



The synchrotron maser emission from relativistic shocks in Fast Radio Bursts: 1D PIC simulations of cold pair plasmas.

Illya Plotnikov and Lorenzo Sironi³

¹Department of Astrophysical Sciences, Princeton University, 4 Ivy Ln., Princeton, NJ 08544, USA

²Université de Toulouse, UPS-OMP, IRAP, 9 av. Colonel Roche, BP 44346, F-31028 Toulouse Cedex 4, France

³Department of Astronomy, Columbia University, 550 W 120th St., New York, NY 10027, USA

Accepted 2019 March 1. Received 2019 February 25; in original form 2018 December 29

ABSTRACT

The emission process of Fast Radio Bursts (FRBs) remains unknown. We investigate whether the synchrotron maser emission from relativistic shocks in a magnetar wind can explain the observed FRB properties. We perform particle-in-cell (PIC) simulations of perpendicular shocks in cold pair plasmas, checking our results for consistency among three PIC codes. We confirm that a linearly polarized X-mode wave is self-consistently generated by the shock and propagates back upstream as a precursor wave. We find that at magnetizations $\sigma \gtrsim 1$ (i.e. ratio of Poynting flux to particle energy flux of the pre-shock flow) the shock converts a fraction $f'_\xi \approx 7 \times 10^{-4}/\sigma^2$ of the total incoming energy into the precursor wave, as measured in the shock frame. The wave spectrum is narrow-band (fractional width $\lesssim 1-3$), with apparent but not dominant line-like features as many resonances concurrently contribute. The peak frequency in the pre-shock (observer) frame is $\omega''_{\text{peak}} \approx 3\gamma_{s|u}\omega_p$, where $\gamma_{s|u}$ is the shock Lorentz factor in the upstream frame and ω_p the plasma frequency. At $\sigma \gtrsim 1$, where our estimated ω''_{peak} differs from previous works, the shock structure presents two solitons separated by a cavity, and the peak frequency corresponds to an eigenmode of the cavity. Our results provide physically grounded inputs for FRB emission models within the magnetar scenario.

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

1 INTRODUCTION

Synchrotron masers are known to produce strong decametric radio emission in the Jovian magnetosphere and kilometric emission in the terrestrial magnetosphere (Auroral Kilometric Emission). They are driven by a ‘population inversion’ of energetic electrons gyrating in an intense magnetic field. The driver of the emission is either a loss-cone or ring-like electron distribution function (Treumann 2006; Melrose 2017). Such a population inversion occurs also in strongly magnetized relativistic perpendicular shocks, where a coherent cold-ring distribution of particles is self-consistently produced as part of the shock evolution (Gallant et al. 1992; Amato & Arons 2006). This distribution is unstable to the synchrotron maser instability (Hoshino & Arons 1991) that causes the emission of a train of high amplitude semicoherent electromagnetic waves propagating from the shock front into the unshocked (upstream) medium (Gallant et al. 1992; Hoshino et al. 1992). The possible importance of the synchrotron maser in astrophysical sources was anticipated long ago (Sazonov 1973). Yet, to our present knowledge,

there is no firm demonstration that shock-powered synchrotron maser emission is dominant in any non-heliospheric astrophysical environment, even though a few potential scenarios have been proposed (e.g. Sagiv & Waxman 2002).

Recently, however, the discovery of Fast Radio Bursts (FRBs; Lorimer et al. 2007; Keane et al. 2012; Thornton et al. 2013; Spitler et al. 2014; Marcote et al. 2017) has revived the interest in this mechanism. These events are bright (~ 1 Jy) pulses of millisecond duration detected in the \sim GHz band. Their extremely high brightness temperature, $T_B \sim 10^{37}$ K, requires a coherent emission mechanism (Katz 2016; Popov, Postnov & Pshirkov 2018). Young magnetars have emerged as one of the most likely progenitors of FRBs, at least for the repeating class (e.g. Popov & Postnov 2013; Murase, Kashiyama & Mészáros 2016; Kashiyama & Murase 2017; Lyutikov 2017; Metzger, Berger & Margalit 2017; Long & Pe'er 2018; Margalit et al. 2018; Margalit & Metzger 2018). Magnetars can naturally explain the short FRB durations, large energy requirements, and ordered magnetic fields needed for coherent emission. Thus far, most works on the FRB emission mechanism are based on mere considerations of energetics and time-scales, in which it is assumed that a fraction of the free energy of the system is radiated away at GHz frequencies by coherent charge

* E-mail: illyap@astro.princeton.edu

‘bunches’ via, e.g. curvature or synchrotron maser processes. In the case of curvature radiation, the emission is postulated to be a product of magnetic reconnection close to the magnetar surface (e.g. Lyutikov 2002; Kumar, Lu & Bhattacharya 2017; Ghisellini & Locatelli 2018; Katz 2018; Lu & Kumar 2018). In the case of the synchrotron maser, the emission is thought to occur at relativistic shocks propagating in the magnetar wind or nebula (Lyubarsky 2014; Murase et al. 2016; Beloborodov 2017) or inside the ultra-relativistic shell ejected from the central compact object (Waxman 2017). Yet, in either case the conditions for coherent emission and the very existence of charge bunches with the required properties are often postulated ad hoc, resulting in models with little predictive power.

The purpose of this work is to demonstrate from first principles that the synchrotron maser at relativistic shocks in the magnetar wind can naturally explain the observed FRB properties. By means of particle-in-cell (PIC) simulations, we investigate how the efficiency and spectrum of the electromagnetic wave emitted by the shock into the pre-shock medium (which we shall call ‘precursor wave’) depend on the physical conditions in the magnetar wind. In this work, the first of a series, we present results from 1D simulations (more precisely, 1D3V, i.e. we employ one spatial dimension, but all three components of velocities and electromagnetic fields are retained), while multidimensional runs will be presented in a future work (Sironi et al., in preparation; see also Appendix B for the precursor energetics in 2D and 3D). We focus on the case of a cold pair-dominated plasma.

There exists extensive literature on PIC modelling of the electromagnetic precursor wave in relativistic perpendicular shocks (e.g. Langdon, Arons & Max 1988; Gallant et al. 1992; Hoshino et al. 1992; Amato & Arons 2006; Hoshino 2008; Sironi & Spitkovsky 2009; Iwamoto et al. 2017, 2018). Our work is motivated by the poor exploration of the extreme regime where the energy content of the plasma is dominated by magnetic fields, as it is supposedly the case in magnetar winds. In other words, we focus on magnetizations $\sigma \gtrsim 1$, where σ is the ratio of upstream Poynting flux to kinetic energy flux. For $\sigma \gtrsim 1$, 1D simulations are adequate, since we find that they agree well with multidimensional results (see Appendix B, for the precursor energetics in 2D and 3D; also Sironi et al., in preparation). We provide an extensive investigation of the dependence on the flow magnetization, from $\sigma = 0.1$ to $\sigma = 30$, with much longer simulations than previously reported, especially in the $\sigma \gtrsim 1$ regime relevant for FRB sources. Previous works arguably never reached a steady state in simulations with $\sigma \gtrsim 1$. We employ several PIC codes (TRISTAN-MP, SMILEI and SHOCKAPIC) to check for consistency, and thus confirm the robustness of our results.

At $\sigma \gtrsim 1$ the shock converts a fraction $f_\xi \approx 2 \times 10^{-3}/\sigma$ of the total incoming energy into the precursor wave, as measured in the post-shock (downstream) frame. In the shock rest frame (SRF), the efficiency is $f'_\xi \approx 7 \times 10^{-4}/\sigma^2$. The spectrum of the precursor wave is narrow-band, $\Delta\omega/\omega_{\text{peak}} \lesssim 1 - 3$, with apparent but not dominant line-like features as many resonances concurrently contribute. The peak frequency scales in the post-shock frame as $\omega_{\text{peak}} \simeq 3\omega_p \max[1, \sqrt{\sigma}]$, where ω_p is the plasma frequency. In the pre-shock frame (which coincides with the observer frame, if the magnetar wind is non-relativistic), this can be recast in a simpler form as $\omega''_{\text{peak}} \approx 3\gamma_{\text{slu}}\omega_p$, where γ_{slu} is the shock Lorentz factor in the upstream frame. At $\sigma \gtrsim 1$, where our estimated ω_{peak} differs from earlier works [that quoted $\omega_{\text{peak}} \propto \sigma\omega_p$, see Gallant et al. (1992), rather than $\omega_{\text{peak}} \propto \sqrt{\sigma}\omega_p$ as we find] we see 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. The efficiency and spectrum of the precursor wave do not depend on the bulk Lorentz factor of the pre-shock flow.

The paper is organized as follows. In Section 2 we present the methods and the numerical set-up. In Section 3 we discuss the main results, in the post-shock (downstream) rest frame. Section 4 discusses the energy content of precursor waves in the frame of the shock front. In Section 5 we present the implications of our results for FRB emission models, and we conclude in Section 6.

2 SIMULATION METHODS AND SET-UP

We use the PIC codes TRISTAN-MP (Spitkovsky 2005; Sironi & Spitkovsky 2009), SMILEI (Derouillat et al. 2018), and SHOCKAPIC (Plotnikov, Grassi & Grech 2018) to perform 1D3V simulations, where we retain one spatial direction, but all three components of velocities and electromagnetic fields. Mainly TRISTAN-MP and SMILEI are employed for large simulations. The pseudo-spectral code SHOCKAPIC is used to check for consistency in shorter simulations. In the main body of the paper, we present only results obtained with TRISTAN-MP, unless stated otherwise. In the Appendix A, we demonstrate the agreement between different codes across the whole range of σ explored in this work.

The use of a reduced 1D spatial geometry is justified in the limit of magnetically dominated plasmas, as we will demonstrate with 2D and 3D simulations in a forthcoming study (Sironi et al., in preparation). For lower magnetizations than explored here, i.e. $\sigma \lesssim 0.5$, Iwamoto et al. (2017) found that the precursor wave energy is reduced by at most an order of magnitude in 2D simulations as compared to 1D. This difference is smaller or even negligible in the high magnetization limit explored here (see Appendix B).

The shock is initialized using the common set-up described in, e.g. Spitkovsky (2008), which we summarize here for completeness. The upstream flow, composed of electrons and positrons, drifts along the $-\hat{x}$ direction with a speed $-\beta_0 c \hat{x}$. The corresponding bulk Lorentz factor is $\gamma_0 = (1 - \beta_0^2)^{-1/2} = 10$, but we have also explored higher values of γ_0 , up to 10^5 . The upstream pair plasma is cold, with thermal spread $k_B T_0/(m_e c^2) = 10^{-4}$. The pre-shock plasma carries a frozen-in magnetic field B_0 oriented along z (so, $B_{z,0} = B_0$), i.e. perpendicular to the flow propagation, and a motional electric field $E_{y,0} = -\beta_0 B_{z,0}$. The flow is reflected at a wall located at $x = 0$. After some time (at least several cyclotron periods ω_c^{-1}), the shock front forms by magnetic reflection and steadily propagates along the $+\hat{x}$ direction with a speed that is in good agreement with the Rankine–Hugoniot conditions. The resulting simulation frame coincides with the frame where the downstream plasma is at rest (downstream rest frame; DRF).

The shock physics is sensitive to the upstream magnetization σ , which we define as the ratio of Poynting to kinetic energy flux

$$\sigma = \frac{B_{z,0}^2}{8\pi\gamma_0 N_0 m_e c^2} = \left(\frac{\omega_c}{\omega_p} \right)^2, \quad (1)$$

and we vary from $\sigma = 0.1$ up to $\sigma = 30$. Here N_0 is the number density of upstream electrons (the overall particle number density is then $2N_0$), m_e is the electron (or positron) mass, and c is the speed of light in vacuum. We also define the typical gyrofrequency as $\omega_c = |q|B_0/(\gamma_0 m_e c)$ and the plasma frequency as $\omega_p = [8\pi N_0 q^2/(\gamma_0 m_e)]^{1/2}$, where q is the elementary electric charge. Both quantities are based on the upstream values of magnetic field and plasma density measured in the simulation frame.

Typical numerical parameters used with TRISTAN-MP are

- (i) The skin depth is well resolved with $c/\omega_p = 100\Delta$, where Δ is the grid size. This ensures that the typical particle gyroradius $\simeq \sigma^{-1/2} c/\omega_p$ is well resolved even for the largest magnetization $\sigma = 30$ explored in this study. A good spatial resolution is also essential to capture the high-frequency part of the spectrum of precursor waves (see also Iwamoto et al. 2017 for a discussion on the required spatial resolution).
- (ii) The simulation time-step is defined as $c\Delta t = 0.5\Delta$ corresponding to a time resolution of $5 \times 10^{-3}\omega_p^{-1}$.
- (iii) The number of particles per cell initialized in the upstream plasma is $N_{\text{ppc}} = 64$ per species (values between 20 and 200 were tested with no appreciable differences).
- (iv) The simulation is evolved up to $T_{\text{sim}} \simeq 1.5 \times 10^3\omega_p^{-1}$ for $\sigma < 1$ and for longer times (up to $2 \times 10^4\omega_p^{-1}$) for $\sigma \gg 1$. This is required in order to reach a steady state in which the precursor emission maintains a constant amplitude (see Section 3.2).

In the Appendix A we also report the typical simulation parameters for the other two PIC codes used in this study. They are not presented here since in the following sections, we mainly discuss the results obtained with TRISTAN-MP.

3 RESULTS

In this section we explore the physics of relativistic highly magnetized electron–positron shocks focusing on the properties of the electromagnetic precursor. In Section 3.1 we discuss the typical shock structure and show the presence of electromagnetic precursor waves. In Sections 3.2 and 3.3 we show the dependence on σ of the precursor wave intensity and spectrum, respectively. The wave strength parameter is discussed in Section 3.4.

3.1 Shock layer structure

As shown in several works using 1D simulations (Langdon et al. 1988; Gallant et al. 1992; Amato & Arons 2006; Lyubarsky 2006; Hoshino 2008), 2D simulations (Sironi & Spitkovsky 2009, 2011; Iwamoto et al. 2017, 2018; Plotnikov et al. 2018), and 3D simulations (Spitkovsky 2005; Sironi, Spitkovsky & Arons 2013), highly magnetized perpendicular relativistic shocks form by magnetic reflection and generate a strong electromagnetic wave propagating from the shock into the upstream region. For high magnetizations (typically, $\sigma \gtrsim 0.01$), the transverse Weibel filamentation instability, which dominates for $\sigma \lesssim 10^{-3}$, plays no significant role in shaping the shock structure. This partly justifies the 1D approach adopted here.

In Fig. 1 we present the structure of the shock transition region for two representative magnetizations. The left-hand column (panels a–e) presents a shock with $\sigma = 0.3$ and the right-hand column (panels f–j) corresponds to $\sigma = 3$. The timespan of the simulations is long enough to reach a stationary state: we show results at $\omega_p t = 540$ for $\sigma = 0.3$ and at $\omega_p t = 1800$ for $\sigma = 3$. From top to bottom we show the electron number density N_e/N_0 (panels a and f), the transverse magnetic field B_z/B_0 (b and g), the transverse electric field E_y/B_0 (c and h), the longitudinal positron phase space $x - u_x$ (d and i), and the transverse positron phase space $x - u_y$ (e and j). Here, we define $u_\alpha = \gamma\beta_\alpha$ as the dimensionless four velocity. The phase space of electrons is identical to the one of positrons in virtue of mass symmetry, except for the opposite sign in variations of u_y . The electrostatic field E_x is not plotted since it is completely negligible in pair plasmas (we have systematically checked this conclusion).

The vertical dashed lines in panels (b) and (g) delimit the region where we have extracted the wave properties, such as amplitude and spectrum, that will be discussed in the sections below. Small insets in the upper right-hand side of panels (d) and (i) show the particle distribution in momentum space $u_x - u_y$ at the location of the shock front.

