

Kinetic beaming in radiative relativistic magnetic reconnection: a mechanism for rapid gamma-ray flares in jets

J. M. Mehlhaff¹,¹★ G. R. Werner¹,¹ D. A. Uzdensky¹ and M. C. Begelman^{2,3}

¹Center for Integrated Plasma Studies, Physics Department, 390 UCB, University of Colorado, Boulder, CO 80309, USA

²JILA, University of Colorado and National Institute of Standards and Technology, 440 UCB, Boulder, CO 80309, USA

³Department of Astrophysical and Planetary Sciences, 391 UCB, Boulder, CO 80309, USA

Accepted 2020 August 2. Received 2020 July 25; in original form 2020 February 17

ABSTRACT

Rapid gamma-ray flares pose an astrophysical puzzle, requiring mechanisms both to accelerate energetic particles and to produce fast observed variability. These dual requirements may be satisfied by collisionless relativistic magnetic reconnection. On the one hand, relativistic reconnection can energize gamma-ray emitting electrons. On the other hand, as previous kinetic simulations have shown, the reconnection acceleration mechanism preferentially focuses high energy particles – and their emitted photons – into beams, which may create rapid blips in flux as they cross a telescope’s line of sight. Using a series of 2D pair-plasma particle-in-cell simulations, we explicitly demonstrate the critical role played by radiative (specifically inverse Compton) cooling in mediating the observable signatures of this ‘kinetic beaming’ effect. Only in our efficiently cooled simulations do we measure kinetic beaming beyond one light crossing time of the reconnection layer. We find a correlation between the cooling strength and the photon energy range across which persistent kinetic beaming occurs: stronger cooling coincides with a wider range of beamed photon energies. We also apply our results to rapid gamma-ray flares in flat-spectrum radio quasars, suggesting that a paradigm of radiatively efficient kinetic beaming constrains relevant emission models. In particular, beaming-produced variability may be more easily realized in two-zone (e.g. spine-sheath) set-ups, with Compton seed photons originating in the jet itself, rather than in one-zone external Compton scenarios.

Key words: acceleration of particles – magnetic reconnection – radiation mechanisms: general – relativistic processes – galaxies: jets – quasars: individual: PKS 1222+21.

1 INTRODUCTION

Powerful and rapid gamma-ray flares are observed in a variety of astrophysical objects, including pulsar wind nebulae (Abdo et al. 2011; Tavani et al. 2011; Buehler et al. 2012), blazars (Aharonian et al. 2007; Albert et al. 2007; Aleksić et al. 2011; Ackermann et al. 2016; Britto et al. 2016; MAGIC Collaboration 2019), and radio galaxies (Acciari et al. 2009; Abramowski et al. 2012; Aleksić et al. 2014). In outbursts from all three types of systems, the gamma-ray flux can change faster than the light crossing time of an important macroscopic scale: the size of the pulsar wind nebula itself, or, for radio galaxies and blazars, the radius of the central black hole’s event horizon. Such extreme flares challenge theories of particle acceleration for two reasons. First, an acceleration mechanism is needed that is capable of energizing particles to emit in the GeV and TeV bands. Secondly, the short observed time-scales require a compact (causally limited) emitting region and an underlying physical process that yields rapid changes in observed flux.

One mechanism that may provide the requisite particle energization is relativistic magnetic reconnection (Blackman & Field 1994; Lyutikov & Uzdensky 2003; Lyubarsky 2005), a plasma process where an abrupt relaxation in the magnetic field topology can power

relativistic single-particle and collective motion. Reconnection can be a promising alternative to shock acceleration, particularly in environments where an inferred high magnetization or relativistic shock front could suppress the latter (Sironi & Spitkovsky 2009, 2011; Sironi, Petropoulou & Giannios 2015). In contrast, it is precisely highly magnetized systems where reconnection may efficiently liberate stored magnetic energy, accelerating relativistic particles that subsequently produce the observed radiation. Partly in view of these considerations, magnetic reconnection has been thought to power emission in many astrophysical objects and phenomena, both flaring and quiescent, including pulsar magnetospheres (Lyubarskii 1996; Chen & Beloborodov 2014; Uzdensky & Spitkovsky 2014; Cerutti et al. 2015; Cerutti, Philippov & Spitkovsky 2016a; Cerutti, Mortier & Philippov 2016b; Cerutti 2018; Philippov & Spitkovsky 2018; Hakobyan, Philippov & Spitkovsky 2019; Schoeffer et al. 2019), pulsar winds (Cerutti & Philippov 2017), pulsar wind nebulae (Uzdensky, Cerutti & Begelman 2011; Clausen-Brown & Lyutikov 2012; Cerutti, Uzdensky & Begelman 2012a; Cerutti et al. 2012b, 2013, 2014a,b; Lyutikov et al. 2018), giant magnetar flares (Lyutikov 2003, 2006b; Uzdensky 2011; Schoeffer et al. 2019), gamma-ray bursts (Drenkhahn & Spruit 2002; Giannios 2006; Lyutikov 2006a; Giannios & Spruit 2007; Uzdensky 2011; McKinney & Uzdensky 2012), black hole accretion disc coronae (Di Matteo 1998; de Gouveia dal Pino & Lazarian 2005; Goodman & Uzdensky 2008; Uzdensky & Goodman 2008; de Gouveia Dal Pino et al.

★ E-mail: john.mehlhaff@colorado.edu

2010; Kadowaki, de Gouveia Dal Pino & Singh 2015; Singh, de Gouveia Dal Pino & Kadowaki 2015; Beloborodov 2017; Sironi & Beloborodov 2020; Werner, Philippov & Uzdensky 2019), blazars (Giannios, Uzdensky & Begelman 2009; Nalewajko et al. 2011, 2012; Giannios 2013; Sironi et al. 2015; Petropoulou, Giannios & Sironi 2016; Werner et al. 2018; Zhang et al. 2018; Christie et al. 2019, 2020; Giannios & Uzdensky 2019; Hakobyan et al. 2020; Ortuño-Macías & Nalewajko 2020), radio galaxies (Giannios, Uzdensky & Begelman 2010), and neutron star merger precursor emission (Crinquad, Cerutti & Dubus 2019; Most & Philippov 2020). Additionally, see the reviews by Hoshino & Lyubarsky (2012) and Uzdensky (2016).

Beyond simply providing a means to access magnetic energy, collisionless relativistic reconnection has been shown to generate non-thermal (power-law) distributions of energetic particles, which are required to produce the non-thermal radiation spectra seen in many of the systems mentioned above. Evidence for reconnection-powered non-thermal particle acceleration stems from a large number of particle-in-cell (PIC) simulation studies of both pair plasmas (Zenitani & Hoshino 2001, 2007, 2008; Jaroschek et al. 2004; Bessho & Bhattacharjee 2007, 2012; Lyubarsky & Liverts 2008; Cerutti et al. 2012b, 2013, 2014a,b; Guo et al. 2014, 2015, 2019; Sironi & Spitkovsky 2014; Sironi, Giannios & Petropoulou 2016; Werner et al. 2016; Werner & Uzdensky 2017; Petropoulou & Sironi 2018; Schoeffler et al. 2019; Werner et al. 2019; Sironi & Beloborodov 2020; Hakobyan et al. 2020; Ortuño-Macías & Nalewajko 2020) and electron-ion plasmas (Melzani et al. 2014a,b; Guo et al. 2016; Ball, Sironi & Özel 2018; Werner et al. 2018).

In this paper, we are specifically interested not in quiescent emission, but in abrupt high- and very high energy flares, where the words ‘high’ and ‘very high’ are used in their technical senses to mean detected photons with energies in the ranges 0.1–10 GeV and $\gtrsim 0.1$ TeV, respectively (e.g. H. E. S. S. Collaboration 2010; Aleksić et al. 2011; Madejski & Sikora 2016). When judging the feasibility of magnetic reconnection to power these extreme events, not only are the preceding energetic and spectral considerations important, but so too is the question of whether this plasma process can facilitate variability consistent with the dramatically short observed time-scales.

Aspects of this question have been addressed in connection with the Crab Nebula synchrotron flares. Cerutti et al. (2012b) showed that, near X-points (points in space where the magnetic field reconnects), particles are both accelerated and collimated, with the particles receiving more energy preferentially focused into tighter beams. This ‘kinetic beaming’ effect naturally leads to rapid changes in the light an observer sees as beams of high energy particles – and their relativistically beamed synchrotron or inverse Compton (IC) emission – sweep across the line of sight. Moreover, kinetic beaming predicts faster variability at higher energies, and is therefore observationally distinguishable from Doppler beaming (e.g. Rees 1966), where relativistic bulk motion focuses emission achronically. For the most energetic radiation, the light curve can change on time-scales as short as one-tenth the light crossing time of the reconnection layer (Cerutti et al. 2013, 2014a,b).

Blazars, though very different from pulsar wind nebulae, also display evidence for kinetic beaming during rapid gamma-ray IC flares. For a 2010 outburst from PKS 1222+21 (Aleksić et al. 2011; Tanaka et al. 2011), kinetic beaming relaxes the energy density required to feed the emission region to a level accessible in the blazar’s jet (Nalewajko et al. 2012). On the observational side, there have been at least two TeV events with approximately symmetric rise and decay times (consistent with a sweeping beam of light) and for which the

TeV variability increases with photon energy (Albert et al. 2007; H. E. S. S. Collaboration 2010). Yet another line of evidence, from numerical simulations, suggests that kinetic beaming at reconnecting current sheets may naturally occur within the turbulent environment of a magnetized blazar jet (Zhdankin et al. 2020).

Although compelling as an explanation for a diverse set of extreme flares, the kinetic beaming paradigm should perhaps be invoked with some caution. In particular, Kagan, Nakar & Piran (2016) suggested that strong radiative cooling may be an important factor. They argued that weakly cooled particles will radiate most of their energy well after their momentum distribution – initially collimated during X-point acceleration – has been isotropized by the magnetic field external to the acceleration site. However, when strongly cooled, particles will dump their reconnection-acquired energy into energetic photons before dispersing. As one expects from this logic, the large simulations of Sironi & Spitkovsky (2014) and Sironi et al. (2016), which, importantly, did not incorporate radiative losses, produced little to no anisotropy in the momentum distributions of high energy particles confined to the largest magnetic islands. Additionally, Yuan et al. (2016) observed beaming to cease at late times in radiatively inefficient simulations. These studies do not cast doubt on kinetic beaming as an explanation for Crab Nebula flares – for which this concept was originally proposed – because in that case the gamma-ray emitting particles are subject to strong synchrotron losses. However, when considering other flaring systems, it seems quite likely that the observable signatures of kinetic beaming may be limited if the emitting particles do not radiate efficiently.

To summarize, kinetic beaming appears to be a generic by-product of the reconnection acceleration mechanism (Uzdensky et al. 2011; Cerutti et al. 2012a,b, 2013, 2014a,b), but it may not always be observationally relevant (Sironi & Spitkovsky 2014; Kagan et al. 2016; Sironi et al. 2016; Yuan et al. 2016). The particles emitting, either via synchrotron or IC processes, at the energies of interest must do so efficiently. Otherwise, the observable signatures of kinetic beaming (i.e. rapid light-curve variability) may become washed out by particle isotropization. This simple picture – despite rapidly progressing numerical work on radiative relativistic reconnection (e.g. Jaroschek & Hoshino 2009; Cerutti et al. 2013, 2014a,b; Yuan et al. 2016; Nalewajko, Yuan & Chruślińska 2018; Hakobyan et al. 2019; Schoeffler et al. 2019; Werner et al. 2019; Sironi & Beloborodov 2020; Ortuño-Macías & Nalewajko 2020) – has yet to be made rigorous by a systematic study of radiatively cooled kinetic beaming. The need for such a study is underscored by its astrophysical implications: it would reveal whether and under what conditions a kinetic beaming scenario may viably explain fast gamma-ray flares observed from many types of sources.

Therefore, in this work, we perform the first systematic investigation into the impact of radiative cooling on observable kinetic beaming. By analysing a series of 2D PIC simulations of relativistic pair-plasma reconnection with varying IC cooling strength, we answer the following related questions:

- (i) For weak radiative cooling, does observable kinetic beaming disappear as the reconnection layer evolves? That is, does particle anisotropy vanish as reported by Sironi & Spitkovsky (2014), Sironi et al. (2016), and Yuan et al. (2016), just as one would expect from Kagan et al. (2016)?
- (ii) If so, does strong cooling restore observable kinetic beaming?

We note that our specialization to IC cooling is mostly for definiteness. We expect our results to carry over, in a qualitative sense, to systems where other emission mechanisms (most notably synchrotron) dominate the radiative output.

Our findings indicate affirmative answers to both questions (i) and (ii). In our simulations where radiative cooling is extremely weak or absent, no significant anisotropy is retained by the distribution of particles beyond a single light crossing time of the reconnection layer. However, as cooling becomes more efficient, a persistent beaming effect emerges across an increasing range of particle (and photon) energies. When persistent, kinetic beaming occurs only among the highest energy particles – those roughly within a decade of their radiatively imposed maximum possible energy.

This last result translates to a powerful constraint on rapid astrophysical gamma-ray flares, adding a radiative efficiency requirement to models attributing the observed variability to reconnection-driven kinetic beaming. As an illustration of the potential applicability of this result, after presenting our numerical findings, we specialize to the case of rapid flares in flat-spectrum radio quasars (FSRQs). Via simple analytical estimates, we show that invoking kinetic beaming constrains possible models for the PKS 1222+21 TeV flare (Aleksić et al. 2011). We find that the most appropriate radiative scenario may be inherently two-zone (for example, spine-sheath; Ghisellini, Tavecchio & Chiaberge 2005; Sikora, Rutkowski & Begelman 2016), where, as opposed to more traditional one-zone external Compton models, the photons seeding TeV IC emission come from inside the jet.

In Section 2, we describe our simulations, including our self-consistent incorporation of IC cooling and how it limits particle acceleration. In Section 3, we develop a quantitative language to describe beaming as manifested in distributions of particles and radiation. Section 4 applies this language to our simulations to answer questions (i) and (ii). In Section 5, we demonstrate the utility of our findings by analysing flaring FSRQs. We conclude in Section 6.

2 SIMULATIONS

2.1 Set-up

We present relativistic pair-plasma simulations run using the radiative electromagnetic PIC code ZELTRON (Cerutti et al. 2013). The simulation domain is a 2D box of size $L_x \times L_y$ with $L_y = 2L_x = 2L$ and periodic boundary conditions enforced in both dimensions. Spatial dependence is limited to the (x, y) dimensions, but vectorial quantities, including velocities and field components, are fully three-dimensional.

We initialize the simulations with standard double Harris current sheets (Kirk & Skjæraasen 2003) of half-thickness δ carrying anti-aligned currents in the $\pm z$ -directions and centred on the planes $y_1 = L_y/4$ and $y_2 = 3L_y/4$. This double-sheet configuration is chosen because it is consistent with the periodic box boundaries. Namely, the currents establish an in-plane magnetic field $B_x(y) = \pm B_0 \tanh[(y - y_{1,2})/\delta]$ that reverses twice – once at each sheet – and is therefore periodic in y . Additionally, we add a uniform initial guide field $\mathbf{B}_g = (B_0/4)\hat{\mathbf{z}}$. This serves chiefly as a numerical device to support magnetic islands against radiative cooling-induced contraction to the point where the Debye length becomes unresolved. The value $B_g = B_0/4$ does not substantially alter non-thermal particle acceleration (Werner & Uzdensky 2017; Werner et al. 2019).

The current in the Harris layers is carried by a drifting plasma component. In each layer, drifting electrons and positrons begin the simulation counterstreaming with bulk velocities $\pm c\beta_d \hat{\mathbf{z}}$ and a combined drifting lab-frame number density $n_d(y) = n_{d,0} \text{sech}^2[(y - y_{1,2})/\delta]$. The resultant current profile $J_z(y) = \pm c\beta_d e n_d(y)$ is precisely that necessary to generate the field $B_x(y)$ as dictated by Ampère’s Law. In addition to the current-governing mean velocity β_d , we

initialize the counterstreaming species with relativistic comoving temperature $\theta_d = T_d/m_e c^2 = 1050 \gg 1$ to support the current layers against the upstream magnetic pressure.

