protons are less affected by the waves, instead proceeding close to their initial velocity. The resulting difference in bulk velocity between the two species generates a longitudinal electric ‘wakefield,’ which boosts electrons towards the shock. At relatively early times, this results in a substantial increase in the efficiency of precursor emission (Iwamoto et al. 2019). However, the nonlinear evolution of the wakefield inevitably leads to appreciable heating of the upstream electrons, which is likely to deteriorate the precursor efficiency in the long term. Future work will help assess the steady-state efficiency of precursor waves from relativistic electron–proton shocks.

Our work can help unveil the importance of the synchrotron maser as a source of coherent emission in astrophysical plasmas.

In particular, our results have implications for models of FRBs that rely on maser emission at ultrarelativistic shocks generated by magnetized pulses launched by an active magnetar (Lyubarsky 2014; Murase et al. 2016; Beloborodov 2017; Waxman 2017; Metzger et al. 2019; Beloborodov 2019; Margalit et al. 2020a,b). Most directly, our study is relevant for scenarios where the shock propagates in a strongly magnetized electron–positron wind (e.g. Beloborodov 2019).

Our findings imply that efficient synchrotron maser emission from electron–positron shocks requires the upstream electrons to have non-relativistic temperatures ($kT/mc^2 \lesssim 0.03$). This poses a direct constraint on both the properties of the pristine medium around the magnetar, as well as on the repetition rate of subsequent FRB-generating shocks, since even an initially cold upstream will be heated by the passage of the shock, raising the temperature of the medium into which the next pulse would collide (Metzger et al. 2019; Beloborodov 2019; Margalit et al. 2020b). This implies that FRBs from two successive shocks should be separated by a minimum lag time, as required for the plasma shocked by the first one to adiabatically cool before the arrival of the second shock. If this were not to be the case, the second shock would not be able to produce strong precursor waves (unless the second shock outruns the plasma shocked by the first one). For the first repeater FRB 121102 the minimum observed wait time seems to peak around 100 s, which in turn can be used to place limits on the shock radius (Metzger et al. 2019).

Even in the case of a cold upstream plasma, some energy may be transferred to the pre-shock flow by either the propagating precursor waves, or by incoherent synchrotron emission of shock-heated electrons (Metzger et al. 2019; Beloborodov 2019). As regards to the former effect, strictly speaking it should not be genuinely considered ‘heating,’ since at any given point ahead of the shock the particle distribution has negligible dispersion. However, when integrating over a few wavelengths, the particles will display some ‘effective’ longitudinal and transverse momentum spread.⁵ Since the wave-induced motions are mostly transverse to the direction of wave propagation, the transverse dispersion is significantly larger than the longitudinal one (fig. 8 in Margalit et al. 2020b). In Appendix B, we show that it is the longitudinal dispersion that primarily controls the maser efficiency. For the FRB conditions envisioned by Margalit et al. (2020b), the self-heating of upstream electrons by the precursor waves is not expected to suppress the maser efficiency.

Concerning the Compton heating of upstream electrons by the γ -rays/X-rays generated via incoherent synchrotron emission at the shock, the requirement that this would not raise the pre-shock temperature above $0.03 mc^2/k$ can be used, if γ -rays/X-ray emission is observed in coincidence with the FRB (as for the recent Galactic FRB, Mereghetti et al. 2020a,b), to place important constraints on the shock radius (see equation 66 in Metzger et al. 2019).

We have also shown that for large magnetizations, the precursor waves are beamed within a narrow cone around the shock normal, with half-opening angle $\simeq 0.7 \sigma^{-1/2}$. The limited range of emission angles implies a reduced range of Doppler factors for the frequency transformation from the post-shock frame to the observer frame (Beloborodov 2019), so the Doppler transformation of the precursor spectrum to the observer frame will not smear out narrow features. Our results demonstrate that for warm temperatures ($10^{-3} \lesssim \Delta\gamma \lesssim$

$10^{-1.5}$), the low-frequency end of the spectrum shows pronounced line-like features with fractional width ~ 0.2 . Given that such narrow features will not be smeared out in high- σ shocks, our spectra may be consistent with the observed complex, and sometimes narrow-band spectral energy distributions of observed FRBs (e.g. Ravi et al. 2016; Law et al. 2017; Macquart et al. 2019).