The shock front is located at $x - x_{\text{front}} = 0$ in Fig. 1. The upstream flow is on the positive side ($x - x_{\text{front}} > 0$) and the downstream plasma is on the negative side ($x - x_{\text{front}} < 0$). The existence of a well developed shock is confirmed by the jump in the electron number density and in the B_z field at the front location. The shock front itself exhibits a soliton-like structure (see e.g. Alsop & Arons 1988), where the particle distribution forms a semicoherent cold ring in momentum space (see insets in panels d and i). The presence of a large amplitude electromagnetic precursor wave is evidenced in the upstream region of the B_z/B_0 and E_y/B_0 panels for both magnetizations (see the ripples in the $x - x_{\text{front}} > 0$ region). The precursor wave amplitude is larger for $\sigma = 0.3$ than for $\sigma = 3$. This wave is steadily emitted from the shock front and is *linearly polarized*. The wave vector \mathbf{k} lies along the shock direction of propagation (i.e. along x), the fluctuating magnetic field is along z (i.e. along the same direction as the upstream field $\mathbf{B}_0 = B_0\hat{z}$), and the fluctuating electric field is perpendicular to both \mathbf{k} and \mathbf{B}_0 . The wave is then identified with the extraordinary mode (X-mode). The phase velocity of the wave is slightly superluminal, as expected for X-mode propagation in a plasma, while its electromagnetic nature is confirmed by the fact that the space-averaged $\langle \delta B_z^2 \rangle = \langle \delta B_z \delta E_y \rangle$.

We note that the field-aligned component of the particle momentum u_z is not affected by the shock. The incoming particles are efficiently isotropized in the xy plane perpendicular to the field, but the post-shock particle distribution remains largely confined to this plane. It follows that the downstream effective adiabatic index corresponds to a 2D relativistically hot gas, $\Gamma_{\text{ad}} = 3/2$, instead of $4/3$ if the downstream plasma was isotropic in all momentum directions. The lack of isotropization is due to the fact that in a $\sigma \gg 1$ flow (with downstream plasma magnetically dominated) it will be harder for the plasma to exceed the threshold for velocity–space instabilities that feed off the particle temperature anisotropy. For example, the plasma will go unstable via the mirror mode if the temperature anisotropy is above a threshold that scales as $\propto \sigma$, which is harder to exceed at higher magnetizations. In addition, the 1D spatial geometry employed here will further suppress the growth of field-aligned modes leading to momentum isotropization.

The downstream particle energy spectrum (not shown) resembles a 2D Maxwell–Jüttner distribution whose temperature is slightly lower than the one expected from the Rankine–Hugoniot jump conditions (the difference is due to the energy transferred to the precursor waves). No non-thermal tail is observed for the runs presented in this work. This is in agreement with the inefficiency of particle acceleration expected at relativistic strongly magnetized perpendicular shocks (Sironi & Spitkovsky 2009; Lemoine & Pelletier 2010; Sironi et al. 2013; Sironi, Keshet & Lemoine 2015; Iwamoto et al. 2017; Pelletier et al. 2017; Plotnikov et al. 2018).

3.2 Precursor wave energy

We now focus on the dependence of the precursor properties, and specifically of its amplitude, on the upstream magnetization σ . The dependence on the upstream bulk Lorentz factor γ_0 will be discussed in the last part of this subsection.

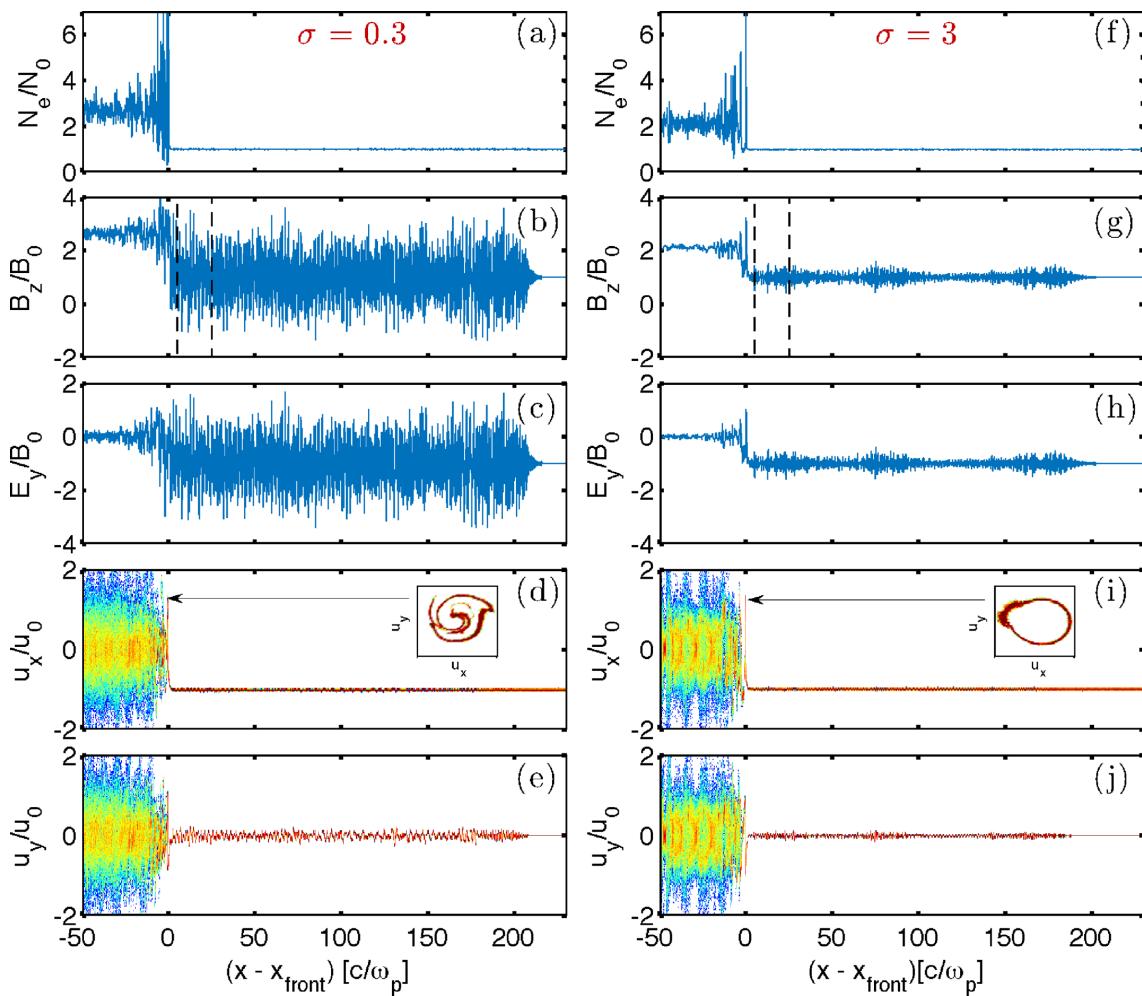


Figure 1. Structure of the shock layer from 1D PIC simulations of $\sigma = 0.3$ (left-hand column) and $\sigma = 3$ (right-hand column) shocks. From top to bottom it is shown: the electron number density N_e/N_0 (panels a and f), the transverse magnetic field B_z/B_0 (panels b and g), the motional transverse electric field E_y/B_0 (panels c and h), the longitudinal phase space of positrons $x - u_x$ (panels d and i), and the transverse phase space of positrons $x - u_y$ (panels e and j), respectively. Here, $u_x = \gamma \beta_x$ and $u_y = \gamma \beta_y$ are the components of the dimensionless four velocity. Small insets in the upper right-hand side of panels (d) and (i) show the particle distribution in momentum space $u_x - u_y$ at the location of the shock front, as indicated by the arrows. The shock front is located at $x - x_{\text{front}} = 0$; it propagates in the $+\hat{x}$ direction. The upstream flow is on the positive $x - x_{\text{front}} > 0$ side and the downstream plasma is on the negative $x - x_{\text{front}} < 0$ side.

3.2.1 Temporal evolution of the precursor wave

After an initial transient—whose duration depends on the upstream magnetization, as we show below—the intensity of the precursor wave settles to its asymptotic value. We measure the wave intensity in a region between 5 and $25 c/\omega_p$ ahead of the shock front:¹ $5 c/\omega_p < x - x_{\text{front}} < 25 c/\omega_p$. 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

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

¹This region is delimited by vertical black dashed lines in panels (b) and (g) of Fig. 1.

In Fig. 2 we show for different magnetizations the time evolution of the normalized wave intensity defined as

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

where $\delta E_y = E_y - E_{y,0} = E_y + \beta_0 B_0$. Different solid lines correspond to different values of the magnetization, from $\sigma = 0.1$ (blue line) to $\sigma = 10$ (black line), as indicated in the legend. By computing the temporal variation of the wave intensity, we can assess when the precursor wave has reached a steady state. Fig. 2 shows that:

- (i) With increasing σ , a longer time is required for the precursor to settle at its time-asymptotic state. This is due to the combination of two effects. First, the shock velocity increases from $\beta_{\text{sd}} = 0.476$ for $\sigma = 0.1$ to $\beta_{\text{sd}} = 0.963$ for $\sigma = 10$, so at higher magnetizations it takes more time for the precursor wave, propagating at $\beta_{\text{wave}} \approx 1$, to detach from the shock front. Second, there is some interaction

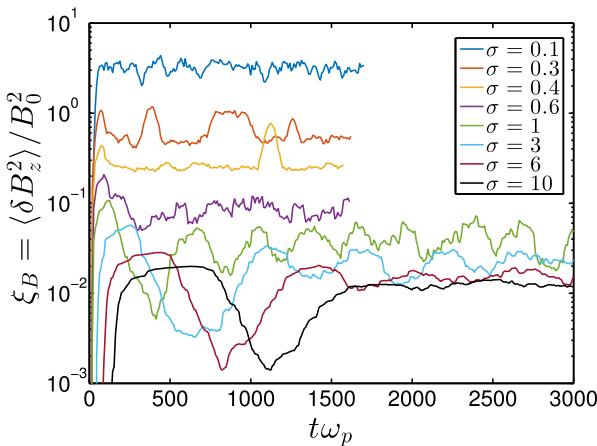


Figure 2. Time evolution of the normalized precursor wave energy, $\xi_B = \langle \delta B_z^2 \rangle / B_0^2$, for different values of the upstream magnetization σ . Lines of different colour correspond to a given σ going from 0.1 (blue line) to 10 (black line). The precursor wave energy was extracted from a $20 c/\omega_p$ wide slab located at $5 c/\omega_p < x - x_{\text{front}} < 25 c/\omega_p$.

occurring between the wave and the upstream plasma, which initially causes a drop in wave efficiency (e.g. at $\omega_p t \sim 1100$ in the black line of Fig. 2). The time required for the wave to self-regulate and settle to a steady state, following this drop, is longer for higher σ (e.g. compare green and black lines in Fig. 2).

(ii) The steady-state value of the normalized wave energy ξ_B decreases with increasing magnetization and for $\sigma \gg 1$ (cyan, brown, and black lines) it approaches a constant value $\xi_B \simeq 0.01$.

(iii) For $\sigma = 0.3$ (orange) and $\sigma = 0.4$ (yellow), which will be called ‘transition cases’ in the following, the wave intensity varies between periods of high efficiency and phases of low efficiency.

(iv) All the simulations have been evolved for long enough to reach a quasi-stationary state. For the largest explored magnetization $\sigma = 30$, the simulation was advanced beyond $2 \times 10^4 \omega_p^{-1}$ (this case is not shown in the figure but reported in subsequent figures).

Once the wave intensity has settled to a steady state, we have extracted a number of wave properties, such as the energy, the spectrum (peak frequency, low-frequency cutoff, and spectral width), and the wave strength parameter, as we now describe.

3.2.2 Dependence on the upstream magnetization

In the previous section we have defined the energy fraction in upstream field fluctuations ξ_B (see equation 3), as the ratio of the precursor wave magnetic energy to the background field energy. In order to get a global idea of the energetics, we also need to complement it with a parameter that quantifies the energy fraction of the incoming plasma (including both kinetic and electromagnetic content) that is radiated from the shock front in the form of precursor waves. When the electromagnetic field fluctuations induced by the precursor are taken into account in the jump conditions across the shock, the energy conservation equation expressed in the simulation frame is (Gallant et al. 1992; Plotnikov et al. 2018):

$$\begin{aligned} & \gamma_0^2 (\beta_0 + \beta_{\text{sl}}) \left(w_u + \frac{b_{0,u}^2}{4\pi} \right) - (1 - \beta_{\text{sl}}) \frac{\delta B_u^2}{4\pi} \\ &= \beta_{\text{sl}} \left(w_d - p_d + \frac{b_{0,d}^2}{8\pi} + \frac{\delta B_d^2}{4\pi} \right), \end{aligned} \quad (4)$$

where the subscripts ‘*u*’ and ‘*d*’ refer to the upstream and downstream regions, respectively. Here, w_i and $b_{0,i}$ are, respectively, the fluid enthalpy density and mean magnetic field, both measured in the fluid rest frame. As above, β_{sl} is the velocity of the shock front as measured in the downstream frame of the simulations. The fluctuating components δB_u and δB_d are measured in the simulation frame. We have made the approximation of negligible thermal pressure upstream (strong shock limit) and we have assumed that electrostatic effects are negligible both upstream and downstream ($\delta E_x \approx 0$). The latter approximation is fully supported by the simulations and, more fundamentally, by the fact that space-charge effects are expected to be negligible in pair plasmas. The strong shock limit means that the upstream plasma pressure can be neglected and the upstream fluid enthalpy density (in the fluid rest frame) is then $w_u = n_u m_e c^2$, where n_u is the upstream plasma proper density. The mean upstream magnetic field $b_{0,u}$ in the upstream frame is related to the pre-shock magnetic field $B_{z,0}$ in the simulation frame via a Lorentz boost: $B_{z,0} = \gamma_0 b_{0,u}$. Hence, the magnetization parameter can be rewritten as $\sigma = b_{0,u}^2 / (4\pi w_u)$ and we note that $\delta B_u^2 / b_{0,u}^2 = \gamma_0^2 \delta B_u^2 / B_{z,0}^2 = \gamma_0^2 \xi_B$, because the fluctuating part is measured directly in the simulation frame.

As we focus on the precursor wave propagating upstream, here we only consider the left-hand side of equation (4).² We find that the fraction of total incoming energy (including both particle and electromagnetic contributions) that is channelled into the precursor wave can be expressed as

$$f_\xi = \xi_B \left(\frac{\sigma}{1 + \sigma} \right) \left(\frac{1 - \beta_{\text{sl}}}{\beta_0 + \beta_{\text{sl}}} \right), \quad (5)$$

as seen from the DRF. In the following, f_ξ will be identified as the ‘energy fraction parameter’. It is also convenient to define the fraction of incoming particle kinetic energy that is converted into precursor emission

$$g_\xi = f_\xi (1 + \sigma). \quad (6)$$

The time-asymptotic values of ξ_B , β_{sl} , and f_ξ measured in our simulations are presented in panels (a), (b), and (d) of Fig. 3, as a function of magnetization. Error bars indicate the standard deviation of our time measurements. As regard to ξ_B , we observe a rapid decrease from $\xi_B \simeq 3.53$ at $\sigma = 0.1$ down to $\xi_B \simeq 0.04$ at $\sigma = 1$. The inflection point of the transition occurs at $\sigma \simeq 0.35$. This decrease accompanies a change in the shock front structure that for $\sigma > 1$ presents a coherent soliton-like shape (compare left-hand and right-hand columns in Fig. 1 at the shock). For $\sigma > 1$, ξ_B slowly decreases and eventually approaches a constant value $\xi_B \simeq 10^{-2}$.