With the Harris equilibrium satisfied by the drifting particles, each simulation also contains an initially stationary uniform background plasma of combined density n_b and relativistically hot temperature $\theta_b = T_b/m_e c^2 = 25 \gg 1$ that provides the inflow material for the reconnection layer. Two important dimensionless quantities associated with the background plasma are the cold and hot upstream magnetizations, $\sigma \equiv B_0^2/4\pi n_b m_e c^2$ and $\sigma_h \equiv B_0^2/4\pi w$, respectively. Here, the enthalpy density w is given in the $\theta_b \gg 1$ limit by $w = 4n_b T_b = 4\theta_b n_b m_e c^2$, and, as a result, the hot magnetization becomes one-half inverse plasma-beta: $\sigma_h = B_0^2/16\pi\theta_b n_b m_e c^2 = 1/(2\beta_{\text{plasma}})$. The two σ -parameters have the following physical interpretations. The cold magnetization characterizes the available magnetic energy per upstream particle; for $\sigma \gg 1$, individual particles may acquire energy far in excess of their rest mass, and we set $\sigma = 10^4$. The hot magnetization, on the other hand, decides whether the energy inflow to the reconnection layer is dominated by the magnetic field or by the particles. Taking $\sigma_h = 10^2$, we operate in the magnetically dominant $\sigma_h \gg 1$ limit, which also implies a relativistic Alfvén speed and places us in the relativistic regime of reconnection.

To minimize system-size effects, the simulation box must considerably exceed the largest kinetic scale in the problem: the typical Larmor radius of energetic particles (see Werner et al. 2016). Because the magnetic energy per upstream particle is $\sigma m_e c^2/2$, an average particle accelerated through the reconnection layer emerges with Larmor radius of order $\sigma \rho_0$ where $\rho_0 = m_e c^2/eB_0$ is a nominal Larmor radius. We conduct simulations in the large system regime $L_x \gg 40\sigma \rho_0$ identified by Werner et al. (2016) and set $L_x/\sigma \rho_0 = 320$. To confirm the insensitivity of our results to our simulation box size, we also run a series of simulations at different L_x .

At the small-scale end, our cell size is $\Delta x = \Delta y = \sigma \rho_0/24$, which is just smaller than the Debye length: $\lambda_D = 1.2\Delta x$. With the cell size set, we employ a corresponding time-step satisfying the Courant–Friedrichs–Lewy condition $c\Delta t = 0.7\Delta x < \Delta x/\sqrt{2}$. Our initial number of simulation particles per grid cell is 80.

We incorporate radiative cooling into our simulations via inverse Compton scattering of a background (‘seed’) radiation field that is static, homogeneous, and isotropic. The photons comprising this field are not tracked simulation entities, but, in the Thomson limit, give rise to a continuous radiative drag force $\mathbf{f}_{\text{IC}}$ that enters self-consistently into the PIC particle push (Tamburini et al. 2010). For a particle of 4-velocity $(c\gamma, c\gamma\boldsymbol{\beta})$, the expression for the drag force is $\mathbf{f}_{\text{IC}} = -(4/3)\sigma_T U_{\text{ph}} \gamma^2 \boldsymbol{\beta}$, where $\sigma_T = 8\pi e^4/3m_e^2 c^4$ is the Thomson cross-section (cf. Blumenthal & Gould 1970; Rybicki & Lightman 1979; Pozdnyakov, Sobol & Syunyaev 1983; Uzdensky 2016; Werner et al. 2019; Sironi & Beloborodov 2020). From this, one sees that particle cooling depends only on the total energy density of background photons $U_{\text{ph}} \equiv \int d\epsilon u(\epsilon)$, not on their spectral distribution. We therefore adopt, without loss of generality, a simple monochromatic spectral energy density $u(\epsilon) = U_{\text{ph}} \delta(\epsilon - \epsilon_{\text{ph}})$. The Thomson limit is satisfied if the photon energies encountered by a particle of Lorentz factor γ in its rest frame are small compared with its mass:

$$\frac{\gamma \epsilon_{\text{ph}}}{m_e c^2} \ll 1. \quad (1)$$

This condition allows the recoil in any single scattering event to be neglected and justifies our continuous treatment of radiative losses (Blumenthal & Gould 1970). When equation (1) is not satisfied, IC cooling transitions to the discrete Klein–Nishina regime, with particles delivering an order-unity fraction of their energies to single photons.

Table 1. Simulation parameters used in this study. Note that we do not scan across all combinations of L_x and γ_{rad} . Instead, we conduct a series of simulations exploring all γ_{rad} values in the table at a fixed system size $L_x = 320\sigma\rho_0$ and a second series across all L_x values at a fixed radiation-reaction strength $\gamma_{\text{rad}} = 4\sigma$.

Parameter	Symbol (=definition)	Value
Upstream magnetic field	B_0	
Nominal gyroradius	$\rho_0 = m_e c^2 / e B_0$	
Radiation-limited Lorentz factor	γ_{rad}	$(1/2, 1, 2, 4, 6, 8, 16, 32, 64, \infty) \times \sigma$
System size	$L_x = L$	$(80, 120, 160, 240, 320) \times \sigma\rho_0$
‘Cold’ magnetization	$\sigma = B_0^2 / 4\pi n_b m_e c^2$	10^4
‘Hot’ magnetization	$\sigma_h = B_0^2 / 16\pi n_b T_b$	10^2
Background temp.	$\theta_b = T_b / m_e c^2$	25
Guide field	B_z	$B_0/4$
Peak drift over background density	$\eta = n_d / n_b$	5
Harris layer drift velocity	$\beta_d c$	$0.3c$
Harris layer (comoving) temp.	$\theta_d = T_d / m_e c^2$	1050
Harris layer half-thickness	$\delta = \sigma\rho_0 / \beta_d \eta$	$0.67\sigma\rho_0$
Cell size	$\Delta x, \Delta y$	$\sigma\rho_0/24$
Time step	Δt	$0.7\Delta x/c$
Macroparticles per cell		80

A convenient dimensionless parameter that quantifies the IC cooling strength is γ_{rad} , the Lorentz factor of a particle whose acceleration force due to the reconnection electric field matches its radiation-reaction force (cf. Uzdensky 2016; Werner et al. 2019; Sironi & Beloborodov 2020)

$$eE_{\text{rec}} = \frac{4}{3}\sigma_T U_{\text{ph}} \gamma_{\text{rad}}^2. \quad (2)$$

(Here we have assumed $\beta \simeq 1$.) If one takes the reconnection electric field as $E_{\text{rec}} = \beta_{\text{rec}} B_0$ with a standard value of the relativistic reconnection rate $\beta_{\text{rec}} = 0.1$, then γ_{rad} can be defined (again, see Uzdensky 2016; Werner et al. 2019) as

$$\gamma_{\text{rad}} \equiv \sqrt{3(0.1)eB_0/(4\sigma_T U_{\text{ph}})}. \quad (3)$$

By construction, γ_{rad} is an upper bound on the particle energy distribution, because particles of higher energy would radiate more power than is delivered them via the reconnection electric field. This suggests the quantity $\gamma_{\text{rad}}/\sigma$ as a quantitative measure of the IC cooling strength. For $\gamma_{\text{rad}}/\sigma \gg 1$, the particle energy distribution may develop a hard power-law tail as in the non-radiative regime (e.g. Guo et al. 2014; Sironi & Spitkovsky 2014; Werner et al. 2016; Petropoulou & Sironi 2018; Werner et al. 2019), with non-thermal particles having typical Lorentz factors of order σ and subject to relatively weak cooling. However, in the case $\gamma_{\text{rad}}/\sigma \lesssim 1$, particles with energy $\sigma m_e c^2$ are strongly cooled and near their upper limit, controlled now by radiative cooling rather than by the available magnetic energy.

In this work, our main results come from a series of simulations scanning across $\gamma_{\text{rad}}/\sigma \in [1/2, 1, 2, 4, 6, 8, 16, 32, 64, \infty]$ at the fiducial box size $L_x/\sigma\rho_0 = 320$. Note that *higher* $\gamma_{\text{rad}}/\sigma$ corresponds to *weaker* IC cooling. In particular, the case $\gamma_{\text{rad}}/\sigma = \infty$ implies $U_{\text{ph}} = 0$: no cooling. In addition, we run a few simulations with a uniform radiative efficiency $\gamma_{\text{rad}}/\sigma = 4$ but differing $L_x/\sigma\rho_0 \in [80, 120, 160, 240, 320]$ as a first step toward characterizing the system-size dependence of our findings. Table 1 summarizes these values and those of the other parameters discussed so far.

Our wide-ranging scan in γ_{rad} is limited on the strong-cooling end by radiative losses in the upstream region, which can cause the background plasma feeding the reconnection layer to change in time. In order to avoid this effect, one should require that the upstream IC cooling time t_{cool} exceed the duration of our simulations $3L_x/c$. Using the average Lorentz factor $\langle\gamma\rangle = 3\theta_b$ of the background plasma, the

ratio $t_{\text{cool}}/(3L_x/c)$ can be written as

$$\begin{aligned} \frac{t_{\text{cool}}}{3L_x/c} &= \frac{\langle\gamma\rangle m_e c^2 / P_{\text{IC}}(\langle\gamma\rangle)}{3L_x/c} = \frac{10}{3} \frac{\sigma\rho_0}{L_x} \frac{\gamma_{\text{rad}}}{\langle\gamma\rangle} \frac{\gamma_{\text{rad}}}{\sigma} \\ &= \frac{40\sigma_h}{9} \frac{\sigma\rho_0}{L_x} \left(\frac{\gamma_{\text{rad}}}{\sigma}\right)^2 \\ &\simeq 1.4 \left(\frac{\sigma_h}{100}\right) \left(\frac{L_x/\sigma\rho_0}{320}\right)^{-1} \left(\frac{\gamma_{\text{rad}}}{\sigma}\right)^2. \end{aligned} \quad (4)$$

Here, we used $P_{\text{IC}}(\gamma) = |c\boldsymbol{\beta} \cdot \mathbf{f}_{\text{IC}}| = (4/3)c\sigma_T U_{\text{ph}} \gamma^2 \beta^2$ and the relativistic limits $\beta = 1$ and $1 \ll \theta_b = \sigma/4\sigma_h$. Let us go one step farther, employing equation (4) along with $d\gamma/dt = -P_{\text{IC}}(\gamma)/m_e c^2$ to estimate the amount by which the upstream plasma cools during a simulation. The temperature $\theta_{b,f}$ reached at $t = 3L_x/c$ may be as low as

$$\frac{\theta_{b,f}}{\theta_b} \simeq \frac{1}{1 + (3L_x/c)/t_{\text{cool}}}. \quad (5)$$

Evidently, our simulation with $\gamma_{\text{rad}} = \sigma/2$ is problematic, with the upstream plasma cooling in time $t_{\text{cool}} \sim 1L_x/c$ and potentially falling to $\simeq 26$ per cent of its initial temperature by the end of the run. This simulation also exhibits the worst energy conservation, with the energy error peaking at about 3.6 per cent (all our other simulations have per cent level or better error). Thus, the results of this most strongly radiative run should not be taken as definitive on their own. Fortunately, the conclusions we draw from our γ_{rad} scan (Section 4.3, Figs 19–21), do not depend on whether we include or exclude this simulation from our quantitative analysis. We have therefore chosen to include it as a tentative endpoint on the data generated by the rest of our γ_{rad} -varying (L_x -constant) simulation series. We also note that our benchmark radiative case $\gamma_{\text{rad}} = \sigma$ is not completely free of the upstream cooling issue. However, according to equation (5), the background temperature may decrease by less than a factor of 2 by the end of that simulation. Because such a discrepancy is within the error bars on our main kinetic beaming quantities measured in Section 4.3, we view our $\gamma_{\text{rad}} = \sigma$ simulation as marginally acceptable.

Finally, we would like to point out that upstream radiative losses only increase our effective σ_h . Because σ_h is quite large to begin with, raising it by order-unity factors preserves (indeed, enhances) the asymptotically large- σ_h limit. Hence, we do not expect the acceleration and beaming of high energy particles to be significantly




Figure 1. Total (drift plus background) electron number density pictured in the lower reconnection layer at key moments during our $\gamma_{\text{rad}} = \sigma$ simulation. The full simulation width ($L = L_x = 320\sigma\rho_0$) and a restricted height range ($L_y/4 \pm 80\sigma\rho_0$) are displayed. See text for a description of the various phases of the time evolution.

impacted, even in our most strongly cooled ($\gamma_{\text{rad}}/\sigma = 1/2, 1$) simulations.

2.2 Evolution of the reconnection layer

Having described the set-up of our simulations, let us now move on to how they evolve in time. In every run, we trigger magnetic reconnection with a small (1 percent) perturbation to the initial magnetic field. The current sheet then tears into a number of magnetic islands or ‘plasmoids’ which begin to merge with one another (Fig. 1, $t = 0.4L/c = 0.4L/c$). Initially, the plasmoids are all about the same size, but eventually – in Fig. 1 at about $t = 1.2L/c$ – a single largest plasmoid begins to dominate the reconnection layer. This primary plasmoid proceeds to consume the others that have also accumulated to considerable but smaller sizes, culminating in a

spectacular merger between the largest and next-to-largest plasmoids (Fig. 1, $t = 2.0L/c$). After this most dramatic merger, additional small plasmoids are continually born from the main X-point and venture across the box to be consumed by the large primary plasmoid (Fig. 1, $t = 2.6L/c$).

We observe significant non-thermal particle acceleration during our simulations. At late times, this energization is bursty: merging plasmoids sporadically punctuate ongoing reconnection from the main X-point with short intense episodes of particle acceleration. In our radiatively efficient runs, following these episodes, the high energy particles rapidly cool, leading to a steepening of their non-thermal energy distribution. This effect was analysed by Werner et al. (2019) (see also Sironi & Beloborodov 2020), and we illustrate it in Fig. 2, which presents time-dependent particle energy distributions from the lower reconnection layer in a subset of our simulations.




Figure 2. Lower layer electron energy distribution as a function of time between L/c and $3L/c$ for a subset of the simulations presented in this work. Coloured solid lines denote the distribution at different times. Black dashed lines represent the time-averaged (over $1 \leq ct/L \leq 3$) distribution. Vertical black dot-dashed lines label γ_{rad} . Red bars indicate power-law segments in the time-averaged particle distributions. The qualitative features are consistent with those reported in Werner et al. (2019). Namely, the simulations with strongest cooling ($\gamma_{\text{rad}} = 1, 4\sigma$) exhibit only a steep power law $f(\gamma) \propto \gamma^{-p}$ with variable index $p \gtrsim 3$; the simulation with intermediate cooling ($\gamma_{\text{rad}} = 8\sigma$) exhibits both soft variable ($p \gtrsim 3$) and hard steady ($p \lesssim 2$) power-law segments; and the weakly cooled simulation ($\gamma_{\text{rad}} = 16\sigma$) only contains a hard steady ($p \lesssim 2$) power law.

In the limit of weak cooling (e.g. $\gamma_{\text{rad}}/\sigma = 16$ in Fig. 2), the distribution develops a shallow power law. However, due to long periods of continuous IC losses interrupted by bursts of plasmoid merger-initiated magnetic reconnection, the particle distributions for the simulations with stronger cooling (e.g. those with $\gamma_{\text{rad}}/\sigma \leq 4$) all exhibit steeper, more variable power laws at late times. Additionally, for all the displayed simulations, the cut-off particle energy is well approximated by γ_{rad} , indicating that radiative losses control this limit (even when too weak to steepen the non-thermal power-law tail).

2.3 A view in angular space

Up until now, we have described the evolution of our magnetic reconnection simulations from spatial and energetic viewpoints. We displayed several snapshots of the electron number density $n(x, y)$ in Fig. 1. Then, in Fig. 2, we described the electron energy distribution $f(\gamma)$ and how its evolution is impacted by radiative cooling. These pictures represent different ways of viewing the master distribution function in phase space $f(x, y, \gamma, \mathbf{\Omega}; t)$. At a given time t , this master distribution is five-dimensional, containing two spatial and three velocity dimensions, the latter of which we decompose into a Lorentz

factor $\gamma = E/mc^2$ and a direction labelled by the solid angle $\mathbf{\Omega}$. In terms of the master distribution, the number density and energy distribution are $n(x, y; t) = \int d\gamma d\mathbf{\Omega} f(x, y, \gamma, \mathbf{\Omega}; t)$ and $f(\gamma; t) = \int dx dy d\mathbf{\Omega} f(x, y, \gamma, \mathbf{\Omega}; t)$, respectively.

