Simulations of Dynamic Relativistic Magnetospheres

Kyle Patrick Parfrey

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2012

ABSTRACT

Simulations of Dynamic Relativistic Magnetospheres

Kyle Parfrey

Neutron stars and black holes are generally surrounded by magnetospheres of highly conducting plasma in which the magnetic flux density is so high that hydrodynamic forces are irrelevant. In this vanishing-inertia—or ultra-relativistic—limit, magnetohydrodynamics becomes force-free electrodynamics, a system of equations comprising only the magnetic and electric fields, and in which the plasma response is effected by a nonlinear current density term. In this dissertation I describe a new pseudospectral simulation code, designed for studying the dynamic magnetospheres of compact objects. A detailed description of the code and several numerical test problems are given. I first apply the code to the aligned rotator problem, in which a star with a dipole magnetic field is set rotating about its magnetic axis. The solution evolves to a steady state, which is nearly ideal and dissipationless everywhere except in a current sheet, or magnetic field discontinuity, at the equator, into which electromagnetic energy flows and is dissipated.

Magnetars are believed to have twisted magnetospheres, due to internal magnetic evolution which deforms the crust, dragging the footpoints of external magnetic field lines. This twisting may be able to explain both magnetars' persistent hard X-ray emission and their energetic bursts and flares. Using the new code, I simulate the evolution of relativistic magnetospheres subjected

to slow twisting through large angles. The field lines expand outward, forming a strong current layer; eventually the configuration loses equilibrium and a dynamic rearrangement occurs, involving large-scale rapid magnetic reconnection and dissipation of the free energy of the twisted magnetic field. When the star is rotating, the magnetospheric twisting leads to a large increase in the stellar spin-down rate, which may take place on the long twisting timescale or in brief explosive events, depending on where the twisting is applied and the history of the system. One such explosive field-expansion and reconnection event may have been responsible for the 27 August 1998 giant flare from SGR 1900+14, and the coincident sudden increase in spin period, or "braking glitch."

The inner magnetospheres of relativistic compact objects are in strongly curved spacetimes. I describe the extension of the code to general-relativistic simulations, including the hypersurface foliation method and the 3 + 1 equations of force-free electrodynamics in curved, evolving spacetimes. A simple test problem for dynamical behavior in the Schwarzschild metric is presented, and the evolutions of the magnetospheres surrounding neutron stars and black holes, in vacuum and in force-free plasma, are compared.

(This page left intentionally blank.)

Contents

List of Figures

(This page left intentionally blank.)

List of Tables

(This page left intentionally blank.)

Acknowledgments

There are many people without whose advice, help, and support I would never have finished this dissertation. I am particularly grateful to my advisors, Andrei Beloborodov and Lam Hui, for their generosity with their time and endless patience.

I'm also indebted to Greg Bryan, for making time for many enlightening discussions about numerics and simulations, and for serving on my thesis and defense committees.

Thank you to Zoltan Haiman for serving on my thesis committee over the years, and for providing some welcome outside perspective.

Thanks also to Jules Halpern and Anatoloy Spitkovsky for taking the time to be on my defense committee.

I'm indebted to my friends, who have kept me sane throughout, especially Josh, Catherine, Kathryn, Cameron, Erika, Taka, and Lia.

Finally, thank you to my parents, Nollaig and Anne, for always encouraging my interests in physics and astronomy, no matter how far away they took me, and to Jessie who was always there for me.

Chapter 1

Introduction

Relativistic compact objects are responsible for a wide variety of disparate phenomena—from megahertz radio pulses to kiloparsec-scale jets to explosions with peak luminosities ten thousand times those of core-collapse supernovae—many of which appear to be directly related to the structure and behavior of magnetic fields in their highly conducting plasma magnetospheres. The dynamics of these magnetospheres poses a complex global nonlinear problem, only amenable to analytic solution in highly symmetric or otherwise restricted configurations. A thorough exploration of the steady-state configurations and time-dependent evolution of these systems requires a numerical approach, which is the subject of this dissertation.