Concerning the shock front speed and its corresponding bulk Lorentz factor, β_{sl} and γ_{sl} , panels (b) and (c) demonstrate an excellent agreement between our measured values, plotted with blue symbols, and the predictions of ideal MHD jump conditions (e.g. Appendix B of Plotnikov et al. 2018), as indicated by the red solid line. The front speed increases from $\beta_{\text{sl}} = 0.476c$ for $\sigma = 0.1$ up to $\beta_{\text{sl}} = 0.987c$ for $\sigma = 30$. The Lorentz factor of the shock front tends asymptotically to $\gamma_{\text{sl}} = \sqrt{\sigma}$, for $\sigma \gg 1$. We remark that the MHD equations used here to derive the jump conditions do not incorporate modifications due to the precursor wave. The accurate agreement of our results with ideal MHD jump conditions for $\sigma \gg 0.1$ is then due to the fact that at high magnetizations the precursor wave

²Since we forego the discussion of the downstream part of the energy conservation equation, an interested reader will find details in the aforementioned works (Gallant et al. 1992; Plotnikov et al. 2018).

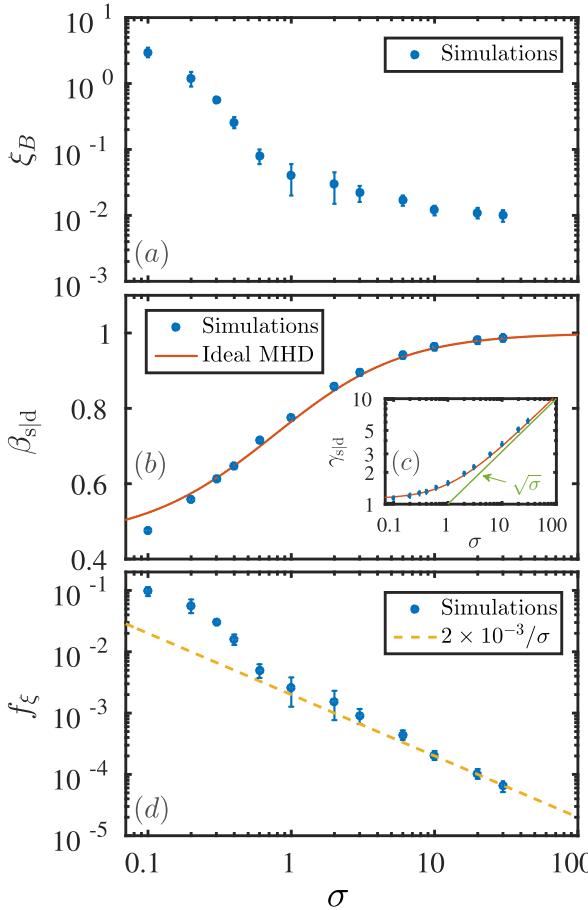


Figure 3. Dependence on σ of the time-asymptotic values of the normalized wave intensity ξ_B (panel a), of the shock front speed β_{sld} (panel b), of the corresponding Lorentz factor γ_{sld} (panel c), and of the energy fraction parameter f_ξ (panel d). The red lines in panels (b) and (c) show the theoretical expectation based on ideal MHD jump conditions (which do not include effects from the precursor waves). The green line in panel (c) shows the asymptotic scaling $\gamma_{sld} = \sqrt{\sigma}$ expected for $\sigma \gg 1$. The yellow dashed line in panel (d) follows the scaling $f_\xi = 2 \times 10^{-3}/\sigma$.

is relatively weak, and it does not have an appreciable dynamical effect on the shock. In contrast, in the case when the precursor wave is the strongest, $\sigma = 0.1$, the agreement is the worst, because the emission of the large amplitude wave can slow down the shock front, as compared to the ideal MHD prediction.

The dependence on σ of the energy fraction parameter f_ξ is presented in panel (d) of Fig. 3. It was calculated by plugging the values from panels (a) and (b) into equation (5). It shows that the energy fraction in the precursor wave decreases from 10 per cent for $\sigma = 0.1$ down to 0.0065 per cent for $\sigma = 30$. The dashed orange line follows the empirical scaling $f_\xi = 2 \times 10^{-3}/\sigma$ that satisfactorily fits our measured values in the $\sigma > 1$ range. The most noticeable result of this panel is that we observe a well-defined scaling $f_\xi \propto \sigma^{-1}$. This result arises from the fact that for $\sigma \gg 1$, the normalized wave intensity ξ_B is roughly constant and $\beta_{sld} \simeq 1 - 1/(2\sigma)$. It follows that in the limit $\sigma \gg 1$ the precursor wave carries a constant fraction of the incoming particle kinetic energy, i.e. $g_\xi \simeq 2 \times 10^{-3}$.

Let us emphasize, however, that this σ -dependence of f_ξ and g_ξ is derived in the DRF (simulation frame). This dependence will be different in the shock front rest frame (SRF), since the front moves

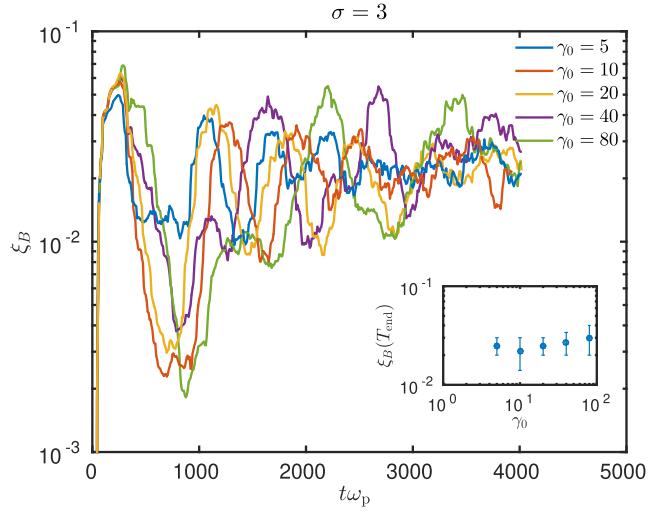


Figure 4. Time evolution of the normalized wave intensity ξ_B for different values of the upstream flow Lorentz factor γ_0 at fixed $\sigma = 3$. The blue, red, orange, violet, and green lines correspond to $\gamma_0 = 5, 10, 20, 40$, and 80 , respectively. The time-asymptotic values of ξ_B , with corresponding error bars, are plotted in the inset at the lower right-hand corner of the figure. Within the error bars, ξ_B is nearly independent from γ_0 .

with ultrarelativistic speeds for $\sigma \gg 1$. This point will be further discussed in Section 4.

3.2.3 Dependence on the upstream bulk Lorentz factor

So far, we have investigated the dependence of the precursor intensity on σ , for a fixed choice of the upstream flow Lorentz factor $\gamma_0 = 10$. Here, we demonstrate that ξ_B is essentially independent from γ_0 for any value of σ . Let us first consider the dependence at a fixed σ . In Fig. 4 we show the time evolution of the precursor wave energy for $\sigma = 3$, when varying γ_0 from 5 to 80. Lines of different colour correspond to different values of γ_0 . Despite large oscillations in time, it appears that ξ_B converges to the same value regardless of γ_0 . The time-asymptotic values of ξ_B , with corresponding error bars, are plotted in the figure inset. Within the error bars, we can assert that there is no obvious dependence on γ_0 .

In order to generalize this conclusion to any σ , it is worth noting that in the seminal study of Gallant et al. (1992), two very different values of the bulk Lorentz factor ($\gamma_0 = 40$ and 10^6) were used, for a range of $\sigma \in [10^{-3}, 5]$. The authors did not notice any dependence on γ_0 . Also, Iwamoto et al. (2017) performed 1D simulations with $\gamma_0 = 40$ and explored σ values between 10^{-3} and 0.5, finding similar values as in Gallant et al. (1992).

In the Appendix A, Fig. A2 shows the values of ξ_B obtained for $\sigma \in [10^{-3}, 1]$ (horizontal axis) and for γ_0 ranging from 10 to 10^6 (different data sets). This figure shows that in the low magnetization regime $\sigma \in [10^{-3}, 0.3]$, the normalized wave intensity ξ_B is nearly independent from γ_0 . In the range $\sigma \in [0.3, 1]$ there is a larger scatter among different data sets (which employ different γ_0). This range of magnetizations corresponds to the transition cases (see Fig. 2). The most plausible reason for the discrepancy among different data sets is that the simulations from earlier studies were not evolved long enough in order to reach the asymptotic state of the transition cases, so the value of ξ_B was not yet stabilized (see Fig. 2). In fact, Fig. 4 shows that even at $\sigma > 1$ the time-asymptotic value of ξ_B is insensitive to the flow Lorentz factor.

3.3 Precursor spectrum

After discussing the wave energy, we now address the dependence on σ of the precursor spectrum and of the typical wavelength of the emission. Our results will be presented in the downstream frame of the simulations. It is important to note, however, that the wave propagates in the upstream plasma and that its emitter is the shock front. Both move with respect to the simulation frame. Hence, when comparing simulation results with the expected scalings, we need to consider the wave dispersion relation first in the upstream frame, and then transform it to the DRF. Also, the typical emission frequency is most naturally estimated in the SRF, and then it should be transformed to the DRF in order to compare with simulation results.

In this section, we first present basic analytical considerations and then we compare them with our simulation results. A special feature of $\sigma > 1$ shocks, where a density and magnetic field cavity is observed in the front structure, is discussed at the end of this section. As we argue, the cavity is instrumental in setting the precursor power and determining its dominant frequency.

3.3.1 Basic considerations

As discussed above, the precursor wave possesses X-mode (extraordinary-mode) polarization, such that its wave vector is perpendicular to \mathbf{B}_0 , its fluctuating magnetic field is parallel to \mathbf{B}_0 , and its fluctuating electric field is perpendicular to both \mathbf{k} and \mathbf{B}_0 . Some basic properties of the extraordinary mode in the context of the shock emission were derived by Gallant et al. (1992) and Iwamoto et al. (2017). We reproduce here their estimations for completeness.

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

$$\frac{k'^2 c^2}{\omega'^2} = 1 - \frac{\omega_p'^2}{\omega'^2 - \omega_c'^2} = 1 - \frac{\omega_p'^2}{\omega'^2 - \sigma \omega_p'^2}, \quad (7)$$

where double primed quantities are measured in the upstream rest frame (URF). Using Lorentz transformations for ω and k and in the limit $\gamma_0^2 \gg \sigma$, the dispersion relation in the DRF becomes

$$k^2 c^2 \simeq \omega^2 - \omega_p^2. \quad (8)$$

Interestingly, as long as $\gamma_0^2 \gg \sigma$, this is identical to the dispersion relation of a simple electromagnetic wave propagating in an unmagnetized plasma.

The motion of the shock front imposes a cutoff frequency below which the wave cannot escape into the upstream medium. It follows that little or no power should be observed in the upstream precursor spectrum below the cutoff frequency. This cutoff frequency is obtained by equating the group velocity of the wave, $d\omega/dk$, with the shock front velocity as:

$$c \sqrt{1 - \frac{\omega_p^2}{\omega^2}} = \beta_{s|d} c. \quad (9)$$

This relation leads to the cutoff frequency and wavelength

$$\omega_{\text{cutoff}} = \gamma_{s|d} \omega_p \quad (10)$$

$$\lambda_{\text{cutoff}} = \frac{2\pi c}{\gamma_{s|d} \beta_{s|d} \omega_p}. \quad (11)$$

As regard to the characteristic frequency of the precursor wave, the most natural assumption is that it corresponds to the collective cyclotron motion of the bunching particles at the shock front, which we now evaluate. First, the magnetic field at the shock can be roughly estimated by assuming that, in the shock frame, all the momentum of the incoming particles is stored in the magnetic field at that point (Alsop & Arons 1988):

$$\frac{B'_{\text{sh}}}{B'_0} \approx \sqrt{1 + \frac{2}{\sigma}}, \quad (12)$$

where primed quantities are measured in the SRF. More detailed considerations on the soliton structure of the shock, as presented by Alsop & Arons (1988), give a similar expression for B'_{sh} . For particles with Lorentz factors comparable to the upstream bulk Lorentz factor, the ratio of the expected emission frequency (which we label ‘sol’ since it is emitted by the soliton at the shock) to the upstream cyclotron frequency is then equal to the magnetic field enhancement ratio, $\omega'_{c,\text{sol}}/\omega_c = B'_{\text{sh}}/B'_0$.³ Lorentz transforming to the DRF ($\omega'_{c,\text{sol}} \rightarrow \omega_{c,\text{sol}}$) and using the dispersion relation in equation (8) leads to

$$\omega_{c,\text{sol}} \approx \left(\sqrt{\sigma + 2} + \sqrt{\sigma + 2 - \beta_{s|d}^2} \right) \gamma_{s|d} \omega_p. \quad (13)$$

Based on these arguments, we expect the precursor spectrum to exhibit a low-frequency cutoff at ω_{cutoff} and prominent line-like features at $\omega_{c,\text{sol}}$ and its harmonics. As we show below, where we compare these scalings with our simulation results, for $\sigma > 1$ the predicted $\omega_{c,\text{sol}}$ systematically overestimates the observed peak frequency ω_{peak} . In Section 3.3.3, we propose a new model for the precursor peak frequency in the high-magnetization regime, and we show that it is in good agreement with our simulation results.

3.3.2 Spectrum dependence on the upstream magnetization

To characterize the spectrum of the precursor wave, we have employed two complementary diagnostics, one spatial and one temporal. They were used to construct the wavenumber spectrum (k -spectrum) and the frequency spectrum (ω -spectrum), respectively.

The wavenumber spectrum was calculated by extracting the spatial profile of $B_z(x) - B_{z,0}$ in the region located at $5c/\omega_p < x - x_{\text{front}} < 105c/\omega_p$, at a time when the precursor has reached the steady state, and then computing its Fourier transform. The frequency spectrum was constructed by recording the temporal variation of $B_z(t) - B_{z,0}$ at one selected grid point in the upstream region, during a time interval of $100\omega_p^{-1}$, and then calculating its Fourier transform. The spatial window for the k -spectrum and the time interval for the ω -spectrum are chosen so that roughly the same segment of the precursor wave was analysed in the two cases.

In Fig. 5 we present the spectrum of the precursor wave for different σ . Five representative cases are shown from top to bottom, $\sigma = 0.1, 0.3, 1, 3$, and 10 , respectively. Each panel contains the k -spectrum, plotted using blue solid lines, and the ω -spectrum, plotted using red dashed lines. For the ω -spectrum, the horizontal axis shows ω/ω_p , whereas for the k -spectrum we take $(k^2 c^2/\omega_p^2 + 1)^{1/2}$. Due to this choice, and given the dispersion relation in equation (8), each wavenumber spectrum should nearly overlap with the corresponding frequency spectrum, as

³There is no prime on the upstream cyclotron frequency as it is Lorentz-invariant for perpendicular shocks.