In this work, we are also interested in how particle momenta (and emitted photons) are distributed directionally. As a result, we must keep the angular information in the distribution function $f(x, y, \gamma, \mathbf{\Omega}; t)$, as was first done by Cerutti et al. (2012b). Furthermore, because we are interested in *kinetic* beaming – beaming as a function of particle or photon energy – we must preserve correlations between γ and $\mathbf{\Omega}$. To visualize all three velocity dimensions of the distribution function, we separate the energy information from the angular information, viewing the entire angular distribution at a single energy. Examples of this view are the angular maps (also ‘intensity maps’ or ‘heatmaps’) of Figs 3 and 4. These display the spatially integrated angular particle distribution $dN_t/d\gamma d\mathbf{\Omega} = \int_{y < L_y/2} dx dy f(x, y, \gamma, \mathbf{\Omega}; t)$ at fixed γ and t using the Aitoff projection. A particle contributes to latitude $\varphi \in [-90^\circ, 90^\circ]$ and longitude $\lambda \in [-180^\circ, 180^\circ]$ on a map if its velocity vector parallels the unit vector

$$\hat{\mathbf{n}} = \cos(\lambda) \cos(\varphi) \hat{\mathbf{z}} + \sin(\lambda) \cos(\varphi) \hat{\mathbf{x}} + \sin(\varphi) \hat{\mathbf{y}}. \quad (6)$$




Figure 3. An angular map for the $\gamma_{\text{rad}} = \sigma$ simulation displaying the angular intensity $I = dN_t/d\gamma d\Omega$ of lower-layer electrons at $2.0L/c$. This is a low-energy map – in the sense that the electron Lorentz factors γ are a decade below γ_{rad} – and exhibits only mild beaming.




Figure 4. An angular map for the $\gamma_{\text{rad}} = \sigma$ simulation displaying the angular intensity $I = dN_t/d\gamma d\Omega$ of lower layer electrons at $2.0L/c$. This is a higher energy map relative to Fig. 3 (the electron Lorentz factors are closer to γ_{rad}) and, as a result, exhibits more intense beaming patterns.

To isolate a single reconnection layer, Figs 3, 4, and all subsequent angular plots are generated using only particles (or, later, photons emitted from particles) located in the lower half of the simulation box.

Let us now describe the physical origins of the basic features in Figs 3 and 4. These features can be neatly decomposed into two broad categories: mild horizontal beaming stemming from bulk motion along the primary reconnection current sheet and extreme beaming arising near X-points (not only in the main current layer, but also in secondary ones between merging plasmoids). In the primary current layer near the main X-point, the reconnection electric field points in the $+z$ -direction ($\varphi = 0^\circ$, $\lambda = 0^\circ$), resulting in electron acceleration in the $-z$ -direction ($\varphi = 0^\circ$, $\lambda = \pm 180^\circ$). As electrons are ejected toward $-z$, they begin to be deflected by lines of reconnected magnetic field, which causes them to disperse toward $\pm x$ ($\varphi = 0^\circ$, $\lambda = \pm 90^\circ$). This results in the mild concentration of particles along the equator in Fig. 4. As these particles radiatively cool, they are simultaneously deposited into plasmoids, and plasmoid bulk motion along the reconnection layer induces a gentle low-energy momentum anisotropy along the $\pm x$ -directions as in Fig. 3.




Figure 5. A positron intensity map drawn for the same energy band and at the same time as the electron map of Fig. 4. Positrons feel the same reconnection electric field as electrons, but experience acceleration in the opposite direction as a result.




Figure 6. Electron density for the same time ($2.0L/c$) used to generate the angular maps of Figs 3, 4, and 5. In the top (bottom) panel, white ovals indicate the locations of a random subset of the electrons contributing to the angular map in Fig. 3 (Fig. 4) with Lorentz factors in the range $0.093\sigma < \gamma < 0.12\sigma$ ($0.47\sigma < \gamma < 0.60\sigma$). The higher energy particles reside in the hearts of the primary and inter-plasmoid current sheets – near X-points – and the lower energy particles in plasmoids.

We move now to the more extreme beaming. Generally, such pronounced anisotropy occurs only among the higher energy particles, a trend that Fig. 4 illustrates nicely. Less universal, but still common, is the fact that the strong beaming patterns in that figure result from plasmoid mergers, evidenced by the prominent vertical swaths in the angular particle distribution. This comes about because, between merging plasmoids, a secondary reconnection layer forms approximately parallel to the zy -plane. The reconnection electric field in this secondary layer points along $-z$ and accelerates electrons along $+z$. Owing to the rotated orientation of the reconnecting magnetic field, these electrons begin to fan out towards $\pm y$ – towards the poles – rather than $\pm x$ as in the primary current sheet.

The fan shapes in Fig. 4 are not perfectly vertical because plasmoids with unequal sizes and speeds are merging. At this time, there are actually two ongoing mergers – one on either flank of

the large primary plasmoid (Fig. 6) – both actively accelerating particles. In the merger on the right-hand side, a smaller and faster left-moving plasmoid creates a secondary current sheet that bends and moves to the left, biasing the accelerated particles towards the $-x$ -direction. The opposite is true for the merger on the left involving a small/fast right-moving plasmoid. The combined result is that the swaths of high energy particles shown in Fig. 4 do not extend along a single meridional plane running through $\lambda = 0^\circ$, but through two slightly offset planes intersecting $\lambda \simeq \pm 30^\circ$: one for each current sheet created at asymmetric plasmoid mergers.

The features in Figs 3 and 4 are nicely mirrored by those in the corresponding positron angular maps. As an example, Fig. 5 displays the angular distribution of positrons in the same energy band and at the same time as the electrons in Fig. 4. Owing to their opposite response to the reconnection electric field, the positrons yield the same X-point-generated beaming configurations found among the electrons in Fig. 4 but reflected about the xy -plane.

To tie the beaming features on angular maps to their spatial origins, Fig. 6 shows the spatial electron number density at the time for which Figs 3–5 were drawn. Also shown are the locations of randomly chosen subsets of electrons with Lorentz factors in the low-energy range of Fig. 3 and the high energy range of Figs 4 and 5. The positional clustering explicitly demonstrates the dichotomy described above: low-energy particles are confined to plasmoids whose bulk motion governs their momentum anisotropy while high energy particles exhibit more extreme beaming shaped by reconnection X-points. Fig. 6 also shows that X-point acceleration and collimation is ongoing in the primary reconnection layer even after the plasmoid chain has fully developed. As a result, it is not always the case (as it is in the angular maps shown previously) that the strongest beaming signatures are vertical. In fact, vertical fan shapes tend to be only intermittently prominent: when plasmoids – particularly large ones like in Fig. 6 – are actively merging. In between these episodes, strong beaming arising in the primary current sheet can still be significant (e.g. Fig. 17).

At this point, we would like to clarify that, of the two beaming origins discussed in this section, it is only the one operating near X-points that was associated with kinetic beaming in the works that originally introduced the concept (Uzdensky et al. 2011; Cerutti et al. 2012a,b). As discussed by those authors, the configuration of electromagnetic fields at these locations is particularly suited to accelerate and collimate high energy particles: the reconnection electric field delivers energy while the reconnecting magnetic field focuses particles into beams. The particles remaining near X-points longer are consequently more energized and more focused. As we have seen already (Fig. 4), this mechanism is responsible for the most severe beaming at the highest particle energies; as we shall see later, it also yields the most energy-*dependent* beaming. By comparison, plasmoid motion-generated anisotropy (Fig. 3) is milder and tends, because it derives from fluid level motion, to be more achromatic. (Plasmoids, along with their associated bulk motion and Doppler beaming, provide the basis for the ‘minijets’ model of Giannios et al. 2009, 2010.)

Thus, when we use the term ‘kinetic beaming’, we are not referring to just any energy-dependence in the degree of particle or photon collimation. We refer specifically to the most extreme and energy-dependent anisotropy at the highest energies generated near X-points. In this sense, the very different signatures of beaming evident in Figs 3 and 4, despite occurring at different particle energies, do not illustrate kinetic beaming. Instead, they portray two separate beaming mechanisms that merely dominate at different energy scales. We will illustrate kinetic beaming as we and previous authors apply the term – which involves energy-dependent anisotropy sourced only by the X-

point mechanism – after we develop a more quantitative description of beaming in the next section.

3 QUANTIFYING BEAMING

In this section, we present two quantitative notions of ‘beaming’ as manifested in angular maps such as Figs 3–5. Before proceeding, it will be helpful to introduce some additional terminology with which to describe the information on these maps: the angular distribution of particles $dN/d\Omega$, of the instantaneous radiated power $dP/d\Omega$, and quantities derived from these.¹ The angular distribution $dP/d\Omega$ is ordinarily called ‘intensity’ I , and the power P radiated into a finite solid angle is $P = \int I d\Omega$. In analogy with light, we shall frequently call the angular distribution of particles $dN/d\Omega$ by the name ‘intensity’, as well as borrow the symbol I . Furthermore, we will use the word ‘power’ to refer to the total number of particles travelling within a finite angular patch. This language enables us to describe beaming in generic terms. Whether we mean a beam of particles or a beam of photons will be clear from the context.

3.1 Two notions of beaming

As demonstrated by the intensity maps of Figs 3 and 4, the ZELTRON particle distributions do not necessarily exhibit what one typically imagines as a beam: a spot of high intensity that is nearly symmetric about some axis. Rather, the high-intensity regions on angular maps can be quite extended and complicated in shape, particularly at higher energies. Any quantitative definition of beaming one adopts must therefore be sufficiently versatile to handle the diverse set of momentum-space configurations attained by the particles (or photons, but for concreteness this section confines the discussion to the particle distribution).

To meet this challenge, we employ two complementary measures of beaming. The first was originally introduced by Cerutti et al. (2012b), who parametrized beaming by Ω_{50} : the smallest total (possibly non-contiguous) solid angle containing half of the power on an angular map (within the given energy bin). This quantity is illustrated in Figs 7 and 8. A smaller value indicates more extreme beaming because a smaller fraction of the sphere contains an order unity fraction of the power.

The second measure of beaming characterizes the angular regions where the intensity exceeds three times the angle-averaged intensity (denoted $\langle I \rangle_\Omega$ in the figures; again within a single energy bin). We define the ‘beamed fraction’, which we abbreviate as bf , to be the fraction of the total power contained within these regions. Extraction of the beamed fraction from the angular maps in Figs 3 and 4 is demonstrated in Figs 9 and 10.

Both measures of beaming – Ω_{50} and bf – have advantages and disadvantages. The beamed fraction does not rely on regions of extreme intensity being confined to small fractions of the sphere. On the other hand, Ω_{50} , when small, is perhaps a more convincing indicator of beaming because it means that the corresponding angular map region contains high power and occupies a small solid angle; large beamed fraction indicates only high power. We will use both tools in order to give a more compelling account of kinetic beaming.

One thing that these metrics have in common is that they are insensitive to the shapes and continuity of beams. Although one may

¹ We temporarily omit to explicitly write the dependence on time t , as well as the γ -dependence of $dN/d\Omega$ and the spectral dependence of $dP/d\Omega$, while establishing our nomenclature.




Figure 7. The same angular map as Fig. 3, but with a blue contour outlining the smallest solid angle that contains half of the total power – i.e. Ω_{50} , which in this case is 20 per cent of 4π .




Figure 8. The same intensity map as in Fig. 4, but with the Ω_{50} contour labelled. In this case Ω_{50} is 5 per cent of 4π .

conceive of more detailed and observer-centric measures of beaming, perhaps characterizing the morphologies of individual contiguous beams, this would greatly complicate the analysis. Leaving that for a future work, we find that our more coarse-grained measures are sufficient to illustrate a number of intriguing properties of the global system-wide beaming produced by magnetic reconnection.

Having developed two notions of beaming, we are now in a position to analyse *kinetic* beaming, which necessarily involves many maps across the particle energy spectrum. In this effort, the chief utility of the Ω_{50} and bf measures is to enable a reduction of the data contained on any given heatmap to two meaningful numbers, which we may then plot as a function of particle or photon energy. This procedure is illustrated in Figs 11 and 12. In the first figure, we display a collection of electron intensity maps spanning a decade in particle energy at a given instant in our $\gamma_{\text{rad}} = \sigma$ simulation (strongly radiative). Each map in that figure is distilled to two numbers, its Ω_{50} and its beamed fraction, which are then plotted as a function of particle energy in Fig. 12. The latter figure depicts the pronounced energy-dependence of beaming more concisely and dramatically, and we will make use of many similar plots throughout the remainder of this work.




Figure 9. The same angular map as in Fig. 3, but with a white contour outlining the high-intensity region (where the intensity exceeds three times its angle-average). The fraction of the heatmap power contained in this contour is the ‘beamed fraction’ and in this case is equal to 11 per cent. The solid angle footprint of the high-intensity region is 3 per cent of the sphere (note that the solid angle Ω_{bf} enclosed by the contour is *not* Ω_{50}). The beamed fraction provides an alternative measure of beaming to Ω_{50} .




Figure 10. The same map as Fig. 4, but with a white contour outlining the high-intensity region ($I \geq 3\langle I \rangle_\Omega$). Here the beamed fraction is equal to 60 per cent and is contained in 7 per cent of the available solid angle.

Before moving on to concentrate more exclusively on the succinct energy-centric view of beaming afforded by plots like Fig. 12, we would like to pause and emphasize, once more, the connection between the angular configurations realized in Fig. 11 and their underlying physical mechanisms (discussed previously in Section 2.3). Namely, Fig. 11 demonstrates: (1) mild beaming in the $\pm x$ -directions ($\varphi = 0^\circ$, $\lambda = \pm 90^\circ$) due to bulk plasmoid motion; and (2) dramatic beaming originating near X-points – in this case, X-points between merging plasmoids – and extending from the $+z$ -direction ($\varphi = 0^\circ$, $\lambda = 0^\circ$) towards the poles ($\varphi = \pm 90^\circ$). The former mechanism is most prominent at lower energies but the latter takes precedence at higher energies and gives rise to the steepest energy-dependence in Fig. 12. As a reminder, it is this more extreme beaming that we call ‘kinetic beaming’ and to which we shall devote the majority of our analysis in the next section.




Figure 11. Electron intensity maps from our $\gamma_{\text{rad}} = \sigma$ simulation, with white contours outlining the ‘beamed fraction’ (where the intensity is more than three times the average), for a series of particle energy bins at a single time $t = 2.0L/c$. Higher energy particles are more strongly beamed and, in this case, originate from reconnection sites between merging plasmoids.




Figure 12. Electron beamed fraction (solid blue), Ω_{50} (dotted black), and energy distribution (solid red) as a function of electron energy at $t = 2.0L/c$ in the $\gamma_{\text{rad}} = \sigma$ simulation. The Lorentz factor $\gamma = \gamma_{\text{rad}} = \sigma$ is denoted by a vertical dashed line. The Ω_{50} curve is normalized such that a perfectly isotropic angular map produces the value $\Omega_{50}/(4\pi \times 0.5) = 1$. In the same limit (attained at low energies), the beamed fraction tends to zero because the intensity is everywhere less than three times the isotropic intensity. The shaded region indicates the energy range shown in Fig. 11. The coincident sharp rise in the beamed fraction and precipitous drop in Ω_{50} demonstrate pronounced kinetic beaming at the highest energies.

4 KINETIC BEAMING AND RADIATIVE COOLING

In this section, we apply the quantitative measures of beaming described above – Ω_{50} and beamed fraction (bf) – to answer questions (i) and (ii) posed in the Introduction. First, we consider the question of observable kinetic beaming for two extreme cases: no radiative cooling and strong radiative cooling. After examining these

scenarios in detail, we conduct a higher level analysis that makes use of our full parameter scan in γ_{rad} to create a more complete picture of the dependence of kinetic beaming on cooling efficiency.

4.1 No cooling: $\gamma_{\text{rad}}/\sigma = \infty$

For our simulation without IC cooling ($\gamma_{\text{rad}}/\sigma = \infty$), the time evolution of three quantities as a function of particle Lorentz factor is displayed in Fig. 13. From top to bottom, these are the electron energy distribution, electron beamed fraction, and electron Ω_{50} .

In the figure, transient behaviour in all three quantities persists through about one light crossing time. During this early stage, beaming is both present and energy-dependent, with bf rising sharply and Ω_{50} falling steeply at high Lorentz factors: we observe clear kinetic beaming.