The infinite-magnetization, or vanishing-particle-inertia, limit is the appropriate one for the magnetospheres of pulsars (Goldreich & Julian 1969; Contopoulos, Kazanas, & Fendt 1999; Gruzinov 2006; Spitkovsky 2006), and magnetars, whose persistent and transient high-energy emission may be due to the distortion, reconnection, and dissipation of strong magnetic fields (Thompson & Duncan 1995a; Lyutikov 2006; Beloborodov 2009a). In this limit, relativistic magnetohydro-

dynamics becomes "force-free electrodynamics," the standard tool for studying the extraction of rotational energy from black holes (Blandford & Znajek 1977; MacDonald & Thorne 1982), where the magnetic field is thought to be supplied by a conducting accretion disc. The natural self-collimation of electromagnetic fields make them attractive candidates for explaining relativistic jets in quasars and active galactic nuclei, whose high Lorentz factors suggest low baryon loading and electromagnetic dominance (Blandford 1976); similar Poynting jets may be responsible for gamma-ray bursts (Mészáros & Rees 1997). An argument can be made that all ultra-relativistic outflows are essentially electromagnetic, rather than gas dynamical (Blandford 2002).

These astrophysical objects are surrounded by plasma whose resistivity is so low that it may be considered to be effectively dissipationless or ideal almost everywhere. One must therefore explain how the observed radiation is produced, since dissipation is required to convert rotational or magnetic energy to either kinetic energy of accelerated particles or plasma thermal energy. The dissipation may occur in specific regions where the ideal plasma approximation is not valid, which may be localized volumes or extended sheet-like surfaces. In either case, one must explain the creation of these non-ideal plasma regions within the larger nearly ideal magnetosphere, and the initiation and regulation of the dissipative process must be consistent with the global configuration, which provides its boundary conditions.

Strong gravitational fields are present in the inner magnetospheres of both neutron stars and black holes, leading to a significant frame-dragging effect in rapidly rotating systems. The launching of relativistic jets, including the collimation and acceleration of the outflowing plasma, is usually ascribed to the dragging of magnetic field in the ergospheres of nearly maximally rotating black holes (e.g. Blandford & Znajek 1977).

In this dissertation we focus on relativistic magnetospheres in the context of isolated (non-accreting) neutron stars. In the following sections we outline the history and physics of pulsars

(in particular their gamma-ray emission, since, unlike the radio pulses, this is both a significant contribution to the overall energy budget and heavily influenced by the global magnetospheric configuration) and magnetars (concentrating on giant flares from soft gamma repeaters, and variable spin-down torque effects). Finally we present a short summary of each chapter.

1.1 Pulsars

1.1.1 Discovery

A source of regular pulsed radio emission, detected at 81.5 MHz, was reported by Hewish et al. (1968); using the pulse duration, they put an upper limit of about 5×10^3 km on the source size and conjectured that the emission may be related to radial pulsations of a white dwarf or neutron star. It was soon realized that the newly discovered sources were more likely to be rotating neutron stars, in particular because the near constancy of the pulse periods argued for the involvement of massive objects rather than tenuous plasma configurations, and that braking torques from a magnetosphere with $B \sim 10^{12}$ G would cause the rotation period to slowly increase over time (Gold 1968; Pacini 1968). The rotational energy loss rate can be calculated from measurements of the pulsar period and period derivative,

$$L = -I\Omega\dot{\Omega} = 4\pi I \frac{\dot{P}}{P^3} \approx 3.95 \times 10^{46} \frac{\dot{P}}{P^3} \text{ erg s}^{-1}, \tag{1.1}$$

where we estimate the moment of inertia to be $I = kMr_\star^2 \approx 10^{45}$ g cm^2, using $M = 1.4\ M_\odot$, $r_\star = 10^6$ cm, and $k = 1/2$ (between a solid sphere, $k = 2/5$, and a thin shell, $k = 2/3$). The spin-down luminosity of the Crab pulsar was found to agree closely with the power required by the surrounding pulsar wind nebula, and a model was proposed in which the neutron star has a

corotating magnetosphere out to the light cylinder, that cylindrical radius $R_{\rm LC}$ at which corotation with the stellar surface requires motion at the speed of light,

$$R_{\rm LC} = \frac{c}{\Omega}, \tag{1.2}$$

and beyond which outflowing relativistic gas would be forced out of corotation, exerting a torque on the star via the magnetosphere (Gold 1969).

