

Particle acceleration and non-thermal emission at the intrabinary shock of spider pulsars – I. Non-radiative simulations

Jorge Cortés¹★ and Lorenzo Sironi^{1,2}★

¹Department of Astronomy, Columbia University, 550 W 120th Street, New York, NY 10027, USA

²Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

Accepted 2024 September 20. Received 2024 September 20; in original form 2024 April 2

ABSTRACT

Spider pulsars are compact binary systems composed of a millisecond pulsar and a low-mass companion. Their X-ray emission – modulated on the orbital period – is interpreted as synchrotron radiation from high-energy electrons accelerated at the intrabinary shock. We perform global two-dimensional particle-in-cell simulations of the intrabinary shock, assuming that the shock wraps around the companion star. When the pulsar spin axis is nearly aligned with the orbital angular momentum, we find that the magnetic energy of the relativistic pulsar wind – composed of magnetic stripes of alternating field polarity – efficiently converts to particle energy at the intrabinary shock, via shock-driven reconnection. The highest energy particles accelerated by reconnection can stream ahead of the shock and be further accelerated by the upstream motional electric field. In the downstream, further energization is governed by stochastic interactions with the plasmoids/magnetic islands generated by reconnection. We also extend our earlier work by performing simulations that have a larger (and more realistic) companion size and a more strongly magnetized pulsar wind. We confirm that our first-principles synchrotron spectra and light curves are in good agreement with X-ray observations.

Key words: acceleration of particles – magnetic reconnection – radiation mechanisms: non-thermal – shock waves – pulsars: general.

1 INTRODUCTION

Relativistic collisionless shocks are efficient sites of particle acceleration, producing non-thermal power-law energy spectra (Blandford & Eichler 1987) via the first-order Fermi process (Fermi 1949). Of particular interest is the intrabinary shock (IBS) in spider pulsars, compact binary systems harbouring a millisecond pulsar and a low-mass companion. Spider pulsars are classified as redbacks (RBs) if the companion is a non-degenerate star with mass $\sim 0.2\text{--}0.4 M_{\odot}$, or as black widows (BWs) if the companion is a degenerate star with mass $\ll 0.1 M_{\odot}$. The relativistic magnetically dominated pulsar wind impacts on to the companion (Phinney et al. 1988), ablating it and slowly ‘devouring’ its atmosphere – hence, the evocative name of these systems. The interaction between the pulsar wind and the companion’s stellar wind (or its magnetosphere) forms the IBS.

Non-thermal emission from the IBS dominates the X-ray band. Many rotation-powered BWs and RBs show evidence for orbitally modulated X-ray emission (Huang et al. 2012; Bogdanov et al. 2015, 2021; Roberts et al. 2015), likely powered by synchrotron cooling of relativistic electrons and positrons accelerated at the IBS (Harding & Gaisser 1990; Arons & Tavani 1993). The observed flux peaks just before and after the pulsar eclipse, which has been attributed to Doppler effects caused by the fast post-shock flow

(Romani & Sanchez 2016; Sanchez & Romani 2017; Wadiasingh et al. 2017, 2018; Kandel, Romani & An 2019; van der Merwe et al. 2020; Kandel, Romani & An 2021). The X-ray spectrum is markedly non-thermal, and has a relatively flat photon index, $\Gamma_X \simeq 1\text{--}1.5$ (Cheung et al. 2012; Romani, Filippenko & Cenko 2014; Arumugasamy, Pavlov & Garmire 2015; Roberts et al. 2015; Swihart et al. 2022). This implies a rather hard electron energy spectrum, with power-law slope $p = 2\Gamma_X - 1 \simeq 1 - 2$ (such that the electron distribution differential in Lorentz factor is $dN/d\gamma \propto \gamma^{-p}$). Such hard particle spectra are generally not expected from Fermi acceleration at relativistic shocks, which normally yields slopes $p > 2$ (Sironi, Keshet & Lemoine 2015 for a review). This apparent contradiction can be resolved by accounting for the fact that the pulsar wind is composed of toroidal stripes of alternating field polarity, separated by current sheets of hot plasma (Bogovalov 1999; Pétri & Lyubarsky 2007). When compressed at the shock, the oppositely directed fields annihilate via shock-driven reconnection. Fully kinetic particle-in-cell (PIC) simulations of the termination shock of a relativistic striped wind have shown that shock-driven reconnection produces hard power-law particle spectra, with a slope as hard as $p = 1$ (Sironi & Spitkovsky 2011; Lu et al. 2021). While useful to elucidate the microphysics of particle acceleration in striped shocks, the *local* approach adopted by these studies did not allow to capture the *global* IBS dynamics, which is typically investigated with fluid-type simulations (Bogovalov et al. 2008, 2012, 2019; Bosch-Ramon et al. 2012; Lamberts et al. 2013; Bosch-Ramon, Barkov & Perucho 2015; Huber et al. 2021).

* E-mails: jorgecortes@astro.columbia.edu (JC); lsironi@astro.columbia.edu (LS)

In Cortés & Sironi (2022; hereafter CS22), we presented the first *global* PIC simulations of the IBS in spider pulsars, assuming that the shock wraps around the companion star as inferred in some BWs. Spider pulsars offer a unique opportunity for global PIC simulations. In fact, the ratio of shock curvature radius R_{curv} to wavelength of the striped wind $\lambda = 2\pi c/\Omega$ (here, Ω is the pulsar spin frequency) is

$$\frac{R_{\text{curv}}}{\lambda} \sim 5 \times 10^1 \left(\frac{R_{\text{curv}}}{10^{10} \text{ cm}} \right) \left(\frac{\Omega}{10^3 \text{ s}^{-1}} \right). \quad (1)$$

The ratio of stripe wavelength to the typical post-shock Larmor radius $r_{\text{L,hot}}$ is (Sironi & Spitkovsky 2011; hereafter SS11)

$$\frac{\lambda}{r_{\text{L,hot}}} \sim 3 \times 10^1 \left(\frac{\kappa}{10^4} \right) \left(\frac{10^{11} \text{ cm}}{d_{\text{IBS}}} \right) \left(\frac{10^3 \text{ s}^{-1}}{\Omega} \right) \quad (2)$$

assuming a wind multiplicity (Goldreich & Julian 1969) of $\kappa \sim 10^4$ (Harding & Muslimov 2011; Timokhin & Harding 2015) and a distance between the shock and the pulsar of $d_{\text{IBS}} \sim 10^{11} \text{ cm}$ (see CS22 for more details). We note that a multiplicity of $\kappa \sim 10^4$ may be an overestimate for millisecond pulsars (Harding & Muslimov 2011), so the ratio in equation (2) might be closer to unity. Equations (1) and (2) imply that less than 3 order-of-magnitude separate global scales (R_{curv}) from plasma scales ($r_{\text{L,hot}}$). In CS22, we adopted realistic values for $\lambda/r_{\text{L,hot}} \sim 30$, but the curvature radius was just a factor of two larger than the stripe wavelength.

In order to bring kinetic global models of the IBS closer to realistic spider pulsar systems, in this work we extend the study by CS22 in three main directions. First, we investigate the case of larger companion stars (i.e. we increase the ratio R_{curv}/λ towards more realistic values). Second, we explore the case of a more magnetized pulsar wind. In both cases, we confirm that the results obtained in CS22 can be reliably applied to realistic spider pulsar systems. Third, we investigate in more detail the physics of particle acceleration following the stage governed by shock-driven reconnection. We find that the highest energy particles accelerated by reconnection can stream ahead of the shock and be further accelerated by the upstream motional electric field via the pick-up process, that is well known in space physics (e.g. Möbius et al. 1985; Iwamoto et al. 2022). In the downstream, further energization is governed by stochastic interactions with reconnection plasmoids/magnetic islands.

The paper is organized as follows: In Section 2, we describe the setup of our simulations. In Section 3, we present our results on the global flow dynamics, the physics of particle acceleration, and the spectrum and light curve of synchrotron emission. We summarize in Section 4 and discuss future prospects.

2 SIMULATION SETUP

We use the 3D electromagnetic PIC code TRISTAN-MP (Buneman 1993; Spitkovsky 2005). We employ a 2D spatial domain in the $x - y$ plane, but we track all three components of velocity and electromagnetic fields. Our setup parallels very close what we employed in CS22, which we summarize here for completeness.

Since the distance between the pulsar and the intrabinary shock is $d_{\text{IBS}} \gtrsim R_{\text{curv}}$, we assume for simplicity that the pulsar wind can be modelled as a sequence of plane-parallel stripes (in Section 4 we discuss the limitations of this assumption). The magnetically dominated electron–positron pulsar wind propagates along $-\hat{x}$. It is injected from a moving boundary, that starts just to the right of the companion and moves along $+\hat{x}$ at the speed of light c . This allows to save memory and computing time, while retaining all the regions that are causally connected with the initial setup (e.g. SS11; Sironi, Spitkovsky & Arons 2013). An absorbing layer for particles and

fields is placed at $x = 0$ (left-most boundary). Periodic boundaries are used along the y direction. The magnetic field in the pulsar wind is initialized as

$$B_y(x, t) = B_0 \tanh \left\{ \frac{1}{\Delta} \left[\alpha + \cos \left(\frac{2\pi(x + \beta_0 c t)}{\lambda} \right) \right] \right\}, \quad (3)$$

where $\beta_0 = (1 - 1/\gamma_0^2)^{1/2}$ is the wind velocity and γ_0 the bulk Lorentz factor. We present results for $\gamma_0 = 3$, but we have verified that a choice of $\gamma_0 = 10$ leads to the same conclusions, apart from an overall shift in the energy scale (see also SS11). The magnetic field flips across current sheets of hot plasma, having a thickness $\sim \Delta\lambda$. The field strength B_0 is parametrized via the magnetization $\sigma \equiv B_0^2/4\pi\gamma_0 mn_0 c^2$ (i.e. the ratio of Poynting to kinetic energy flux). Here, m is the electron (or positron) mass and n_0 the density of particles in the ‘cold wind’ (i.e. the region outside of current sheets, which instead have peak density of $4n_0$). We employ a fiducial magnetization of $\sigma = 10$, but we also discuss how our results change for a higher magnetization of $\sigma = 40$. Finally, α is a measure of the magnetic field averaged over one wavelength, such that $\langle B_y \rangle_\lambda / B_0 = \alpha/(2 - |\alpha|)$. A value of $\alpha = 0$ – or equivalently, ‘positive’ and ‘negative’ stripes of comparable width – indicates the equatorial plane of the pulsar, whereas $|\alpha|$ increases when moving away from the mid-plane (SS11). If ξ is the pulsar obliquity angle (i.e. the angle between the pulsar rotation and magnetic axes), the value of α at a given latitude ℓ above the pulsar equatorial plane is

$$\alpha = \frac{\tan \ell}{\tan \xi}. \quad (4)$$

For BWs, the spin axis of the pulsar is believed to be well aligned with the orbital angular momentum (Clark et al. 2023), so we will only explore small values of α : $\alpha = 0, 0.1$, and 0.3 .

The relativistic skin depth in the cold wind $c/\omega_p \equiv (\gamma_0 mc^2/4\pi e^2 n_0)^{1/2}$ is resolved with 10 cells, where e is the positron charge. It follows that the pre-shock Larmor radius $r_L \equiv \gamma_0 mc^2/eB_0 = (c/\omega_p)/\sqrt{\sigma}$ is resolved with three cells for $\sigma = 10$ and 1.6 cells for $\sigma = 40$. The post-shock Larmor radius, assuming complete field dissipation, is defined as $r_{\text{L,hot}} = \sigma r_L = (\gamma_0 \sigma) mc^2/eB_0 \simeq \gamma_\sigma mc^2/eB_0$ (below, we define $\gamma_\sigma = \gamma_0(1 + \sigma)$ as the mean particle Lorentz factor assuming full dissipation). The numerical speed of light is 0.45 cells/time-step. Within the cold wind, each computational cell is initialized with two pairs of cold ($kT/mc^2 = 10^{-4}$) electrons and positrons, but we have checked that higher numbers of particles per cell lead to the same results. The temperature in the overdense current sheets is set by pressure balance, which yields a thermal spread $kT_h/mc^2 = \Theta_h = \sigma/2\eta$, where we choose that current sheets are denser than the striped wind by a factor of $\eta = 3$.

Unless otherwise noted, our computational domain is $9600 c/\omega_p$ wide (in the y direction), or equivalently 96 000 cells wide. This is twice as large as compared to what we employed in CS22. The centre of the companion is placed at $(x_c, y_c) = (3000, 4800) c/\omega_p$, with a companion radius of $R_* = 140 c/\omega_p$. The companion surface (a cylinder, for our 2D geometry) is a conducting boundary for fields and a reflecting boundary for particles. The value for R_* is chosen such that the companion wind (see below) is stopped by the pulsar wind at $R_{\text{curv}} \simeq 400 c/\omega_p$, which then gives the characteristic shock curvature radius. The large width of our domain in the y direction is required such that the shock wraps around the companion until the $x = 0$ open boundary, i.e. the shock surface does not cross the periodic y boundaries. We set the stripe wavelength to be $\lambda = 100 c/\omega_p$ for $\sigma = 10$ and $\lambda = 200 c/\omega_p$ for $\sigma = 40$, so that the ratio $\lambda/r_{\text{L,hot}} \simeq 30$ in both cases. For $\sigma = 10$, we then have

$R_{\text{curv}}/\lambda \simeq 4$, i.e. twice as large as in our earlier work. Instead, for $\sigma = 40$ we have $R_{\text{curv}}/\lambda \simeq 2$ as in CS22.