In addition, in the decelerating shock scenario, the sub-pulses observed within a given burst (Hessels et al. 2019; CHIME/FRB Collaboration 2019) may be interpreted as due to individual peaks in the maser spectrum, as they drift downwards across the observing band (Metzger et al. 2019). Our spectra for warm flows provide some ground for this interpretation. However, we caution that in the recently detected Galactic FRB (Mereghetti et al. 2020a,b; Ridnaia et al. 2020), there are two distinct X-ray peaks accompanying the two radio peaks, which rather suggests that the two radio peaks are associated with two separate shocks, with each X-ray burst coming from incoherent synchrotron emission of the corresponding shock-heated electrons (Metzger et al. 2019; Beloborodov 2019; Margalit et al. 2020a; but see e.g. Lu, Kumar & Zhang 2020, for caveats). In this regard, the results presented in this work require the second shock to outrun the plasma shocked by the first one, otherwise the shock generating the first peak would heat the gas too much to get efficient maser emission from the second shock.

ACKNOWLEDGEMENTS

We are grateful to A. Beloborodov, B. Margalit, B. Metzger, J. Nättilä, I. Plotnikov, E. Sobacchi, N. Sridhar, and A. Tran for inspiring discussions and comments on the manuscript. This work is supported by NASA ATP 80NSSC18K1104 and NSF AST-1716567. The simulations have been performed at Columbia (Habanero and Terremoto), and with NERSC (Cori) and NASA (Pleiades) resources.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Alsop D., Arons J., 1988, *Phys. Fluids*, 31, 839
 Amato E., Arons J., 2006, *ApJ*, 653, 325
 Begelman M. C., Kirk J. G., 1990, *ApJ*, 353, 66
 Beloborodov A. M., 2017, *ApJ*, 843, L26
 Beloborodov A. M., 2019, *ApJ*, 896, 142
 Belyaev M. A., 2015, *New Astron.*, 36, 37
 Blandford R., Eichler D., 1987, *Phys. Rep.*, 154, 1
 Bochenek C., Kulkarni S., Ravi V., McKenna D., Hallinan G., Belov K., 2020, *Astron. Telegram*, 13684, 1
 CHIME/FRB Collaboration, 2019, *ApJ*, 885, L24
 Chiueh T., Lai T.-C., 1991, *Phys. Rev. A*, 44, 6944
 Cordes J. M., Chatterjee S., 2019, *ARA&A*, 57, 417
 Dieckmann M. E., Frederiksen J. T., Bret A., Shukla P. K., 2006, *Phys. Plasmas*, 13, 112110
 Drake J. F., Kaw P. K., Lee Y. C., Schmid G., Liu C. S., Rosenbluth M. N., 1974, *Phys. Fluids*, 17, 778
 Gallant Y. A., Hoshino M., Langdon A. B., Arons J., Max C. E., 1992, *ApJ*, 391, 73
 Greenwood A. D., Cartwright K. L., Luginsland J. W., Baca E. A., 2004, *J. Comput. Phys.*, 201, 665
 Hessels J. W. T. et al., 2019, *ApJ*, 876, L23
 Hoshino M., 2001, *Prog. Theor. Phys. Suppl.*, 143, 149
 Hoshino M., 2008, *ApJ*, 672, 940

⁵Here, ‘longitudinal’ is along the shock direction of propagation (x -direction), while ‘transverse’ is along the wave electric field (y -direction).

- Hoshino M., Arons J., 1991, *Phys. Fluids B*, 3, 818
- Hoshino M., Arons J., Gallant Y. A., Langdon A. B., 1992, *ApJ*, 390, 454
- Iwamoto M., Amano T., Hoshino M., Matsumoto Y., 2017, *ApJ*, 840, 52
- Iwamoto M., Amano T., Hoshino M., Matsumoto Y., 2018, *ApJ*, 858, 93
- Iwamoto M., Amano T., Hoshino M., Matsumoto Y., Niemiec J., Ligorini A., Kobzar O., Pohl M., 2019, *ApJ*, 883, L35
- Katz J. I., 2016, *Mod. Phys. Lett. A*, 31, 1630013
- Langdon A. B., Arons J., Max C. E., 1988, *Phys. Rev. Lett.*, 61, 779
- Law C. J. et al., 2017, *ApJ*, 850, 76
- Lu W., Kumar P., Zhang B., 2020, *MNRAS*, 498, 1397
- Lyubarsky Y., 2006, *ApJ*, 652, 1297
- Lyubarsky Y., 2014, *MNRAS*, 442, L9
- Macquart J. P., Shannon R. M., Bannister K. W., James C. W., Ekers R. D., Bunton J. D., 2019, *ApJ*, 872, L19
- Margalit B., Beniamini P., Sridhar N., Metzger B. D., 2020a, *ApJ*, 899, L27
- Margalit B., Metzger B. D., Sironi L., 2020b, *MNRAS*, 494, 4627
- Max C. E., Arons J., Langdon A. B., 1974, *Phys. Rev. Lett.*, 33, 209
- Mereghetti S., Savchenko V., Gotz D. et al., 2020a, GRB Coordinates Network, 27668, 1
- Mereghetti S. et al., 2020b, *ApJ*, 898, L29
- Metzger B. D., Margalit B., Sironi L., 2019, *MNRAS*, 485, 4091
- Murase K., Kashiyama K., Mészáros P., 2016, *MNRAS*, 461, 1498
- Pétri J., Lyubarsky Y., 2007, *A&A*, 473, 683
- Petroff E., Hessels J. W. T., Lorimer D. R., 2019, *A&AR*, 27, 4
- Platts E., Weltman A., Walters A., Tendulkar S. P., Gordin J. E. B., Kandhai S., 2019, *Phys. Rep.*, 821, 1
- Plotnikov I., Sironi L., 2019, *MNRAS*, 485, 3816
- Plotnikov I., Grassi A., Grech M., 2018, *MNRAS*, 477, 5238
- Ravi V. et al., 2016, *Science*, 354, 1249
- Ridnaia A. et al., 2020, preprint (arXiv:2005.11178)
- Scholz P., CHIME/FRB Collaboration, 2020, *Astron. Telegram*, 13681, 1
- 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. Proc. Vol. 801, Astrophysical Sources of High Energy Particles and Radiation*. Am. Inst. Phys., New York, p. 345
- Stockem A., Fiúza F., Fonseca R. A., Silva L. O., 2012, *ApJ*, 755, 68
- Waxman E., 2017, *ApJ*, 842, 34