Although these early studies recognized that the magnetosphere would be filled with plasma, the dominant model of the magnetic field structure and rate of energy loss would remain for many years the vacuum retarded dipole solution (Deutsch 1955). In this model the spin-down luminosity of the star is,

$$L_{\rm vac} = \frac{2}{3}\frac{\mu^2\Omega^4}{c^3}\sin^2\chi, \tag{1.3}$$

where μ is the star's magnetic dipole moment, Ω its angular velocity, and χ the angle between the spin and magnetic axes. This equation is still the standard for estimating a pulsar's magnetic field strength from the observed period and period derivative; the field at the magnetic equator is given by

$$B_{\rm eq} = 3.2 \times 10^{19}\sqrt{P\dot{P}}\ \ {\rm G}, \tag{1.4}$$

for a system in which the spin and magnetic axes are orthogonal ($\chi = \pi/2$). In vacuum, no energy is lost by a star rotating about its magnetic dipole axis (an "aligned rotator.")

1.1.2 Force-free plasma-filled magnetospheres

A consistent plasma-filled model for the magnetosphere of the aligned rotator was developed by Goldreich & Julian (1969). They argued that the radial electrostatic force, due to the electric field

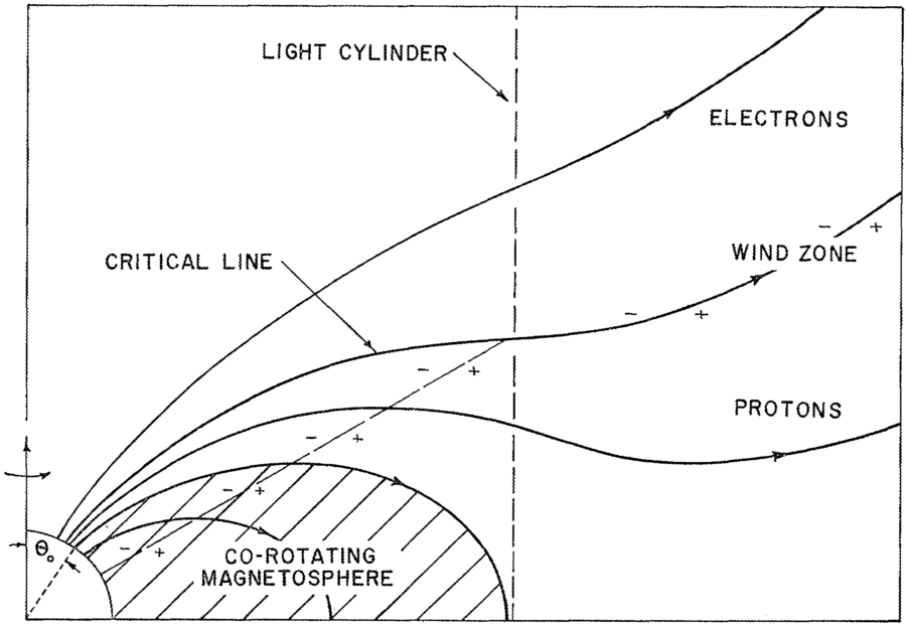

Figure 1.1: Schematic diagram of the force-free aligned rotator magnetosphere, showing open and closed field lines. From Goldreich & Julian (1969).

induced by rotation of the magnetosphere, would be orders of magnitude greater than the gravitational force on electrons and protons for realistic pulsar magnetic fields $B \sim 10^{12}$ G, and therefore that the particles would ripped from the surface, populating the surrounding magnetosphere with plasma. This conclusion has stood up to the later inclusion of effects due to the iron crystal lattice at the stellar surface. The magnetospheric plasma is in a "force-free" state, in the sense that the magnetic field strength is so high ($B \gtrsim 10^{12}$ G) that hydrodynamic forces can be neglected (i.e. the plasma's magnetization is approximately infinite), implying vanishing Lorentz force density.