In our setup, the pulsar wind is stopped by a companion wind launched isotropically from its surface. We do not aim to reproduce the realistic properties of the companion wind. Given that our focus is on pulsar wind particles, their acceleration, and emission, the companion wind will merely serve to halt the pulsar wind. We initialize an unmagnetized companion wind with realistic values of the radial momentum flux (twice larger than the momentum flux of the pulsar wind), but with artificially smaller particle density (and so, artificially higher wind velocity) to make the problem computationally tractable. In the remainder of this work we will only consider acceleration and emission of pulsar wind particles.

3 RESULTS

In this section, we present our main results on the global flow dynamics, the physics of particle acceleration, and the resulting synchrotron spectra and light curves. Our fiducial simulations in Sections 3.1–3.4 have $\sigma = 10$, $\lambda = 100 c/\omega_p$, and $R_* = 140 c/\omega_p$. When investigating the dependence on magnetization in Section 3.5, we will compare a simulation having $\sigma = 40$, $\lambda = 200 c/\omega_p$, and $R_* = 140 c/\omega_p$ with a simulation having $\sigma = 10$, $\lambda = 100 c/\omega_p$, and $R_* = 70 c/\omega_p$ (the same as in CS22), so that the two runs have the same R_*/λ and $\lambda/r_{\text{L,hot}}$. A detailed assessment of the dependence on R_* for $\sigma = 10$ is presented in Appendix A.

3.1 Flow dynamics

Fig. 1 shows the typical global morphology of our system once a quasi-steady state¹ has been reached. All panels in Fig. 1 and all the subsequent figures in this section refer to a time of $\omega_p t = 6524$. Given the absence of orbital motion in our simulations, the fluid structure is roughly symmetric with respect to the $y = 4800 c/\omega_p$ axis (which passes through the centre of the companion star), and so it is sufficient to show just the top half of our simulation domain. We present, from top to bottom, the number density of pulsar wind particles, the flow velocity of the pulsar wind, and the magnetic energy density. We remark that the top and middle rows only account for pulsar wind particles (the grey region around and to the left of the companion star is populated by companion wind particles). We vary α from 0 to 0.3, increasing from left to right columns.

The pulsar wind travels in the $-\hat{x}$ direction at relativistic speeds and comes into contact with the companion’s isotropic, unmagnetized wind. The interaction between the two winds sends a fast magnetohydrodynamic (MHD) shock back into the striped pulsar wind. The fast MHD shock is visible in the bulk flow velocity panels of Fig. 1 (middle row) as the right-most arc. The fast MHD shock compresses the incoming current sheets and initiates magnetic reconnection. As the flow propagates downstream from the fast MHD shock, reconnection progressively erases the striped structure of the pulsar wind, and the main shock eventually forms.² This is visible as a sharp arc in Fig. 1, and from now on we refer to this shock as the IBS. The IBS compresses and slows down the incoming flow (top and middle row in Fig. 1, respectively). Regardless of α , the middle row of Fig. 1 shows that the flow becomes partially stagnant near

¹Refer to Appendix B for an assessment of the quasi-steady state.

²Local PIC simulations have demonstrated that the distance between the fast MHD shock and the main shock is set by the time needed for reconnection islands to grow up to a size $\sim \lambda/2$ (SS11).

the apex of the IBS, transitioning to greater speeds $\sim 0.8c$ at higher latitudes, farther downstream. The trans-relativistic velocity of the downstream flow at high latitudes is expected to leave an imprint in the synchrotron light curves, via Doppler boosting (Arons & Tavani 1993; Romani & Sanchez 2016; Wadiasingh et al. 2017). We postpone this discussion to Section 3.3, where we explore the effect of Doppler boosting on synchrotron emission.

The development of reconnection from the fast shock to the IBS causes the formation of magnetic islands/plasmoids, easily distinguishable as the overdense, magnetized, quasi-circular blobs seen in the number density (top row), and magnetic energy density (bottom row) in Fig. 1 (in particular, for $\alpha = 0$). The typical plasmoid size at the IBS is set by the separation between two consecutive current sheets in the pulsar wind, see SS11. Adopting the same formalism as in SS11, the stripe-averaged field can be expressed as $\langle B_y \rangle_\lambda / B_0 = (\lambda_+ - \lambda_-)/\lambda$, where λ_+ is the width of the ‘positive’ field region ($B_y = +B_0$), λ_- is the width of the ‘negative’ field region ($B_y = -B_0$), and $\lambda_+ + \lambda_- = \lambda$. The spacing between successive current sheets alternates between λ_+ and λ_- , the minimum of which is responsible for inhibiting plasmoid growth, and thus setting the maximum plasmoid size at the IBS. In fact, the distance between the fast MHD shock and the IBS is set by the requirement that, as the flow propagates downstream from the fast MHD shock, reconnection plasmoids grow up to a size $\sim \min(\lambda_+, \lambda_-)$. This condition identifies the location of the IBS. There, the ordered striped structure of the magnetically dominated pulsar wind transitions to a more disordered, turbulent medium with weaker fields, since some of the field energy has been converted to particle energy (at the IBS, B^2/B_0^2 turns from cyan to purple, see bottom row in Fig. 1).

The argument we have just presented sets the typical plasmoid size at the IBS. As Fig. 1 shows, the plasmoid properties downstream from the IBS significantly depend on α . For $\alpha = 0$, it is apparent that plasmoids merge with each other and grow in size while flowing away from the IBS, with the largest plasmoids reaching a size that is comparable to global scales (e.g. to the companion radius, or even the shock curvature radius). As α increases, the distance behind the IBS where plasmoids survive progressively shrinks. This is a consequence of the fact that, for $\alpha \neq 0$, the downstream flow retains a net magnetic field in the y direction (for $\alpha > 0$, a net positive field). When a plasmoid merges with a region threaded by the net field, it will necessarily be engulfed and disappear, since the region with net field contains much more magnetic flux than the plasmoid itself. This explains why only few plasmoids survive in the downstream region of the $\alpha = 0.1$ case, and none for $\alpha = 0.3$.

We conclude this subsection by commenting on the fact that, while the shape and location of the IBS is nearly the same for the three values of α that we explore,³ the location of the contact discontinuity between the shocked companion wind and the shocked pulsar wind is significantly different. The contact discontinuity, seen as the upper boundary of the grey regions in the number density and flow velocity panels of Fig. 1, wraps around the companion more tightly as α increases. This is ultimately related to the fact that for larger α , a smaller fraction of the incoming Poynting flux is available for dissipation, since the stripe-averaged field $\langle B_y \rangle_\lambda / B_0 = \alpha/(2 - |\alpha|)$

³We point out, however, that the IBS apex is farther from the companion for increasing α . This is due to the fact that the net stripe-averaged field accumulates secularly in front of the companion, for our 2D geometry. We expect the accumulation of magnetic energy/pressure to be partially alleviated in 3D, as the field lines can pass above/below the spherical companion (unlike for the cylindrical companion in 2D).

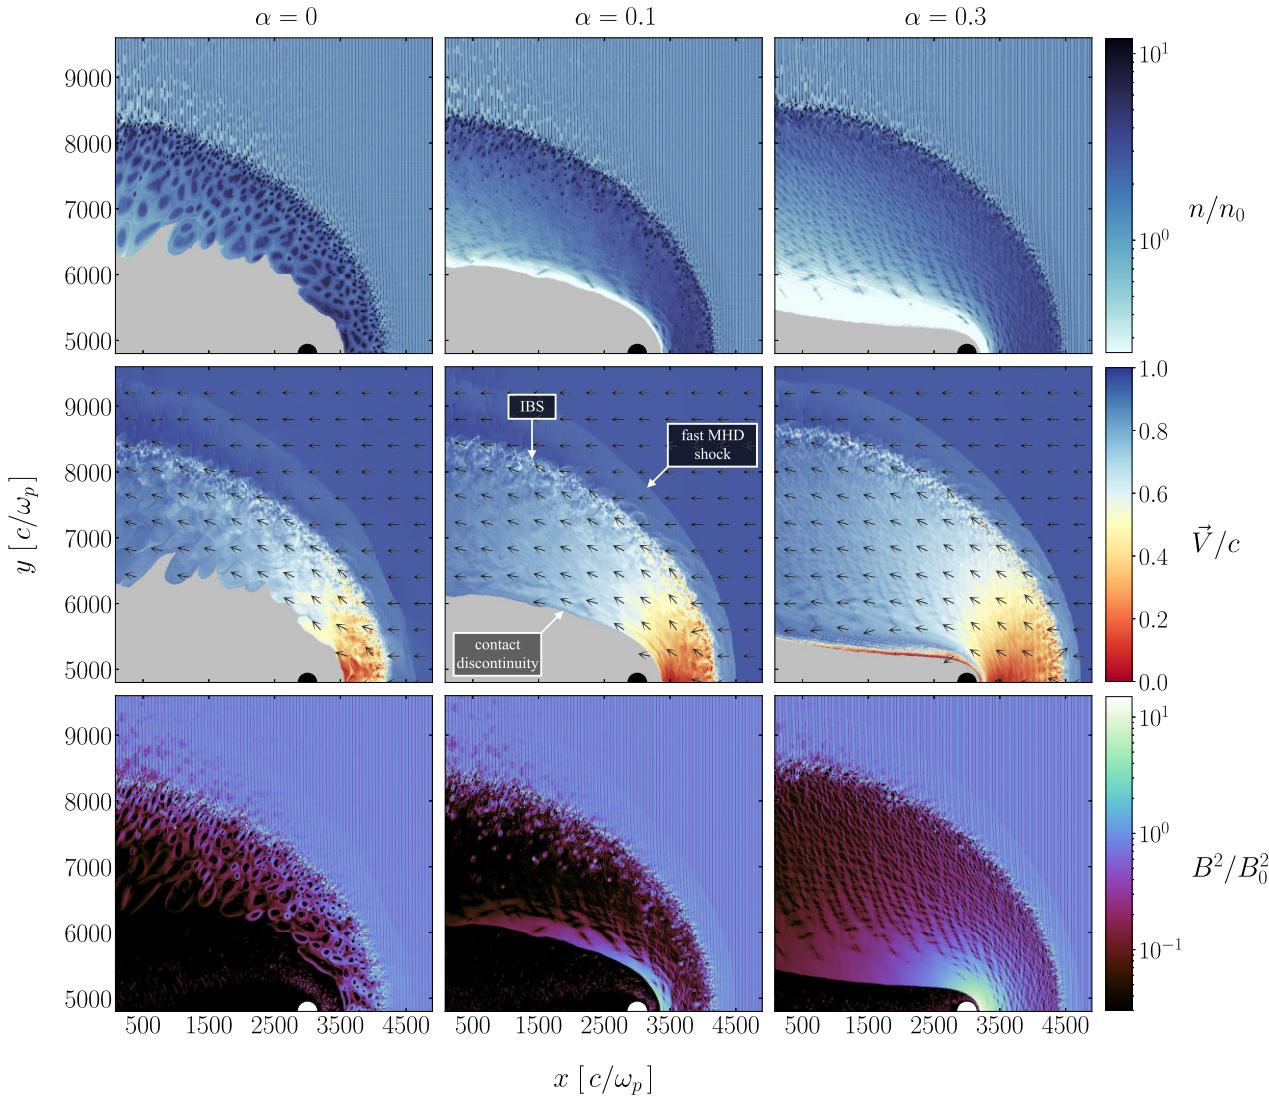


Figure 1. 2D plots of the top half simulation domain for $\alpha = 0, 0.1$ and 0.3 (left to right columns). We adopt a magnetization of $\sigma = 10$ and a companion size of $R_* = 140 c/\omega_p$, and we show results for $\omega_p t = 6524$. *Top*: number density of pulsar wind particles in units of n_0 . *Middle*: flow velocity of the pulsar wind in units of c , with arrows of unit length depicting flow direction. *Bottom*: magnetic energy density in units of the upstream value $B_0^2/8\pi$. Either a black circle (top two rows) or a white circle (bottom row) represents the companion star. In the top and middle rows, the grey region around and to the left of the companion star is populated by companion wind particles, which we exclude from our analysis. The particle density (top) and the mean velocity (middle) are obtained in each cell by averaging over the neighbouring 5×5 cells (here, and in the following).

is preserved (and compressed) across the shock. The mean energy per particle in the incoming flow (in units of the rest-mass energy) is $\gamma_\sigma = \gamma_0(1 + \sigma)$. Assume that a fraction ζ is converted to particle energy, while $1 - \zeta$ stays in magnetic fields, with ζ decreasing for larger α given that the stripe-averaged field will not dissipate. While the energy per particle in the downstream is still γ_σ , the pressure that each particle can provide (more precisely, the pressure per unit rest-mass energy) is $[(\hat{\gamma} - 1)\zeta + (1 - \zeta)]\gamma_\sigma$, where $\hat{\gamma}$ is the adiabatic index ($\hat{\gamma} = 4/3$ for an ultrarelativistic gas). It follows that, at fixed γ_σ , the post-shock pressure is higher for smaller ζ (and so, for larger α). The larger post-shock pressure in the higher α cases results in greater confinement of the companion's wind, forcing the contact discontinuity to be closer to the companion.

The properties of the flow dynamics illustrated here are not much different from the results that we presented in CS22. We refer the

reader to Appendix A for a detailed comparison between the results of our previous simulations (having $R_* = 70 c/\omega_p$) and those of this work (where $R_* = 140 c/\omega_p$).