APPENDIX A: DEPENDENCE ON NUMERICAL PARAMETERS

In this Appendix, we describe the dependence of the power spectrum of precursor waves on numerical parameters, for a representative set of 2D simulations with $\sigma = 1$ and $\Delta\gamma = 10^{-4}$. In Fig. A1, we show how the spectrum changes when varying the number of particles per cell N_0 between 4 and 16, and the number of iterations of the low-pass filter for electric currents between $N_{sm} = 10$ and $N_{sm} = 20$. Our reference simulations presented in the main body of the paper have $N_0 = 16$ and $N_{sm} = 20$ (the red curve in Fig. A1).

The figure shows that the spectral shape and normalization at $k \lesssim 20 \omega_p/c$ (to the left of the vertical dashed line) are robust to our choice of numerical parameters. On the other hand, the highest wavenumber part at $k \gtrsim 20 \omega_p/c$ varies. A larger number of either particles per cell or filter passes reduces the power at high wavenumbers. It appears that a comparable level of high-wavenumber suppression is achieved with the combinations ($N_0 = 4, N_{sm} = 20$) and ($N_0 = 16, N_{sm} = 10$), compare green and yellow lines. In the main body of the paper, we have used the threshold at $k_{max} = 20 \omega_p/c$ indicated by the vertical dashed line as the upper limit of the wavenumber range where our spectra are robust against numerical artefacts.

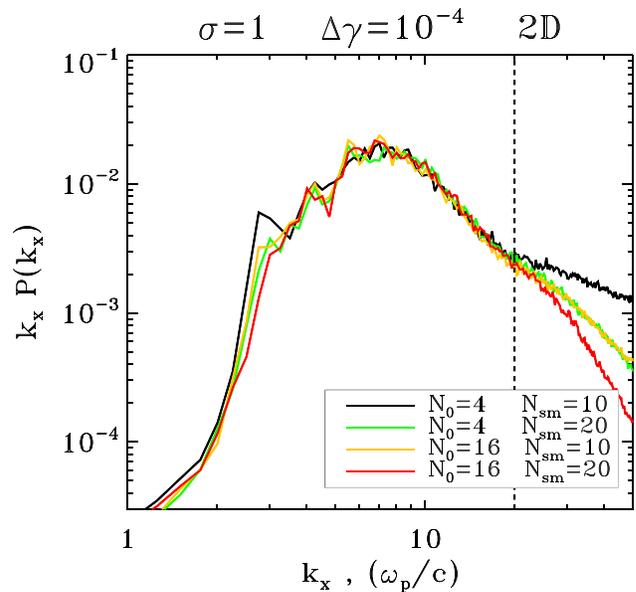


Figure A1. 1D wavenumber power spectrum $P(k_x) = \int P_{2D}(k_x, k_y) dk_y$ for 2D simulations with $\sigma = 1$ and different choices of particles per cell N_0 and number of passes of the smoothing filter of electric currents N_{sm} (see legend). The spectrum is extracted from the same region ahead of the shock where ξ_B is computed, i.e. $5c/\omega_p < x - x_{sh} < 30c/\omega_p$. At each time, the spectrum is normalized such that its integral equals ξ_B . Here, each curve is obtained by averaging the spectra in the time interval $2000 \leq \omega_p t \leq 2500$. The vertical dashed line indicates the boundary $k_{max} = 20 \omega_p/c$ beyond which our spectra differ because of numerical effects.