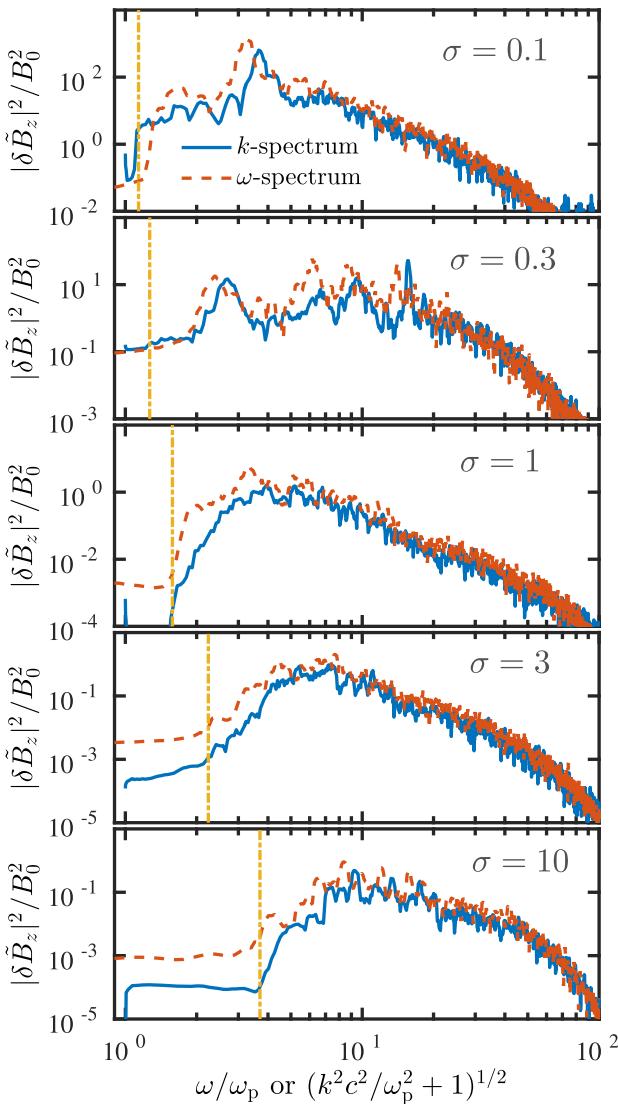


Figure 5. Spectrum of the precursor wave for different σ . Five representative cases are shown from top to bottom, $\sigma = 0.1, 0.3, 1, 3$, and 10 , respectively. The blue solid lines in each panel show the spectrum in k -space and red dashed lines show the same in ω -space. The horizontal axis shows ω/ω_p for the ω -spectrum. For the k -spectrum, the choice for the horizontal axis is motivated by the dispersion relation in equation (8), so that the k -spectrum should nearly overlap with the corresponding ω -spectrum. The method to compute the spectra is described in the main text. All the spectra are normalized as $\int |\delta\tilde{B}_z(k)|^2/B_0^2 dk = \int |\delta\tilde{B}_z(\omega)|^2/B_0^2 d\omega = \xi_B$. The orange vertical lines mark the position of the expected low-frequency cutoff, as given by equation (10).

it is indeed the case. The spectra were normalized such that $\int |\delta\tilde{B}_z(k)|^2/B_0^2 dk = \int |\delta\tilde{B}_z(\omega)|^2/B_0^2 d\omega = \xi_B$.

In each spectrum, the power drops rapidly below the cutoff frequency given by equation (10), which is indicated by an orange vertical line in each panel. This is expected, since for lower frequencies (or wavenumbers) the group velocity is smaller than the shock speed, so the wave cannot propagate ahead of the shock.

The spectra are narrow-band, but they are not consistent with a unique line, as it would be expected for cyclotron emission. This is due to the fact that the ring-like particle distribution at the shock front possesses ultrarelativistic energies. The emission is then controlled not by the non-relativistic cyclotron maser, but rather

by the ultrarelativistic synchrotron maser instability, that generates a large number of harmonics with comparable growth rate to the fundamental (Hoshino & Arons 1991).

Prominent line-like features are observed at $\sigma < 1$, with the fundamental at $\omega = \omega_{c, \text{sol}}$ or the second harmonic dominating the spectrum at low magnetizations (see the peak at $\omega \simeq 4\omega_p$ for $\sigma = 0.1$). In the transition cases with $0.1 < \sigma < 1$, we observe the generation of very strong harmonics up to $N = 5$, where $N = \omega/\omega_{c, \text{sol}}$, with high-order harmonics producing stronger lines than the fundamental (see the case with $\sigma = 0.3$). For $\sigma > 1$ the spectrum shows much less prominent lines. As we will argue later, supplementary amplification mechanisms operate in this regime, and the characteristic frequency $\omega_{c, \text{sol}}$ given by equation (13) no longer controls the location of the spectral peak.

The dependence of the relevant wavelengths and frequencies on the magnetization is presented in Fig. 6. The top row refers to wavelengths, the bottom row to frequencies. The left-hand column shows the variation with σ of the cutoff wavelength λ_{cutoff} (panel a) and cutoff frequency ω_{cutoff} (panel b). The values derived from our simulations are plotted using blue circles, and they are in very good agreement with the analytical predictions of equations (10) and (11), indicated by the black dashed lines. The only exception is the transition case $\sigma \simeq 0.3$, where the low-frequency cutoff is non-stationary.

The central column (panels c and d) presents the variation with σ of the peak wavelength and frequency (red squares are the results of our simulations), defined as the location where the precursor spectrum peaks (see Fig. 5). The black dashed lines indicate the expectation for soliton emission (equation 13). It is apparent that the peak values obtained in the simulations do not agree with the analytical estimate given by equation (13) for any $\sigma > 0.1$.⁴ To understand the disagreement we define two regimes: (i) the transition cases ($0.1 < \sigma < 1$) and (ii) the magnetically dominated cases ($\sigma > 1$).

In case (i), high-order harmonics in the precursor spectrum are stronger than the fundamental, and the spectral peak is not at the fundamental frequency. If we artificially select the lowest frequency corresponding to a local maximum in the spectrum, we find that its location is in reasonable agreement with the expected fundamental frequency $\omega_{c, \text{sol}}$ (see top two panels in Fig. 5). In case (ii), we do not find evidence of any strong line at the expected $\omega_{c, \text{sol}}$ or its harmonics, but rather we observe less prominent lines at frequencies that have no clear connection with $\omega_{c, \text{sol}}$. The measured peak frequency scales as $\omega_{\text{peak}} \approx 3\sigma^{1/2}\omega_p$. In contrast, from equation (13) we would expect a stronger scaling with σ , since $\omega_{c, \text{sol}} \rightarrow \sigma\omega_p$ in the limit $\sigma \gg 1$. As discussed in the next subsection, we attribute the observed scaling to the presence of a resonant cavity in the shock structure that builds up only for $\sigma > 1$. We show below that the peak wavelength in case (ii) corresponds to an eigenmode of the cavity, and it is roughly three times shorter than the cavity width (see the green dot-dashed lines in panels c and d).

The right-hand column (panels e and f, respectively) presents the dependence on σ of the fractional spectral width in wavelength and frequency space ($\Delta\lambda/\lambda_{\text{peak}}$ and $\Delta\omega/\omega_{\text{peak}}$, respectively). The width $\Delta\omega$ is the difference between the two frequencies (one above the peak frequency ω_{peak} and one below) where the power drops by a factor of 30 below the peak. The width $\Delta\lambda$ is defined in

⁴We note, however, that in simulations with $\sigma < 0.1$, not presented here, we have obtained a very good agreement between the measured ω_{peak} and $\omega_{c, \text{sol}}$ given in equation (13) (see also Gallant et al. 1992).

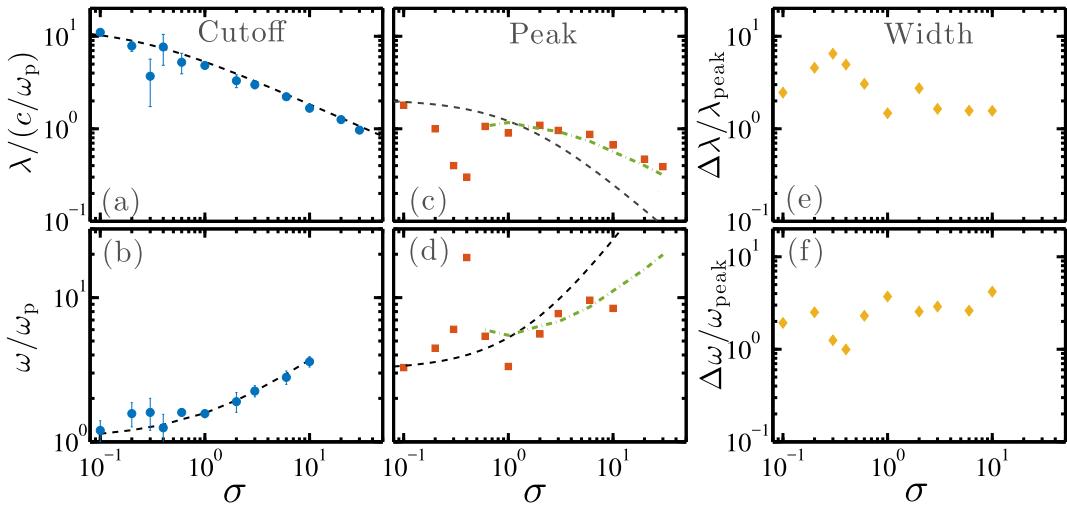


Figure 6. Characteristic wavelengths and frequencies as a function of σ . Cutoff wavelengths λ_{cutoff} and frequencies ω_{cutoff} are presented in the left-hand column (blue circles in panels a and b, respectively). The dashed black lines show the analytical predictions in equations (10) and (11). The wavelengths and frequencies at the peak of the spectrum (λ_{peak} and ω_{peak}) are plotted in the central column (panels c and d) using red squares. The black dashed lines in panels (c) and (d) show the expected $\omega_{\text{c, sol}}$ from equation (13). The green dot-dashed line in panel (c) is the width of the density cavity at the shock front divided by three: $L_{\text{cav}}/3$ (see the text, subsection 3.3.3). The right-hand column presents the dependence on σ of the fractional spectral width $\Delta\lambda/\lambda_{\text{peak}}$ and $\Delta\omega/\omega_{\text{peak}}$ (panels e and f, respectively).

an analogous way. This shows quantitatively that the spectrum is narrow, with $\Delta\omega/\omega_{\text{peak}} \lesssim 3$ nearly independently of σ . The spectra of the cases with $\sigma < 1$, that show pronounced line-like features, are even narrower, with line widths of $\Delta\omega/\omega_{\text{peak}} \lesssim 1$ (see the top two panels in Fig. 5).

3.3.3 Resonating cavity in the shock structure at $\sigma > 1$

In the previous subsection we have found that in the magnetically dominated regime $\sigma > 1$, the peak frequency in our simulations does not scale as the expected gyration frequency in the soliton, $\omega_{\text{c, sol}}$. The physical picture that led to the estimate of $\omega_{\text{c, sol}}$ must then be revised, since the shock structure for $\sigma > 1$ appears to be different than for lower magnetizations. In fact, instead of one density peak defining the shock front, as it is the case in the $\sigma \lesssim 1$ regime, we observe for $\sigma \gtrsim 1$ the build-up of two density peaks separated by a cavity.⁵ As we now argue, it appears that the density cavity plays an essential role in amplifying the precursor emission and in selecting a well-defined wavelength for the precursor waves that corresponds to an eigenmode of the cavity.

In Fig. 7 we illustrate the structure of the shock transition region for $\sigma = 10$ at a well-advanced stage of the simulation when the precursor power has reached a steady state. Panel (a) of this figure shows the profile of the electron density (blue line) and of the transverse magnetic field B_z/B_0 (red line). The shock front is located at $x - x_{\text{shock}} = 0$ and it propagates in the $+\hat{x}$ direction. The two density peaks near the shock are separated by a cavity of width $L_{\text{cav}} \simeq 1.6 c/\omega_p$, just behind the shock front. The magnetic field profile peaks at the positions of the two density spikes, but in addition it exhibits a wave-like pattern within the density cavity.

⁵We believe that the structure of the shocks studied here is controlled by wave dispersion (rather than dissipation), given the importance of the precursor emission from the shock. For $\sigma > 1$, the amount of dispersion provided by the leading soliton becomes insufficient to sustain the shock structure, and a secondary soliton forms to provide additional dispersion.

For this particular snapshot, only a mode with wavelength $\lambda = L_{\text{cav}}/2$ is clearly seen in the cavity. However, the cavity is dynamic in nature, and different eigenmodes are distinctly seen at different times.

In panel (b) of Fig. 7 we demonstrate the role of the cavity in shaping the precursor spectrum by showing the wavenumber spectrum as a function of $\lambda^{-1} = k/(2\pi)$. Some characteristic emission wavelengths are easily identified. For instance, the cutoff wavelength at $\lambda_{\text{cutoff}} \simeq 1.6 c/\omega_p$ seems to be closely related to the width of the density cavity L_{cav} , which is indicated by a vertical dashed blue line. The other two vertical lines (red and orange, respectively) correspond to wavelengths equal to $L_{\text{cav}}/2$ and $L_{\text{cav}}/3$, respectively. The latter matches well the position of the strongest emission line. As discussed below, this holds for all $\sigma \gtrsim 1$.

To assess the connection between the cavity size and the precursor efficiency we show in panel (c) the time evolution of L_{cav} (blue line) and of the precursor wave energy ξ_B multiplied by a factor of 100 (red line). The value of ξ_B was computed in a region closer to the shock front than we have done before (here, between 1 and $5 c/\omega_p$ ahead of the front), which allows to probe more directly the causal connection between the precursor efficiency and the instantaneous shock structure. This panel shows that the cavity width (blue line) initially increases, then it decreases, and finally settles to a steady state. The time evolution of the precursor efficiency appears to be anticorrelated to the cavity width: when the cavity size is larger the emitted precursor is weaker (no amplification), and the wave intensity settles to a steady state at the same time ($\omega_p t \sim 1000$) as the cavity width. We interpret this behaviour as a self-regulation in the shock structure, such that the cavity width self-tunes to the value where it can efficiently channel the precursor emission into the upstream, i.e. L_{cav} has to be roughly equal to λ_{cutoff} (see also panel b). When this condition is met, the wave is amplified and its efficiency settles to the steady state. The critical role of the cavity for efficient wave emission is also revealed by inspecting the shock profile at the time when the precursor intensity sharply increases, right before settling to a steady state ($\omega_p t \sim 1000$): we see that large B_z fluctuations are first amplified in the cavity, and the emission of

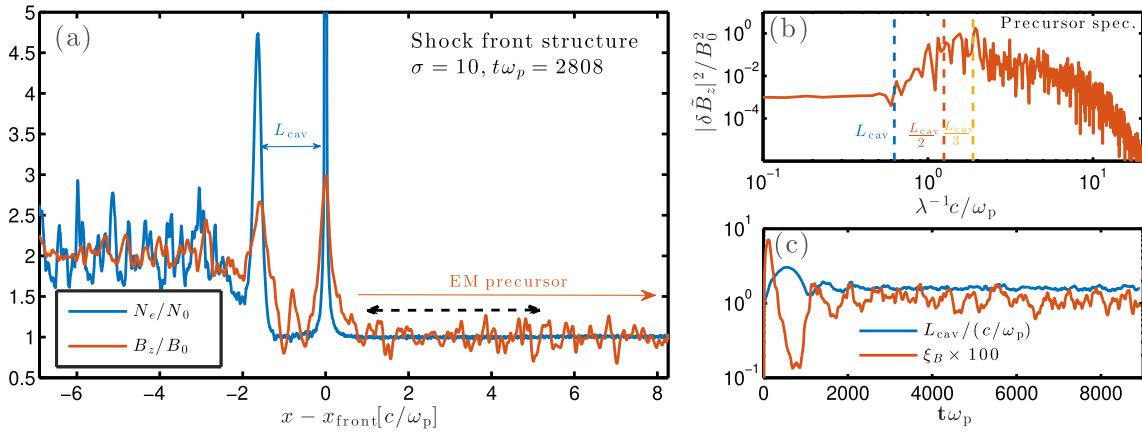


Figure 7. Structure of the shock transition region for $\sigma = 10$. Panel (a): profile of the electron density N_e/N_0 (blue line) and of the transverse magnetic field B_z/B_0 (red line). Panel (b): precursor spectrum in k -space, $|\delta B_z(k)|^2/B_0^2$, as a function of the inverse wavelength $\lambda^{-1} = k/(2\pi)$. The three vertical dashed lines correspond to the cavity width L_{cav} (blue), to $L_{\text{cav}}/2$ (orange), and to $L_{\text{cav}}/3$ (red). Panel (c): time evolution of L_{cav} (blue line) and of the precursor wave energy ξ_B multiplied by 100 (red line). The value of ξ_B was derived in the region closer to the shock front than previously (between 1 and $5c/\omega_p$ ahead of the front). The efficiency settles to a steady state at the same time as the cavity length does.

a strong precursor propagating upstream is then the consequence of partial transmission of these waves from the cavity through the leading soliton.