3.2 Particle energy spectra and synchrotron spectra

The particle energy spectra and the synchrotron spectra extracted from our simulations are presented in Fig. 2 for $\omega_p t = 6524$ (same time as in Fig. 1, when the system has achieved a quasi-steady state). For both particle energy spectra and synchrotron spectra, we compute them in spatial patches of 100×100 cells, and we consider only the contribution from pulsar wind particles in the post-shock region. We identify the shock downstream as the region where the mean particle Lorentz factor $\langle \gamma \rangle > 0.3 \gamma_\sigma$. The top panel of Fig. 2 shows the downstream particle spectra $dN/d\ln(\gamma - 1)$ for the three

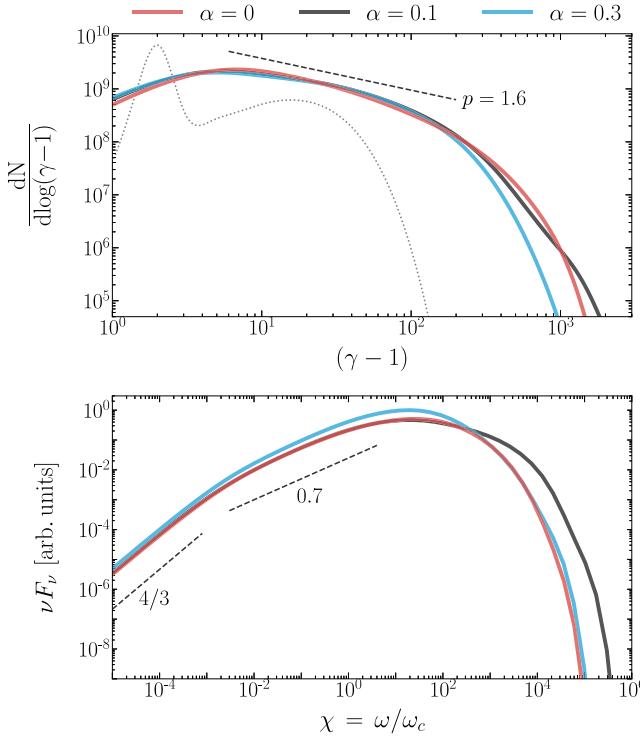


Figure 2. *Top:* downstream particle energy spectra for $\alpha = 0$ (red solid), 0.1 (black solid), and 0.3 (cyan solid) at $\omega_p t = 6524$, the same time as in Fig. 1. For comparison, the dotted grey line is the pre-shock spectrum. *Bottom:* angle-integrated synchrotron spectra vF_v (same colour coding as in the top panel). We define $\chi = \omega/\omega_c$, where the characteristic synchrotron frequency $\omega_c = \gamma_\sigma^2 e B_0 / mc$ is calculated for the average post-shock Lorentz factor $\gamma_\sigma = \gamma_0(1 + \sigma)$ assuming complete field dissipation.

values of α used in this work. For comparison, we also show the energy spectrum of the pre-shock pulsar wind, integrated over an area extending for $\Delta x = 400 c/\omega_p$ along the x direction and as wide as the simulation domain in the y direction. The pre-shock spectrum (dotted grey line) is a combination of the ‘cold’ plasma in the striped wind (with $\gamma \simeq \gamma_0$) and the ‘hot’ Maxwellian of particles initialized in the current sheets (with $\gamma \simeq 3\gamma_0\Theta_h \simeq 15$, where Θ_h is the thermal spread of the hot particles in the current sheets).

Downstream from the IBS, the particle spectrum can be described as a single broad distribution. In the range between $\gamma \sim \gamma_0 = 3$ (the initial bulk flow Lorentz factor) and $\gamma \sim \gamma_\sigma \sim 30$ (the Lorentz factor that particles acquire in the case of complete field dissipation), the particle spectrum can be modelled as a power-law $dN/dy \propto (\gamma - 1)^{-p}$ with a slope $p \simeq 1.6$ independent of α ($\alpha = 0$ in red, $\alpha = 0.1$ in black and $\alpha = 0.3$ in blue). In fact, this slope persists up to a few times γ_σ . As we further discuss below, particle acceleration in this energy range is governed by shock-driven reconnection at the IBS. Local PIC studies by SS11 showed that the power-law range exists as long as $\lambda/r_{L,hot} \gtrsim 10$, while for lower values of this ratio the spectrum resembles a relativistic Maxwellian.

The main differences among the three α values are seen in the high-energy tail (i.e. $\gamma \gtrsim 200$). The cases $\alpha = 0$ and 0.1 have very similar cutoffs. As we describe in Section 3.4, for $\alpha = 0$ additional channels of energization can operate both upstream and downstream, which extend the spectrum well beyond the energy range controlled by shock-driven reconnection. Similar arguments apply to $\alpha = 0.1$. In particular, following SS11, we attribute the high-energy bump seen for $\alpha = 0.1$ at $\gamma \gtrsim 800$ to shock-drift acceleration in the stripe-

averaged motional electric field $\langle E_z \rangle_\lambda = \beta_0 \langle B_y \rangle_\lambda$, as particles pre-energized by shock-driven reconnection gyrate around the shock front, before eventually advecting downstream. As $\langle E_z \rangle_\lambda = 0$ for $\alpha = 0$, shock-drift acceleration cannot operate, and in fact we find that particle energization ahead of the shock is governed by scattering off self-generated magnetic turbulence (more on this in Section 3.4). While the $\alpha = 0.3$ case possesses a non-zero $\langle E_z \rangle_\lambda$, here shock-drift acceleration appears inhibited (the spectral cutoff lies at lower energies than for smaller α). As α increases, the average energy per particle in the downstream drops below the reference value $\gamma_\sigma mc^2$, since the stripe-averaged magnetic field is preserved, and so a smaller fraction of the incoming Poynting flux is available for dissipation (see SS11, where values of α up to unity were explored). It follows that the highest energy particles resulting from shock-driven reconnection have lower energies for higher α , and so they may not be able to propagate back into the upstream far enough to sample the motional electric field and be energized by shock-drift acceleration.

The synchrotron spectra in the bottom panel of Fig. 2 roughly mirror the properties of the particle energy spectra. Synchrotron spectra are calculated following Reville & Kirk (2010), by summing over the angle-integrated synchrotron emission from every particle in the downstream region.⁴ We normalize the synchrotron frequency on the horizontal axis with respect to the characteristic frequency $\omega_c = \gamma_\sigma^2 e B_0 / mc$ of particles with $\gamma = \gamma_\sigma$. At low frequencies, all spectra display the expected $vF_v \propto v^{4/3}$ scaling. At $\chi = \omega/\omega_c \sim 10^{-2}$, all three spectra transition to a power-law scaling $vF_v \propto v^{0.7}$ that extends slightly beyond $\chi \sim 1$. This frequency range corresponds to the energy range where the particle spectrum can be modelled as $dN/dy \propto (\gamma - 1)^{-1.6}$. In fact, a power-law index of $p \simeq 1.6$ in the particle energy spectrum results in a synchrotron spectrum $vF_v \propto v^{(3-p)/2} = v^{0.7}$, as indeed observed. Since the power-law energy range in the particle spectrum extends from $\gamma \sim \gamma_0$ up to $\gamma \gtrsim \gamma_\sigma$, the corresponding power-law synchrotron spectrum will range from $\chi \sim 10^{-2}$ up to $\chi \gtrsim 1$, as indeed confirmed by the bottom panel in Fig. 2. The high-frequency cutoff of the synchrotron spectrum for $\alpha = 0$ is nearly the same as for $\alpha = 0.3$, even though the particle energy spectrum of $\alpha = 0$ cuts off at higher energies than for $\alpha = 0.3$. While a smaller fraction of the incoming Poynting flux is converted to particle energy for increasing α , the residual post-shock magnetic field is stronger at higher α (see bottom row in Fig. 1), so the two opposite effects nearly cancel.

To identify how the synchrotron spectral hardness in the $\chi = 10^{-2} - 10^0$ range depends on spatial location, in Fig. 3 we create spectral hardness maps for $\alpha = 0, 0.1$ and 0.3 (from left to right). We divide the simulation domain into patches of 100×100 cells, compute the isotropic synchrotron spectrum within each patch, and perform a linear fit (in log-log space) to obtain the slope Γ of the vF_v spectrum in the range $\chi = 10^{-2} - 10^0$. If the synchrotron spectrum is $vF_v \propto v^\Gamma$, then the spectral hardness Γ is related to the particle power-law slope p as $\Gamma = (3 - p)/2$. Regardless of α , we find harder spectra with $\Gamma \simeq 0.75$ near the apex of the IBS, transitioning to $\Gamma \simeq 0.6$ farther downstream and at higher latitudes, where the upstream flow enters the shock at a more oblique angle. For $\alpha = 0$, the spectrum in plasmoid cores appears softer, with $\Gamma \simeq 0.3$. Plasmoid cores are mostly populated by particles that initially resided in the hot current sheets of the striped wind. They got locked in plasmoid cores without experiencing further energization via shock-

⁴As in our previous work (CS22), we do not include radiative cooling losses in the particle equations of motion. We defer the exploration of the effects of synchrotron cooling to an upcoming work.

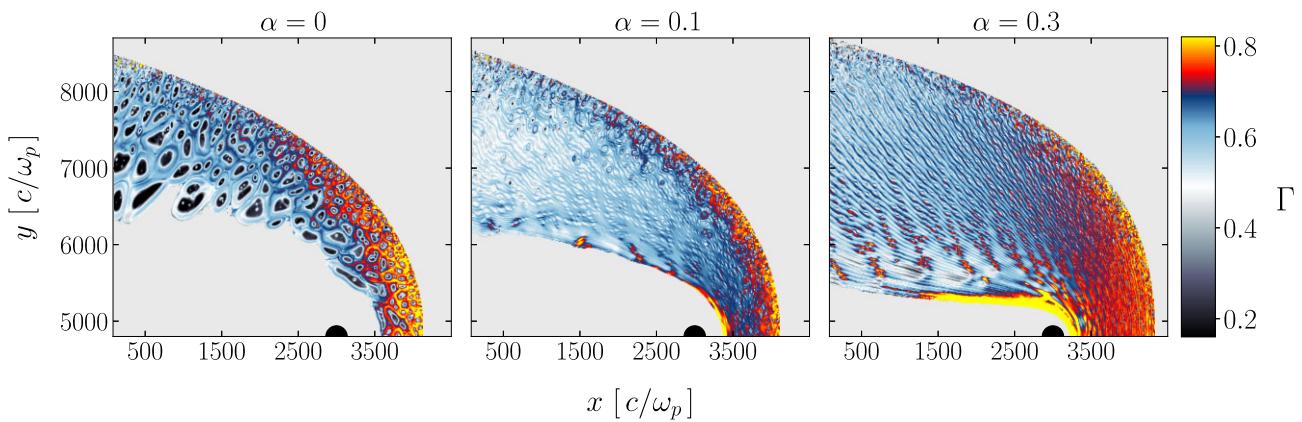


Figure 3. 2D plots of the spectral hardness Γ , defined as the best-fitting power-law slope $\nu F_\nu \propto \nu^\Gamma$ in the range $\chi = 10^{-2} - 10^0$, where $\chi = \omega/\omega_c$. Both the pulsar wind particles in the upstream region and the companion wind particles are omitted (these regions are depicted in light grey). From left to right, we present results for $\alpha = 0, 0.1$ and 0.3 at $\omega_p t = 6524$, the same time as in Figs 1 and 2.

driven reconnection, which explains why their synchrotron spectrum is softer. From the bottom panel of Fig. 2, plasmoid cores do not appear to play a significant role in setting the hardness of the space-integrated synchrotron spectrum.

3.3 Synchrotron light curves

Phase-resolved light curves for $\chi = 1$ are shown in the top row of Fig. 4.⁵ All curves are normalized to the peak of the $\alpha = 0.3$ case. The $\alpha = 0.3$ light curve (right-most panel) is strongly peaked at superior conjunction ($\phi = 0$, when the pulsar is eclipsed by the companion); the peak is shallower for $\alpha = 0.1$ (top-centre panel); interestingly, for $\alpha = 0$ (left-most panel) the light curve peaks at $\phi \simeq \pm\pi/8$, i.e. just before and after superior conjunction. For $\alpha = 0$, the two peaks are separated by $\simeq 0.8$ rad (equivalently, ~ 0.13 in normalized phase), and the flux at $\phi = 0$ drops by ~ 10 per cent below the peak flux. The flux at inferior conjunction ($\phi = \pm\pi$) is nearly a factor of 10 lower than the peak. As we argued in CS22, all these properties hold throughout the spectral range $\chi = 10^{-2} - 10^1$ where $\nu F_\nu \propto \nu^{0.7}$, and are in good agreement with X-ray data (e.g. Bogdanov, Grindlay & van den Berg 2005; Bogdanov et al. 2011, 2021; Huang et al. 2012; Roberts et al. 2014, 2015).