Here, we are restricting our discussion to the highest particle energies, ignoring the non-monotonic behaviour in bf and Ω_{50} that takes place at lower energies (and primarily at early times). This behaviour stems chiefly from a competition between the two sources of anisotropy in the particle distribution discussed previously in Section 2.3: plasmoid motion, which induces mild beaming among the low-energy particles, and collimation near reconnection X-points, which has a much more dramatic beaming effect primarily at high particle energies. At intermediate energies, the contributions from both plasmoids and X-points to the global (spatially integrated) distribution of particles can be approximately equal, causing bright regions to cover a larger portion of the angular map and, hence, making it appear more isotropic. By focusing on the highest particle (and, later, photon) energies, where beaming is monotonically increasing [$d(bf)/d\gamma > 0$ and $d\Omega_{50}/d\gamma < 0$], we isolate the contribution from X-points, the true underlying agents of ‘kinetic beaming’ as defined here and in previous works (Uzdensky et al. 2011; Cerutti et al. 2012a,b, see also Section 2.3).




Figure 13. The evolution of the electron energy distribution (top panel), beamed fraction (middle panel), and Ω_{50} (bottom panel) for our simulation with $\gamma_{\text{rad}}/\sigma = \infty$ (no cooling). Dashed lines indicate data prior to the first light crossing time and solid lines data between 1 and 3 light crossings. The solid lines show that, at late times, beaming vanishes at all but the highest (recently populated) energies: the beamed fraction tends to zero and Ω_{50} to $0.5 \times 4\pi$.

Returning now to Fig. 13, one sees that at later times, beaming is quenched. After one light crossing, both the bf and Ω_{50} curves approach their isotropic values – 0 and $0.5 \times 4\pi$, respectively – across nearly all particle energies. As discussed below, this occurs because, after their initial energization, particles quickly isotropize due to gyration about reconnected magnetic field lines. At first glance, it may appear that the highest Lorentz factors – those near the cut-off in the particle distribution – are exceptions to this rule, with dramatic beaming occurring even at late times. This is not really a persistent effect, however, because beaming lasts only temporarily at any fixed Lorentz factor, beginning when the high energy cut-off crosses (from below to above) that particular energy and ending shortly thereafter. Evidently, high energy bands retain their beaming only until they may be populated by a significant number of particles. The reason for this is illustrated in Fig. 14, and we discuss it here.

High energy bands in the particle distribution begin the simulation empty. At some point, X-point energization begins to populate such a band with particles. These ‘young’ (recently accelerated) particles are beamed in the same way that particles emerge beamed from the reconnection layer in the early part of the simulation ($t < L/c$, as observed previously by Cerutti et al. 2012b, who confined their analysis to early times). In the absence of radiative cooling, however, particles essentially remain in the energy band to which they are first accelerated. Meanwhile, they settle into plasmoids, where magnetic gyromotion isotropizes their momentum distribution. After a while, a given high energy band is dominated by ‘old’ particles whose angular spread has lost the beaming imprint left by X-points. Though a few young particles may still be injected into the energy band, their contribution to the angular intensity is washed out by the large number of older particles that has already accrued there. As a result, the high energy parts of the particle distribution contain brief, intense




Figure 14. Intensity maps for an exemplary high energy bin across several times for the $\gamma_{\text{rad}}/\sigma = \infty$ simulation. Beamed fraction contours are in white and beamed fraction values are labelled below each map. Initially, the energy bin contains no particles; only around $0.36L/c$ are particles first accelerated into the bin. Although initially beamed, in the absence of cooling, the particles maintain their energies and isotropize over time. Later, newly accelerated and collimated particles continue to appear in this energy band, but their contribution is drowned out by older, isotropized particles. The generic behaviour of the high energy uncooled particle distribution may be summarized as follows: high energy bands start empty, they become briefly beamed when they first acquire particles, but they then relax to an isotropic state.




Figure 15. The same as Fig. 13 but for our simulation with $\gamma_{\text{rad}}/\sigma = 1$ (strongly cooled). Unlike the non-radiative case, beaming is present across a moderate range of energies and persists well beyond $t = L/c$. The envelope of the late-time beamed fraction curves indicates at least mild kinetic beaming across a decade in particle energies. The strongest beaming occurs over a somewhat smaller range, where the beamed fraction curves begin to rise steeply and the Ω_{50} curves begin to turn over.

beaming when they first acquire particles, but subsequently isotropize as older, increasingly isotropic particles begin to pile up.

We have simplified this explanation by considering particle acceleration to be dominated by the impulsive X-point mechanism. Alternative slower and more isotropic acceleration channels have been studied, for example, by Petropoulou & Sironi (2018), Guo et al. (2019), and Hakobyan et al. (2020). However, these should be suppressed in the strongly radiative regime, unable to keep pace with the rapid cooling of the most energetic particles. Moreover, even when radiative losses are weaker (as in the present case $\gamma_{\text{rad}}/\sigma = \infty$) and these slower mechanisms are more likely to operate, they can only serve to reduce the amount of beaming we measure, tending to swamp the highly anisotropic angular signatures produced at X-points. Practically speaking, this means that, while an isotropic angular map may be the combined result of dispersing beams and intrinsically isotropic energization mechanisms, maps indicating strong kinetic beaming can only be attributed to X-point acceleration.

4.2 Strong cooling: $\gamma_{\text{rad}}/\sigma = 1$

We now turn to the case of strong radiative cooling $\gamma_{\text{rad}}/\sigma \lesssim 1$, analysing our $\gamma_{\text{rad}}/\sigma = 1$ simulation in detail. In this regime, the radiative cut-off γ_{rad} is not far above the typical energy (σ) – and well below the maximum energy (e.g. several σ , Werner et al. 2016) – that an energized particle would have in the absence of cooling. Thus, γ_{rad} chops off the part of the non-thermal power-law tail that could otherwise extend to energies above γ_{rad} (see Fig. 2 and the surrounding discussion).

Paralleling our treatment of the non-radiative case, Fig. 15 plots for the simulation with $\gamma_{\text{rad}}/\sigma = 1$ what Fig. 13 plots for the simulation with $\gamma_{\text{rad}}/\sigma = \infty$. Here, strong cooling causes the electron




Figure 16. The time-averaged electron distribution (red solid line), median electron beamed fraction (blue solid line), and median electron Ω_{50} (black dotted line) as a function of particle energy for times between L/c and $3L/c$ in our $\gamma_{\text{rad}}/\sigma = 1$ simulation. The shaded beamed fraction and Ω_{50} envelopes indicate the middle 68 percent of the time series data at each particle energy. As discussed in Fig. 15, but somewhat more obvious here, the beamed fraction indicates at least weak beaming across a decade on the horizontal axis. Both beamed fraction and Ω_{50} indicate strong beaming over a slightly narrower energetic range.

distribution, beamed fraction, and Ω_{50} to depart from their previous behaviour, where they essentially grew monotonically in horizontal or vertical extent. Because this makes it hard to discern the time evolution in Fig. 15, we also supply Fig. 16, which presents the post-one-light-crossing time-averaged particle distribution, median beamed fraction, and median Ω_{50} . As evident from Fig. 16, beaming persists, when cooling is strong, to late times across almost a decade in particle energy. In contrast, as we saw in the previous section, kinetic beaming is only transient when cooling is weak; it is present across a wide range of energies at early times and thereafter relegated to energies near the cut-off of the particle distribution.

As is the case without radiative losses, acceleration from reconnection X-points preferentially collimates the more energetic particles into beams. The crucial difference with strong cooling is that particles radiate away most of their energy before they have had time to isotropize: the most energetic particles are always ‘young’ (recently accelerated). As a result, the high energy part of the particle distribution remains beamed at much later times (see Fig. 17). Beaming falls off with decreasing particle energy, however, because particles that have been cooling longer have also been isotropizing longer (they are ‘older’).

In principle, whether kinetic beaming persists at a given particle energy comes down to whether the isotropization time-scale for those particles is longer or shorter than their cooling time-scale. Suppose, for illustration, that the particle isotropization time-scale is the gyration period $t_{\text{iso}} \sim \gamma/\omega_0$ where the nominal Larmor frequency is $\omega_0 = eB_0/m_e c$. The cooling time-scale for the same particles is $t_{\text{IC}} \sim \gamma m_e c^2 / P_{\text{IC}}(\gamma) \sim 10\gamma_{\text{rad}}^2 / \omega_0 \gamma$. One expects the smallest Lorentz factor γ_{iso} for which kinetic beaming persists to late times to be that for which these time-scales are equal: $\gamma_{\text{iso}} \sim \sqrt{10\gamma_{\text{rad}}}$. This picture is oversimplified, for it predicts γ_{iso} to exceed (somewhat) the radiation-reaction limit γ_{rad} , and it predicts $\gamma_{\text{rad}}/\gamma_{\text{iso}} \sim \text{constant}$, while we find (see Section 4.3) a non-trivial scaling of $\gamma_{\text{rad}}/\gamma_{\text{iso}}$ with γ_{rad} . What can be said for certain in the case $\gamma_{\text{rad}}/\sigma = 1$ is that particles radiate more quickly than they isotropize over a considerable range of energies.




Figure 17. The time-dependence of a high energy heatmap for the $\gamma_{\text{rad}}/\sigma = 1$ simulation. In the uncooled simulation, a similar set of figures (Fig. 14) demonstrated the transience of beaming in that energy bin. But in this strongly cooled case, spectacular beaming patterns persist to times well past $t = L/c$. Note that sometimes (particularly $t = 1.6L/c$) X-points in the primary reconnection layer dominate the beaming, producing horizontal swaths of high-intensity regions; at other times, the prominent vertical swaths indicate beaming produced at X-points between merging plasmoids (cf. Section 2.3 discussion).

4.3 Kinetic beaming as a function of radiative efficiency

Having demonstrated that strong kinetic beaming persists in the presence of efficient radiative cooling but disappears after about $t = L/c$ when cooling is negligible, we now analyse the transition between these regimes. In particular, we examine how sustained kinetic beaming weakens as the result of decreasing IC radiative efficiency. We also shift our focus from the angular particle distribution $dN_i/d\gamma d\Omega$ to the IC emission spectrum $dP_i/d\epsilon d\Omega$, where ϵ and Ω are the energy and direction, respectively, of IC photons.² This presents no challenge from an analysis standpoint, since the diagnostics we have been using so far (angular maps, beamed fraction, Ω_{50}) apply as well to photons as to particles. In fact, analysing the photons themselves rather than the emitting particles enables a more precise measurement of kinetic beaming, a point that we now briefly elaborate.

In this section, we shall be interested in the energetic extent of kinetic beaming: the range of (high) photon or particle energies across which a strong energy-dependent anisotropy is evident. Now, in the Thomson regime, photons Comptonized by a particle of Lorentz factor γ attain increased energies by the factor γ^2 and are emitted along the particle's velocity vector within a cone of half-opening angle $1/\gamma$. So, in the ultrarelativistic limit ($\gamma \gg 1$), not only is the angular distribution of emission nearly identical to that of the radiating particles, but kinetic beaming actually encompasses a wider range of energies in the former than in the latter. Therefore, treating the IC emission rather than the radiating particles directly allows us to measure more precisely the energetic extent of kinetic beaming as it becomes small.

In Fig. 18, we illustrate beaming as a function of photon energy for three of our radiative runs ($\gamma_{\text{rad}}/\sigma = 1, 4$, and 16). In that figure, one sees that kinetic beaming – marked by rising bf and declining Ω_{50}

at the highest energies – persists well beyond the first light crossing time in all simulations. Also, the energy range across which beaming is kinetic widens for the more strongly radiative simulations. Fig. 18 quantifies these observations by displaying two characteristic photon energies. The first is the energy ϵ_{iso} above which beaming acquires pronounced spectral dependence – where the Ω_{50} curves begin to turn downward. The second is ϵ_c , the cut-off in the IC emission spectrum. The ratio of the cut-off ϵ_c to the ‘isotropic’ energy ϵ_{iso} characterizes the beamed range of photon energies.

These energy scales and their ratio will be critical to our eventual quantitative portrait of kinetic beaming as a function of radiative efficiency. Therefore, we will here expound upon the techniques we use to measure them as well as describe the trends in our measurements across our series of simulations. Let us begin with ϵ_c . Following the method of equation 37 in Bai et al. (2015) (see also Sironi et al. 2016; Hakobyan et al. 2019), we take

$$\epsilon_c = \frac{\int d\epsilon \epsilon^n F_{\text{IC}}(\epsilon)}{\int d\epsilon \epsilon^{n-1} F_{\text{IC}}(\epsilon)}, \quad (7)$$

where $F_{\text{IC}}(\epsilon) = \int d\Omega dP_i/d\epsilon d\Omega$ and n is empirically determined. We use $n = 4$ (higher values do not change the power-law scaling of ϵ_c with cooling strength in Fig. 19). The result of this calculation is displayed in Fig. 18 for three reference simulations, and the functional dependence of ϵ_c on γ_{rad} is displayed, for our entire series of simulations, in Fig. 19. (Note in that figure, as well as in Figs 20 and 21, the horizontal axis is $\sigma/\gamma_{\text{rad}}$, which *increases* with stronger cooling.) We omit $\gamma_{\text{rad}}/\sigma = 64, \infty$ from Figs 19–21 because those simulations exhibit secular growth in ϵ_c throughout our analysis interval $1 < ct/L \leq 3$, never reaching a quasi-steady state.

In Fig. 19, the large- γ_{rad} (weak radiation-reaction) scaling of ϵ_c with γ_{rad} significantly deviates from $\epsilon_c = 4\gamma_{\text{rad}}^2 \epsilon_{\text{ph}}$ (recall here that ϵ_{ph} is the monochromatic energy of IC seed photons). This scaling is a special case of the more general result $\epsilon_c = 4\gamma_c^2 \epsilon_{\text{ph}}$, which equals the maximum emitted photon energy from a particle at the cut-off Lorentz factor γ_c in the particle distribution. At large γ_{rad} , we

²The quantity $dP_i/d\epsilon d\Omega$ is the instantaneous (at time t), lower-layer, volume-integrated IC emission coefficient j_{IC} , with j_{IC} as defined by Rybicki & Lightman (1979): $dP_i/d\epsilon d\Omega = \int_{\gamma < L_{\gamma}/2} dx dy j_{\text{IC}}(x, y, \epsilon, \Omega; t)$.




Figure 18. The time-averaged (over $1 \leq ct/L \leq 3$) electron IC emission spectrum (solid red), median beamed fraction (bf ; solid blue), and median Ω_{50} curve (dotted black) for the simulations with $\gamma_{\text{rad}}/\sigma = 1$ (left), 4 (middle), and 16 (right). Shaded regions indicate, for Ω_{50} and bf only, the middle 68 percent of data. In each panel, the Ω_{50} data (dotted black lines with shaded regions serving as error bars) are fit using a smoothly broken power law parametrized as in equation (10). To avoid needing a more complicated fitting formula, only data following the last local maximum in the dotted Ω_{50} curve are fit. The fit is drawn as a solid black line on top of the data that were used. Solid black vertical lines indicate the onset of the best-fitting spectral break (ϵ_{iso} in the text); dot-dashed red vertical lines show the spectral cut-off (ϵ_c in the text) in the displayed emission spectrum. The horizontal axis is normalized to the maximum photon energy $4\sigma^2\epsilon_{\text{ph}}$ to which a particle of Lorentz factor σ can upscatter ϵ_{ph} -energy photons.




Figure 19. The cut-off photon energy displayed as a function of IC cooling strength using two different normalizations. In the upper panel, ϵ_c is normalized to $4\sigma^2\epsilon_{\text{ph}}$; in the bottom panel, the normalization is $4\gamma_{\text{rad}}^2\epsilon_{\text{ph}}$. Each data point represents the median cut-off computed as a function of time (over $1 \leq ct/L \leq 3$) for a given simulation. Error bars indicate the middle 68 percent of data. A power-law fit to the data is presented in blue, with shaded blue envelope indicating the uncertainty in the fit. The expected low- γ_{rad} scaling $\epsilon_{c1} \sim \gamma_{\text{rad}}^2$ is displayed in dashed magenta and the high- γ_{rad} scaling $\epsilon_{c2} \sim \gamma_{\text{rad}}^{4/3}$ in dotted magenta. Additionally, the lower panel shows the empirical fitting formula (9). A χ^2 goodness-of-fit test does not reject equations (8) and (9), or the single power law.

measure $\epsilon_c < 4\gamma_{\text{rad}}^2\epsilon_{\text{ph}}$. This means that $\gamma_c < \gamma_{\text{rad}}$ and, perhaps, that diminished radiative efficiency allows a slower particle acceleration mechanism to dominate the highest energies.³

³We have verified that $\gamma_c < \gamma_{\text{rad}}$, but do not present a corresponding plot.