In the Goldreich-Julian model, the magnetosphere is divided into two parts: the closed zone extending out to the light cylinder, and a bundle of open field lines, extending from the polar regions, through the light cylinder, and out to infinity; a schematic diagram is shown in Figure 1.1. The closed magnetosphere is in corotation which the star, $B_\phi = 0$ on all closed field lines. The open poloidal field lines are asymptotically radial at large distances from the star, in the "wind zone,"

where the toroidal field is dominant: $B_\phi \gg B_r \gg B_\theta$. Particles leaving the system along open field lines are most likely responsible for the observed pulsed radio emission, although the plasma processes that produce the low-frequency coherent radiation are not at all clear (see Ruderman & Sutherland 1975).

Unlike in the Deutsch vacuum model, a plasma-filled magnetosphere exerts a spin-down torque even when the spin and magnetic axes are aligned, due to the sweeping back of the open magnetic field lines beyond the light cylinder. Far out in the wind zone the poloidal magnetic field appears essentially monopolar,

$$B_\mathrm{p} \approx \frac{\Theta(\theta)}{r^2}, \tag{1.5}$$

with this approximation the total torque on the star can be estimated, by integrating the r–ϕ component of the Maxwell stress tensor over a spherical surface centered on the star,

$$T = \frac{\Omega}{c} \int_0^{\pi/2} \sin^3\theta \, [\Theta(\theta)]^2 \, \mathrm{d}\theta. \tag{1.6}$$

In the idealized Goldreich-Julian model the closed and open field lines are separated by an infinitely thin layer of infinite current density, or "current sheet." Another current sheet extends outward along the equator from the "Y-point" (where the closed zone touches the light cylinder), separating magnetic flux with radial component of opposite signs. In real pulsars, the high current densities present in these regions are expected to generate resistivity via collective plasma interactions and instabilities, which may be related to the observed high-energy emission (see e.g. Arons 2011).

It would be many years before a numerical construction of the Goldreich-Julian magnetosphere was achieved. The chief difficulty was ensuring continuity and smoothness of the solution across the light cylinder. The problem can be specified, in cylindrical coordinates (R, Z, ϕ), in what

has become known as the pulsar equation (Michel 1973; Scharlemann & Wagoner 1973),

$$(1-x^2)\left(\frac{\partial^2\Psi}{\partial x^2}-\frac{1}{x}\frac{\partial\Psi}{\partial x}+\frac{\partial^2\Psi}{\partial z^2}\right)-2x\frac{\partial\Psi}{\partial x}=-R_{\rm LC}^2A\frac{{\rm d}A}{{\rm d}\Psi}, \tag{1.7}$$

where Ψ is the flux function generating the poloidal magnetic field,

$$\boldsymbol{B}_{\rm p}=\frac{\nabla\Psi\times\hat{\boldsymbol{e}}_\phi}{R}, \tag{1.8}$$

A is the current function related to the toroidal magnetic field,

$$B_\phi=\frac{A(\Psi)}{R}, \tag{1.9}$$

$x=R/R_{\rm LC}$, and $z=Z/R_{\rm LC}$. This elliptic equation was solved by Contopoulos et al. (1999), who used an iterative relaxation procedure to find a solution that was approximately smooth through the light cylinder. More accurate solutions, with the Y-point located variously at and inside the light cylinder, were found by Timokhin (2006).