Orbital modulations in the X-ray band have been attributed to Doppler boosting in the post-shock flow (Arons & Tavani 1993; Romani & Sanchez 2016; Wadiasingh et al. 2017). The light curves extracted from our PIC simulations self-consistently include Doppler effects. Still, we can further investigate the role of Doppler boosting, by artificially overemphasizing or de-emphasizing it. We now concentrate on the bottom row of Fig. 4. We divide the simulation domain into patches of 100×100 cells. In each patch, we compute the mean bulk fluid velocity β_b , by averaging the velocities of individual particles in that patch. We then compute the local Doppler factor $\delta = [\Gamma_b(1 - \hat{n} \cdot \beta_b)]^{-1}$, which, for patch (i, j) , we call δ_{ij} . Here, $\Gamma_b = 1/(1 - \beta_b^2)^{1/2}$ and \hat{n} identifies the line of sight. We then weigh the local synchrotron flux $F_{ij}(\chi = 1)$ by different powers s of the Doppler factor δ_{ij} , as indicated by the legend below the

plot, and finally sum over the patches, $\sum_{i,j} F_{ij}(\chi = 1) \delta_{ij}^s$. Positive powers ($s > 0$) artificially enhance the role of Doppler boosting, while negative powers tend to remove Doppler effects (i.e. they approach the case of emission computed in the fluid comoving frame). The lines for $s = 0$ (shown in black) are the same as in the top three panels, i.e. they show the light curves extracted from our simulations, that properly include Doppler effects.

For $\alpha = 0$, we find that the double-peaked nature of the light curve is emphasized even more with $s \geq 1$, suggesting that Doppler effects have a key role in shaping the light curve. In fact, the same double-peaked structure appears also for $\alpha = 0.1$ if $s \geq 1$. In contrast, the $\alpha = 0.3$ light curves peak at superior conjunction regardless of s . We then regard $\alpha = 0.1$ as a borderline case (with a light curve ‘almost’ double-peaked), and conclude that $\alpha \lesssim 0.1$ orientations are generally conducive to the double-peaked light curves often detected from spider pulsar systems.

3.4 The physics of particle acceleration

As we already mentioned, shock-driven reconnection can energize the incoming particles up to a typical Lorentz factor $\gamma \sim \gamma_\sigma = \gamma_0(1 + \sigma)$. Further acceleration is governed by other mechanisms, which we now investigate. The dependence on α of the properties of high-energy particles ($\gamma \gg \gamma_\sigma$) is illustrated in Fig. 5. First of all, we note that the top and bottom rows cover the same region of space, i.e. the lower half of the simulation domain ($y = 0 - 4800 c/\omega_p$). The vertical axis of the top row is ‘inverted’, with y increasing downward; this is done to more easily correlate top and bottom panels for a given α . The top row shows the spatial structure of the maximum Lorentz factor γ_{max} , defined as follows. We divide the simulation domain in patches of 100×100 cells and in each patch we compute the particle energy spectrum $dN/d\gamma$ (including electrons and positrons). Then γ_{max} is calculated in each patch as

$$\gamma_{max} = \frac{\int d\gamma (\gamma - 1)^{n+1} dN/d\gamma}{\int d\gamma (\gamma - 1)^n dN/d\gamma}. \quad (5)$$

We choose $n = 5$ which is sufficiently large that $(\gamma - 1)^n dN/d\gamma$ peaks near the high-energy cutoff of the particle spectrum, for which γ_{max} is then a good proxy. In the bottom row of Fig. 5, we present the 2D structure of the out-of-plane magnetic field component B_z in units of B_0 . We overlay a random sample of high-energy particles having $\gamma > 700$ (electrons in pink, positrons in green).

⁵We remark that, if the orbital plane does not coincide with the pulsar equatorial plane, the value of α should depend on orbital phase. Given that the most pronounced differences in light curves emerge at/near superior conjunction, one should interpret the value of α that we quote below as the instantaneous value at/near superior conjunction.

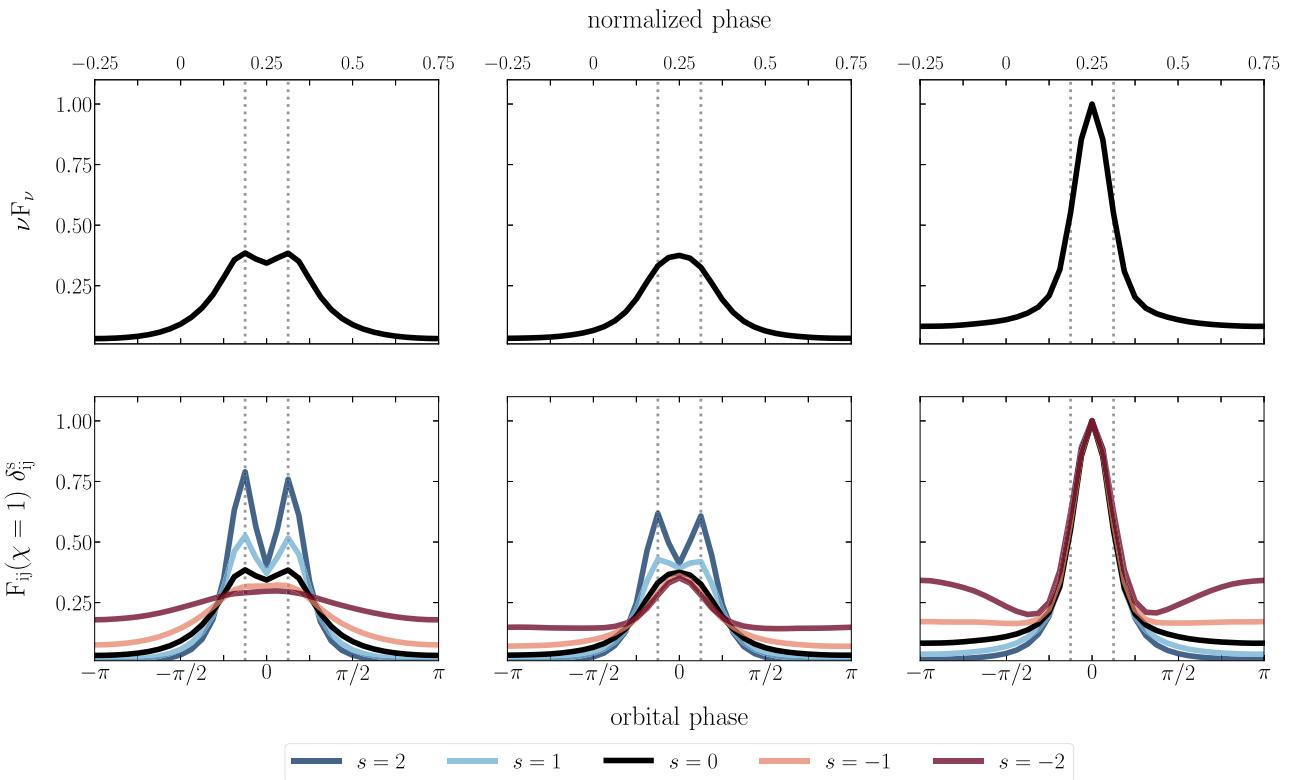


Figure 4. Phase-resolved synchrotron light curves. *Top:* light curves for $\alpha = 0, 0.1$, and 0.3 at $\omega_p t = 6524$, as a function of orbital phase ϕ (bottom axis). We choose $\phi = 0$ to correspond to superior conjunction (when the pulsar is eclipsed), while $\phi = \pm\pi$ to inferior conjunction. The top axis is the normalized phase, with 0.25 at the pulsar superior conjunction, as typical for X-ray observations. The light curves are shown for $\chi = 1$, where $\chi = \omega/\omega_c$. Grey dotted lines at $|\phi| = \pi/8$ indicate the location of the two peaks for the $\alpha = 0$ case. *Bottom:* In every patch of the domain (of 100×100 cells), we weigh the corresponding synchrotron flux $F_{ij}(\chi = 1)$ by different powers s of the local Doppler factor δ_{ij} , as indicated by the legend, and then we compute the sum $\sum_{i,j} F_{ij}(\chi = 1) \delta_{ij}^s$. For each s , the light curves are normalized to the peak value of the corresponding $\alpha = 0.3$ light curve.

As the top row shows, the 2D distribution of γ_{\max} is very sensitive to α . For $\alpha = 0.3$, the downstream region is nearly uniform, with $\gamma_{\max} \sim 400$. Higher values of $\gamma_{\max} \sim 500 - 600$ are seen in the downstream regions of $\alpha = 0$ and $\alpha = 0.1$. This trend mirrors the hierarchy of high-energy cutoffs in the particle spectra of Fig. 2, which are integrated over the whole downstream. The difference among the three α cases is even more dramatic in the upstream region near the shock. The case with $\alpha = 0.3$ does not show any evidence of high-energy particles residing in the upstream. In contrast, the near-upstream region of the $\alpha = 0.1$ case displays large values of $\gamma_{\max} \sim 700 - 1000$, modulated on the periodicity of the striped wind. We find that $\gamma_{\max} \sim 700$ in front of the IBS apex, and it reaches $\gamma_{\max} \sim 1000$ at higher latitudes. We attribute the enhancement of γ_{\max} near the shock for $\alpha = 0.1$ to the effect of shock-drift acceleration, as already discussed in Section 3.2. The increase of γ_{\max} with latitude can then be understood from the fact that particles energized via shock-drift acceleration at higher latitudes can spend more time in the upstream before being advected downstream.

The main differences between $\alpha = 0.1$ and $\alpha = 0.3$ are confined to the near upstream region. Further ahead of the shock, panels (b) and (c) are nearly indistinguishable. In contrast, the case $\alpha = 0$ in panel (a) shows that γ_{\max} is enhanced throughout the upstream region. This reveals that for $\alpha = 0$ high-energy particles can stream far ahead of the shock. The counter-streaming between the population of high-energy particles propagating back upstream and the incoming flow generates magnetic filaments in B_z (see panel d) via the Weibel

instability (Fried 1959; Weibel 1959), as already observed in local PIC simulations (SS11). In panel (d), we overlay on the 2D map of B_z a sample of high-energy particles having $\gamma > 700$. We find that both electrons (pink) and positrons (green) tend to cluster in regions of small B_z (i.e. at the transition between yellow and blue). High-energy electrons are mostly confined in between a filament with $B_z > 0$ lying below and a filament with $B_z < 0$ lying above, while the opposite holds for high-energy positrons. In order to propagate back upstream, the ∇B drift velocity in the B_z field needs to be oriented away from the shock; for back-streaming electrons, this requires doing the lower half of their Larmor gyration with $B_z > 0$, while the upper half with $B_z < 0$ (the opposite holds for positrons). Thus, their spatial clustering is enforced by the condition to propagate far into the upstream. The non-uniformity in the spatial distribution of high-energy particles in panel (d) is mirrored by the filamentary structures seen in γ_{\max} in panel (a).

Going forward, we restrict our analysis to the $\alpha = 0$ case, which displays the most pronounced evidence for particle acceleration to $\gamma \gg \gamma_\sigma$. We track a sample of $\sim 10^5$ particles with $\gamma_{\text{end}} \geq 200$, where γ_{end} is the Lorentz factor at the final time of our simulation ($\omega_p t = 6524$). The tracked particles are saved with an output cadence of $\Delta t = 450 \omega_p^{-1}$.

In Fig. 6, we show the 2D histogram of the tracked particles in the $\gamma - \mathcal{R}$ space as well as two representative particle trajectories overlaid. The quantity \mathcal{R} is defined as follows. At each time, we fit the shape of the IBS with an ellipse, whose centre is at $(x_0, y_0 =$

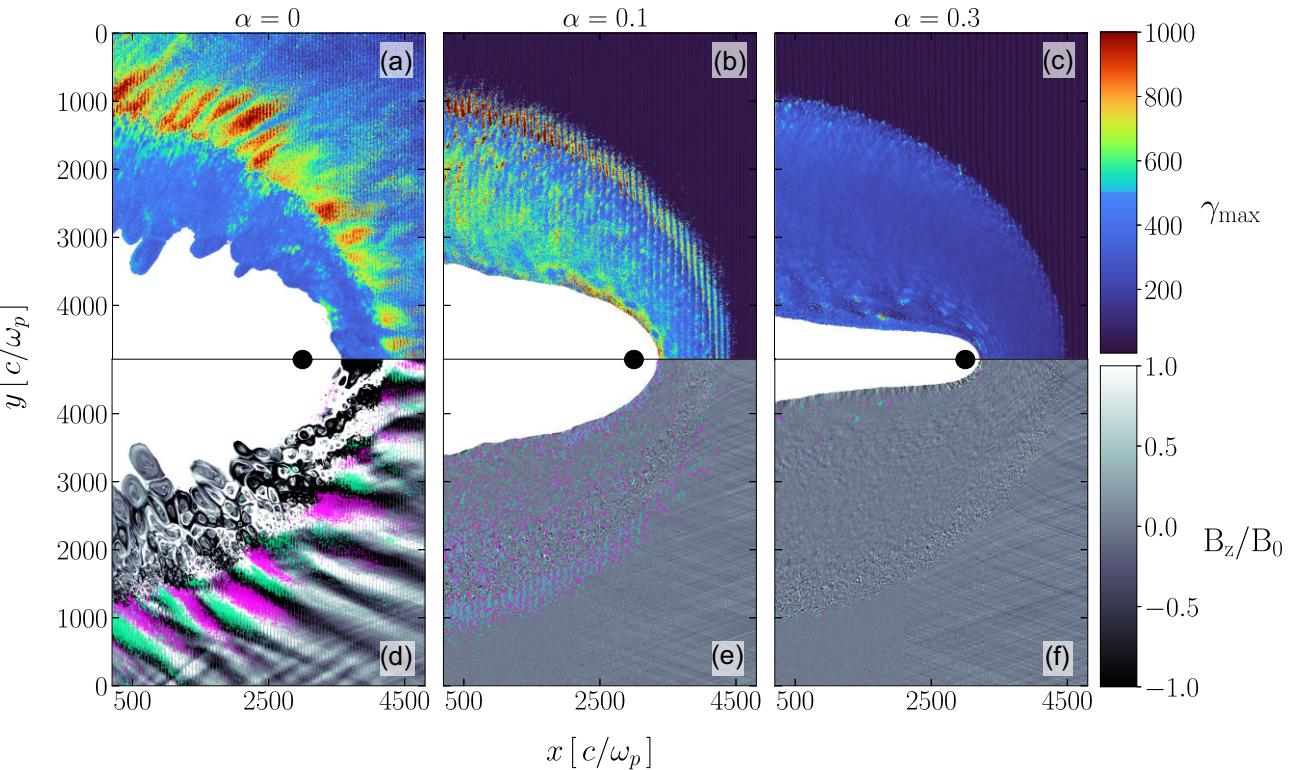


Figure 5. 2D plots of the spectral cutoff γ_{\max} (top; as defined in the text and in equation 5) and of the out-of-plane field B_z/B_0 (bottom) for $\alpha = 0, 0.1$ and 0.3 (from left to right) at $\omega_p t = 6524$. The vertical axis of the top row has the y coordinate increasing downward, to more easily correlate top and bottom panels for a given α . A subset of particles with $\gamma > 700$ are overlaid on the B_z/B_0 plots. Electrons are shown in pink, while positrons in green. The spatial locations of the two species are roughly coincident in panel (e), whereas they are rather distinct in panel (d). The white region around and to the left of the companion star is populated by companion wind particles, which we exclude from our analysis.