Figure 20. The isotropic photon energy ϵ_{iso} , which marks the transition to strong kinetic beaming in the Ω_{50} curve, as a function of IC cooling strength. Error bars indicate 68 percent confidence intervals given by Markov chain Monte Carlo fits (Foreman-Mackey et al. 2013) to each simulation's Ω_{50} curve (see Fig. 18) using a smoothly broken power law [as parametrized in equation (10)]. In this figure, the ϵ_{iso} data are fit with unbroken power laws across both the full range of data $1/2 \leq \gamma_{\text{rad}}/\sigma \leq 32$ (solid blue with shaded error envelope) and a restricted range $1/2 \leq \gamma_{\text{rad}}/\sigma \leq 16$ (dashed green with shaded error envelope). The restricted fit excludes the non-radiative asymptotic behaviour where ϵ_{iso} is expected to clamp to ϵ_c .

To see how this might work, let us suppose that such a mechanism operates and that the associated acceleration time t_{slow} for a particle to double its Lorentz factor $\gamma \rightarrow 2\gamma$ scales as $t_{\text{slow}} \propto \gamma^\zeta$. Generally, we expect $\zeta > 1$ since the linear case $\zeta = 1$ corresponds to the fast time-scale $t_X \propto \gamma$ associated with direct X-point acceleration [to be exact, $t_X(\gamma) = \gamma m_e c^2 / l c e \beta_{\text{rec}} B_0 \simeq 10 \gamma / \omega_0$]. Furthermore, for self-consistency, this picture requires that the fast acceleration scaling $t_X \propto \gamma$ does not continue indefinitely to high energies; otherwise there would never be a large- γ regime where $t_{\text{slow}} < t_X$. We will therefore assume that X-point acceleration can only deliver particles up to a threshold Lorentz factor γ_X and beyond that ceases to operate. (Formally, $t_X \propto \gamma$ for $\gamma < \gamma_X$ and $t_X = \infty$ otherwise.)

In this scenario, a good candidate Lorentz factor for γ_X is $\gamma_X \sim 4\sigma$. This has already been suggested by Werner et al. (2016) as a natural limit set by the size of elementary current layers in the plasmoid




Figure 21. The ‘range’ of kinetic beaming as a function of $\sigma/\gamma_{\text{rad}}$, expressed as the ratio of cut-off photon energy to photon energy of beaming onset: $\epsilon_c/\epsilon_{\text{iso}}$. For each simulation, the photon energy ϵ_{iso} is measured as the break energy in a broken power-law fit to the Ω_{50} curve, as in Fig. 18. Error bars presented are ‘worst case’, given by computing the ratio $\epsilon_c/\epsilon_{\text{iso}}$ at the extreme ends of the confidence intervals in Figs 19 and 20. The solid blue power-law fit uses the entire range of data $1/2 \leq \gamma_{\text{rad}}/\sigma \leq 32$ and the dashed green power law uses the restricted range $1/2 \leq \gamma_{\text{rad}}/\sigma \leq 16$. (As in Figs 19 and 20, shaded blue and green envelopes indicate fit uncertainties.) Fits omitting additional data from the weakly radiative end (high γ_{rad} , low $\sigma/\gamma_{\text{rad}}$) yield power-law scalings consistent with the dashed-green line, suggesting that $\epsilon_c/\epsilon_{\text{iso}}$ may asymptote to a non-radiative limit of order unity for $\gamma_{\text{rad}} \gtrsim 16\sigma$.

hierarchy. Moreover, $\gamma_{\text{rad}} = 4\sigma$ is, quite suggestively, where the scaling $\epsilon_c \propto \gamma_{\text{rad}}^2$ appears to break down in Fig. 19. Let us therefore tentatively assign $\gamma_X = 4\sigma$ (appropriate for our simulations, but, as discussed in Section 4.5, not necessarily the case in all astrophysical instances of reconnection).

Assuming radiative losses are weak enough that particles reach Lorentz factors exceeding $\gamma_X = 4\sigma$ (i.e. $\gamma_{\text{rad}} > 4\sigma$), we are in a regime where $t_{\text{slow}} < t_X$ at the highest energies. Equating, therefore, t_{slow} to the IC cooling time $t_{\text{IC}}(\gamma) = 10\gamma_{\text{rad}}^2/\gamma\omega_0 \propto \gamma_{\text{rad}}^2/\gamma$ gives an expected cut-off in the particle distribution $\gamma_c \propto \gamma_{\text{rad}}^{2/(\zeta+1)}$ decided not by the competition between radiative losses and X-point acceleration, but by that between radiation and the putative slower acceleration mechanism. Importantly, the corresponding photon energy cut-off $\epsilon_c = 4\gamma_c^2\epsilon_{\text{ph}} \propto \gamma_{\text{rad}}^{4/(\zeta+1)}$ scales more gently than $\propto \gamma_{\text{rad}}^2$, which is expected only if $\gamma_{\text{rad}} \leq \gamma_X = 4\sigma$.

The above considerations are more than just an abstract hypothetical exercise. In fact, a slower energization process with $\zeta = 2$ – in which the Lorentz factors of high energy particles generally follow $\gamma(t) \propto \sqrt{t}$ – has been identified by Petropoulou & Sironi (2018) and recently elaborated by Hakobyan et al. (2020). In contrast to direct acceleration at reconnection X-points, this mechanism operates on particles inside plasmoids. There, particles are accelerated gradually due to conservation of their magnetic moments in the presence of a slowly growing magnetic field. The resultant scaling $\epsilon_c \propto \gamma_{\text{rad}}^{4/3}$ is not far from the apparent weak radiation-reaction scaling in Fig. 19.

Let us now use these ideas to construct a theoretical model that explains all of our ϵ_c measurements. In this effort, we regard the cut-offs γ_c and ϵ_c as dependent functions of the independent variable γ_{rad} . The particular value $\gamma_{\text{rad}} = \gamma_X = 4\sigma$ is special, because we assume that, for $\gamma_{\text{rad}} > \gamma_X$, X-point acceleration is subdominant.

Instead, particles are primarily accelerated by a process similar to that of Petropoulou & Sironi (2018) and Hakobyan et al. (2020), which operates on a time-scale $t_{\text{slow}} = C\gamma^2$. To fix the proportionality constant C , we require that the slow and fast acceleration time-scales, t_{slow} and t_X , give equal cut-off Lorentz factors $\gamma_c(\gamma_{\text{rad}}) = \gamma_{\text{rad}}$ at the transition value $\gamma_{\text{rad}} = \gamma_X$. This can be expressed as the condition $C\gamma_c^2 = t_{\text{slow}} = t_X = t_{\text{IC}} = 10\gamma_{\text{rad}}^2/\gamma_c\omega_0$, and yields, upon inserting $\gamma_c = \gamma_{\text{rad}} = 4\sigma$, the result $t_{\text{slow}} = 5\gamma^2/2\sigma\omega_0$ ($C = 5/2\sigma\omega_0$). For $\gamma_{\text{rad}} < \gamma_X$, the cut-off γ_c is set by the competition between X-point acceleration and radiative cooling and is given by $t_X(\gamma_c) = t_{\text{IC}}(\gamma_c)$; for larger γ_{rad} , cooling balances the slower acceleration mechanism and γ_c can be found from the condition $t_{\text{slow}}(\gamma_c) = t_{\text{IC}}(\gamma_c)$. The IC photon cut-off energy $\epsilon_c = 4\gamma_c^2\epsilon_{\text{ph}}$ is then a broken power law in γ_{rad} :

$$\epsilon_c = \begin{cases} \epsilon_{c1} = 4\gamma_{\text{rad}}^2\epsilon_{\text{ph}} & \gamma_{\text{rad}} \leq \gamma_X = 4\sigma \\ \epsilon_{c2} = 4\gamma_{\text{rad}}^2\epsilon_{\text{ph}}(4\sigma/\gamma_{\text{rad}})^{2/3} & \text{otherwise} \end{cases} \quad (8)$$

Should one wish to smooth the transition between $\epsilon_c = \epsilon_{c1}$ and $\epsilon_c = \epsilon_{c2}$, we find that the empirical formula

$$\frac{1}{\epsilon_c^2} = \frac{1}{\epsilon_{c1}^2} + \frac{1}{\epsilon_{c2}^2} \quad (9)$$

describes our $\epsilon_c(\gamma_{\text{rad}})$ data quite well. Both this smoothed form and ϵ_{c1} and ϵ_{c2} individually are displayed in Fig. 19. Also shown is a power-law fit $\epsilon_c \propto \gamma_{\text{rad}}^{1.6}$ exhibiting a scaling intermediate between ϵ_{c1} and ϵ_{c2} . Formally, equations (8) and (9), and a single power law all acceptably reproduce our $\epsilon_c(\gamma_{\text{rad}})$ data, but, based on our theoretical considerations, we suspect that a broken power law more accurately reflects the underlying physics.

Let us now move on to describe our other important beaming-related energy scale ϵ_{iso} . As we did with ϵ_c , we will first describe how we measure this quantity, interpreting our measurements thereafter. To calculate ϵ_{iso} , the first step is to fit a smoothly broken power law of the form

$$\Omega_{50}(\epsilon) = A \left(\frac{\epsilon}{\epsilon_{\text{br}}} \right)^{-p_1} \left\{ \frac{1}{2} \left[1 + \left(\frac{\epsilon}{\epsilon_{\text{br}}} \right)^{1/\Delta} \right] \right\}^{(p_1-p_2)\Delta} \quad (10)$$

to the Ω_{50} curve (astropy.modeling.powerlaws.SmoothlyBrokenPowerLaw1D; Astropy Collaboration 2018, see Fig. 18). The parameters A , p_1 , p_2 , and ϵ_{br} are the scale, power-law indices, and spectral break of the fit. The parameter Δ controls the width of the break in the sense that equation (10) constitutes a pure power law with index $-p_1$ ($-p_2$) at energies below $\epsilon_{\text{br}}/10^\Delta$ (above $\epsilon_{\text{br}}10^\Delta$).

Once the parameters in (10) are determined, we take $\epsilon_{\text{iso}} = \epsilon_{\text{br}}/10^\Delta$. This definition is empirically motivated from two observations. First, our fits generally produce p_1 nearly flat and p_2 steep ($|p_1| \leq 0.2$ and $p_2 \geq 1.5$ across all fits). Second, in the (low-energy) p_1 segment of the curve, Ω_{50} hovers near isotropy [$\Omega_{50}(\epsilon) \simeq 0.5 \times 4\pi$]. Hence, at energies below $\epsilon_{\text{br}}/10^\Delta$, Ω_{50} is both energy-independent and isotropic (i.e. kinetic beaming is absent), but above $\epsilon_{\text{br}}/10^\Delta$, Ω_{50} begins to turn over, eventually declining precipitously with photon energy. Thus, our choice $\epsilon_{\text{iso}} = \epsilon_{\text{br}}/10^\Delta$ provides a good description for when beaming starts becoming kinetic, as intended.⁴

In order to build confidence in our ϵ_{iso} -extraction method, and to illustrate the utility of having two metrics of beaming (Ω_{50} and b_f), we will discuss one subtlety associated with our procedure. Namely,

⁴Because we wish to flag the onset of kinetic beaming, we choose $\epsilon_{\text{iso}} = \epsilon_{\text{br}}/10^\Delta$ rather than $\epsilon_{\text{iso}} = \epsilon_{\text{br}}$. This means that, intentionally, the measured ϵ_{iso} values indicated in Fig. 18 are often just before the Ω_{50} curves turn over, rather than in the middle of the spectral break.

because 10^A lies between 2 and 8 across all our fits (except for one broad transition $10^A = 20$ in our $\gamma_{\text{rad}}/\sigma = 2$ simulation), and because ϵ_{br} is rather large to begin with, the energy $\epsilon_{\text{br}} 10^A$ signalling the end of the spectral break often falls near the rightmost edge of the Ω_{50} data (or, for $\gamma_{\text{rad}}/\sigma = 2$, well beyond it). This means that p_2 and, to some extent, Δ and ϵ_{br} are not necessarily well constrained. We deal with this difficulty in two ways. First, we do not rigorously study the p_2 measurements. We only report the lowest value $p_2 = 1.5$ (see above) to generally indicate the pronounced energy-dependence acquired by Ω_{50} beyond ϵ_{iso} . Secondly, for our fitted ϵ_{br} and Δ values, which are used directly in our definition of ϵ_{iso} , we provide the following sanity check using our second metric of beaming: the beamed fraction. In particular, for each simulation, the location $\epsilon_{\text{iso}} = \epsilon_{\text{br}}/10^A$ – despite being entirely determined from the Ω_{50} data – roughly coincides with photon energies where the beamed fraction slope increases most rapidly [where $d^2(bf)/d\epsilon^2$ is peaked]. This qualitative agreement between beaming metrics suggests that ϵ_{iso} flags a real feature in the Ω_{50} curves, and is not merely an artefact of truncated high energy Ω_{50} information. Fig. 20 displays the dependence of ϵ_{iso} on $\sigma/\gamma_{\text{rad}}$.

The kinetic beaming range $\epsilon_c/\epsilon_{\text{iso}}$, measured from actual PIC simulations using the above techniques, allows us to quantify the energetic extent of kinetic beaming versus cooling strength. This is done in Fig. 21, where, for each of our radiative simulations, we present the value of $\epsilon_c/\epsilon_{\text{iso}}$ as a function of $\sigma/\gamma_{\text{rad}}$. In both Figs 20 and 21, we supply power-law fits to our data. This is not meant to indicate a robust theoretical description, but merely to characterize how quickly these quantities change from the weakly radiative ($\gamma_{\text{rad}} \gg \sigma$) to the strongly radiative ($\gamma_{\text{rad}} \lesssim \sigma$) regime.

Fig. 21 demonstrates a clear dependence of the kinetic beaming range on IC cooling strength. Moreover, this dependence may be stronger than is suggested by naively fitting a single power law to our entire data set $1/2 \leq \gamma_{\text{rad}}/\sigma \leq 32$. This is because ϵ_{iso} can never exceed ϵ_c : particles and photons can be neither isotropic nor beamed at energies where none of them exist. Thus, the ratio $\epsilon_c/\epsilon_{\text{iso}}$, although strongly dependent on γ_{rad} when γ_{rad} approaches the strong cooling regime, must ultimately asymptote to unity in the non-radiative limit. According to our data, this occurs closer to $\gamma_{\text{rad}} = 16\sigma$ than to $\gamma_{\text{rad}} = 32\sigma$. To demonstrate this, we conduct a series of fits to the data in Fig. 21, of which we display only the first two. In each successive iteration, we remove the most weakly radiative simulation (that is, we keep all the data for the first fit, omit $\gamma_{\text{rad}} = 32\sigma$ for the second, omit $\gamma_{\text{rad}} = 32\sigma$ and $\gamma_{\text{rad}} = 16\sigma$ for the third, etc.). The power law becomes insensitive to this procedure once we are restricted to $\gamma_{\text{rad}} \leq 16\sigma$ and at that point exhibits nearly linear scaling.

This suggests that our simulation series captures an important transition in the range of beamed photon energies. For the mildest radiative cooling ($\gamma_{\text{rad}} \gtrsim 16\sigma$), kinetic beaming is unobservable, manifesting itself nowhere in the distribution of emitted photons ($\epsilon_{\text{iso}} \sim \epsilon_c$). However, once the radiative efficiency is increased, kinetic beaming suddenly appears, and persists at late times across a sizeable range of photon energies. This range increases throughout the entire set of γ_{rad} that is numerically accessible to us, and even surpasses one decade when $\gamma_{\text{rad}} \lesssim 4\sigma$.

Finally, we note that the simulation with $\gamma_{\text{rad}} = \sigma/2$, which (as noted in Section 2) is problematic from the standpoint of cooling in the upstream plasma, is not essential to any of the findings in this section. In particular, excluding it leaves all best-fitting power-law scalings essentially unchanged. Because this simulation does not modify any of the overall trends, and indeed appears to fall in line with those trends, we have included it in Figs 19–21.