The validity of our ‘resonating cavity’ interpretation is tested in Fig. 6 (panels c and d), where we show that the peak wavelength of the precursor emission (red squares) is consistent with $L_{\text{cav}}/3$ (green dot-dashed lines in panel c), for all $\sigma \gtrsim 1$. In other words, for magnetically dominated plasmas the wave amplification inside the cavity plays an important role in selecting the dominant wavelength of the emitted precursor as an eigenmode of the cavity. It follows that the peak frequency for $\sigma \gtrsim 1$ scales as $\omega_{\text{peak}} \simeq 3\omega_{\text{cutoff}} \simeq 3\sqrt{\sigma}\omega_p \simeq 3\omega_c$ in the simulation frame, where we have used that $\gamma_{\text{std}} \simeq \sqrt{\sigma}$ for $\sigma \gg 1$. This should be contrasted with equation (13), whose scaling ($\propto \sigma\omega_p$ in the $\sigma \gg 1$ limit) is not supported by our simulations.

3.4 Wave strength parameter

The wave strength parameter (also known as ‘wiggler’) measures the dynamical effect of the propagating wave on the background plasma. It is defined through the equation of motion of particles in a high-amplitude wave (Lyubarsky 2006; Iwamoto et al. 2017):

$$u_x = \gamma\beta_x = -\gamma_0\beta_0 + \frac{a^2}{2} \cos^2[\omega(x/c + t)], \quad (14)$$

$$u_y = \gamma\beta_y = a \cos[\omega(x/c + t)], \quad (15)$$

where

$$a = \frac{e\delta E_y}{m_e c \omega} \quad (16)$$

is the strength parameter of the wave. δE_y is the electric field of the wave and ω is the wave frequency; e and m_e are the particle charge and mass, respectively. When $a > 1$, the particle quiver motion becomes relativistic and the plasma back-reacts strongly on to the wave. We have employed two measures for the wave strength parameter: either from the maximum excursion in u_y , $a_{\text{max}} = \max(u_y)$; or from the root mean square value, a_{std} . These choices are motivated by the form of equation (15), where the wiggler parameter controls the y -oscillations of the particle 4-velocity. In either case,

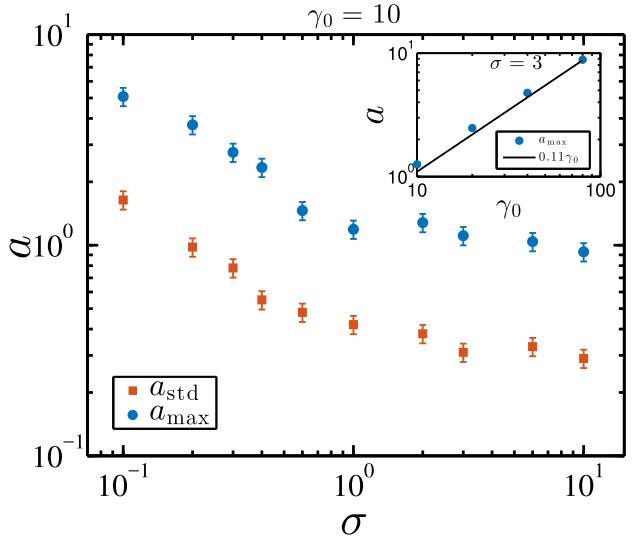


Figure 8. Dependence on σ of the wave strength parameter a . The two colours refer to different measurements of a : the blue symbols correspond to the maximum value of u_y in the precursor region (between 5 and $105 c/\omega_p$ ahead of the front) and red symbols correspond to the root mean square value of u_y in the same region. The inset presents the dependence of a_{max} on γ_0 for $\sigma = 3$. The solid black line shows a linear scaling.

we have extracted the measurement from the region between 5 and $105 c/\omega_p$ ahead of the front at the final time of the simulations.

Fig. 8 presents the dependence on σ of the wave strength parameter derived from our simulations. The maximal value a_{max} decreases from $\simeq 5$ for $\sigma = 0.1$ down to $\simeq 1$ for $\sigma = 10$, while the root mean square value has the same dependence on σ but it is three times smaller, $a_{\text{std}} \simeq a_{\text{max}}/3$. The subpanel of this figure shows the dependence on γ_0 . Supplementary simulations were performed for this purpose, where we fixed $\sigma = 3$. There is a clear linear dependence of a on γ_0 , as already suggested by Iwamoto et al. (2017).

The linear dependence on γ_0 arises naturally from the fact that ξ_B does not depend on γ_0 , combined with the fact that the typical frequency of the precursor wave is $\simeq 3\omega_p$ for $\sigma \ll 0.1$ and $\simeq 3\omega_c$

for $\sigma > 1$. It follows from equation (16) that $a \approx \gamma_0 \sqrt{\xi_B \sigma}$ for $\sigma \ll 0.1$ and $a \approx \gamma_0 \sqrt{\xi_B}/3$ for $\sigma > 1$, which justifies the linear scaling with γ_0 shown in the inset of Fig. 8.

The wiggler parameter is Lorentz-invariant under transformations along the shock propagation direction, in virtue of equation (15). Alternatively, one can note that the electric field of the wave transforms in the same way as its frequency. Values presented in Fig. 8 will then be the same in the SRF and in the URF. This is in contrast to the precursor normalized energy ξ_B and the precursor spectrum, which are frame-dependent.

4 ENERGETICS IN THE SHOCK FRONT REST FRAME

The SRF is, by definition, the frame where the shock is stationary. In this frame the upstream plasma flows along the shock normal with a negative velocity in the x direction (whose magnitude is larger than in the DRF). The downstream plasma recedes from the front along the negative x direction. This frame can be naturally employed to quantify the incoming (and outgoing) momentum and energy, and so to derive the energy fraction channelled into the precursor wave.

4.1 From the simulation frame to the shock rest frame

So far, all the quantities related to the precursor waves have been given in the DRF, so we need to Lorentz transform them to the SRF. We will employ primed variables for the SRF. The amplitude of the mean magnetic field transforms as

$$B'_0 = \gamma_{s|d} (B_0 + \beta_{s|d} E_0) = \gamma_{s|d} B_0 (1 + \beta_{s|d} \beta_0), \quad (17)$$

where we have used the shortcut notations $B_0 = B_{z,0}$ and $E_0 = E_{y,0}$. Since by transforming into the shock frame we are ‘catching up’ with the precursor wave, the precursor amplitude will decrease as

$$\delta B'_z = \gamma_{s|d} (\delta B_z - \beta_{s|d} \delta E_y) \simeq \gamma_{s|d} \delta B_z (1 - \beta_{s|d}). \quad (18)$$

The ξ_B parameter then transforms as (Gallant et al. 1992)

$$\begin{aligned} \xi'_B &= \xi_{B|sh} = \frac{\langle \delta B_z'^2 \rangle}{B_0'^2} = \left(\frac{1 - \beta_{s|d}}{1 + \beta_0 \beta_{s|d}} \right)^2 \frac{\langle \delta B_z'^2 \rangle}{B_0'^2} \\ &= \left(\frac{1 - \beta_{s|d}}{1 + \beta_0 \beta_{s|d}} \right)^2 \xi_B. \end{aligned} \quad (19)$$

We compute directly ξ'_B with the following procedure. The values of $\gamma_{s|d}$ and $\beta_{s|d}$ obtained from our simulations (see Fig. 3, panels b and c) are used to Lorentz transform the electromagnetic fields into the SRF at a given snapshot of the simulation. Then, ξ'_B is computed directly, by averaging between 5 and 105 c/ω_p ahead of the front (the distance is still measured in the simulation frame).

In Fig. 9 (panel a) we present the dependence of $\xi'_B = \xi_{B|sh}$ on σ , obtained independently with the three PIC codes used in this study: orange squares for SHOCKAPIC, red circles for SMILEI, and blue diamonds for TRISTAN-MP. First, the figure demonstrates excellent agreement between the three codes. Secondly, it shows that, beyond the transition cases with $0.1 < \sigma < 1$, where ξ'_B attains the largest values, the normalized wave energy in the SRF scales as $\xi'_B \simeq 7 \times 10^{-4} \sigma^{-2}$ for $\sigma > 1$. This scaling is plotted with a dashed black line, and it can be easily justified. In fact, in Section 3 we have shown that for $\sigma \gg 1$ the wave amplitude in the DRF converges to a constant (i.e. σ -independent) value, $\xi_B \simeq 10^{-2}$. In addition, the asymptotic shock velocity in the DRF is $\beta_{s|d} \simeq 1 - 1/(2\sigma)$ for $\sigma \gg 1$. Plugging these two scalings into equation (19) leads to $\xi'_B = 6.3 \times 10^{-4} \sigma^{-2}$, which is very close to the measured scaling.

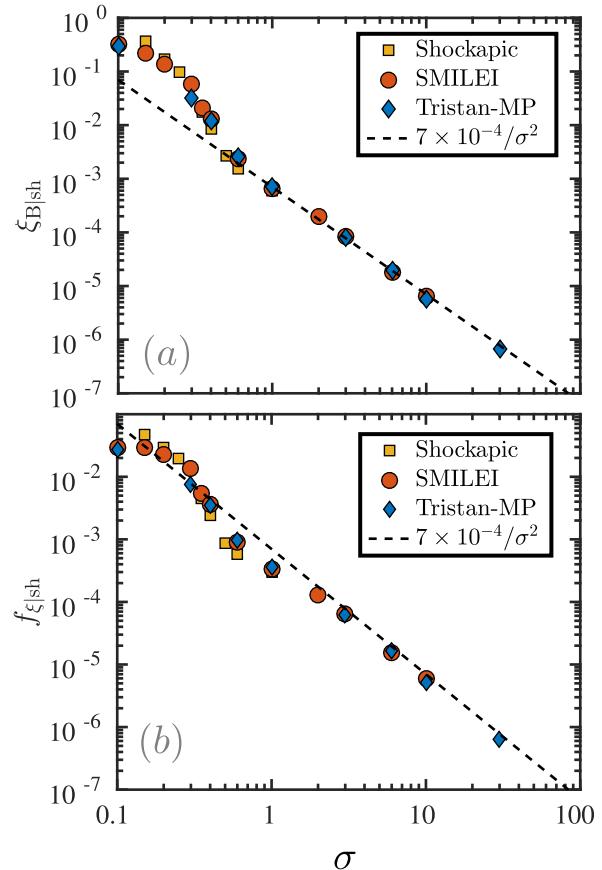


Figure 9. Energetics in the shock front rest frame: dependence on σ of the normalized precursor wave energy $\xi_{B|sh} = \xi'_B$ (top panel) and of the energy fraction in the precursor wave relative to the total incoming energy, $f_{\xi|sh} = f'_\xi$ (bottom panel). Different symbols refer to a different code: the red circles for SMILEI, blue diamonds for TRISTAN-MP, and orange squares for SHOCKAPIC. The dashed black lines indicate the $\propto \sigma^{-2}$ scaling.

4.2 Energy budget in the precursor

Let us also discuss the global energy budget as seen from the SRF (i.e. the fraction of total incoming energy channelled into the precursor). In the SRF, the energy conservation equation including wave contributions can be written as

$$\gamma_u'^2 \beta_u' W_u - \frac{\delta B_u'^2}{4\pi} = \gamma_d'^2 \beta_d' W_d + \frac{\delta B_d'^2}{4\pi}, \quad (20)$$

where $W_i = w_i + b_{0,i}^2/(4\pi)$ is the generalized enthalpy expressed in the proper frame of the fluid. The left-hand side corresponds to the upstream total energy content and the right-hand side to the downstream energy content. We have used that the upstream and downstream electromagnetic wave energies can be expressed as

$$\frac{\delta B_u'^2}{4\pi} = \frac{\langle \delta B_{z,u}^2 + \delta E_{y,u}^2 \rangle}{8\pi}, \quad (21)$$

$$\frac{\delta B_d'^2}{4\pi} = \frac{\langle \delta B_{z,d}^2 + \delta E_{y,d}^2 \rangle}{8\pi}, \quad (22)$$

respectively. The brackets represent either space averages at a given time or equivalently time averages at one spatial position. We have neglected the contribution from electrostatic waves, since it is largely subdominant in pair plasmas.

Reminding that the upstream magnetization is Lorentz invariant, we use $\sigma = \sigma' = b_{0,u}^2/(4\pi w_u)$. The fraction of total incoming energy channelled into the precursor wave is then

$$f'_\xi = f_{\xi|\text{sh}} = \frac{\text{EM wave energy}}{\text{Total incoming energy}} = \frac{\xi'_B}{\beta'_u} \frac{\sigma}{1 + \sigma}, \quad (23)$$

where β'_u is the upstream flow velocity measured in SRF. Equivalently, it is the shock front speed in the URF. One can also give the fraction of incoming particle kinetic energy channelled into the precursor wave:

$$g'_\xi = \frac{\text{EM wave energy}}{\text{Incoming kinetic energy}} = \frac{\delta B_u'^2}{4\pi \gamma_u'^2 \beta'_u n_u m c^2} = \frac{\xi'_B}{\beta'_u} \sigma. \quad (24)$$

In the latter equation n_u is the upstream plasma proper density, including both species (so, $n_u = 2N_0/\gamma_0$).

Getting back to the simulation results, in Fig. 9 (panel b) we present the dependence on σ of the energy fraction $f_{\xi|\text{sh}} = f'_\xi$, as measured in the SRF. The maximum value is reached at $\sigma \sim 0.1$, where the precursor carries up to 5 per cent of the incoming energy. For $\sigma > 0.3$ the energy content in the wave rapidly drops. Similarly to the ξ'_B scaling, there is a clear dependence as $\propto \sigma^{-2}$ for $\sigma > 1$ (more precisely, $f'_\xi \simeq 7 \times 10^{-4} \sigma^{-2}$). The similarity comes from the fact that the upstream velocity is $\beta'_u \rightarrow 1$ and the upstream energy content is dominated by the magnetic field (i.e. $\sigma \gg 1$). This implies from equation (23) that $f'_\xi \simeq \xi'_B$.

In the limit $\sigma \gg 1$, the conversion efficiency of incoming particle kinetic energy into wave energy scales as $g'_\xi \simeq \xi'_B \sigma \simeq 7 \times 10^{-4} \sigma^{-1}$. This should be contrasted with what we have obtained in the DRF, where this quantity became constant in the $\sigma \gg 1$ limit.

We remark that the scalings reported so far have been obtained from 1D runs. While we expect that the dependence on σ will remain unchanged in 2D and 3D, we speculate that the normalizations of f_ξ and g_ξ will decrease due to transverse effects that cannot be captured in 1D, e.g. wave filamentation and self-focusing through interaction with the upstream plasma. In fact, the 2D simulations of Iwamoto et al. (2017, 2018), performed in the low magnetization regime $\sigma < 0.5$, demonstrated that the wave energy is reduced typically by a factor of 3 (and up to 10), when going from 1D to 2D. However, we expect that the efficiency drop from 1D to 2D (and 3D) will be much less severe in the magnetically dominated regime ($\sigma > 1$) of interest for our work, given the rapid decrease of the wave strength parameter with magnetization (see Fig. 8), and so of the wave feedback on to the upstream plasma. This point will be addressed in a forthcoming study (Sironi et al., in preparation).