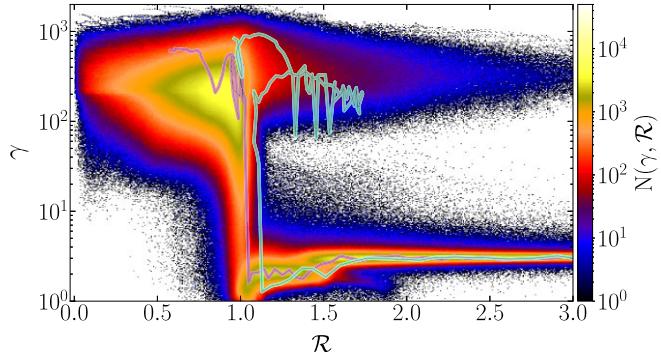


Figure 6. 2D histogram of a sample of $\sim 10^5$ particles tracked in the $\gamma - R$ space for $\alpha = 0$, with two representative trajectories overlaid. The particles are selected such that at the final time their Lorentz factor is $\gamma_{\text{end}} \geq 200$ (this threshold is visible as a horizontal discontinuity on the downstream side at $R \lesssim 0.7$). Each particle appears as many times as it is saved.

y_c). At each time, the best-fitting values for the semimajor (a_p) and semiminor (b_p) axes are then used to compute

$$R = \frac{(x - x_0)^2}{a_p^2} + \frac{(y - y_0)^2}{b_p^2} \quad (6)$$

for a particle having coordinates (x, y) . It follows that R serves as a proxy for the particle position, at any point in time, relative to the IBS. A particle with $R = 1$ is on the ellipse (i.e. at the IBS). Based off of its R , a particle can thus be classified as residing in one of

three regions: (i) downstream, $R \lesssim 0.9$; (ii) IBS, $0.9 \lesssim R \lesssim 1.1$; or (iii) upstream, $R \gtrsim 1.1$.

Fig. 6 describes the typical trajectory of high-energy particles ($\gamma_{\text{end}} \geq 200$) as a 2D histogram in the $\gamma - R$ space, with two representative trajectories overlaid (cyan and purple). The pulsar wind comes towards the shock with a typical Lorentz factor $\gamma \simeq \gamma_0 = 3$. The flow is slightly decelerated at the fast MHD shock ($R \sim 1.4$), where its typical Lorentz factor decreases down to $\gamma \simeq 2$. Upon interaction with the IBS ($R \sim 1.0$), the incoming particles are rapidly energized by shock-driven reconnection up to $\gamma \sim \gamma_\sigma \simeq 30$. Relatively few particles populate the area delimited by $2 \lesssim \gamma \lesssim 30$ and $0.9 \lesssim R \lesssim 1.1$ where particle energization is governed by shock-driven reconnection, since it is a fast acceleration mechanism, and so the residence time in this portion of phase space is short. Further energization to $\gamma \gg \gamma_\sigma$ can occur both in the upstream ($R \gtrsim 1.1$, see the cyan trajectory) as well as in the downstream ($R \lesssim 0.9$, see the purple trajectory). Eventually, most high-energy particles end up in the downstream at $0.5 \lesssim R \lesssim 1.0$.

The histogram shows that particles energized by shock-driven reconnection can propagate back into the upstream only if their Lorentz factor is $\gamma \gtrsim 100$, as apparent by the lack of particles having $10 \lesssim \gamma \lesssim 100$. In order to propagate far ahead of the shock, a particle must be able to traverse half of the stripe wavelength (and so, reverse its sense of gyration in the upstream alternating B_y field) before being overtaken by the shock. This criterion is best phrased in the upstream frame, where the time to cross half of the stripe wavelength is $\gamma_0 \lambda / 2c$. For a relativistic shock seen in the upstream frame, particles returning upstream are caught up by the

shock after completing a fraction $\sim \gamma_0^{-1}$ of their Larmor orbit, and so after a time $\gamma_0^{-1}(2\pi/\omega_L)$, where $\omega_L = eB_0/\gamma_0^2\gamma mc$ is the Larmor frequency – measured in the upstream frame – for a particle having Lorentz factor γ in the downstream frame. Writing $\gamma = \xi\gamma_\sigma$ and assuming $\sigma \gg 1$, the condition for propagating back upstream can be cast as

$$\xi \gtrsim \frac{1}{4\pi} \frac{\lambda}{r_{L,\text{hot}}}, \quad (7)$$

where $r_{L,\text{hot}} = \sqrt{\sigma}c/\omega_p$ is the typical post-shock Larmor radius assuming complete field dissipation. For our parameters $\lambda/r_{L,\text{hot}} \simeq 30$, which implies $\xi \gtrsim 2.4$ and so a lower limit on the Lorentz factor of $\gamma \gtrsim 70$, in good agreement with Fig. 6.

The acceleration physics of high-energy particles is further investigated in Fig. 7, where we show the trajectories of two representative electrons: on the left, an electron that acquires most of its energy while in the upstream; on the right, an electron that undergoes a number of stochastic scatterings while in the downstream. We present their spatial tracks in the top panels, and the time evolution of their Lorentz factor in the bottom panels. The early stages of energization are the same for the two particles: they interact with the shock (at $\omega_p t \sim 2500$ on the left, at $\omega_p t \sim 1500$ on the right) and get accelerated by shock-driven reconnection up to $\gamma \gtrsim 100$.

After interacting with the IBS, the electron in the left column propagates back upstream, and eventually undergoes an episode of significant energization after $\omega_p t \sim 5800$ (marked by the green circle). Here, the electron is kicked back towards the shock and its Lorentz factor increases up to $\gamma \sim 1000$. We identify the acceleration mechanism as akin to the pick-up process widely discussed in space physics (e.g. Möbius et al. 1985; Iwamoto et al. 2022). The electron is efficiently accelerated by the motional electric field $E_z = \beta_0 B_y$ while gyrating around the upstream field B_y . Since the sign of E_z switches every half wavelength, efficient acceleration requires that the particle moves back towards the shock with speed $\beta_x \simeq -\beta_0$, so the electron always ‘sees’ the same sign of the upstream motional electric field. We find that this mechanism governs the energization of most of the particles that propagate back upstream. In Fig. 8, we show the $x - \gamma$ trajectories of several electrons undergoing pick-up acceleration. We manually identify the beginning of the energization episode, and record the location x_{kick} and the energy γ_{kick} at this time. The figure shows that the tracks of various particles nearly overlap in the $[x - x_{\text{kick}}, (\gamma - \gamma_{\text{kick}})/\gamma_0]$ space. In the inset of Fig. 8, we separate the contributions of various electric field components. We define $\Delta\gamma_j = -\int_{t_{\text{kick}}}^t eE_j \beta_j dt/mc$ (here, β_j is the j -component of the particle dimensionless velocity), such that $\gamma - \gamma_{\text{kick}} = \Delta\gamma_x + \Delta\gamma_y + \Delta\gamma_z$. The inset in Fig. 8 demonstrates that most of the energization is to be attributed to E_z , i.e. to the motional field of the striped wind.

The right column in Fig. 7 shows the trajectory of an electron that, after interacting with the IBS, remains in the post-shock region. There, it interacts with the plasmoids produced by shock-driven reconnection, that flow away from the shock towards the left boundary of the simulation. The electron energy increases during head-on collisions (e.g. at $\omega_p t \sim 6000$), while it decreases for tail-on collisions (e.g. at $\omega_p t \sim 3900$, marked by the green circle). This should ultimately lead to a net energy gain, since head-on collisions are more frequent than tail-on collisions. However, the stochastic nature of the process results in a slower energization rate, in a time-averaged sense, as compared to the dramatic events of pick-up acceleration enjoyed by upstream particles.

Fig. 9 provides a statistical assessment of which region (upstream or downstream) mostly contributes to the energy acquired beyond the

reconnection-driven stage. For each high-energy particle, we determine the ‘injection’ time t_{inj} when it interacted with the IBS (and so, got accelerated by shock-driven reconnection) as the time of largest $d \log \gamma/dt$. For each electron, we then compute $\Delta\gamma_{\text{tot}} = \gamma_{\text{end}} - \gamma_{\text{inj}}$, where γ_{end} is the particle Lorentz factor right before it exits the box (or at the end of the simulation, whatever happens earlier) and γ_{inj} is the particle Lorentz factor at the injection time t_{inj} . For each particle, we also compute the amount of post-injection energy acquired while upstream ($\Delta\gamma_{\text{up}}$, using $\mathcal{R} > 1$) and downstream ($\Delta\gamma_{\text{down}}$, using $\mathcal{R} < 1$), such that $\Delta\gamma_{\text{tot}} = \Delta\gamma_{\text{up}} + \Delta\gamma_{\text{down}}$. We remark that $\Delta\gamma_{\text{up}}$ and $\Delta\gamma_{\text{down}}$ may be negative. Fig. 9 shows the fractional amount $\Delta\gamma_{\text{up}}/\Delta\gamma_{\text{tot}}$ of post-injection energy acquired while upstream, as a function of the overall post-injection energy gain $\Delta\gamma_{\text{tot}}$. It shows that most particles acquire an additional $\Delta\gamma_{\text{tot}} \sim 200$ beyond the injection stage, and the energy gain occurs primarily in the downstream for $\Delta\gamma_{\text{tot}} \lesssim 1000$. In contrast, the highest energy particles, having $\Delta\gamma_{\text{tot}} \gtrsim 1000$, gain most of their post-injection energy via the pick-up process in the upstream.

3.5 Dependence on the pulsar wind magnetization

So far, we have considered a pulsar wind with magnetization $\sigma = 10$. Here, we investigate the dependence of our results – namely, the flow properties, the particle energy/synchrotron spectra, and the light curves – on the pulsar wind magnetization, for the specific case $\alpha = 0$. We consider a magnetization of $\sigma = 40$, paired with a stripe wavelength of $\lambda = 200c/\omega_p$ and a companion radius of $R_* = 140c/\omega_p$, so that the ratios $\lambda/r_{L,\text{hot}}$ and $R_*/r_{L,\text{hot}}$ are the same as in CS22. There, we used $\sigma = 10$, $\lambda = 100c/\omega_p$, and $R_* = 70c/\omega_p$.

In Fig. 10, we compare the flow dynamics of $\sigma = 10$ (left column) and $\sigma = 40$ (right column). Only minor differences are seen, with the primary one being the strength and location of the fast MHD shock. The case with higher magnetization produces a stronger MHD shock, resulting in more efficient disruption of current sheets in the upstream region. This is clearly seen in the upper right-hand side of the magnetic energy density plot (bottom panel of the second column). None the less, the main properties of the downstream region (plasmoid size distribution, pattern of flow velocity, magnetic energy structure) are nearly independent of magnetization.

The downstream energy spectra $dN/d\ln(\gamma - 1)$ of pulsar wind particles are presented in the top panel of Fig. 11, where $\sigma = 40$ is shown in red and $\sigma = 10$ in blue. We normalize the spectra such that their integral is the same. Given that R_* is twice larger for $\sigma = 40$ than for $\sigma = 10$, the post-shock area is four times greater in the higher magnetization case. For a fixed n_0 , we then expect the two integrals to differ by roughly a factor of four. To facilitate comparison, we also rescale the horizontal axis by $10/\sigma$, i.e. we align the two cases as regard to the typical Lorentz factor $\gamma_\sigma \simeq \gamma_0\sigma$ of post-shock particles. With this rescaling, the high-energy parts of the two spectra, $(\gamma - 1)(10/\sigma) \gtrsim 2$, nearly overlap. Some difference is seen at the low-energy end. In fact, the low-energy break $\gamma \sim \gamma_0$ of the post-shock particle spectrum should occur at a value of $(\gamma - 1)(10/\sigma)$ that is four times lower for $\sigma = 40$ than for $\sigma = 10$, in good agreement with the top panel of Fig. 11.