Figure 22. Top panel: ϵ_c and ϵ_{iso} measured using the same techniques as for Figs 19 and 20 but for a series of simulations at uniform $\gamma_{\text{rad}} = 4\sigma$ and varying $L/\sigma\rho_0 = [80, 120, 160, 240, 320]$. Bottom panel: the kinetic beaming range $\epsilon_c/\epsilon_{\text{iso}}$ for the same series of simulations. Solid lines indicate constant fits to each quantity, with shaded regions displaying 68 percent confidence intervals. A χ^2 goodness-of-fit test does not reject any of these flat lines, showing that, given the error bars, the data are consistent with L -independence.

4.4 System size dependence

Before placing our numerical results in an astrophysical context, let us briefly explore how those results depend on the size of our computational box $L_x = L$. To do so, we report on a small series of simulations with fixed radiation-reaction strength $\gamma_{\text{rad}} = 4\sigma$ and varying $L/\sigma\rho_0 \in [80, 120, 160, 240, 320]$. All other parameters are the same as described in Section 2 and summarized in Table 1.

We have chosen $\gamma_{\text{rad}} = 4\sigma$ for these simulations because, on the one hand, our results in the previous section indicate that this radiative efficiency is strong enough to yield a substantial range of beamed photon energies. On the other hand, the radiation is weak enough that the upstream plasma does not cool at all during the simulation [$t_{\text{cool}}/(3L/c) = 22$ in equation 4]. In our γ_{rad} scan, we tolerated a larger degree of upstream cooling because the expected error imparted to our main measured quantities – ϵ_{iso} , ϵ_c , and $\epsilon_c/\epsilon_{\text{iso}}$ – was both within our measurement error and swamped by the pronounced observed γ_{rad} -dependence. However, here we need to be more strict. Our goal is to demonstrate system-size insensitivity of the same beaming quantities, and, hence, we need to eliminate any upstream cooling effects that could selectively come into play at larger L .

For this series of simulations, Fig. 22 displays the values of ϵ_c , ϵ_{iso} , and the kinetic beaming range $\epsilon_c/\epsilon_{\text{iso}}$ calculated by the methods described in Section 4.3. If kinetic beaming were to weaken with system size, presumably because of a diminished importance of kinetic scale phenomena, then one would expect the kinetic beaming range $\epsilon_c/\epsilon_{\text{iso}}$ to exhibit a downward trend with L . Instead, our data

show that this quantity is consistent with being constant with L . This is in fact true not only of the ratio $\epsilon_c/\epsilon_{\text{iso}}$ but also of ϵ_c and ϵ_{iso} individually. If there is any non-constant trend at all, although this is not statistically significant, the kinetic beaming range *increases* with L (owing primarily to a decrease in ϵ_{iso}).

In this system-size scan, we have not explored larger L because (1) our existing results establish a compelling trend and (2) the cost of larger boxes is prohibitive: a factor of 2 increase leads to a factor of 2^3 additional computational expense. Pending a future study that more thoroughly elucidates system-size effects – both by going to larger L and by trialing additional values of γ_{rad} – we will move forward assuming that the basic picture of kinetic beaming is captured by our main simulation series (with constant $L = 320\sigma\rho_0$ and varying γ_{rad}). We will, in particular, assume that those results can be extrapolated to astrophysically large systems, as suggested by the L -invariance apparent in this system-size investigation.

4.5 Summary of kinetic beaming

In this section, we would like to collect and summarize what we have learned so far, particularly from Section 4.3 (the fundamental features of which, as suggested in Section 4.4, may apply even to astrophysically large systems), with an eye towards extrapolating these findings to astrophysical situations. These remarks will be mostly at a general level, independent of any particular class of astrophysical objects, but they will set the stage for our specific application to TeV blazar flares in Section 5.

Assuming the connection between beaming and rapid light-curve variability (described in the earlier works of Cerutti et al. 2012b, 2013, 2014a,b), our main result is that kinetic beaming requires a high degree of radiative efficiency in order to leave an observational signature. From our analysis of the kinetic beaming range $\epsilon_c/\epsilon_{\text{iso}}$ (Fig. 21), we have quantitatively discovered that beaming-induced variability should only exist in the highest energy spectral bands, where the emitting particles are near their radiatively imposed cut-off energy. An implication in the context of gamma-ray (e.g. blazar) flares is that, if high-cadence observations could be made in both bands, TeV flares should not have similar time-scale GeV counterparts (unless multiple radiative processes enable strongly cooled particles to emit in several bands): kinetic beaming does not appear to be sustained over such a broad energy range.

Our findings further suggest that, when radiative cooling is weak, kinetic beaming does not necessarily explain the total duration of rapid flares. While the transient initial phase of beaming seen in our weakly cooled simulations could influence the rising part of a flare, subsequent isotropization would likely prevent, in that case, the fast-rise fast-decay pattern characteristic of a collimated beam crossing the line of sight. On the other hand, kinetic beaming may well shape the entire temporal profile of rapid outbursts when radiative cooling is efficient. Then the energetic particle beams emit corresponding photon beams before diverging. Near the spectral cut-off, beaming is pronounced and kinetic, and the light curves in this band are expected to exhibit increasingly dramatic variability at higher and higher energies.

One point worth emphasizing is that our usage of the terms ‘strong’ and ‘efficient’ cooling is not necessarily the same as that in other works. Often in astrophysics, radiative cooling is said to be efficient if particles cool faster than some macroscopic system time-scale. Here, the definition of efficient radiative cooling is at least as strong, but (and depending on the particular system) often much stronger than these more conventional notions. For us, strong cooling is *microscopically* strong. A particle with Lorentz factor close

to its radiatively imposed limit γ_{rad} has a cooling time matching its *acceleration* time through the reconnection layer, or, equivalently, has a cooling length (the distance it travels in one cooling time) of order its Larmor radius (see Uzdensky 2016). Either scale may potentially be much smaller than any macroscopic system scale.

Finally, we would like to state a conjecture that may broaden the scope of our results, enhancing the potential variety of astrophysical sources for which kinetic beaming may explain rapid flares. This conjecture concerns the range of $\gamma_{\text{rad}}/\sigma$ for which kinetic beaming extends across an appreciable span of energies (for which $\epsilon_c/\epsilon_{\text{iso}}$ is sizable). Fig. 21 suggests that this range only exceeds about a decade when $\gamma_{\text{rad}} \leq 4\sigma$. This, however, does not necessarily mean that sources for which $\gamma_{\text{rad}} \gg \sigma$ should not exhibit kinetically beamed emission. In particular, we believe that the figure of merit for a kinetic beaming scenario is not whether γ_{rad} is of order several σ or less, but whether the actual energies achieved by particles are close to γ_{rad} . As suggested by our results (Fig. 19), particles are only expected to reach Lorentz factors $\sim \gamma_{\text{rad}}$ if they are accelerated by the fast X-point mechanism. Slower acceleration channels radiatively saturate at energies less than γ_{rad} . Hence, if particles are somehow able to reach γ_{rad} even when $\gamma_{\text{rad}} \gg \sigma$, then they must have been accelerated and, consequently, beamed near an X-point.

Thus, what kinetic beaming really depends on is not whether $\gamma_{\text{rad}} \lesssim (\text{several})\sigma$, but whether γ_{rad} is comparable to or less than the maximum Lorentz factor γ_X achievable due to X-point acceleration. As hypothesized by this and prior studies (Werner et al. 2016; Petropoulou & Sironi 2018; Hakobyan et al. 2020), γ_X may be of the order of 4σ in reconnection set-ups like the one employed by us in this work, meaning that the requirement $\gamma_{\text{rad}} \leq \gamma_X$ in our case simplifies to $\gamma_{\text{rad}} \leq 4\sigma$. However, in alternative and more astrophysical situations, it may be possible for γ_X to circumvent this 4σ limit. For example, many of the sources (e.g. pulsar wind nebulae and active galactic nuclei) for which kinetic beaming nicely explains a number of aspects of observed flares are also highly non-thermal emitters, even in their quiescent states. This suggests that the upstream plasma is itself non-thermal, possessing a long tail of already high energy particles – very different from the thermal upstream conditions in our (and almost all other) simulations. Injected into the reconnection layer near an X-point, these high energy particles may not be limited to Lorentz factors 4σ , and could indeed already exceed those Lorentz factors *before* even experiencing the reconnection electric field.

Such particles could, in principle, reach Lorentz factors all the way up to the Hillas limit $eE_{\text{rec}}/m_e c^2$ where $E_{\text{rec}} \simeq 0.1B_0$ is the reconnection electric field, and l is its potentially macroscopic coherence length. Importantly, the effective l should be larger for particles that are more energetic upon entering the reconnection region. Indeed, Werner et al. (2016) found that the characteristic $\gamma_X \simeq 4\sigma$ limit arises because cold thermal particles only experience direct linear acceleration in elementary current layers between the smallest-scale plasmoids. Before their energy can grow too large, these particles become magnetized and trapped inside small plasmoids flanking the elementary layer where they were originally accelerated. However, higher energy particles have much larger Larmor radii and therefore sample larger scale fields. They may potentially traverse multiple acceleration regions (spanning several elementary current layers) before finally becoming trapped inside a necessarily large (and therefore rare) plasmoid (cf. Cerutti et al. 2013).

These remarks motivate a future systematic study of the maximum Lorentz factors achievable by X-point acceleration, and of kinetic beaming, in the presence of alternative upstream conditions. However, when it comes to our astrophysical discussion below, we will simply assume that it is possible to achieve Lorentz factors $\gamma \sim$

γ_{rad} (i.e. $\gamma_{\text{rad}} \leq \gamma_X$) even if $\gamma_{\text{rad}} \gg \sigma$, and we will not require $\gamma_{\text{rad}} \leq 4\sigma$ as a necessary condition for kinetic beaming.

At this point, it is clear that the astrophysical relevance of our findings is predicated on whether they survive under a number of non-trivial generalizations (e.g. to larger systems, to more realistic upstream conditions, and even to 3D), most of which are beyond the scope of this work. With that in mind, one should read our specific astrophysical remarks below not as predictions made by a robust and fully-fledged theory, but as provocative inferences that can be made should the fundamental character of our findings be preserved in real astrophysical systems. It is to those inferences that we now turn. Focusing specifically on TeV blazar flares, we examine whether kinetic beaming – as understood within the simplified framework of this study – can reasonably explain the extreme variability observed in these events in a manner that is consistent with, and possibly constrains, blazar radiative environments.

5 RAPID TEV FLARES IN FSRQS

We now shift our discussion towards a concrete astrophysical application of our numerical results: rapid TeV blazar (specifically FSRQ) flares. Before our analysis, we provide some brief background concerning blazars.

Blazars comprise a class of active galactic nuclei (AGNs) with a relativistic jet pointed towards us. Their observed spectra are generally quite broad, extending from the radio band to the gamma-rays, and characterized by two non-thermal humps. In models where the emission is leptonic, the lower energy (optical/UV/X-ray) hump is thought to stem from synchrotron radiation and the higher energy (gamma-ray) component from the IC process, whereby soft ambient photons are upscattered by relativistic particles (Begelman, Fabian & Rees 2008; Böttcher et al. 2013; Madejski & Sikora 2016). The photons seeding IC scattering are typically supplied either by synchrotron emission from within the jet itself (synchrotron self-Compton, or SSC, models; e.g. Maraschi, Ghisellini & Celotti 1992; Bloom & Marscher 1996) or by various external sources (external IC models; e.g. Begelman et al. 1987; Melia & Konigl 1989; Sikora, Begelman & Rees 1994). The most common view is that internal synchrotron emission seeds Compton scattering in BL Lacs – blazars characterized by a lack of strong emission lines – whereas external photons dominate the ambient radiation field in FSRQs, which are more luminous, showing strong emission lines and thermal radiation attributed to an accretion disc (Tavecchio et al. 2011; Madejski & Sikora 2016).

Here, we will follow in the footsteps of many prior studies (e.g. Giannios et al. 2009; Nalewajko et al. 2011; Giannios 2013; Sironi et al. 2015; Petropoulou et al. 2016; Werner et al. 2018; Christie et al. 2019, 2020; Ortuño-Macías & Nalewajko 2020), positing relativistic magnetic reconnection as the driving mechanism behind blazar flares. Because it is difficult for a reconnection layer to sustain internal radiation energy density larger than the upstream (unreconnected) magnetic energy density (Beloborodov 2017; Christie et al. 2019), the most natural emission model for reconnection-powered Compton-dominated flares – for which the IC spectral component dominates the synchrotron emission – is external IC (e.g. Christie et al. 2020). Conveniently, the objects with the most extreme observed Compton dominance, FSRQs, are also those objects that come pre-equipped with rich external radiation environments.⁵ Not only is an externally

illuminated reconnection region precisely the set-up addressed in our simulations, but in the case of the first sub-hour TeV FSRQ flare ever observed – that from PKS 1222+21 on 2010 June 17 (Aleksić et al. 2011; Tanaka et al. 2011) – a strong case has already been made by Nalewajko et al. (2012) that kinetic beaming was at play. For these reasons, we will devote the main part of our analysis to understanding rapid TeV FSRQ flares, concentrating on the prototypical PKS 1222+21 outburst.

Two prominent sources of external background radiation in FSRQs like PKS 1222+21 are the broad-line region (BLR) and dusty torus (also called the hot dust region; HDR). The BLR contains gas subject to ionizing radiation from the AGN accretion disc, and it reprocesses this light into UV line emission (most prominently Ly α ; Tavecchio & Ghisellini 2008). The HDR is made of dust clouds radiantly heated by the AGN and producing thermal emission predominantly in the IR (Nenkova et al. 2008a,b).

As discussed by Nalewajko et al. (2012) (see also Aleksić et al. 2011; Tavecchio et al. 2011), the very high energy (VHE; $\gtrsim 0.1$ TeV) radiation detected from PKS 1222+21 must have been produced beyond the BLR, at least ~ 0.5 pc from the AGN. Otherwise, it would have been absorbed while traversing the intense BLR radiation fields. At that distance, the extremely rapid variability time-scale requires the VHE flare to be fed by an unrealistically high energy density packed into a small fraction of the jet’s cross-section. However, via kinetic beaming, magnetic reconnection can achieve the same variability time-scale in a much larger space (Cerutti et al. 2013, 2014a,b). This, combined with the highly collimated emitting particles (this study and, originally, Cerutti et al. 2012b), relaxes the necessary energy density, enabling the flare to be fuelled at the parsec-scale on a reasonable energy budget (the full details of this argument are presented by Nalewajko et al. 2012).

Our numerical results may be used to constrain further this general picture of kinetic beaming in reconnection-powered VHE FSRQ flares. Namely, one may stipulate that the putative beaming operates in the regime of strong radiative cooling (as defined in this work), and examine what new astrophysical insight may be derived from this requirement. Let us make this idea more quantitative. In our simulations, kinetic beaming was apparent in the late-time distribution of particles only for strong cooling ($\gamma_{\text{rad}} \lesssim 4\sigma$) and for emitting particle Lorentz factors γ_{emit} above the isotropization threshold γ_{iso} . In the case of particularly efficient radiation ($\gamma_{\text{rad}} = \sigma$), γ_{iso} was nearly an order of magnitude smaller than the radiative cut-off, well-approximated by γ_{rad} (see Figs 15 and 16 and surrounding discussion). Thus, a necessary condition for efficiently cooled kinetic beaming is that the emitting particles bear energies within a fairly narrow-band given by

$$\gamma_{\text{rad}} \geq \gamma_{\text{emit}} \geq \gamma_{\text{iso}} \sim \frac{\gamma_{\text{rad}}}{10}. \quad (11)$$

As already discussed (Section 4.5), despite that it was only in our simulations with $\gamma_{\text{rad}} \lesssim 4\sigma$ that we measured an appreciable range of beamed particle and photon energies, we refrain from employing $\gamma_{\text{rad}} \lesssim 4\sigma$ as a requirement for kinetic beaming in addition to (11). Rather, we assume that it is possible for X-points to accelerate particles up to the radiative limit γ_{rad} even if $\gamma_{\text{rad}} \gg \sigma$, provided some particles in the reconnection inflow already possess relatively high Lorentz factors. This could be the case, for example, if the upstream plasma is highly non-thermal, as may reasonably be expected from non-thermal quiescent blazar spectra.