APPENDIX B: EFFECT OF LONGITUDINAL AND TRANSVERSE DISPERSION

In this Appendix, we present dedicated 1D simulations to clarify whether it is the longitudinal or the transverse dispersion that is most detrimental for the efficiency of the precursor emission. We employ $N_0 = 400$ particles per cell, a spatial resolution of $c/\omega_p = 112$ cells and a numerical speed of light of 0.5 cells/time-step. The magnetic field initialization is the same as in the 2D simulations presented in the main body of the paper, and we focus on $\sigma = 1$.

Fig. B1 shows that the efficiency drops by nearly two orders of magnitude from $\Delta\gamma = 10^{-6}$ (black) to $\Delta\gamma = 10^{-0.5}$ (red). Indeed, as Fig. 8 shows, the drop occurs between $\Delta\gamma = 10^{-1.5}$ and $\Delta\gamma = 10^{-1}$. We perform two additional simulations, in which we start with a hot upstream flow ($\Delta\gamma = 10^{-0.5}$, as for the red line), but for $\omega_p t \geq 700$ we artificially change its properties right ahead of the shock, in the region $4c/\omega_p \leq x - x_{sh} \leq 22c/\omega_p$. For the yellow line, we suppress the transverse momentum dispersion (i.e. along y), while for the green line we suppress the longitudinal momentum dispersion (i.e. along x). If the region where we enforce the suppression were to be far ahead of the shock, gyration around the upstream magnetic field would interchange the transverse and longitudinal motions on a time-scale $(\pi/2)\gamma_0 \sigma^{-1/2} \omega_p^{-1}$, or equivalently on a distance $(\pi/2)\gamma_0 \sigma^{-1/2} c/\omega_p$. For $\gamma_0 = 10$ and $\sigma = 1$, the choice of suppressing one momentum component between 4 and 22 skin depths is then sufficient to guarantee that the flow entering the shock preserves our imposed temperature anisotropy.

Fig. B1 shows that, if we suppress the transverse dispersion (the yellow line), the shock retains the low efficiency of the hot case (the red curve). In contrast, if we suppress the longitudinal dispersion (the green line), the efficiency increases and eventually settles to

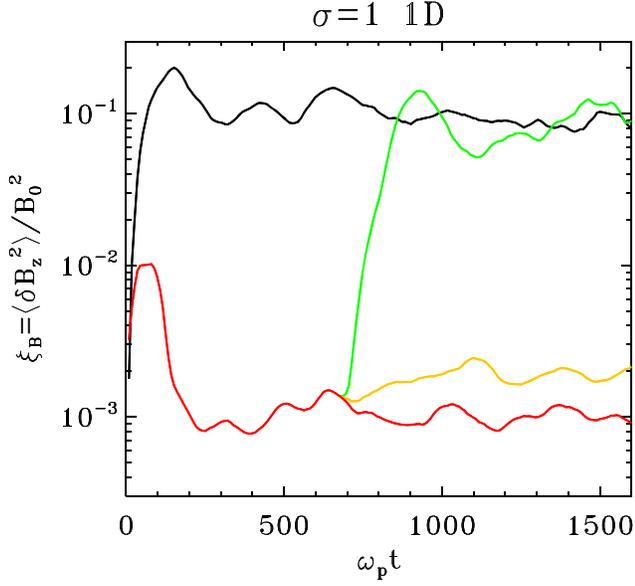


Figure B1. Time evolution of the normalized precursor wave energy, $\xi_B = \langle \delta B_z^2 \rangle / B_0^2$ for 1D simulations with $\sigma = 1$. The black and red lines, respectively, correspond to a cold case ($\Delta\gamma = 10^{-6}$) and a hot case ($\Delta\gamma = 10^{-0.5}$). The yellow and green lines, respectively, correspond to cases where we start with a hot thermal spread ($\Delta\gamma = 10^{-0.5}$), but for $\omega_p t \geq 700$ we suppress by hand either the longitudinal momentum dispersion (i.e. along x) or the transverse one (i.e. along y) just ahead of the shock. Here, ‘longitudinal’ and ‘transverse’ refer to the shock direction of propagation.

the level corresponding to the cold case (the black line). We have also verified that at late times the power spectra corresponding to the black and green lines are similar (and the same applies to the spectral comparison of the yellow and red lines).

This demonstrates that, if the upstream plasma were to be continuously heated in an anisotropic way, it is the longitudinal dispersion that determines the efficiency of the precursor emission. Such anisotropic heating may result from the propagation of the precursor waves themselves since they preferentially induce particle motions in the transverse direction.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.