Similar conclusions hold for the angle-integrated synchrotron spectra in the bottom panel of Fig. 11. Since the synchrotron frequency ω is normalized to ω_c , the location of the spectral peak should be independent of σ , as indeed confirmed by Fig. 11. However, the extent of the power-law range $\nu F_\nu \propto \nu^{0.7}$ is different, with the $\sigma = 40$ case extending down to lower $\chi = \omega/\omega_c$ than the $\sigma = 10$ case. This mirrors the different extent in the power-law energy range (having $p \simeq 1.6$) of the particle spectrum discussed above. Particles

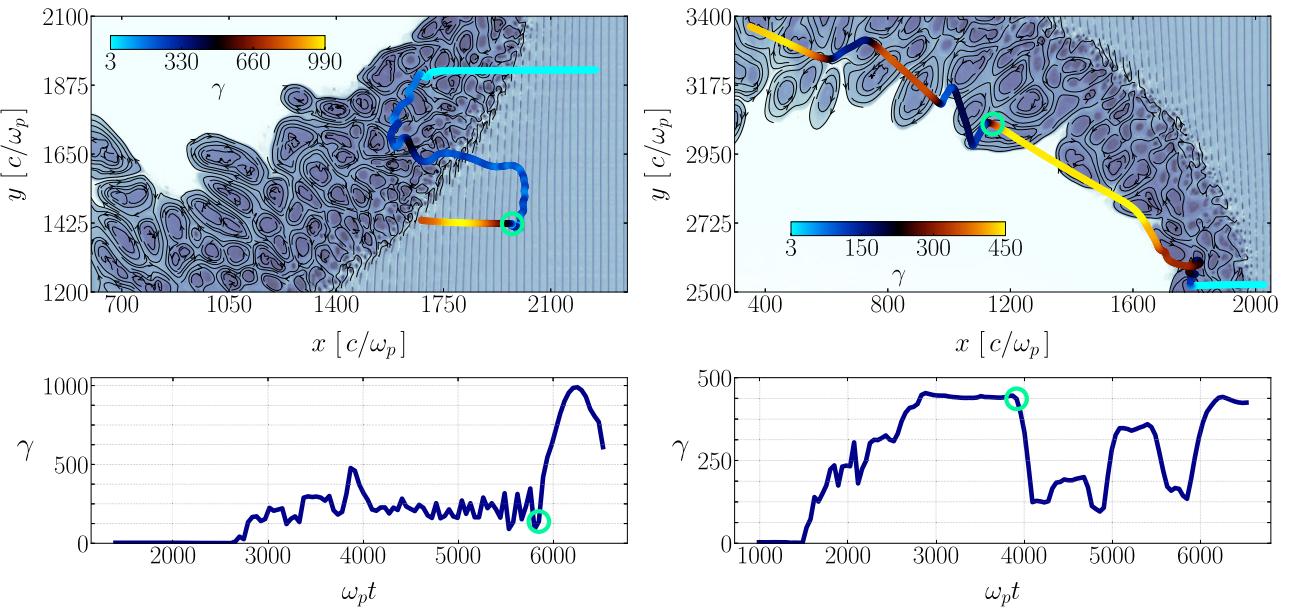


Figure 7. Two representative electron trajectories for $\alpha = 0$, showing: (left) Fermi-type acceleration in the upstream, where the green circle is highlighting the beginning of the main acceleration episode at $\omega_p t \sim 5800$; (right) stochastic acceleration in the downstream, where the green circle is highlighting a deceleration episode at $\omega_p t \sim 3900$. These are not the same electrons as in Fig. 6. We present the (x, y) particle trajectories in the top panels (the colour of the trajectory indicates the instantaneous Lorentz factor, see the colourbar), and their energy histories in the bottom panels. The underlying 2D map in the top panels shows the particle density at the same time as the green circle (so, at $\omega_p t \sim 5800$ in the left panel and at $\omega_p t \sim 3900$ in the right panel) with magnetic field lines overlaid in black. An animated version of this figure can be found online.

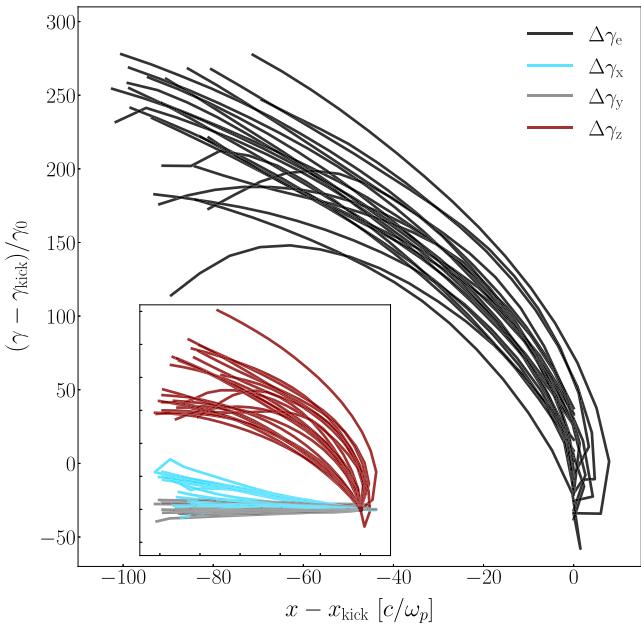


Figure 8. Trajectories in the $x - \gamma$ space for a sample of high-energy electrons energized via the pickup process. We identify the time t_{kick} when a given electron gets ‘picked up’ by the upstream flow, and we define its x -coordinate at that time as x_{kick} and its Lorentz factor as γ_{kick} . The main panel shows the energy gain $(\gamma - \gamma_{\text{kick}})/\gamma_0$ as a function of the spatial displacement $x - x_{\text{kick}}$. The inset shows the work done by E_x ($\Delta\gamma_x$ – cyan), E_y ($\Delta\gamma_y$ – grey), and E_z ($\Delta\gamma_z$ – red), see text for details. The inset has the same axis range as the main panel.

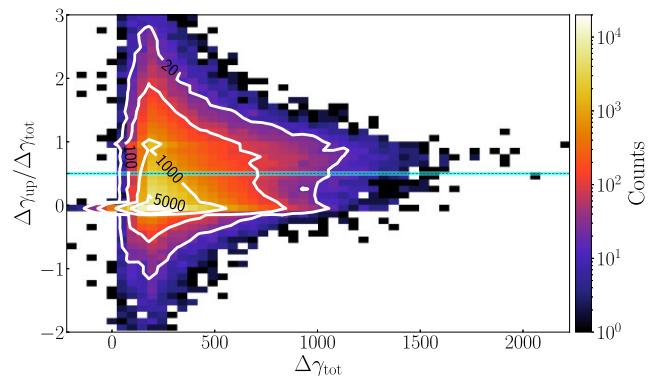


Figure 9. 2D histogram of the fractional amount $\Delta\gamma_{\text{up}}/\Delta\gamma_{\text{tot}}$ of post-injection energy acquired while upstream, as a function of the overall post-injection energy gain $\Delta\gamma_{\text{tot}}$ (i.e. following the interaction with the IBS), with white contours overlaid on top showing counts of 20, 100, 1000, and 5000. Each particle appears only once. The histogram refers to electrons, but the one of positrons is nearly identical. A black dashed line, highlighted in cyan, is placed at $\Delta\gamma_{\text{up}}/\Delta\gamma_{\text{tot}} = 0.5$ (where a particle obtains equal contributions in upstream and downstream).

with $\gamma \simeq \gamma_0$ (i.e. near the low-energy break of the particle spectrum) will emit at $\chi \sim \sigma^{-2}$. It follows that the break between $\nu F_{\nu} \propto \nu^{4/3}$ and $\nu F_{\nu} \propto \nu^{0.7}$ in the synchrotron spectrum should occur at a value of χ that is four times smaller for $\sigma = 40$ than for $\sigma = 10$. This is in good agreement with the synchrotron spectra in Fig. 11.

Local studies of plane-parallel shocks in striped pulsar winds have demonstrated that, at fixed $\lambda/r_{\text{L,hot}}$, the power-law range of

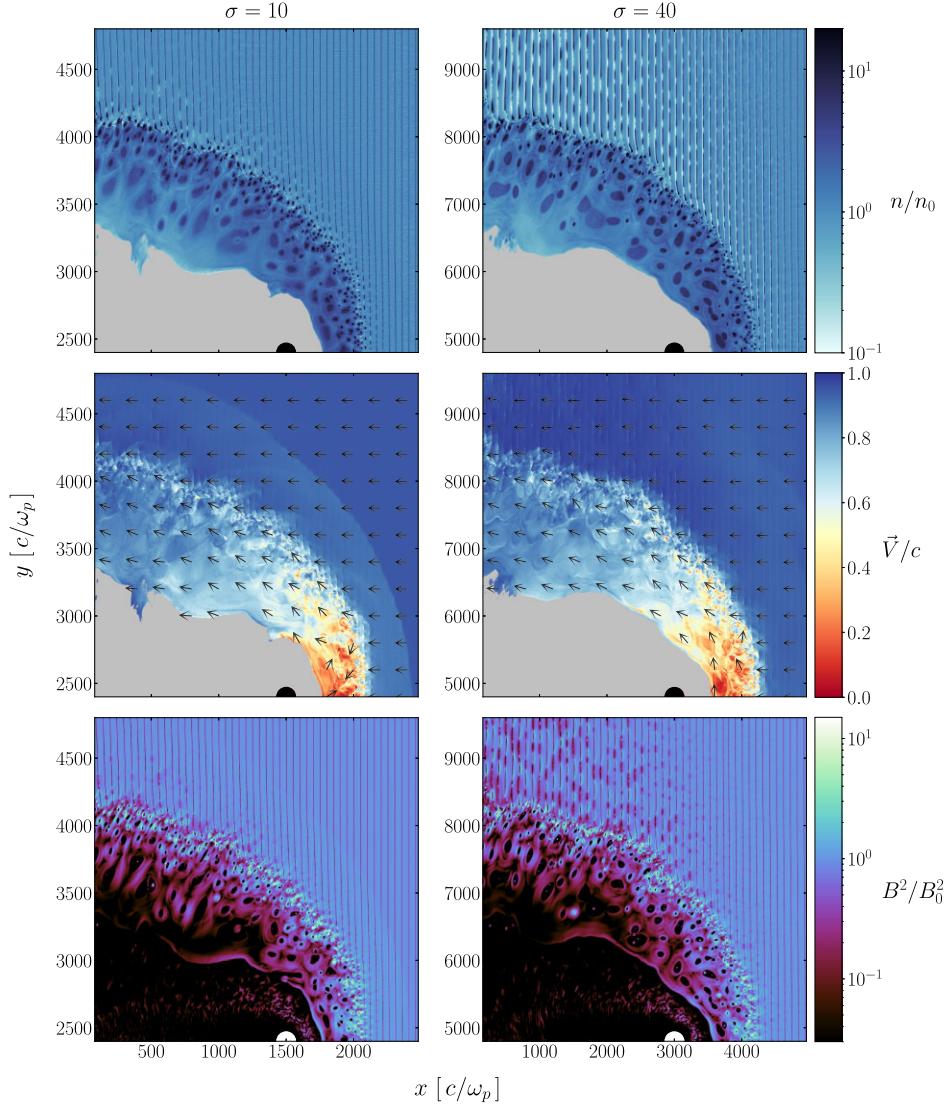


Figure 10. 2D plots comparing $\sigma = 10$ (left column) and $\sigma = 40$ (right column) for $\alpha = 0$. Both cases have the same $\lambda/r_{\text{L,hot}}$ and $R_*/r_{\text{L,hot}}$. We only show the top half of the simulation domain. *Top*: number density of pulsar wind particles in units of n_0 . *Middle*: flow velocity of pulsar wind in units of c , with arrows of unit length depicting flow direction. *Bottom*: magnetic energy density in units of the upstream value $B_0^2/8\pi$. The companion star is represented by either a black circle (top two rows) or a white circle (bottom row). In the top and middle rows, the grey region around and to the left of the companion star is populated by companion wind particles, which we exclude from our analysis.

the particle spectrum tends to be harder for higher σ , with $p \rightarrow 1$ in the $\sigma \gg 1$ limit (SS11). This is confirmed by our global simulations in Fig. 11. In the limit of very high magnetization, we expect a power-law particle spectrum with $p \simeq 1$ between $\gamma/\gamma_\sigma \sim \sigma^{-1}$ and $\gamma/\gamma_\sigma \sim 1$. Correspondingly, the synchrotron spectrum displays a power-law $\nu F_\nu \propto \nu^1$ extending between $\chi \sim \sigma^{-2}$ and $\chi \sim 1$.

We conclude our comparison of cases with different magnetization by showing the phase-resolved light curves measured at $\chi = 1$ (Fig. 12). Apart from an overall (arbitrary) normalization, the light curves are nearly identical. In particular, we confirm that for $\alpha = 0$ the light curve displays two peaks, just before and after the pulsar eclipse, regardless of σ .

4 SUMMARY AND DISCUSSION

We have performed global two-dimensional PIC simulations of the intrabinary shock in spider pulsars, assuming that the shock wraps

around the companion star. This work extends our previous paper (CS22) in several directions. In order to bring our simulations closer to realistic parameters, we have investigated the cases of larger companion stars and of a more magnetized pulsar wind. In both cases, we confirm that the results obtained in CS22 can be reliably applied to realistic spider pulsar systems. Our first-principles synchrotron spectra and light curves are in good agreement with X-ray observations: (1) The synchrotron spectrum is hard, $\nu F_\nu \propto \nu^{0.7-1}$, with a tendency for harder spectra at higher magnetizations (we expect $\nu F_\nu \propto \nu^1$ in the limit $\sigma \gg 1$); (2) when the pulsar spin axis is nearly aligned with the orbital angular momentum (i.e. $\alpha = 0$), the light curve displays two peaks, just before and after the pulsar eclipse (superior conjunction), separated in phase by ~ 0.8 rad; (3) the peak flux exceeds the one at inferior conjunction by a factor of ten, regardless of σ .