We will now present simple estimates to check whether equation (11) is satisfied by an external IC model for the PKS 1222+21 VHE flare. Here, unprimed quantities are evaluated in the observer’s

⁵A notable exception is the BL Lac PKS 2155–304, which produced a Compton-dominated flare in 2006 (Abramowski et al. 2012).

frame and primed quantities in the frame of the VHE emitting region, which is the assumed rest frame of the reconnection layer. (However, we leave particle Lorentz factors γ unprimed, though they are always evaluated in the reconnection frame.) These frames are connected by the emitting region bulk Lorentz factor Γ . For simplicity, we ignore the source redshift $z \simeq 0.4$, and assume the angle θ_{obs} between the emitting region bulk velocity and the line of sight to be such that the Doppler factor $\delta = \{\Gamma[1 - (1 - 1/\Gamma^2)^{1/2}\cos\theta_{\text{obs}}]\}^{-1}$ is approximately equal to Γ . We follow Nalewajko et al. (2012), adopting $\Gamma = 40$ (sufficient to render external IC radiation more efficient than SSC) and a fiducial comoving (unreconnected) magnetic field strength $B'_0 = 0.1$ G typical at the parsec-scale.

Beyond the BLR, the likely dominant source of external photons illuminating the jet is the dusty torus. In the observer's frame, the torus radiation is approximately uniform and isotropic, with energy density $U_{\text{HDR}} \simeq 9 \times 10^{-5} \text{ erg cm}^{-3}$ and typical photon energy $\epsilon_{\text{HDR}} \simeq 0.3 \text{ eV}$ (Nenkova et al. 2008a,b; Sikora et al. 2009; Malmrose et al. 2011; Tavecchio et al. 2011; Nalewajko et al. 2012). Particles upscattering these photons to the characteristic observed energy $\epsilon_{\text{obs}} = 100 \text{ GeV}$ (Aleksić et al. 2011) have approximate Lorentz factors

$$\gamma_{\text{emit,HDR}} \sim \sqrt{\frac{\epsilon'_{\text{obs}}}{\epsilon'_{\text{HDR}}}} \sim \sqrt{\frac{\epsilon_{\text{obs}}/\Gamma}{\epsilon_{\text{HDR}}\Gamma}} \sim 1 \times 10^4 \left(\frac{\Gamma}{40}\right)^{-1} \left(\frac{\epsilon_{\text{obs}}}{100 \text{ GeV}}\right)^{1/2} \left(\frac{\epsilon_{\text{HDR}}}{0.3 \text{ eV}}\right)^{-1/2}. \quad (12)$$

We note that the Comptonization occurs in the marginal Klein–Nishina regime, since, in the rest frames of the scattering particles, the seed photon energies are close to the electron rest mass [cf. equation (1)]:

$$\frac{\gamma_{\text{emit,HDR}} \epsilon'_{\text{HDR}}}{m_e c^2} \sim \sqrt{\frac{\epsilon_{\text{obs}} \epsilon_{\text{HDR}}}{(m_e c^2)^2}} \sim 0.3 \left(\frac{\epsilon_{\text{obs}}}{100 \text{ GeV}}\right)^{1/2} \left(\frac{\epsilon_{\text{HDR}}}{0.3 \text{ eV}}\right)^{1/2}. \quad (13)$$

Klein–Nishina effects are even more important for hypothetical particles at the much-higher upper-limit Lorentz factor imposed by the HDR, which can be estimated via equation (3) as

$$\gamma_{\text{rad,HDR}} = \sqrt{\frac{0.3eB'_0}{4\sigma_T U'_{\text{HDR}}}} \sim \sqrt{\frac{0.3eB'_0}{4\sigma_T U_{\text{HDR}} \Gamma^2}} \sim 6 \times 10^6 \left(\frac{\Gamma}{40}\right)^{-1} \left(\frac{B'_0}{0.1 \text{ G}}\right)^{1/2} \left(\frac{U_{\text{HDR}}}{9 \times 10^{-5} \text{ erg cm}^{-3}}\right)^{-1/2}. \quad (14)$$

Even though our numerical study was confined to Thomson IC radiation, with our definition of γ_{rad} even relying on that fact, let us momentarily maintain equation (3) as a definition and suppose that our result (11) also holds in the deep Klein–Nishina regime. Then, because $\gamma_{\text{rad,HDR}}$ and $\gamma_{\text{emit,HDR}}$ are widely separated, the emitting particles are far below our expected isotropization threshold:

$$\gamma_{\text{emit}} \sim \frac{\gamma_{\text{rad,HDR}}}{400} \ll \frac{\gamma_{\text{rad,HDR}}}{10} \sim \gamma_{\text{iso,HDR}}. \quad (15)$$

Equation (15) suggests that the IC (HDR) process does not impose sufficient radiative losses for kinetic beaming to imprint itself upon the emitted photons. Rather, the radiating particles are expected to emit isotropically. Admittedly, the fact that the VHE photons are produced in the marginal Klein–Nishina regime challenges the applicability of our numerical results – and, indeed, we plan to study kinetic beaming using a fully Klein–Nishina Compton cross-section in a future work. However, we do not expect this to extend kinetic beaming to a broader range of particle energies. This is because

Klein–Nishina effects suppress radiative cooling, likely lengthening a particle's cooling time relative to its isotropization time.

We therefore see that, on the one hand, radiatively efficient kinetic beaming appears strained to fit into the picture of Compton-dominated flares seeded by dusty torus photons. On the other hand, kinetic beaming solves an important and challenging energy budget problem for parsec-scale FSRQ flares independently of the underlying radiative mechanism (Nalewajko et al. 2012). Rather than abandon the kinetic beaming framework, we submit that the new insight gleaned in this work – that kinetic beaming requires efficient radiative cooling to manifest itself observationally – hints that a more elaborate emission model may be appropriate.

Let us therefore conduct our analysis in the opposite direction. Rather than model the flare's radiative environment, testing afterward whether it is consistent with strongly cooled kinetic beaming, let us start by assuming that beaming and efficient radiation operate together and see what this implies about the background photon population. In that spirit, we consider the properties of a hypothetical radiation field, characterized by its (assumed narrowly distributed) photon energy ϵ_{ph} and energy density U_{ph} , that satisfies our main requirement $\gamma_{\text{emit,ph}} \geq \gamma_{\text{iso,ph}} \sim \gamma_{\text{rad,ph}}/10$ in equation (11). This requirement can be recast, using equations (3) and (12), as the following inequality involving ϵ'_{ph} , U'_{ph} , B'_0 , and ϵ'_{obs} :

$$10 \geq \frac{\gamma_{\text{rad,ph}}}{\gamma_{\text{emit,ph}}} \simeq \sqrt{\frac{9}{40} \frac{B'_0}{B_c} \frac{U_c}{U_{\text{ph}}} \frac{\epsilon'_{\text{ph}}}{\epsilon'_{\text{obs}}}}, \quad (16)$$

where $B_c = 8\pi e/3\sigma_T = 6.0 \times 10^{15} \text{ G}$ is the classical critical field and $U_c = B_c^2/8\pi$. For a fixed B'_0 and ϵ'_{obs} , saturation of this inequality defines a 1D space of radiation fields $U'_{\text{ph}}(\epsilon'_{\text{ph}}) \propto \epsilon'_{\text{ph}}$ for which $\gamma_{\text{emit,HDR}}$ is at the expected isotropization threshold. To pinpoint one candidate combination of ϵ'_{ph} and U'_{ph} , we require that $U'_{\text{ph}} \geq U'_{\text{HDR}}$, necessary for Comptonization of U_{ph} -photons to dominate those from the dusty torus, and implying $\gamma_{\text{rad,ph}} \leq \gamma_{\text{rad,HDR}}$. In turn, this yields a smallest permissible emitting particle Lorentz factor [via (14) and (16)] of $\gamma_{\text{emit,min}} \sim \gamma_{\text{rad,ph}}/10 \leq \gamma_{\text{rad,HDR}}/10 \sim 6 \times 10^5$. The corresponding Compton seed photons have characteristic energies $\epsilon'_{\text{ph}} \sim \epsilon'_{\text{obs}}/\gamma_{\text{emit,min}}^2 \geq 7 \times 10^{-3} \text{ eV}$, which are small enough that the IC emission takes place safely in the Thomson regime: $\gamma_{\text{emit,min}} \epsilon'_{\text{ph}}/m_e c^2 \sim 8 \times 10^{-3}$.

Now that we know what kind of seed photon population (i.e. combination of ϵ_{ph} and U_{ph}) is required for efficiently cooled kinetic beaming, we ask whether such a population can be realized in nature. As an affirmative plausibility argument, we briefly consider the possibility of a structured jet. However, since a detailed global flare model is beyond the scope of our present study, we discuss only a subset of the possible parameters.

In particular, we consider a spine-sheath configuration, where the transverse jet structure consists of two regions: a central, fast-moving spine surrounded by a slower moving sheath (Ghisellini et al. 2005; Sikora et al. 2016; Tavecchio & Ghisellini 2016; Tavecchio 2017). Photons produced in the sheath are blueshifted to the frame of the spine where they seed reconnection-powered Compton radiation.⁶ We suggest that the sheath emission mechanism is synchrotron – perhaps due to a simultaneous but less luminous reconnection event – but, to maintain a simple and general discussion, avoid explicitly invoking this fact. We will merely suppose that the spine, which contains the VHE-producing magnetic reconnection site, inherits the (fast) bulk Lorentz factor $\Gamma_{\text{spine}} = \Gamma = 40$, while the sheath moves at

⁶Strictly speaking, we could equally well consider an emitting blob plowing through an otherwise unstructured jet; the important part is the relative motion.




Figure 23. Radiative phase diagram for the PKS 1222+21 VHE flare. We project an otherwise high-dimensional parameter space on to the $U' - \epsilon'$ plane by fixing the parameters $\Gamma = \Gamma_{>} = 40$, $\Gamma_{<} = 10$, $B'_0 = 0.1$ G, and $\epsilon_{\text{obs}} = 100$ GeV. Individual radiation fields are assumed to be monochromatic, occupying single points. To the right of the dotted vertical line, VHE photons pair-produce with their seed population. This excludes broad-line region illumination (see text), which we illustrate by adopting a characteristic BLR radiation field: $U_{\text{BLR}} = 6 \times 10^{-3}$ erg cm $^{-3}$ and $\epsilon_{\text{BLR}} = 10$ eV (Sikora et al. 2009; Tavecchio et al. 2011; Nalewajko et al. 2012). Right of the dashed red line, particles scatter photons in the Klein–Nishina ($\epsilon' \gamma_{\text{rad}} \geq m_e c^2$) regime with suppressed efficiency. In contrast, the blue band indicates efficient but not unphysical ($\gamma_{\text{rad}} \geq \gamma_{\text{emit}} \geq \gamma_{\text{iso}} \sim \gamma_{\text{rad}}/10$) Thomson ($\epsilon' \gamma_{\text{rad}} < m_e c^2$) IC cooling, and its lower border is given by saturating inequality (16). Lying outside this band, HDR illumination is probably unable to mediate observable kinetic beaming. A spine-sheath radiation field is more viable, and is above the shaded grey zone, where fields more tenuous than U'_{HDR} yield potentially unobservable IC output.

more typical (slower) speeds: $\Gamma_{<} = 10$. The relative Lorentz factor between the two regions is $\Gamma_r \simeq \Gamma_{>}/2\Gamma_{<} = 2$. To generalize our prior convention, primed quantities continue to refer to the reconnection (i.e. spine) rest frame and unprimed quantities (save particle Lorentz factors γ) to the observer’s frame; we will not write anything down in the sheath frame itself.

In this set-up, the photon energy $\epsilon'_{\text{ph}} \sim 7 \times 10^{-3}$ eV corresponds to an observed seed photon energy $\epsilon_{\text{ph}} \sim \epsilon'_{\text{ph}} \Gamma_{<}/\Gamma_r \sim 0.03$ eV. This lies on the part of the broad-band spectral energy distribution presented by Tavecchio et al. (2011) attributed to the dusty torus (Malmrose et al. 2011; Tavecchio et al. 2011), and hence is consistent with observed spectral features. Let us see whether the inferred energy density U'_{ph} – which, again, is expected to be larger than U'_{HDR} in order for its Comptonization to dominate the flare – is also consistent with observations. To that end, we suppose the sheath luminosity peaks at $L_{\text{ph}} = 10^{46}$ erg s $^{-1}$, similar to that observed in the broad-band spectrum near ϵ_{ph} (Tavecchio et al. 2011; however, these data are not simultaneous with the VHE flare). The spine-frame seed photon energy density is then $U'_{\text{ph}} \sim \Gamma_r^2 L_{\text{ph}}/4\pi c \Gamma_{<}^2 R_{\text{sh}}^2 \sim 0.2$ erg cm $^{-3}$, and, importantly, exceeds $U'_{\text{HDR}} \sim \Gamma_{>}^2 U_{\text{HDR}} \sim 0.1$ erg cm $^{-3}$, as required. In this estimate, we have assumed that the transverse size R_{sh} of the sheath photon source is comparable to that of the VHE-emitting region, $R_{\text{sh}} \simeq R' \sim 10 c t_{\text{var}} \Gamma_{>} \simeq 2 \times 10^{-3}$ pc, implied by the TeV variability time-scale $t_{\text{var}} = 10$ min (Aleksić et al. 2011) and enlarged by a factor of 10 due to kinetic beaming (Nalewajko et al. 2012; cf.

Cerutti et al. 2012b; Ackermann et al. 2016). We have checked that the sheath thickness R_{sh} can be relaxed without substantial change to the model (reducing U'_{ph} much less severely than the naive expectation $U'_{\text{ph}} \propto R_{\text{sh}}^{-2}$).

Kinetic beaming is more viable in this spine-sheath model not so much because $U'_{\text{ph}} > U'_{\text{HDR}}$, but because $\epsilon'_{\text{ph}} \ll \epsilon'_{\text{HDR}}$. That is, the sheath photons appear much softer than the HDR photons in the spine frame, requiring higher energy particles for Comptonization to the VHE band – particles that are then quite strongly cooled. This is illustrated in Fig. 23, which presents the main results of this section through a radiative ‘phase diagram’. The blue band in the figure, with lower border given by saturating inequality (16), designates seed photon populations conducive to efficiently radiative kinetic beaming. The sheath photons occupy this band, and the HDR photons do not, mostly because of the large energetic disparity between the two populations. This is accentuated by the sheath motion, which enlarges the energy gap to $\epsilon'_{\text{HDR}}/\epsilon'_{\text{ph}} = (\Gamma_{>}\Gamma_{<}/\Gamma_r)\epsilon_{\text{HDR}}/\epsilon_{\text{ph}} = 200\epsilon_{\text{HDR}}/\epsilon_{\text{ph}}$ in the spine frame.

Hence, purely by invoking relative motion between a VHE-emitting spine and a seed-photon-emitting sheath, one may reconcile the seed photon population required by radiatively efficient kinetic beaming with one that may plausibly be realized during an actual flare. Although we leave a detailed model to future work, we view the above remarks as illustrating the potential utility of a kinetic beaming framework that includes the new ingredient of strong radiative cooling. Whereas kinetic beaming on its own has previously been used to balance the energy budget in VHE FSRQ flares (Nalewajko et al. 2012), we now see that the added radiative requirement may constrain possible emission mechanisms. Surprisingly, the most appropriate radiative model for rapid FSRQ flares may be one that does not rely on external structures at all, and, therefore, presents a potentially universal mechanism for the most rapid TeV flares in all blazars, even BL Lacs.

6 CONCLUSIONS

In this paper, we present the first systematic investigation of the role radiative cooling plays in the kinetic (energy-dependent) beaming of particles and their emission in collisionless relativistic magnetic reconnection. In agreement with prior studies (Cerutti et al. 2012b; Kagan et al. 2016), we measure definite and pronounced kinetic beaming during the early stages of all our simulations, independent of cooling strength (see Figs 13 and 15). When radiation is inefficient, kinetic beaming fades at later times to a nearly isotropic distribution of particles and emission, as anticipated (Kagan et al. 2016; Sironi et al. 2016; Yuan et al. 2016). However, as we demonstrate explicitly, kinetic beaming remains persistently observable when radiative cooling is strong (see again Figs 13 and 15), and may then extend across more than an order of magnitude in photon energies. Moreover, enhanced radiative efficiency increases the beamed range of photon energies (Figs 18 and 21). In every case, late-time kinetic beaming is apparent only when the emitting particles have energies that are moderately close to (within an order of magnitude of) the radiatively imposed cut-off γ_{rad} .