Evolving to a steady-state aligned rotator solution using time-dependent magnetohydrodynamic (MHD) simulations is complicated by two factors. The first is the difficulty of initially prescribing, and continuously regulating, a mass density field, which in a force-free problem is essentially undetermined. The second complicating factor is that a realistic solution requires a high plasma magnetization (ratio of magnetic to inertial energy densities) which is difficult to achieve accurately in MHD simulations. For these reasons dedicated electrodynamic simulation codes were developed, which solved the equations of force-free electrodynamics, the vanishing-inertia limit of MHD. This infinite-magnetization formulation does not include any mass density field; the charge and current densities are determined entirely by the demands of the electromagnetic

configuration.

Steady-state aligned rotator solutions were found using time-dependent electrodynamic simulations by Komissarov (2006); McKinney (2006b); Spitkovsky (2006); these were similar to the configurations generated by solving the elliptic equation. Furthermore, solutions to the non-aligned, or oblique, rotator problem were presented by Spitkovsky (2006), who found that the spin-down luminosity for force-free plasma-filled magnetospheres is given by

$$L_{\rm plasma} = \frac{\mu^2 \Omega^4}{c^3} \left(1 + \sin^2 \chi\right). \tag{1.10}$$

1.1.3 Pulsed gamma-ray emission

Gamma-ray pulses, at energies exceeding 35 MeV, were detected by the SAS-2 satellite (Fichtel et al. 1975); detections included the Crab (Kniffen et al. 1974) and Vela pulsars (Thompson et al. 1975). At these high energies, both sources display gamma-ray pulses at the same period as the radio pulses, a double peak structure with separation 0.4 P, and large pulsed fractions. The double pulses, compared to single pulses at radio wavelengths, suggest that the high-energy emission may be produced under different conditions or in different locations. Soft X-ray pulsations were detected from the gamma-ray source Geminga, indicating that it too is a gamma-ray pulsar (Halpern & Holt 1992); the gamma-ray pulses were then detected by the EGRET instrument on the Compton Gamma Ray Observatory (Bertsch et al. 1992). EGRET expanded the number of known gamma-ray pulsars to seven.

The largest expansion of the population of known high-energy pulsars came from the Fermi space telescope, which had 46 clear detections in the first six months of data taken with the LAT instrument (Abdo et al. 2010). Approximately 75% of these sources have gamma-ray pulse profiles with two peaks separated by $\gtrsim$ 0.2 P. Most interestingly, the gamma-ray conversion efficiency,