5 APPLICATIONS TO FRBS

During a magnetar flare, in response to the motions of the neutron star crust, the above-lying magnetosphere is violently twisted and a strongly magnetized pulse is formed, which propagates away through the magnetar wind. The FRB can be potentially generated at ultrarelativistic shocks resulting from the collision of the magnetized pulse with the steady magnetar wind produced by its spin-down luminosity or by the cumulative effect of earlier flares (Lyubarsky 2014; Beloborodov 2017; Waxman 2017). The train of electromagnetic waves emitted by the shock front via the synchrotron maser is the candidate FRB. Most works up to now assumed empirical values for the conversion efficiency of the shock kinetic energy into the precursor waves. These values were primarily motivated by the work of Gallant et al. (1992) where, however, high- σ simulations were not evolved long enough to reach a stationary state. Here, we use long-term simulations to quantify the steady-

state energetics and spectrum of the precursor waves for a wide range of magnetizations (up to $\sigma \gg 1$). As we now argue, our work can provide a physically grounded model for the origin of coherent emission in FRBs.

First, the synchrotron maser at shocks is a coherent process, which helps explaining the extremely high brightness temperatures of FRBs. In this work, we have derived the fraction of incoming flow energy channelled into the precursor waves. If considered in the ejecta frame (post-shock frame), our simulations show that for $\sigma > 1$ the emitted wave carries a fraction $f_\xi = 2 \times 10^{-3}/\sigma$ of the total energy. This corresponds to a fraction $g_\xi \simeq 2 \times 10^{-3}$ of the incoming particle kinetic energy regardless of σ . If one considers the energy budget in the SRF, the previous scalings become $f'_\xi \simeq 7 \times 10^{-4}/\sigma^2$ and $g'_\xi \simeq 7 \times 10^{-4}/\sigma$, respectively.

Secondly, the precursor emission is linearly polarized, in agreement with the observations of several non-repeating FRBs (Ravi et al. 2016; Petroff et al. 2017; Caleb et al. 2018) and of the repeating FRB 121102 (Gajjar et al. 2018; Michilli et al. 2018). Linear polarization is a natural consequence of the resonance of bunching particles with the extraordinary mode (X-mode). This mode can escape out of the plasma and become a vacuum electromagnetic wave. A contribution from the ordinary mode (O-mode) was also observed in the 2D simulations of Iwamoto et al. (2018), but it was found to be largely subdominant in strongly magnetized plasmas.

Thirdly, the spectral peak can fall in the GHz range for a reasonable choice of parameters. In particular, we have found that in the post-shock frame the emission peak frequency scales as $\omega_{\text{peak}} \simeq 3 \omega_p$ for $\sigma \lesssim 0.1$ and as $\omega_{\text{peak}} \simeq 3 \omega_c$ for $\sigma > 1$. Several high-order harmonics characterize the transition region with $0.1 < \sigma < 1$. Joining the two regimes, and neglecting for simplicity the transition cases, we can cast the peak frequency as $\omega_{\text{peak}} \simeq 3 \omega_p \max[1, \sqrt{\sigma}]$. This can be recast in a simpler form in the pre-shock frame as

$$\omega''_{\text{peak}} \approx 3 \gamma_{s|u} \omega_p, \quad (25)$$

as long as the shock is moving with an ultrarelativistic bulk Lorentz factor $\gamma_{s|u}$ into the upstream medium. The emission frequency for an upstream observer is then

$$\nu'' = \frac{\omega''}{2\pi} \approx 2.7 \times 10^4 \gamma_{s|u} \left(\frac{n_e}{1 \text{ cm}^{-3}} \right)^{1/2} \text{ Hz}, \quad (26)$$

where n_e is the pre-shock electron density. If we assume that the upstream frame corresponds to the observer frame (which is true if the pre-burst wind expands with a non-relativistic velocity), then the combination $\gamma_{s|u} \sqrt{n_e/1 \text{ cm}^{-3}} \approx 4 \times 10^4$ is required for the shock to emit in the GHz band, in rather good agreement with the estimates of Beloborodov (2017). As recently found by Metzger, Margalit & Sironi (2019), this frequency is also consistent with \sim GHz emission from decelerating blast waves produced by flare ejecta in young magnetars.

Finally, the spectrum is narrow-band, $\Delta\omega/\omega_{\text{peak}} \lesssim 1 - 3$ (see Fig. 6), which is again consistent with the observations (e.g. Law et al. 2017; Macquart et al. 2019).

5.1 Comment on criticisms to the synchrotron maser

A number of criticisms have recently been moved against the synchrotron maser emission as a source of the coherent FRB radiation. Lu & Kumar (2018) looked into a wide variety of maser mechanisms operating in either vacuum or plasma and found that none of them can explain the high luminosity of FRBs without invoking unrealistic or fine-tuned plasma conditions. Here, we argue

that the synchrotron maser at relativistic shocks — due to its unique properties — still remains a viable candidate for powering FRBs.

First, it was argued that the synchrotron maser in vacuum requires fine-tuned plasma conditions where the magnetic field is nearly uniform (to within an angle γ^{-1}) and the particles' pitch-angle distribution is narrowly peaked with spread $\lesssim \gamma^{-1}$. Here, γ is the typical Lorentz factor of the emitting particles. This is indeed the natural configuration expected at a relativistic magnetized shock, if the pre-shock particles have non-relativistic temperatures (which is anyway a requirement for efficient synchrotron maser emission). In the shock transition region, the magnetic field is nearly uniform, and the particles coherently rotate in a plane perpendicular to the field (with negligible pitch angle spread).

Secondly, it was argued that it is unclear how the mechanism for the population inversion required by the maser is achieved. Once again, this is naturally realized in the shock transition of a magnetized relativistic shock, where the particles form a ring in momentum space at fixed Lorentz factor $\gamma \sim \gamma_0$, while the inner region of the ring (i.e. at lower γ) is devoid of particles, as indeed required for the existence of a population inversion.

Also, it was argued that during the maser amplification process, high-energy electrons radiate faster than low-energy ones, so the population inversion condition may be quickly destroyed. This is indeed true for each generation of particles passing through the shock, since the synchrotron maser instability relaxes by 'filling up' the hollow ring in momentum space, thus destroying the population inversion. However, while this happens, a new generation of particles is entering into the shock. They establish a new ring in momentum space, and keep sustaining the radiated train of precursor waves. In other words, the continuous passage of plasma through the shock ensures that the population inversion is steadily maintained (yet, at each time by different particles).

Finally, Lu & Kumar (2018) considered more specifically the maser synchrotron emission at shocks, which they named as 'bunching in the gyration phase.' In order to minimize the effect of induced Compton scattering, they estimated that the radiative efficiency of the shock must be extremely small. However, they considered only internal shocks occurring in between two identical consecutive density shells propagating inside the pre-burst wind, and not the leading shock moving directly into the wind. Aside from the limitations of induced Compton scattering, it is anyway hard for internal shocks to be efficient emitters of maser synchrotron radiation, since they propagate into a relativistically hot shocked plasma (the downstream region of the leading shock). The arguments by Lu & Kumar (2018) will not apply to the leading shock. First, this shock is likely to be ultrarelativistic, unlike internal shocks. Secondly, the properties of the shell and of the pre-burst wind (as regard to magnetization, temperature, and composition) are generally different, in contrast to what Lu & Kumar (2018) implicitly assumed. We believe that the quantitative results on precursor energetics and spectrum that we provide in this work will help revisit the estimates provided by Lu & Kumar (2018) for the case of the leading shock.

6 SUMMARY AND CONCLUSIONS

In this work we have investigated by means of 1D PIC simulations the physics of synchrotron maser emission from perpendicular relativistic shocks that propagate in highly magnetized electron-positron plasmas (with magnetization $0.1 \leq \sigma \leq 30$). For strongly magnetized shocks, we expect that multidimensional simulations (to be discussed in a forthcoming work) will not yield very different results than what we present here. We have explored the

efficiency and spectrum of the electromagnetic precursor emission as a function of σ and γ_0 . We have found that:

(i) The shock front emits efficiently and steadily a train of high-amplitude electromagnetic precursor waves for any σ and γ_0 , in the range $0.1 \leq \sigma \leq 30$ and $\gamma_0 \geq 5$ that we have explored. The emission is linearly polarized, with fluctuating magnetic field along the same direction as the upstream mean field.

(ii) Thanks to unprecedently long simulations, we have been able to reach the stage when the precursor emission settles to a steady state, which allows to systematically extract the wave properties (energetics and spectrum). We find that the ratio of the wave energy to the upstream magnetic energy, ξ_B , decreases rapidly from 3.5 at $\sigma = 0.1$ down to 0.04 at $\sigma = 1$, as measured in the post-shock frame of the simulations. For $\sigma \gg 1$, this ratio converges to a constant value $\xi_B \simeq 0.01$. In the SRF, the asymptotic scaling in the limit $\sigma \gg 1$ becomes $\xi'_B \propto \sigma^{-2}$.

(iii) For $\sigma > 1$, the energy output in precursor waves normalized to the total incoming energy scales as $f_\xi \simeq 2 \times 10^{-3} \sigma^{-1}$ in the post-shock frame and as $f'_\xi \simeq 7 \times 10^{-4} \sigma^{-2}$ in the SRF. The former implies that in the downstream frame, $\sigma > 1$ shocks convert a constant fraction of the incoming particle kinetic energy into precursor waves (equal to $g_\xi \simeq 2 \times 10^{-3}$).

(iv) Magnetically dominated shocks with $\sigma > 1$ exhibit a resonating cavity in the shock front structure in between two solitons, instead of the single soliton loop that is observed for $\sigma \ll 1$ shocks. This cavity plays an essential role in amplifying the radiation and selecting the dominant emission frequency as an eigenmode of the cavity. This effect causes the peak emission frequency, as measured in the downstream frame, to scale as $\omega_{\text{peak}} \simeq 3\omega_c = 3\sqrt{\sigma}\omega_p$ for $\sigma > 1$, whereas earlier works (Gallant et al. 1992) quote a stronger scaling with magnetization, $\omega_{\text{peak}} \simeq \sigma\omega_p$.

(v) The characteristic frequency of the emission, as measured in the post-shock frame, is $\omega \simeq 3\omega_p$ for weakly magnetized shocks $\sigma \leq 0.1$, and $\omega \simeq 3\omega_c$ for $\sigma \gg 1$, as we have just discussed. In the transition region $0.1 < \sigma < 1$, prominent high-order harmonics of $\omega_{c,\text{sol}}$ (given in equation 13) were observed along with the fundamental at $\omega_{c,\text{sol}}$. Aside from the transition cases, we can interpolate between the low- and high-magnetization results and state that the peak emission occurs at $\omega_{\text{peak}} \simeq 3\omega_p \max[1, \sqrt{\sigma}]$, as measured in the downstream frame. In the pre-shock frame (which coincides with the observer frame, if the magnetar wind is non-relativistic), this can be recast in a simpler form as $\omega''_{\text{peak}} \approx 3\gamma_{s|u}\omega_p$, where $\gamma_{s|u}$ is the shock Lorentz factor in the upstream frame.

(vi) The spectrum of the precursor is narrow-band, $\Delta\omega/\omega_{\text{peak}} \lesssim 1 - 3$, with a low-frequency cutoff at $\omega_{\text{cutoff}} = \gamma_{s|d}\omega_p$ (here, $\gamma_{s|d}$ is the shock Lorentz factor in the downstream frame) set by the requirement that the group velocity be faster than the shock speed.

(vii) We did not observe any dependence on γ_0 of the energy fraction, ξ_B , and of the characteristic emission frequency, $\omega_{\text{peak}}/\omega_p$, in the post-shock frame.

We conclude with a few caveats. First, we have assumed that the upstream plasma has negligible thermal spread, $k_B T_0/m_e c^2 = 10^{-4}$. Higher temperatures are likely to suppress high-order harmonics and reduce the global energy of the wave. Secondly, we have mostly focused on strongly magnetized ($\sigma > 1$) plasmas, a regime that so far has received little attention. Even though this work only presents 1D simulations, we anticipate that the multidimensional physics of $\sigma > 1$ shocks (Sironi et al., in preparation) will not depart significantly from what we report here. In contrast, for weaker magnetizations ($\sigma \ll 10^{-2}$), transverse effects (e.g. Weibel-driven filamentation)

will significantly reduce the energy carried by the precursor waves (Sironi et al. 2013; Iwamoto et al. 2017). In summary, both higher pre-shock temperatures and multidimensional effects at low σ are expected to degrade the precursor efficiency, which might become too low to explain the FRB emission.

Finally, we have only considered electron–positron shocks. Recently, a very large Faraday Rotation Measure (RM) of $\sim 10^5$ rad m $^{-2}$ was reported from the repeating FRB 121102 (Michilli et al. 2018). This challenges the pure electron–positron composition assumed in this study, since the presence of an appreciable fraction of ions is required to produce non-zero RM (Margalit & Metzger 2018). This urges us to explore the shock physics for electron–proton and electron–positron–proton compositions. Yet, it is still possible that the FRB pulse is produced in localized regions with pristine electron–positron composition, even though most of the magnetar wind (which inflates the surrounding nebula, where the RM accumulates) is proton-dominated.

ACKNOWLEDGEMENTS

IP acknowledges discussions with Anatoly Spitkovsky, Patrick Crumley, and Yuri Cavecchi. LS is grateful to Brian Metzger for many inspiring discussions. IP was supported by NSF grants PHY-1804048 and PHY-1523261. This work was facilitated by the Max-Planck/Princeton Center for Plasma Physics. LS acknowledges support from NASA ATP 80NNSC18K1104. The simulations were performed on Habanero cluster at Columbia University, NERSC (Edison) and NASA (Pleiades) resources, PICSciE-OIT High Performance Computing Center and Visualization Laboratory at Princeton University, and on CALMIP supercomputing resources at Université de Toulouse (France) under the allocation 2016-p1504.

REFERENCES

Alsop D., Arons J., 1988, *Phys. Fluids*, 31, 839
 Amato E., Arons J., 2006, *ApJ*, 653, 325
 Beloborodov A. M., 2017, *ApJ*, 843, L26
 Birdsall C. K., Langdon A. B., 1991, *Plasma Physics via Computer Simulation*. Taylor and Francis, Milton Park, UK
 Caleb M. et al., 2018, *MNRAS*, 478, 2046
 Derouillat J. et al., 2018, *Comput. Phys. Commun.*, 222, 351
 Gajjar V. et al., 2018, *ApJ*, 863, 2
 Gallant Y. A., Hoshino M., Langdon A. B., Arons J., Max C. E., 1992, *ApJ*, 391, 73
 Ghisellini G., Locatelli N., 2018, *A&A*, 613, A61
 Greenwood A. D., Cartwright K. L., Luginsland J. W., Baca E. A., 2004, *J. Comput. Phys.*, 201, 665
 Hoshino M., 2008, *ApJ*, 672, 940
 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
 Kashiyama K., Murase K., 2017, *ApJ*, 839, L3
 Katz J. I., 2016, *Mod. Phys. Lett. A*, 31, 1630013
 Katz J. I., 2018, *MNRAS*, 481, 2946
 Keane E. F., Stappers B. W., Kramer M., Lyne A. G., 2012, *MNRAS*, 425, L71
 Kumar P., Lu W., Bhattacharya M., 2017, *MNRAS*, 468, 2726
 Langdon A. B., Arons J., Max C. E., 1988, *Phys. Rev. Lett.*, 61, 779
 Law C. J. et al., 2017, *ApJ*, 850, 76
 Lemoine M., Pelletier G., 2010, *MNRAS*, 402, 321
 Long K., Pe'er A., 2018, *ApJ*, 864, L12
 Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, *Science*, 318, 777
 Lu W., Kumar P., 2018, *MNRAS*, 477, 2470
 Lyubarsky Y., 2006, *ApJ*, 652, 1297
 Lyubarsky Y., 2014, *MNRAS*, 442, L9
 Lyutikov M., 2002, *ApJ*, 580, L65
 Lyutikov M., 2017, *ApJ*, 838, L13
 Macquart J.-P., Shannon R. M., Bannister K. W., James C. W., Ekers R. D., Bunton J. D., 2019, *ApJ*, 872, L19
 Marcote B. et al., 2017, *ApJ*, 834, L8
 Margalit B., Metzger B. D., 2018, *ApJ*, 868, L4
 Margalit B., Metzger B. D., Berger E., Nicholl M., Eftekhar T., Margutti R., 2018, *MNRAS*, 481, 2407
 Melrose D. B., 2017, *Rev. Mod. Plasma Phys.*, 1, 5
 Metzger B. D., Berger E., Margalit B., 2017, *ApJ*, 841, 14
 Metzger B. D., Margalit B., Sironi L., 2019, preprint ([arXiv:1902.01866](https://arxiv.org/abs/1902.01866))
 Michilli D. et al., 2018, *Nature*, 553, 182
 Murase K., Kashiyama K., Mészáros P., 2016, *MNRAS*, 461, 1498
 Pelletier G., Bykov A., Ellison D., Lemoine M., 2017, *Space Sci. Rev.*, 207, 319
 Petroff E. et al., 2017, *MNRAS*, 469, 4465
 Plotnikov I., Grassi A., Grech M., 2018, *MNRAS*, 477, 5238
 Popov S. B., Postnov K. A., 2013, preprint ([arXiv:1307.4924](https://arxiv.org/abs/1307.4924))
 Popov S. B., Postnov K. A., Pshirkov M. S., 2018, preprint ([arXiv:1806.03628](https://arxiv.org/abs/1806.03628))
 Ravi V. et al., 2016, *Science*, 354, 1249
 Sagiv A., Waxman E., 2002, *ApJ*, 574, 861
 Sazonov V. N., 1973, *Soviet Astron.*, 16, 971
 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
 Sironi L., Keshet U., Lemoine M., 2015, *Space Sci. Rev.*, 191, 519
 Spitkovsky A., 2005, in Bulik T., Rudak B., Madejski G., eds, *AIP Conf. Ser.* Vol. 801, *Astrophysical Sources of High Energy Particles and Radiation*. Astron. Soc. Pac., San Francisco, p. 345
 Spitkovsky A., 2008, *ApJ*, 682, L5
 Spitler L. G. et al., 2014, *ApJ*, 790, 101
 Thornton D. et al., 2013, *Science*, 341, 53
 Treumann R. A., 2006, *A&AR*, 13, 229
 Waxman E., 2017, *ApJ*, 842, 34