The underlying picture is a competition of time-scales: that over which the radiating particles cool and that over which they isotropize. As a generic side effect of impulsive X-point acceleration, particles are always initially beamed (Uzdensky et al. 2011; Cerutti et al. 2012a,b). However, in the limit of weak radiative losses, they produce most of their radiation after they have isotropized, and their initial collimation leaves no observable remnant. When

cooling is strong, the opposite situation occurs: high energy radiation comes only from beamed, recently accelerated particles. Reconnection focuses particles in both circumstances, but only in one is this focusing imprinted on the high energy emission, manifesting itself as rapid variability along a particular observer's line of sight.

Our findings have important consequences for rapid high energy (HE; $\sim$ GeV) and VHE ($\sim$ TeV) astrophysical flares. At a very general level, we predict that rapid flares observed in a given spectral band (e.g. VHE) should not exhibit similar variability at energies that are lower by more than a couple orders of magnitude (e.g. HE), barring counterparts produced by the same particles simultaneously shining via multiple radiative processes. At a more detailed level, a kinetic beaming framework may constrain emission models in specific flaring systems, and we examine TeV FSRQ flares as an example. Analysing the 2010 June 17 flare of PKS 1222+21, we find that a kinetic beaming origin of the rapid variability seems at odds with the picture of IC-scattered dusty torus photons dominating the TeV outburst (we do not consider models invoking Comptonization from inside the broad-line region, which are precluded by pair-production considerations; e.g. Aleksić et al. 2011). We postulate that an alternative spine-sheath model, wherein reconnection-energized particles upscatter photons originating in the outer fringes of the jet, may be viable. Importantly, this could potentially operate in both FSRQs and BL Lacs. Thus, our results hint that the same physical mechanism may underlie rapid TeV flares from all blazars, regardless of their class.

This study opens the door to a wide scope of future work. It remains to be seen, for example, how our numerical results regarding the effects of radiative cooling on kinetic beaming extend to: different magnetizations σ , 3D, larger systems, different guide field strengths, a non-thermal upstream plasma, and electron-ion reconnection. Indeed, our main astrophysical results depend on whether the fundamental picture of kinetic beaming described here survives in the presence of a number of more realistic physical setups, which future studies may test. Additionally, our astrophysical analysis motivates the incorporation of more exotic physics into future simulations. For example, as we have seen in our study of the 2010 June 17 flare of PKS 1222+21, if Comptonization of dusty torus photons were primarily responsible for the outburst, then the IC emission would have taken place in the marginal Klein–Nishina limit. Although we do not expect radiative cooling to be strong enough to mediate kinetic beaming in this case, a dedicated study of kinetic beaming with fully Klein–Nishina Compton cooling may yield surprising results. We intend to carry out such a study in the future.

Finally, we stress that, although we have specialized to rapid FSRQ flares, the general framework presented here of kinetic beaming mediated by radiative cooling and facilitating short time-scale variability is quite general; it requires only that relativistic magnetic reconnection take place in its radiative regime. As illustrated by our particular application to FSRQs, a kinetic beaming hypothesis places powerful constraints on a flaring system without introducing many free parameters, and, if the case of blazars is any indication, may help to refine our understanding of reconnection-powered outbursts from other types of astrophysical systems. There appears, then, great potential for the kinetic beaming paradigm – as has already been applied to pulsar wind nebulae (Cerutti et al. 2012b, 2013, 2014a,b) and now to blazars (see also Nalewajko et al. 2012; Zhdankin et al. 2020) – to find fruitful application in attempts to explain many of the most extreme and diversely sourced flares in astrophysics.

ACKNOWLEDGEMENTS

This work is supported by NASA, NSF, and the DOE, grant numbers NASA ATP NNX16AB28G, NASA ATP NNX17AK57G, NSF AST 1411879, NSF AST 1903335, and DE-SC0008409. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) Stampede2 at the Texas Advanced Computing Center through allocation TG-PHY140041.

DATA AVAILABILITY

The simulation data underlying this article were generated at the XSEDE/TACC Stampede2 supercomputer and are archived at the TACC/Ranch storage facility. As long as the data remain in the archive, they will be shared on reasonable request to the corresponding author.

REFERENCES

- Abdo A. A. et al., 2011, *Science*, 331, 739
 Abramowski A. et al., 2012, *ApJ*, 746, 151
 Acciari V. A. et al., 2009, *Science*, 325, 444
 Ackermann M. et al., 2016, *ApJ*, 824, L20
 Aharonian F. et al., 2007, *ApJ*, 664, L71
 Albert J. et al., 2007, *ApJ*, 669, 862
 Aleksić J. et al., 2011, *ApJ*, 730, L8
 Aleksić J. et al., 2014, *Science*, 346, 1080
 Astropy Collaboration, 2018, *AJ*, 156, 123
 Bai X.-N., Caprioli D., Sironi L., Spitkovsky A., 2015, *ApJ*, 809, 55
 Ball D., Sironi L., Özel F., 2018, *ApJ*, 862, 80
 Begelman M. C., Fabian A. C., Rees M. J., 2008, *MNRAS*, 384, L19
 Begelman M. C. et al., 1987, *ApJ*, 322, 650
 Beloborodov A. M., 2017, *ApJ*, 850, 141
 Bessho N., Bhattacharjee A., 2007, *Phys. Plasmas*, 14, 056503
 Bessho N., Bhattacharjee A., 2012, *ApJ*, 750, 129
 Blackman E. G., Field G. B., 1994, *Phys. Rev. Lett.*, 72, 494
 Bloom S. D., Marscher A. P., 1996, *ApJ*, 461, 657
 Blumenthal G. R., Gould R. J., 1970, *Rev. Mod. Phys.*, 42, 237
 Britto R. J., Bottacini E., Lott B., Razzaque S., Buson S., 2016, *ApJ*, 830, 162
 Buehler R. et al., 2012, *ApJ*, 749, 26
 Böttcher M., Reimer A., Sweeney K., Prakash A., 2013, *ApJ*, 768, 54
 Cerutti B., 2018, A Decade of Agile: Results, Challenges and Prospects of Gamma-Ray Astrophysics. In memory of Giovanni Fabrizio ‘Nanni’ Bignami, preprint ([arXiv:1811.09215](https://arxiv.org/abs/1811.09215))
 Cerutti B., Mortier J., Philippov A. A., 2016b, *MNRAS*, 463, L89
 Cerutti B., Philippov A., Parfrey K., Spitkovsky A., 2015, *MNRAS*, 448, 606
 Cerutti B., Philippov A. A., 2017, *A&A*, 607, A134
 Cerutti B., Philippov A. A., Spitkovsky A., 2016a, *MNRAS*, 457, 2401
 Cerutti B., Uzdensky D. A., Begelman M. C., 2012a, *ApJ*, 746, 148
 Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2012b, *ApJ*, 754, L33
 Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2013, *ApJ*, 770, 147
 Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2014a, *Phys. Plasmas*, 21, 056501
 Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2014b, *ApJ*, 782, 104
 Chen A. Y., Beloborodov A. M., 2014, *ApJ*, 795, L22
 Christie I. M., Petropoulou M., Sironi L., Giannios D., 2019, *MNRAS*, 482, 65
 Christie I. M., Petropoulou M., Sironi L., Giannios D., 2020, *MNRAS*, 492, 549
 Clausen-Brown E., Lyutikov M., 2012, *MNRAS*, 426, 1374
 Crinquand B., Cerutti B., Dubus G., 2019, *A&A*, 622, A161
 de Gouveia dal Pino E. M., Lazarian A., 2005, *A&A*, 441, 845

- de Gouveia Dal Pino E. M., Piovezan P., Kadowaki L., Kowal G., Lazarian A., 2010, *Highlights Astron.*, 15, 247
- Di Matteo T., 1998, *MNRAS*, 299, L15
- Drenkhahn G., Spruit H. C., 2002, *A&A*, 391, 1141
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Ghisellini G., Tavecchio F., Chiaberge M., 2005, *A&A*, 432, 401
- Giannios D., 2006, *A&A*, 457, 763
- Giannios D., 2013, *MNRAS*, 431, 355
- Giannios D., Spruit H. C., 2007, *A&A*, 469, 1
- Giannios D., Uzdensky D. A., 2019, *MNRAS*, 484, 1378
- Giannios D., Uzdensky D. A., Begelman M. C., 2009, *MNRAS*, 395, L29
- Giannios D., Uzdensky D. A., Begelman M. C., 2010, *MNRAS*, 402, 1649
- Goodman J., Uzdensky D., 2008, *ApJ*, 688, 555
- Guo F., Li H., Daughton W., Liu Y.-H., 2014, *Phys. Rev. Lett.*, 113, 155005
- Guo F., Liu Y.-H., Daughton W., Li H., 2015, *ApJ*, 806, 167
- Guo F., Li X., Daughton W., Kilian P., Li H., Liu Y.-H., Yan W., Ma D., 2019, *ApJ*, 879, L23
- Guo F. et al., 2016, *ApJ*, 818, L9
- H. E. S. S. Collaboration, 2010, *A&A*, 520, A83
- Hakobyan H., Petropoulou M., Spitkovsky A., Sironi L., 2020, preprint ([arXiv:2006.12530](https://arxiv.org/abs/2006.12530))
- Hakobyan H., Philippov A., Spitkovsky A., 2019, *ApJ*, 877, 53
- Hoshino M., Lyubarsky Y., 2012, *Space Sci. Rev.*, 173, 521
- Jaroschek C. H., Hoshino M., 2009, *Phys. Rev. Lett.*, 103, 075002
- Jaroschek C. H., Treumann R. A., Lesch H., Scholer M., 2004, *Phys. Plasmas*, 11, 1151
- Kadowaki L. H. S., de Gouveia Dal Pino E. M., Singh C. B., 2015, *ApJ*, 802, 113
- Kagan D., Nakar E., Piran T., 2016, *ApJ*, 826, 221
- Kirk J. G., Skjæraasen O., 2003, *ApJ*, 591, 366
- Lyubarskii Y. E., 1996, *A&A*, 311, 172
- Lyubarsky Y., Liverts M., 2008, *ApJ*, 682, 1436
- Lyubarsky Y. E., 2005, *MNRAS*, 358, 113
- Lytikov M., 2003, *MNRAS*, 346, 540
- Lytikov M., 2006a, *New J. Phys.*, 8, 119
- Lytikov M., 2006b, *MNRAS*, 367, 1594
- Lytikov M., Komissarov S., Sironi L., Porth O., 2018, *J. Plasma Phys.*, 84, 635840201
- Lytikov M., Uzdensky D., 2003, *ApJ*, 589, 893
- Madejski G., Sikora M., 2016, *ARA&A*, 54, 725
- MAGIC Collaboration, 2019, *A&A*, 623, A175
- Malmrose M. P., Marscher A. P., Jorstad S. G., Nikutta R., Elitzur M., 2011, *ApJ*, 732, 116
- Maraschi L., Ghisellini G., Celotti A., 1992, *ApJ*, 397, L5
- McKinney J. C., Uzdensky D. A., 2012, *MNRAS*, 419, 573
- Melia F., Konigl A., 1989, *ApJ*, 340, 162
- Melzani M., Walder R., Folini D., Winisdoerffer C., Favre J. M., 2014a, *A&A*, 570, A111
- Melzani M., Walder R., Folini D., Winisdoerffer C., Favre J. M., 2014b, *A&A*, 570, A112
- Most E. R., Philippov A. A., 2020, *ApJ*, 893, L6
- Nalewajko K., Begelman M. C., Cerutti B., Uzdensky D. A., Sikora M., 2012, *MNRAS*, 425, 2519
- Nalewajko K., Giannios D., Begelman M. C., Uzdensky D. A., Sikora M., 2011, *MNRAS*, 413, 333
- Nalewajko K., Yuan Y., Chruślińska M., 2018, *J. Plasma Phys.*, 84, 755840301
- Nenkova M., Sirocky M. M., Ivezić, Ž., Elitzur M., 2008a, *ApJ*, 685, 147
- Nenkova M., Sirocky M. M., Nikutta R., Ivezić, Ž., Elitzur M., 2008b, *ApJ*, 685, 160
- Ortuño-Macías J., Nalewajko K., 2020, *MNRAS*, 497, 1365
- Petropoulou M., Giannios D., Sironi L., 2016, *MNRAS*, 462, 3325
- Petropoulou M., Sironi L., 2018, *MNRAS*, 481, 5687
- Philippov A. A., Spitkovsky A., 2018, *ApJ*, 855, 94
- Pozdnyakov L. A., Sobol I. M., Syunyaev R. A., 1983, *Astrophys. Space Phys. Res.*, 2, 189
- Rees M. J., 1966, *Nature*, 211, 468
- Rybicki G. B., Lightman A. P., 1979, *Radiative Processes in Astrophysics*, New York, Wiley, 382
- Schoeffler K. M., Grismayer T., Uzdensky D., Fonseca R. A., Silva L. O., 2019, *ApJ*, 870, 49
- Sikora M., Begelman M. C., Rees M. J., 1994, *ApJ*, 421, 153
- Sikora M., Rutkowski M., Begelman M. C., 2016, *MNRAS*, 457, 1352
- Sikora M., Stawarz Ł., Moderski R., Nalewajko K., Madejski G. M., 2009, *ApJ*, 704, 38
- Singh C. B., de Gouveia Dal Pino E. M., Kadowaki L. H. S., 2015, *ApJ*, 799, L20
- Sironi L., Beloborodov A. M., 2020, *ApJ*, 899, 52
- Sironi L., Giannios D., Petropoulou M., 2016, *MNRAS*, 462, 48
- Sironi L., Petropoulou M., Giannios D., 2015, *MNRAS*, 450, 183
- Sironi L., Spitkovsky A., 2009, *ApJ*, 698, 1523
- Sironi L., Spitkovsky A., 2011, *ApJ*, 726, 75
- Sironi L., Spitkovsky A., 2014, *ApJ*, 783, L21
- Tamburini M., Pegoraro F., Di Piazza A., Keitel C. H., Macchi A., 2010, *New J. Phys.*, 12, 123005
- Tanaka Y. T. et al., 2011, *ApJ*, 733, 19
- Tavani M. et al., 2011, *Science*, 331, 736
- Tavecchio F., 2017, *AIP Conf. Proc.*, 1792, 020007
- Tavecchio F., Becerra-Gonzalez J., Ghisellini G., Stamerra A., Bonnoli G., Foschini L., Maraschi L., 2011, *A&A*, 534, A86
- Tavecchio F., Ghisellini G., 2008, *MNRAS*, 386, 945
- Tavecchio F., Ghisellini G., 2016, *MNRAS*, 456, 2374
- Uzdensky D. A., 2011, *Space Sci. Rev.*, 160, 45
- Uzdensky D. A., 2016, in Gonzalez W., Parker E., eds, *Astrophysics and Space Science Library*, Vol. 427, *Magnetic Reconnection: Concepts and Applications*, Springer International Publishing, Cham, p. 473
- Uzdensky D. A., Cerutti B., Begelman M. C., 2011, *ApJ*, 737, L40
- Uzdensky D. A., Goodman J., 2008, *ApJ*, 682, 608
- Uzdensky D. A., Spitkovsky A., 2014, *ApJ*, 780, 3
- Werner G. R., Philippov A. A., Uzdensky D. A., 2019, *MNRAS*, 482, L60
- Werner G. R., Uzdensky D. A., 2017, *ApJ*, 843, L27
- Werner G. R., Uzdensky D. A., Begelman M. C., Cerutti B., Nalewajko K., 2018, *MNRAS*, 473, 4840
- Werner G. R., Uzdensky D. A., Cerutti B., Nalewajko K., Begelman M. C., 2016, *ApJ*, 816, L8
- Yuan Y., Nalewajko K., Zrake J., East W. E., Blandford R. D., 2016, *ApJ*, 828, 92
- Zenitani S., Hoshino M., 2001, *ApJ*, 562, L63
- Zenitani S., Hoshino M., 2007, *ApJ*, 670, 702
- Zenitani S., Hoshino M., 2008, *ApJ*, 677, 530
- Zhang H., Li X., Guo F., Giannios D., 2018, *ApJ*, 862, L25
- Zhdankin V., Uzdensky D. A., Werner G. R., Begelman M. C., 2020, *MNRAS*, 493, 603

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