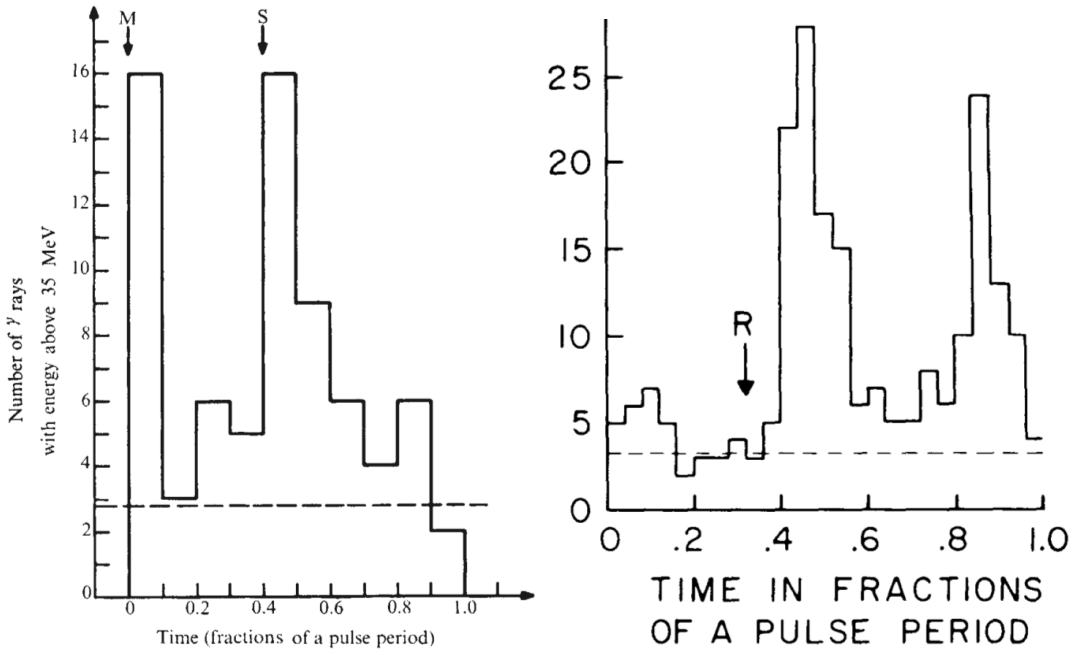

(a) Crab (b) Vela

Figure 1.2: Gamma-ray (> 35 MeV) pulses. (a) The Crab pulsar. "M" and "S" indicate the positions of the main and secondary radio pulses respectively. From Kniffen et al. (1974). (b) The Vela pulsar. The ordinate is number of detected gamma-rays, and "R" indicates the position of the radio pulse. From Thompson et al. (1975).

$L_\gamma/L_{\rm rot}$, varies from $\sim 0.1\%$ to order unity, meaning that high-energy emission can make a large contribution to the total energy budget (Figure 1.3).

There are three classical models for explaining pulsar high-energy emission. In each case the emission is linked to a particular region in which the electric field is not fully screened, allowing $E_\parallel$ to exist along the magnetic field lines, leading to particle acceleration. In such regions the force-free approximation must not apply, because force-free electrodynamics assumes $\boldsymbol{E}\cdot\boldsymbol{B} = 0$. In these models the gamma-rays either are curvature radiation from accelerated electrons (e.g. Sturrock 1971), or come from inverse-Compton up-scattering of thermal emission from the stellar surface, optical synchrotron radiation, or other soft photons (see Morini (1983) for an inverse-Compton model of the Vela high-energy emission). We present these models mainly for comparison to more

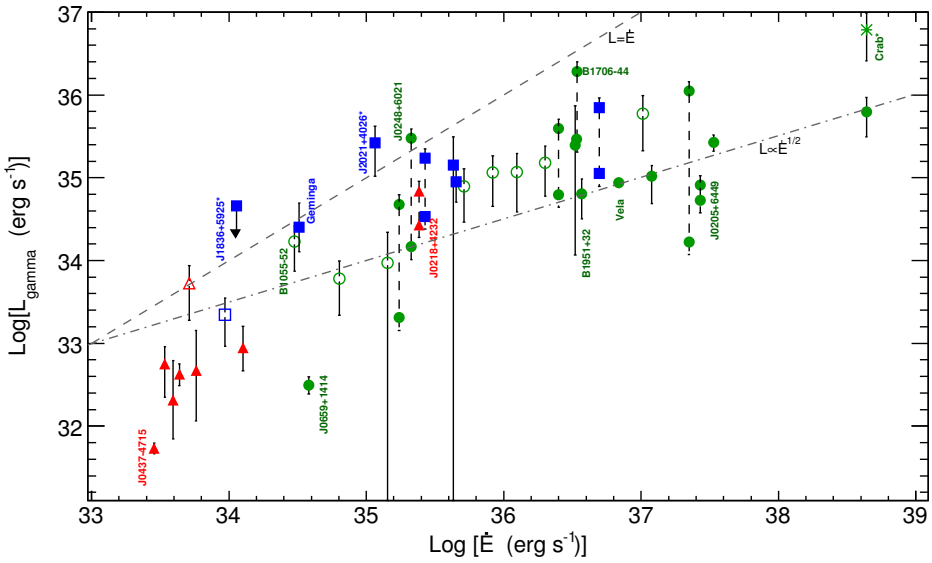

Figure 1.3: gamma-ray luminosity L_γ versus the rotational energy-loss rate $\dot{E}$ for the Fermi gamma-ray pulsars. The dot-dashed line indicates $L_\gamma \propto \sqrt{\dot{E}}$. Blue squares: gamma-ray selected pulsars; red triangles: millisecond gamma-ray pulsars; green circles: all other radio-loud gamma-ray pulsars. From Abdo et al. (2010).

recent work, because several of their underlying assumptions, such as the fine tuning of the current to the Goldreich-Julian current ($c\rho_{\rm GJ}$, see below), are no longer believed to hold (e.g. Arons 2007; Beloborodov 2008).