As compared to CS22, this work investigates in more detail the physics of particle acceleration following the stage governed by

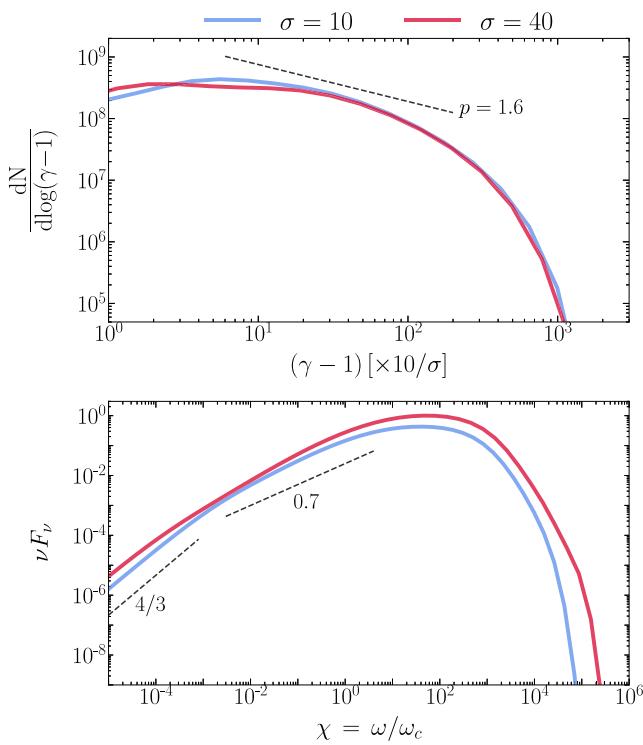


Figure 11. *Top:* downstream particle energy spectra for $\sigma = 10$ (blue) and $\sigma = 40$ (red) at $\omega_{\text{pt}} = 3263$ and $\omega_{\text{pt}} = 6526$, respectively (so, at the same time in units of $r_{\text{L,hot}}/c$). Spectra are shifted along the horizontal axis by $10/\sigma$ to facilitate their comparison, and they are normalized such that their integral is the same. *Bottom:* angle-integrated synchrotron spectra $v F_v$ (same colour coding as in the top panel), with arbitrary normalization.

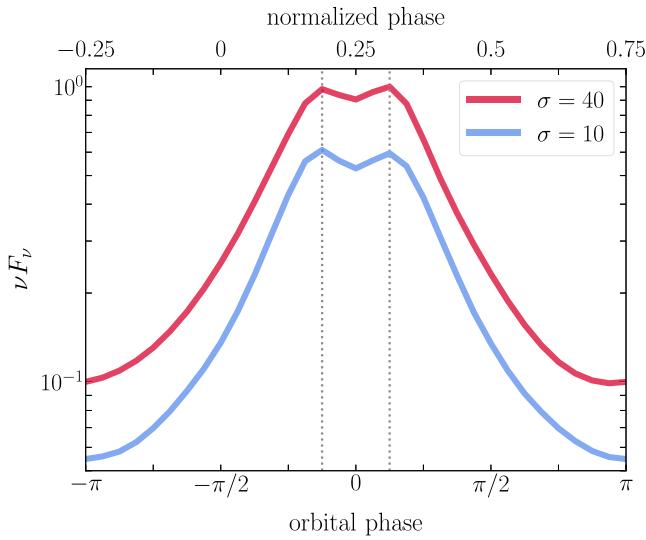


Figure 12. Phase-resolved synchrotron light curves for $\sigma = 10$ (blue) and $\sigma = 40$ (red) at $\omega_{\text{pt}} = 3263$ and $\omega_{\text{pt}} = 6526$, respectively. The light curves are shown for $\chi = 1$, where $\chi = \omega/\omega_c$. Bottom axis is orbital phase ϕ , where $\phi = 0$ corresponds to superior conjunction (when the pulsar is eclipsed), while $\phi = \pm\pi$ to inferior conjunction. The top axis is the normalized phase, with 0.25 at the pulsar superior conjunction, as typical for X-ray observations. The normalization is the same as for the synchrotron spectra in Fig. 11.

shock-driven reconnection. We find that the highest energy particles accelerated by reconnection can stream ahead of the shock and be further accelerated by the upstream motional electric field via the pick-up process. In the downstream, further energization is governed by stochastic interactions with the plasmoids generated by reconnection. Particle acceleration beyond the typical energy provided by shock-driven reconnection could have important implications for synchrotron spectra and light curves beyond the peak frequency.

As already discussed in CS22, realistic values of γ_σ are such that the resulting emission at $\sim \omega_c$ falls naturally in the X-ray band. Conservation of energy along the pulsar wind streamlines implies that $\gamma_0(1 + \sigma)\kappa = \omega_{\text{LC}}/2\Omega$, where $\omega_{\text{LC}} = eB_{\text{LC}}/mc$ is the cyclotron frequency at the light cylinder radius. We extrapolate the field from the pulsar surface to the light cylinder radius with a dipolar scaling, $B_{\text{LC}} \sim B_{\text{P}}(R_{\text{NS}}/R_{\text{LC}})^3$ (here, $R_{\text{NS}} \sim 10$ km is the neutron star radius). We then assume that the post-shock field is of the same order as the pre-shock field $B_0 \sim B_{\text{LC}}(R_{\text{LC}}/R_{\text{TS}})$. The characteristic synchrotron photon energy will be

$$\hbar\omega_c = \hbar(\gamma_0\sigma)^2 \frac{eB_0}{mc} \quad (8)$$

$$\simeq 0.2 \left(\frac{10^4}{\kappa} \right)^2 \left(\frac{10^{11} \text{ cm}}{R_{\text{TS}}} \right)^3 \left(\frac{\Omega}{10^3 \text{ s}^{-1}} \right)^6 \text{ keV}. \quad (9)$$

The peaks of our synchrotron spectra are located at $\hbar\omega_{\text{pk}} \sim 20 - 100 \hbar\omega_c$ (with marginal evidence for higher magnetizations peaking at higher frequencies, see Fig. 11). For realistic parameters, our conclusions on light curve shape and spectral hardness below the peak can be promptly applied to X-ray observations, especially if realistic multiplicities of millisecond pulsars are well below $\kappa \sim 10^4$.

Based on our findings, we propose a simple, order-of-magnitude estimate for the X-ray luminosity emitted by the IBS in spider pulsars. Let L_{sd} be the spin-down luminosity of the millisecond pulsar, and let us assume that the energy flux in the pulsar wind is roughly isotropic (in reality, more energy flows near the pulsar equator, e.g. Tchekhovskoy, Spitkovsky & Li 2013). The fraction of spin-down power dissipated at the IBS is $\sim (\Omega_{\text{IBS}}/4\pi)$, where Ω_{IBS} is the solid angle subtended by the shock, as seen from the pulsar. The fraction of dissipated energy that gets radiated away depends on the ratio $t_{\text{cool}}(\gamma_\sigma)/t_{\text{adv}}$ between the synchrotron cooling time of particles with $\gamma = \gamma_\sigma$ and the advection time $t_{\text{adv}} \sim R_{\text{curv}}/c$. Finally, the observed X-ray luminosity depends on how far below the peak frequency ω_{pk} is the X-ray observing frequency ω_X . Assuming a spectrum $v F_v \propto v^\Gamma$ with $\Gamma \simeq 0.7 - 1$ in the X-ray range, we predict

$$\frac{L_X}{L_{\text{sd}}} \sim \frac{\Omega_{\text{IBS}}}{4\pi} \left(\frac{\omega_X}{\omega_{\text{pk}}} \right)^\Gamma \min \left[\frac{t_{\text{adv}}}{t_{\text{cool}}(\gamma_\sigma)}, 1 \right] \quad (10)$$

(as long as $\omega_X \lesssim \omega_{\text{pk}}$), which should be compared with existing observations of spider pulsars.

We conclude with a few caveats. First, we have employed 2D simulations, and we defer to future work an assessment of 3D effects (e.g. for cases with $\alpha \neq 0$, in 2D field lines artificially accumulate ahead of the companion). A realistic 3D geometry is also known to bring qualitative differences in the physics of particle acceleration in reconnection (Zhang, Sironi & Giannios 2021; Zhang et al. 2023), as compared to 2D counterparts. Second, we have neglected the orbital motion of the system, which has been invoked to explain asymmetries in the light curve, with the two peaks having different heights (Kandel et al. 2021). Third, we have ignored radiative cooling losses in the particle equation of motion. Depending on parameters, the post-shock flow may be slow- or fast-cooling (Wadiasingh et al. 2017), so our results are applicable only to the slow-cooling cases.

Fourth, the thickness Δ of the pre-shock current sheets is chosen such that reconnection does not spontaneously start before the shock; in reality, some degree of magnetic field dissipation is expected to occur since the pulsar light cylinder (Cerutti & Philippov 2017; Cerutti, Philippov & Dubus 2020).

Finally, we have assumed for simplicity that the pulsar wind can be modelled as a sequence of plane-parallel stripes. This is appropriate if $d_{\text{IBS}} \gg R_{\text{curv}}$, whereas for realistic spider systems d_{IBS} is not much larger than the characteristic size R_{curv} of the shock. Future work should properly account for the curvature of the pulsar wind stripes. In a realistic 3D configuration, this also implies that one cannot assume a single value of α . In the case that the pulsar spin is aligned with the orbital angular momentum (i.e. the pulsar equator, where $\alpha = 0$, coincides with the orbital plane), the value of α at a distance d_{IBS} from the pulsar will depend on the height z above the orbital plane as

$$\alpha = \frac{z/d_{\text{IBS}}}{\tan \xi}. \quad (11)$$

It follows that significant deviations from the value of α appropriate for the orbital plane will occur unless $R_{\text{curv}}/d_{\text{IBS}} \ll \tan \xi$. Future work will need to account for the z -dependence of α .

ACKNOWLEDGEMENTS

We thank Elena Amato, Damiano Caprioli, Sasha Philippov, and Andrew Sullivan for useful discussions. We also thank Zorawar Wadiasingh for the careful and insightful review of our manuscript. LS acknowledges support from DE-SC0023015 and NSFAST 2307202. This research was facilitated by the Multimessenger Plasma Physics Center (MPPC), NSF grant PHY-2206609. This work was also supported by a grant from the Simons Foundation (MP-SCMPS-00001470) to LS. The authors acknowledge computing resources from Columbia University's Shared Research Computing Facility project and NASA Pleiades.

DATA AVAILABILITY

The simulated data underlying this paper will be shared upon reasonable request to the corresponding author(s).

REFERENCES

Arons J., Tavani M., 1993, *ApJ*, 403, 249

Arumugam P., Pavlov G. G., Garmire G. P., 2015, *ApJ*, 814, 90

Blandford R., Eichler D., 1987, *Phys. Rep.*, 154, 1

Bogdanov S., Grindlay J. E., van den Berg M., 2005, *ApJ*, 630, 1029

Bogdanov S., Archibald A. M., Hessels J. W. T., Kaspi V. M., Lorimer D., McLaughlin M. A., Ransom S. M., Stairs I. H., 2011, *ApJ*, 742, 97

Bogdanov S. et al., 2015, *ApJ*, 806, 148

Bogdanov S., Bahramian A., Heinke C. O., Freire P. C. C., Hessels J. W. T., Ransom S. M., Stairs I. H., 2021, *ApJ*, 912, 124

Bogovalov S. V., 1999, *A&A*, 349, 1017

Bogovalov S. V., Khangulyan D. V., Koldoba A. V., Ustyugova G. V., Aharonian F. A., 2008, *MNRAS*, 387, 63

Bogovalov S. V., Khangulyan D., Koldoba A. V., Ustyugova G. V., Aharonian F. A., 2012, *MNRAS*, 419, 3426

Bogovalov S. V., Khangulyan D., Koldoba A., Ustyugova G. V., Aharonian F., 2019, *MNRAS*, 490, 3601

Bosch-Ramon V., Barkov M. V., Khangulyan D., Perucho M., 2012, *A&A*, 544, A59

Bosch-Ramon V., Barkov M. V., Perucho M., 2015, *A&A*, 577, A89

Buneman O., 1993, Computer Space Plasma Physics. Terra Scientific, Tokyo, p. 67

Cerutti B., Philippov A. A., 2017, *A&A*, 607, A134

Cerutti B., Philippov A. A., Dubus G., 2020, *A&A*, 642, A204

Cheung C. C., Donato D., Gehrels N., Sokolovsky K. V., Giroletti M., 2012, *ApJ*, 756, 33

Clark C. J. et al., 2023, *Nat. Astron.*, 7, 451

Cortés J., Sironi L., 2022, *ApJ*, 933, 140

Fermi E., 1949, *Phys. Rev.*, 75, 1169

Fried B. D., 1959, *Phys. Fluids*, 2, 337

Goldreich P., Julian W. H., 1969, *ApJ*, 157, 869

Harding A. K., Gaisser T. K., 1990, *ApJ*, 358, 561

Harding A. K., Muslimov A. G., 2011, *ApJ*, 743, 181

Huang R. H. H., Kong A. K. H., Takata J., Hui C. Y., Lin L. C. C., Cheng K. S., 2012, *ApJ*, 760, 92