APPENDIX A: CODES COMPARISON

In this appendix we show how the results from the different codes compare. The synchrotron maser emission occurs through the resonance of the cyclotron harmonics with the X-mode branch. It is not guaranteed that a typical PIC code can capture accurately a large number of these resonances, especially at the high-frequency end of the branch. For instance, in typical Yee-type second-order solvers of Maxwell's equations the numerical speed of electromagnetic waves is known to be artificially suppressed at high ω if the CFL number is smaller than unity (Birdsall & Langdon 1991).

For this reason we undertook an extensive comparison of three different PIC codes: two Finite-Difference Time-Domain (FDTD) codes, TRISTAN-MP and SMILEI, and the pseudo-spectral code SHOCKAPIC. In principle, SHOCKAPIC is the best suited to capture the dispersion relation of waves in a plasma, but it is the least optimized among the three codes, making it challenging to perform long-term simulations. Concerning SMILEI, it is a well-optimized code, but in 1D setups it currently has only a standard Yee solver. TRISTAN-MP is the most optimized for shock setups. Also, it allows us to use a fourth-order scheme to solve Maxwell's equations (Greenwood et al. 2004) that reproduces accurately the dispersion relation of electromagnetic waves even at low CFL numbers. This is the reason why the runs presented in the main body of the paper were performed with TRISTAN-MP. In general, each code employs different algorithms and implementations. The agreement between

Table A1. Typical parameters of the PIC simulations presented in this study: Δt is the time-step in units of the inverse plasma frequency ω_p^{-1} (defined with both species), T_{sim} is the simulation timespan, Δx is the cell size in units of c/ω_p , $k_B T_0$ is the upstream thermal energy in units of $m_e c^2$, N_{ppc} is the number of particles-per-cell for each species, σ_{min} and σ_{max} are the minimal and maximal values of the magnetization explored with a given code.

| PIC code | $\Delta t \omega_p$ | $T_{\text{sim}} \omega_p$ | $\Delta x (c/\omega_p)$ | N_{ppc} | $k_B T_0 (m_e c^2)$ | σ_{min} | σ_{max} | γ_0 |
|------------|---------------------|---------------------------|-------------------------|------------------|---------------------|-----------------------|-----------------------|------------|
| TRISTAN-MP | 1/200 | $\lesssim 2 \times 10^4$ | 1/100 | 64 | 10^{-4} | 10^{-1} | 30 | 10 |
| SMILEI | 1/224 & 1/113 | 6.7×10^3 | 1/112 | 20 | 10^{-4} | 10^{-1} | 30 | 10 |
| SHOCKAPIC | 1/90 | 1.2×10^3 | 1/44.7 | 20 | 10^{-4} | 10^{-3} | 1 | 10 |
| SMILEI (2) | 1/90 | 1×10^3 | 1/44.7 | 20 | 10^{-6} | 10^{-3} | 2 | 10 |
| SMILEI (3) | 1/90 | 1.5×10^3 | 1/44.7 | 20 | 10^{-4} | 10^{-3} | 2 | 160 |

the three codes will then be a strong indication of the physical robustness of our results.

The simulation parameters for each code are presented in Table A1. The table reports the space and time resolution, the simulation timespan, the number of particles per cell, the values of the upstream temperature T_0 and bulk Lorentz factor γ_0 , and the explored range of σ . For better comparison we used comparable space and time resolutions: the skin depth was resolved with 100 cells in TRISTAN-MP simulations, with 112 cells in SMILEI simulations, and with 44.7 cells in SHOCKAPIC simulations. The latter has a twice smaller resolution due to code performance limitations (not parallelized). We noticed that a resolution lower than 20 cells per skin depth affected the results for any σ negatively. The results become stable for any resolution higher than 40 cells per c/ω_p , as long as $\sigma \leq 10$. Similar conclusions were reached by Iwamoto et al. (2017). For this reason a high spatial and time resolution was employed in the simulations presented in the main body of the paper. Only short simulations were affordable with SHOCKAPIC. For this reason, the $\sigma > 1$ regime was not explored with this code (as we have discussed, at high σ it takes longer to reach a steady state). With TRISTAN-MP and SMILEI it was possible to reach the stationary state for σ up to 30. Concerning the number of particles per cell, the results are very weakly dependent on N_{ppc} , as long as at least a dozen of particles per cell are initialized.

In the following we present in more detail the comparison of precursor energy and spectrum as derived from different codes.

A1 Precursor energy

In Fig. A1 we present the normalized wave energy ξ_B as a function of σ obtained with the three codes. Values obtained with SHOCKAPIC, SMILEI, and TRISTAN-MP are plotted using orange squares, red circles, and blue diamonds, respectively. In the overlapping range of σ , we observe good agreement among different codes. For instance, in the $\sigma > 1$ regime SMILEI and TRISTAN-MP give the same values of ξ_B . In the range $0.1 < \sigma \leq 1$, where all codes overlap, the scatter among codes is slightly larger, although the rapid drop in ξ_B is common to all codes, and it happens around the same σ . We note that the transition is more abrupt in SHOCKAPIC than in SMILEI and TRISTAN-MP, but differences remain minor.

In Fig. A2 we extend the comparison to different studies in the literature and to different values of γ_0 from 10 to 10^6 . The range of σ in this figure is from 10^{-3} to 1, since other studies did not explore highly magnetized cases with sufficiently long simulations (i.e. they did not reach a steady state in the regime $\sigma \gg 1$). The blue circles and green diamonds report the values obtained with SMILEI using $\gamma_0 = 10$ and 160, respectively. Both give nearly the same values for any explored σ confirming that ξ_B does not depend on the flow Lorentz factor. The red squares report the values from TRISTAN-MP using $\gamma_0 = 10$ (same as in Fig. A1). The data from the 1D

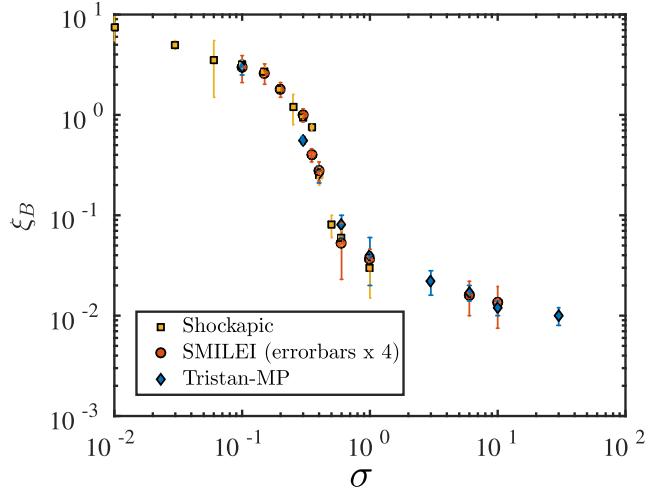


Figure A1. Precursor wave energy ξ_B as a function of σ : comparison between three PIC codes used in this study (TRISTAN-MP, SMILEI, and SHOCKAPIC). Values obtained with SHOCKAPIC, SMILEI, and TRISTAN-MP are plotted using orange squares, red circles, and blue diamonds, respectively.

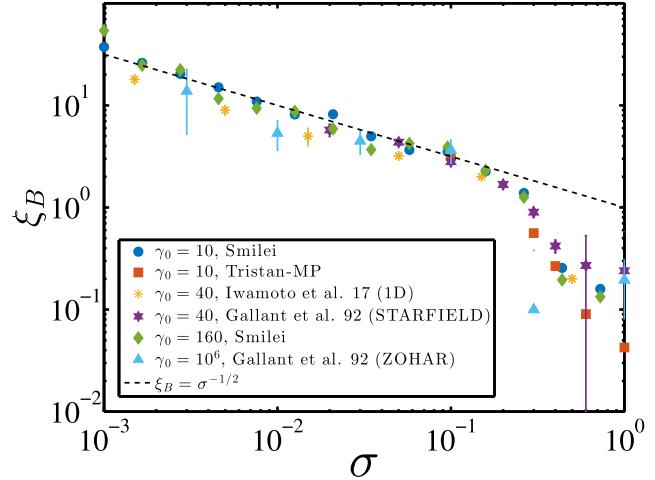


Figure A2. Same as Fig. A1, but comparing the results from PIC codes in this work with earlier studies from the literature. The explored values of σ range here from 10^{-3} to 1, since other studies did not explore highly magnetized cases with sufficiently long simulations. Also, results with different values of γ_0 are shown here. The blue circles correspond to SMILEI simulations with $\gamma_0 = 10$, red squares are from TRISTAN-MP with $\gamma_0 = 10$, orange stars present the data taken from Iwamoto et al. (2017) with $\gamma_0 = 40$ (1D), magenta stars are from Gallant et al. (1992) with $\gamma_0 = 40$, green diamonds present SMILEI results with $\gamma_0 = 160$, and the values presented using light blue triangles are taken from Gallant et al. (1992) with $\gamma_0 = 10^6$. The dashed black line presents the scaling $\xi_B = 1/\sqrt{\sigma}$, which roughly fits the data points in the range $\sigma \in [10^{-3}, 0.1]$.

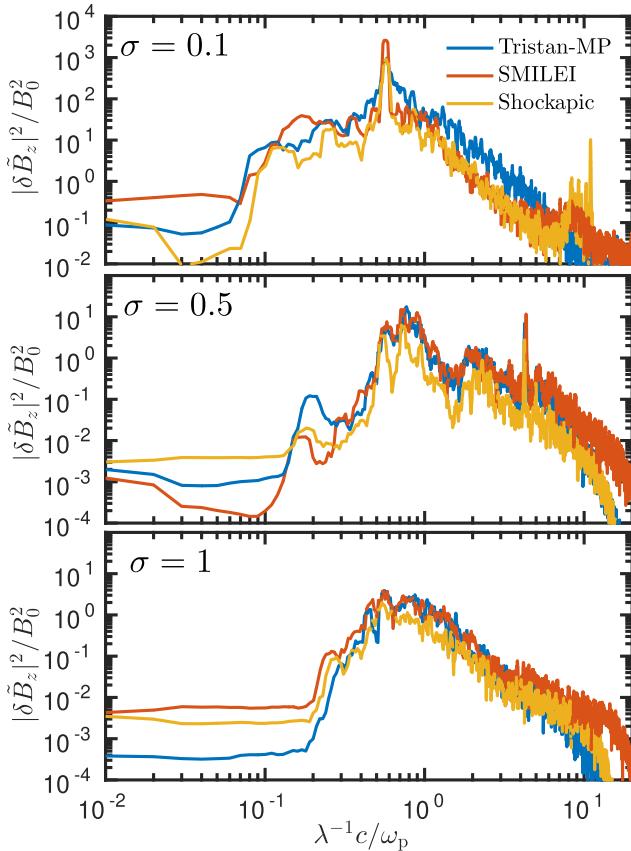


Figure A3. Precursor spectrum in k -space for three values of σ as indicated in the legend. We perform a comparison among the three PIC codes used in this study: the spectrum extracted from TRISTAN-MP simulations is plotted using a solid blue line, whereas red lines are used for SMILEI and orange lines for SHOCKAPIC.

simulations of Iwamoto et al. (2017) using $\gamma_0 = 40$ are plotted with orange stars. Their values are slightly smaller than what is found in this study, though generally in good agreement. The violet stars and light-blue triangles report the values from Gallant et al. (1992) using $\gamma_0 = 40$ and 10^6 , respectively. We notice that all codes provide the same results in the range of $\sigma \in [10^{-3}, 0.3]$, regardless of γ_0 . This demonstrates that the precursor wave normalized energy ξ_B is not dependent on γ_0 , and that our study is in very good agreement with earlier results.

For $\sigma > 0.3$ there is a noticeable scatter between different simulations. The most plausible reason for the discrepancy among different data sets is that the high- σ simulations from earlier studies were not evolved long enough to reach the asymptotic state, so the value of ξ_B was not yet stabilized (see Fig. 2 for the time convergence of the efficiency).

A2 Precursor spectrum

We now compare the precursor k -spectrum among the three codes. Some differences are expected since the numerical schemes for the integration of Maxwell's equations differ among the codes.