The polar cap accelerator model comes in two flavors: the vacuum gap (Ruderman & Sutherland 1975) and the space-charge-limited flow (SCLF) gap (Arons & Scharlemann 1979). In the vacuum gap, the stellar surface temperature T_s is smaller than the electron and ion surface thermionic emission temperatures $T_{e,i}$, and so the full vacuum electric field $E_\parallel \approx \Omega B_0 r_\star$ exists above the surface, causing pair production by vacuum breakdown. In the SCLF model, charges are "boiled off" from the surface and flow along open field lines, emitting curvature radiation which eventually pair-produces off the magnetic field at the pair formation front. Above the pair formation front there are enough charges to screen the parallel electric field, and acceleration ceases.

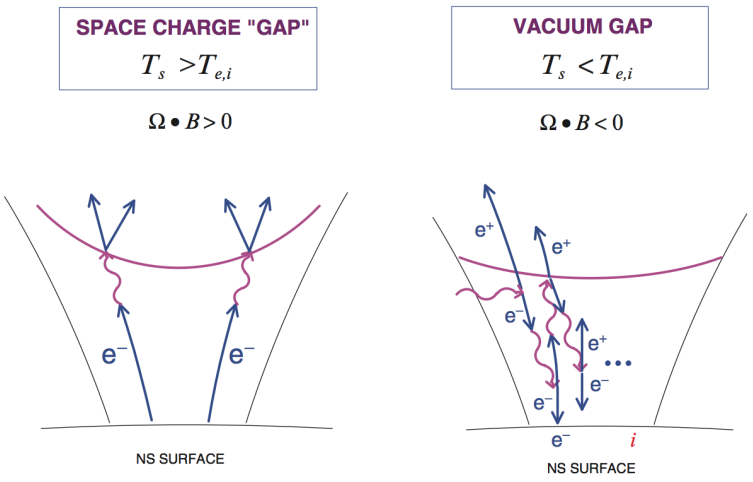

Figure 1.4: Vacuum and space-charge-limited flow polar cap models. From Harding (2009).

These polar cap models are illustrated in Figure 1.4.

The second emission model is the outer gap (Cheng et al. 1986), which is a region bordered by the closed zone and the $\boldsymbol{\Omega} \cdot \boldsymbol{B} = 0$ surface (Figure 1.5a). The "Goldreich-Julian" charge density required for steady-state corotation with $\boldsymbol{E} \cdot \boldsymbol{B} = 0$,

$$\rho_{\rm GJ} = \frac{\nabla \cdot \boldsymbol{E}}{4\pi} = -\frac{\nabla \cdot [(\Omega \times \boldsymbol{r}) \times \boldsymbol{B}]}{4\pi c} \propto -\frac{\boldsymbol{\Omega} \cdot \boldsymbol{B}}{2\pi c}, \tag{1.11}$$

changes sign at the $\boldsymbol{\Omega} \cdot \boldsymbol{B} = 0$ surface[1], and therefore charges from the surface cannot be brought past this surface in a charge-separated flow. If the star is the dominant charge source, unscreened $E_\parallel$ can build up in this outer gap when $\rho_{\rm e} \neq \rho_{\rm GJ}$, leading to large voltages $\sim 10^{15}$ V and vacuum breakdown.

The final classical emission site is the slot gap (Arons & Scharlemann 1979; Muslimov & Harding 2003). This region is an extension of the polar cap acceleration region, and lies just inside the last open field line (Figure 1.5b). The height of the pair formation front varies across the polar

[1] $\rho_{\rm GJ} = -\boldsymbol{\Omega} \cdot \boldsymbol{B} / (2\pi c)$ at $r = r_\star$, and this is a good approximation for $r \ll R_{\rm LC}$.