Huber D., Kissmann R., Reimer A., Reimer O., 2021, *A&A*, 646, A91

Iwamoto M., Amano T., Matsumoto Y., Matsukiyo S., Hoshino M., 2022, *ApJ*, 924, 108

Kandel D., Romani R. W., An H., 2019, *ApJ*, 879, 73

Kandel D., Romani R. W., An H., 2021, *ApJ*, 917, L13

Lamberts A., Fromang S., Dubus G., Teyssier R., 2013, *A&A*, 560, A79

Lu Y., Guo F., Kilian P., Li H., Huang C., Liang E., 2021, *ApJ*, 908, 147

Möbius E., Hovestadt D., Klecker B., Scholer M., Gloeckler G., Ipavich F. M., 1985, *Nature*, 318, 426

Pétri J., Lyubarsky Y., 2007, *A&A*, 473, 683

Phinney E. S., Evans C. R., Blandford R. D., Kulkarni S. R., 1988, *Nature*, 333, 832

Reville B., Kirk J. G., 2010, *ApJ*, 724, 1283

Roberts M. S. E., McLaughlin M. A., Gentile P., Aliu E., Hessels J. W. T., Ransom S. M., Ray P. S., 2014, *Astron. Nachr.*, 335, 313

Romani R. W., Sanchez N., 2016, *ApJ*, 828, 7

Roberts M. S. E., McLaughlin M. A., Gentile P. A., Ray P. S., Ransom S. M., Hessels J. W. T., 2015, preprint ([arXiv:1502.07208](https://arxiv.org/abs/1502.07208))

Romani R. W., Filippenko A. V., Cenko S. B., 2014, *ApJ*, 793, L20

Sanchez N., Romani R. W., 2017, *ApJ*, 845, 42

Sironi L., Spitkovsky A., 2011, *ApJ*, 741, 39

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. Am. Inst. Phys., New York, p. 345

Swihart S. J., Strader J., Chomiuk L., Aydi E., Sokolovsky K. V., Ray P. S., Kerr M., 2022, *ApJ*, 941, 199, 20

Tchekhovskoy A., Spitkovsky A., Li J. G., 2013, *MNRAS*, 435, L1

Timokhin A. N., Harding A. K., 2015, *ApJ*, 810, 144

van der Merwe C. J. T., Wadiasingh Z., Venter C., Harding A. K., Baring M. G., 2020, *ApJ*, 904, 91

Wadiasingh Z., Harding A. K., Venter C., Böttcher M., Baring M. G., 2017, *ApJ*, 839, 80

Wadiasingh Z., Venter C., Harding A. K., Böttcher M., Kilian P., 2018, *ApJ*, 869, 120

Weibel E. S., 1959, *Phys. Rev. Lett.*, 2, 83

Zhang H., Sironi L., Giannios D., 2021, *ApJ*, 922, 261

Zhang H., Sironi L., Giannios D., Petropoulou M., 2023, *ApJ*, 956, L36

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

IBS_accel.mp4

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: DEPENDENCE ON THE COMPANION SIZE

For realistic parameters of spider pulsars, the ratio of shock curvature radius R_{curv} to wavelength of the striped wind $\lambda = 2\pi c/\Omega$ (here, Ω is the pulsar spin frequency) is

$$\frac{R_{\text{curv}}}{\lambda} \sim 5 \times 10^1 \left(\frac{R_{\text{curv}}}{10^{10} \text{ cm}} \right) \left(\frac{\Omega}{10^3 \text{ s}^{-1}} \right). \quad (\text{A1})$$

For $\sigma = 10$, the simulations employed in this work have $R_{\text{curv}}/\lambda \sim 4$, as compared to $R_{\text{curv}}/\lambda \sim 2$ in CS22. In Fig. A1, we compare the

flow properties of the two cases, for a fixed $\alpha = 0$ ($R_{\text{curv}}/\lambda \sim 2$ on the left and $R_{\text{curv}}/\lambda \sim 4$ on the right). We find that the overall flow properties are nearly the same. In particular, in both cases the largest post-shock plasmoids grow to scales comparable to the shock curvature radius. Particle and synchrotron spectra are also nearly indistinguishable, see Fig. A2. The only marginal difference is a trend for higher cutoff energies at larger R_{curv}/λ for $\alpha = 0$ (both in particle energy spectra and in synchrotron spectra).

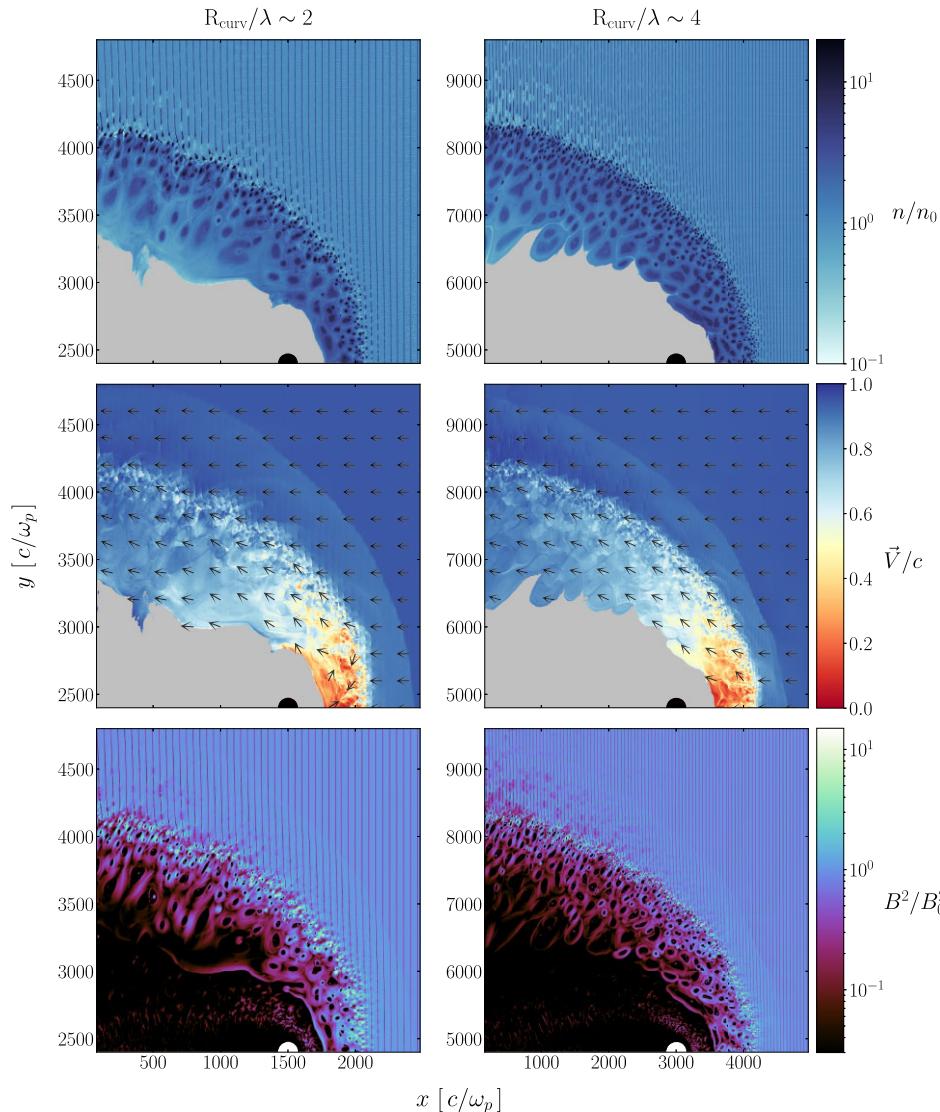


Figure A1. 2D plots comparing $R_{\text{curv}}/\lambda \sim 2$ (left column) and $R_{\text{curv}}/\lambda \sim 4$ (right column) for $\alpha = 0$. Both cases have the same $\lambda/(c/\omega_p)$. We only show the top half of the simulation domain. *Top:* number density of pulsar wind particles in units of n_0 . *Middle:* flow velocity of pulsar wind in units of c , with arrows of unit length depicting flow direction. *Bottom:* magnetic energy density in units of the upstream value $B_0^2/8\pi$. The companion star is represented by either a black circle (top two rows) or a white circle (bottom row). In the top and middle rows, the grey region around and to the left of the companion star is populated by companion wind particles, which we exclude from our analysis.

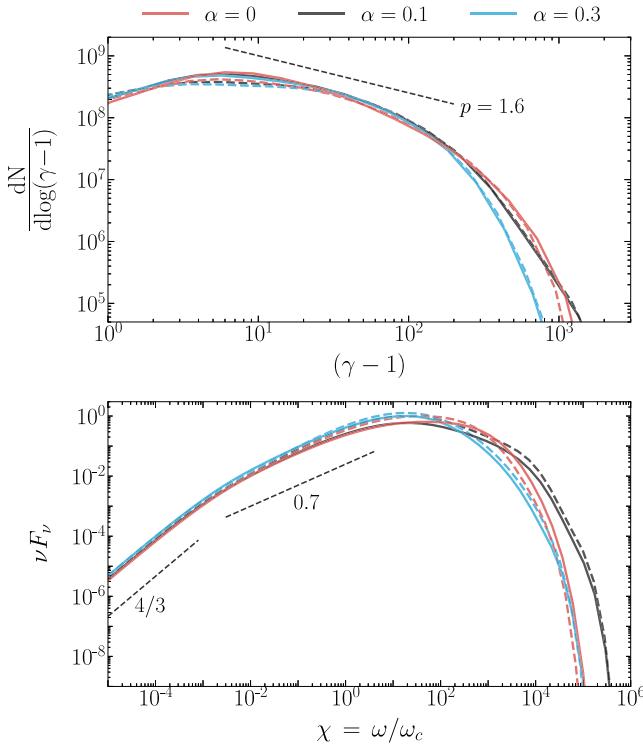


Figure A2. *Top:* downstream particle energy spectra for $\alpha = 0$ (red), 0.1 (black), and 0.3 (cyan). Spectra for $R_{\text{curv}}/\lambda \sim 2$ are at $\omega_p t = 3263$ and are shown with dashed lines; spectra for $R_{\text{curv}}/\lambda \sim 4$ are at $\omega_p t = 6526$ (so, at the same time in R_{curv}/c) and are shown with solid lines. *Bottom:* angle-integrated synchrotron spectra νF_ν (same colour coding and linestyle as in the top panel).

APPENDIX B: ASSESSMENT OF THE QUASI-STEADY STATE

In this section, we demonstrate that our simulations have reached a quasi-steady state. Fig. B1 shows the average particle Lorentz factor, $\langle \gamma \rangle$, for $\alpha = 0$ at $\omega_p t = 6524$. As seen in the main panel, the downstream is characterized by $\langle \gamma \rangle \gtrsim \gamma_\sigma$. We thus mask regions with $\langle \gamma \rangle \leq \gamma_\sigma$, and use the remaining area to track the boundaries of the IBS. We identify three length-scales of interest: Δx_c is the horizontal

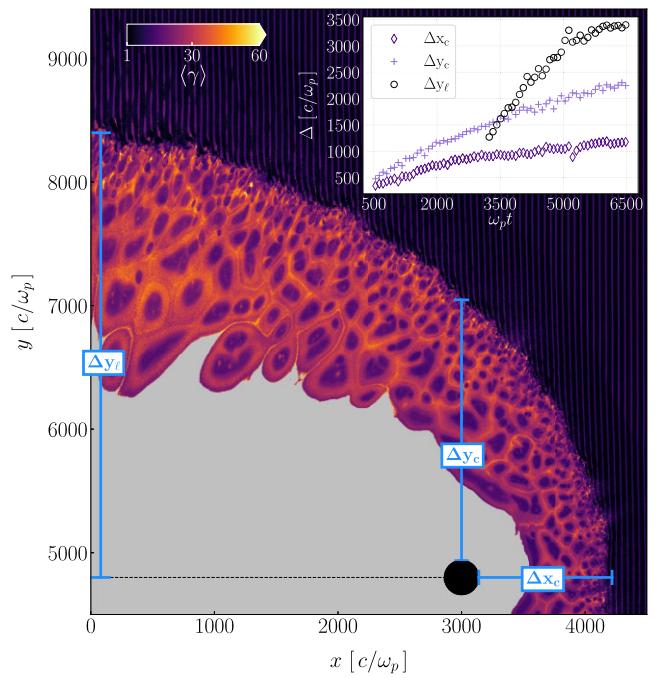


Figure B1. Assessment of the IBS quasi-steady state. The main panel shows a 2D plot of the average particle Lorentz factor, $\langle \gamma \rangle$, for $\alpha = 0$ at $\omega_p t = 6524$. Boundaries of the IBS are tracked via three length-scales: Δx_c , Δy_c , and Δy_ℓ , all labelled in blue in the figure. The inset panel depicts the measured quantities Δx_c (dark purple diamonds), Δy_c (light purple crosses), and Δy_ℓ (black circles) as a function of time.

distance between the apex of the shock and the companion's surface (at $y = y_c$), Δy_c is the vertical distance between the companion's surface (at $x = x_c$) and the IBS boundary, and Δy_ℓ is the vertical distance at the left-most edge of the simulation domain between $y = y_c$ and the IBS boundary. These length-scales are labelled in blue in the main panel. As depicted in the inset panel, the three length-scales approach a nearly constant value at late times ($\omega_p t \gtrsim 5000$), suggesting that we are reaching a quasi-steady state.

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