In Fig. A3 we compare the precursor spectrum extracted from the three codes for a few representative values of magnetization. From top to bottom, the value of σ is 0.1, 0.5, and 1, respectively. We cannot perform any comparison for $\sigma > 1$ as this range was not

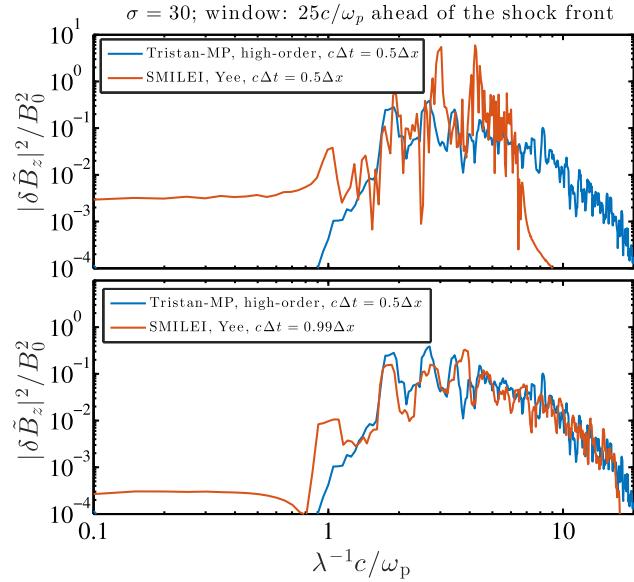


Figure A4. Comparison of the precursor spectrum in k -space obtained from TRISTAN-MP and SMILEI for $\sigma = 30$ (the largest magnetization that we have explored, where differences among codes are most dramatic). The upper panel presents the spectrum from TRISTAN-MP using a fourth-order scheme to solve Maxwell's equations (blue line) and from SMILEI using a Yee-type scheme with CFL number $= c\Delta t/\Delta x = 0.5$ (red line). The latter presents a sharp cutoff at high- k (i.e. for $\lambda < 0.14 c/\omega_p$) and irregular line-like emission features. The lower panel presents the same comparison but with CFL = 0.99 for SMILEI (red line). The high- k cutoff disappears and a very good agreement with TRISTAN-MP is obtained. We note that the spectra presented in this figure are normalized to unity, instead of the previously adopted normalization $\int |\delta \tilde{B}_z(k)|^2 / B_0^2 dk = \int |\delta \tilde{B}_z(\omega)|^2 / B_0^2 d\omega = \xi_B$.

explored with SHOCKAPIC (but see below for a comparison between SMILEI and TRISTAN-MP at $\sigma = 30$). The spectrum extracted from TRISTAN-MP is plotted using a solid blue line. The red and orange lines are used for SMILEI and SHOCKAPIC, respectively. There is generally a good agreement among the codes for all values of σ as regard to the low- k cutoff wavenumber, the high- k slope, and the main peaks in the spectrum. For example, the dominant emission line for $\sigma = 0.1$ and the high-order harmonic line at $\lambda = 0.24 c/\omega_p$ for $\sigma = 0.5$ are exactly at the same wavelength for the three codes. One difference can be noted: the spectral energy density is slightly smaller in SHOCKAPIC than in the two FDTD codes around $\lambda^{-1}c/\omega_p \sim 1$, for $\sigma = 0.5$ and $\sigma = 1$. Yet, this difference is not systematic and the overall energy in the precursor is very close among the three codes.

As an exception and a word of caution, we noticed that the use of a small CFL number with a Yee-type solver of Maxwell's equations (as used in the SMILEI code) has a negative impact on the results for the largest magnetizations explored here, i.e. $\sigma > 10$. In fact, the emission peaks at high frequencies where the light-wave branch is affected by the artificial reduction of the phase speed. The spectrum of the precursor is then sharply cut at high frequencies, affecting the overall energy output in the precursor. This effect is evidenced in Fig. A3 for $\sigma = 30$ (the largest value explored in this work). The upper panel of the figure compares the spectrum from TRISTAN-MP (blue), where a fourth-order scheme was used, with the spectrum from SMILEI (red), which employs a Yee-type scheme with $c\Delta t/\Delta x = 0.5$. There is an artificial suppression in the high- k region in the SMILEI simulation. The bottom panel shows the same comparison, but with $c\Delta t/\Delta x = 0.99$ being used with

SMILEI. In this case, the spectra agree very well, up to details in line-like features. This conveys that the high- k (and so, high- ω) part of the precursor spectrum can be properly captured only when the numerical integrator is capable of reproducing correctly the dispersion relation of electromagnetic waves. This problem does not arise in TRISTAN-MP (with high-order spatial solver) and SHOCKAPIC, since for them the numerical dispersion of the light-wave branch is much closer to the realistic one even for small CFL numbers.

All SMILEI simulations that use a CFL number as close as possible to unity (CFL = 0.99) display spectra that are in very good agreement with the other two codes for any σ .

A3 Concluding remark

We find that our results do not depend on the code that we employ if these three conditions are realized: (i) a high spatial resolution (i.e. large c/ω_p) is employed; (ii) in a Yee-type based code, the CFL number is as close as possible to unity; (iii) the simulations are sufficiently long to reach the steady state.

APPENDIX B: PRECURSOR ENERGETICS: 1D VERSUS MULTIDIMENSIONAL SIMULATIONS

In order to support our claim that the precursor wave energy does not significantly decrease due to multidimensional effects (in the $\sigma \geq 1$ regime of interest for this work), here we present a preliminary analysis of 2D and 3D simulations performed with TRISTAN-MP.⁶ We explore a range of $\sigma \in [0.1, 10]$ in 2D, and $\sigma \in [0.1, 3]$ in 3D. In 2D simulations we focus on the *out-of-plane* configuration: the simulation plane is the xy plane, the shock front propagates in the x -direction, and the upstream magnetic field is along the z -direction. We do not present any *in-plane* 2D simulation results here because we find that 3D simulations are in excellent agreement with 2D out-of-plane results.

In 2D simulations we keep all parameters the same as in 1D, except that the number of particles per cell per species is set to 8 (values between 2 and 32 have been tested with no significant differences). The transverse dimension of the simulation box is set to $14 c/\omega_p$. We find that a transverse width of more than $2 - 3 c/\omega_p$ is sufficient to capture multidimensional effects. In particular, the effects of wave filamentation and self-focusing that lead to efficient pre-heating of the upstream plasma in the longitudinal momentum are properly captured with a box width of a few skin depths.

In 3D simulations we reduce the transverse dimension to $4 c/\omega_p$ (in both y and z directions). The spatial resolution in 3D runs is set to 25 cells per c/ω_p (four times lower than in 1D and 2D) and the number of particles per cell per species is varied between 3 and 18 (again with little differences). This was necessary to produce sufficiently long runs while still capturing the relevant physics. The effect of a lower spatial resolution was only apparent in the $\sigma = 3$ run, since the spectrum extends to higher frequencies, which are not captured properly if the resolution is insufficient.

Fig. B1 presents the time evolution of the precursor energy for $\sigma = 1$ as measured in 1D (blue line), 2D (red line), and 3D (orange line) simulations. The direct comparison between 1D and multidimensional simulations shows that the asymptotic value of ξ_B

⁶The code accuracy and stability in multidimensional simulations of relativistic shocks was assessed in several studies (Spitkovsky 2005, 2008; Sironi & Spitkovsky 2009, 2011; Sironi et al. 2013).

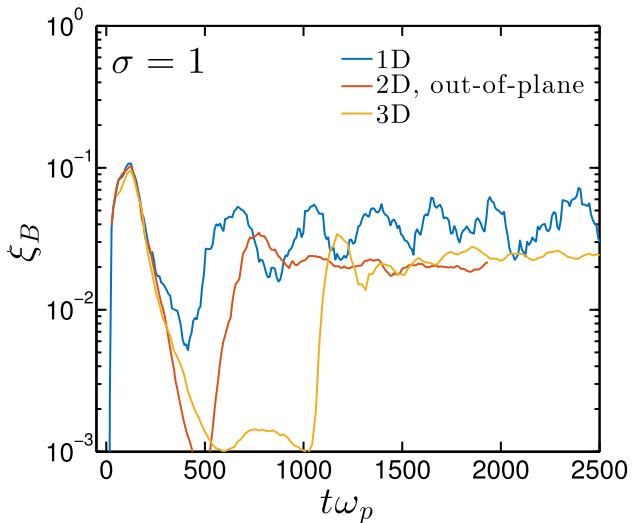


Figure B1. Comparison of the time-evolution of the escaping Poynting flux along the shock-normal direction in 1D, 2D, and 3D simulations for $\sigma = 1$. The blue solid line presents 1D, red for 2D (out-of-plane $B_{z,0}$), and orange for 3D. 1D results are the same as in Fig. 2, as the Poynting flux is nearly equal to ξ_B in 1D (see the main text for details on numerical parameters in 2D and 3D simulations).

in 2D and 3D simulations is only a factor of two smaller than in 1D, while for $\sigma < 0.5$ – as we will show below, and see also Iwamoto et al. (2017) – the energy of the wave decreases by a factor of about 3–10 when going from 1D to 2D and 3D configurations. It also shows that the 2D and 3D energetics are in very good agreement. The only difference between 2D and 3D is that it takes more time in 3D to settle into the steady state (see the rise of the red line after $t\omega_p = 500$ and of the orange line after $t\omega_p = 1000$). So, we can confidently state that the decrease in precursor efficiency due to multidimensional effects is much less severe in the high-magnetization case $\sigma = 1$ than for $\sigma < 0.5$.

Let us note that in this appendix we have redefined the ξ_B parameter. Here, ξ_B corresponds to the normalized Poynting flux in the x -direction, $\xi_B = \langle \delta E_y \delta B_z - \delta E_z \delta B_y \rangle / B_0^2$. The average is done over the region between 5 and $25 c/\omega_p$ ahead of the shock front, for consistency with our 1D results, and over all the transverse directions (y in 2D; y and z in 3D). In 1D we have systematically verified that $\langle \delta B_z^2 \rangle = \langle \delta E_y \delta B_z - \delta E_z \delta B_y \rangle = \langle \delta E_y \delta B_z \rangle$, but this equality is not obviously satisfied in multidimensional simulations with $\sigma \leq 0.6$. The choice of defining ξ_B as the precursor Poynting flux is due to the fact that the most relevant measure of the electromagnetic energy output of the shock is the Poynting flux of the escaping wave in the shock-normal direction.

In Fig. B2, using a suite of 1D, 2D, and 3D simulations, we show the dependence on σ of the normalized Poynting flux of the precursor wave ξ_B (panel a), of the energy fraction parameter as measured in the simulation frame f_ξ (panel b), and of the energy fraction parameter as measured in the SRF $f'_\xi = f_{\xi|\text{sh}}$ (panel c). The definition of the latter two is given in the main body of the article: equation (5) and equation (23), respectively. Values from 1D, 2D, and 3D simulations are plotted using blue circles, red squares, and green stars, respectively. The results of 2D out-of-plane simulations of Iwamoto et al. (2017) are plotted using orange triangles in panel (a). The measurement of ξ_B in 2D and 3D simulations was done by considering the asymptotic values in the time evolution for each σ , as shown in Fig. B1 for the particular case of $\sigma = 1$. Error

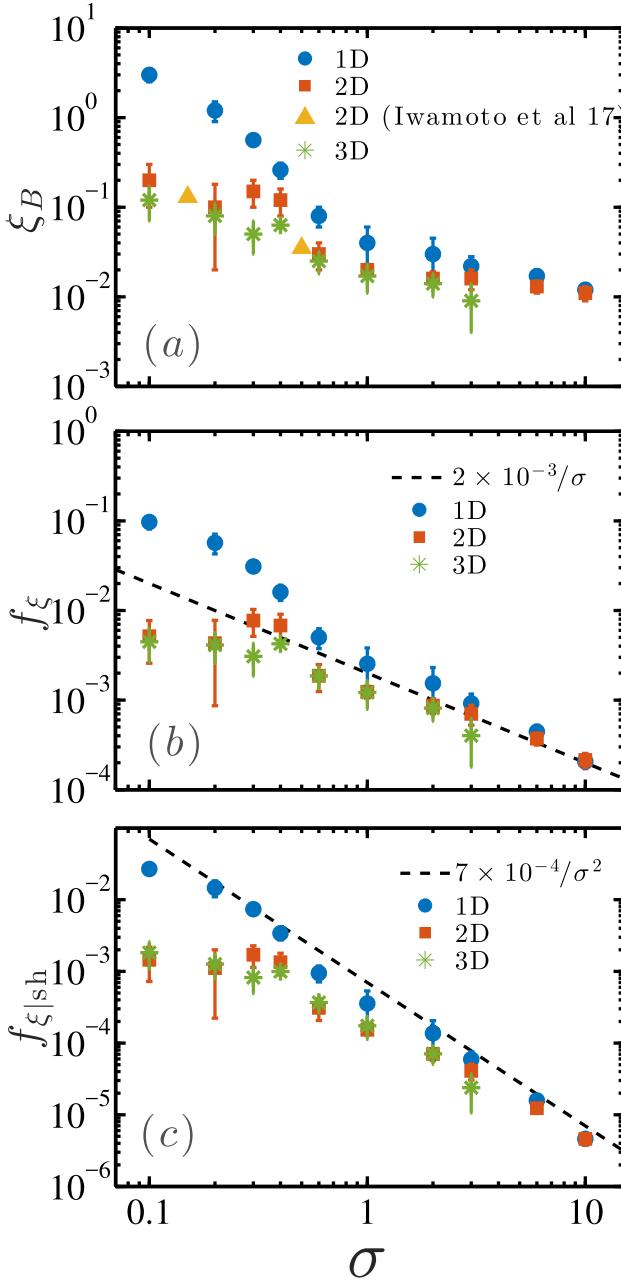


Figure B2. Dependence on the upstream magnetization σ of the normalized Poynting flux of the precursor wave ξ_B (panel a), of the energy fraction parameter measured in the simulation frame f_ξ (panel b), and of the energy fraction measured in the shock front rest frame $f'_\xi = f_\xi|_{\text{sh}}$ (panel c). The dashed line in panel (b) follows the scaling $f_\xi = 2 \times 10^{-3}/\sigma$, in the same way as in Fig. 3. The dashed line in panel (c) follows $f'_\xi = 7 \times 10^{-4}/\sigma^2$, identical to Fig. 9.

bars quantify uncertainties due to temporal oscillations of the time-evolution curves. Knowing ξ_B and measuring directly the shock front velocities from simulations, the values in panels (b) and (c) were produced using equations (5) and (23), respectively.

Fig. B2 shows that:

(i) In 2D and 3D (red and green symbols), for $\sigma = 0.1$ the Poynting flux of the precursor wave ξ_B , the energy fractions f_ξ and $f_{\xi|_{\text{sh}}}$ are reduced by a factor of $\approx 10-20$ as compared to 1D (blue circles). This is in agreement with Iwamoto et al. (2017).

(ii) The suppression in efficiency becomes gradually smaller when σ increases from 0.1 to 3. For $\sigma \gtrsim 1$, the difference between 1D and multidimensional results becomes negligible.

(iii) Values from 2D out-of-plane and 3D simulations are generally in very good agreement, except for $\sigma = 0.3$ and 0.4 (which we have called ‘transition cases’ in the main body of the text).

(iv) If the precursor energy fraction is cast in the SRF, panel (c) shows that $f'_\xi \approx 10^{-3}$ for $\sigma \sim 0.1-0.4$, instead of ~ 0.01 in 1D. For $\sigma > 1$, multidimensional simulations converge towards 1D values and follow the scaling $f'_\xi \approx 5 \times 10^{-4}/\sigma^2$, only slightly lower than reported in the main text for 1D simulations only.

Using 3D simulations we can address other aspects of the precursor physics, such as the importance of the O-mode (δB_y component, since $\delta \mathbf{B} \perp \mathbf{B}_0$ for this mode) versus X-mode (δB_z component, since $\delta \mathbf{B} \parallel \mathbf{B}_0$ for this mode) and beaming of the emitted precursor wave. By extracting systematically the values of $\langle \delta B_y^2 \rangle$ and $\langle \delta B_z^2 \rangle$ in 3D simulations, we find that the O-mode is subdominant for all magnetizations explored here, i.e. $\langle \delta B_y^2 \rangle / \langle \delta B_z^2 \rangle \sim 0.2-0.5$. This implies that the precursor wave retains (at the 99 per cent level, or more) the linear polarization of the X-mode, with magnetic field of the wave lying in the same direction as the upstream background field.

Concerning the beaming of the precursor wave in 3D, we considered the components of the Poynting vector in different directions. We find that the Poynting flux along the y -direction (and z -direction) is largely subdominant as compared to the shock-normal direction. The ratio is $|\Pi_y|/\Pi_x \sim 5 \times 10^{-2}$ for any $\sigma \in [0.1, 3]$, where the Poynting vector of the wave is defined as $\pi = \delta \mathbf{E} \times \delta \mathbf{B} / B_0^2$. This shows that the emitted wave is strongly beamed in the shock-normal direction. For an external observer the beaming will be further enhanced by Lorentz transformation from the simulation frame to the observer frame (in the case of shocks in magnetar winds from the post-shock frame to the pre-shock frame).

This preliminary analysis of multidimensional runs demonstrates that 1D simulations provide accurate numbers in the $\sigma \gg 1$ regime, in agreement with 2D out-of-plane and 3D simulations.

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