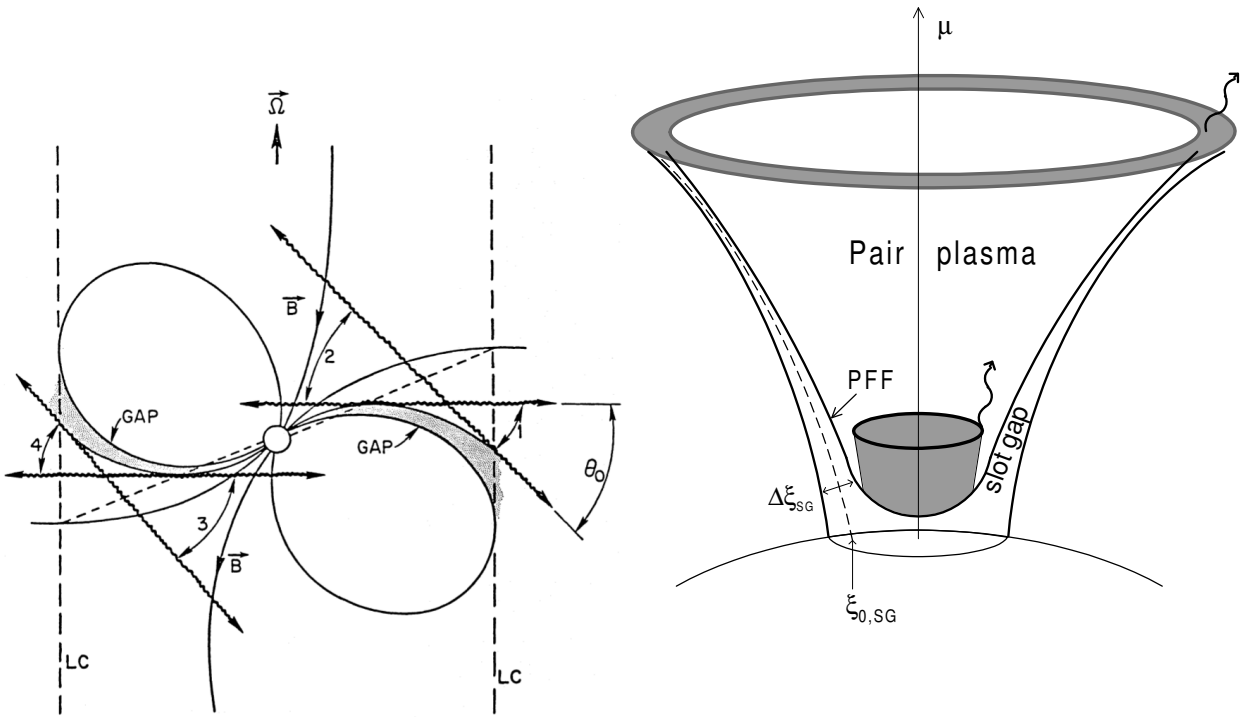

(a) Outer gap geometry

(b) Slot gap geometry

Figure 1.5: (a) The outer gaps (shaded) exist between the last closed field lines and the $\boldsymbol{\Omega} \cdot \boldsymbol{B} = 0$ surface (dashed line). Four trajectories of emitted curvature photons are shown. Note two outer gap regions are not indicated. From Cheng et al. (1986). (b) The slot gap lies in the polar cap region, just inside the last open field line. The pair formation front asymptotically approaches the last open field line. From Muslimov & Harding (2003).

cap because the $E_\parallel$ applied by rotation is not constant: it is strong well inside the polar cap and decreases to zero at the polar cap boundary since $E_\parallel = 0$ on closed field lines. When the electric field is weaker, electrons must be accelerated over a greater distance to achieve the Lorentz factor required to emit photons energetic enough to pair produce off the magnetic field; at the same time, the required Lorentz factor increases with increasing height because the magnetic field becomes weaker. Therefore the height above the surface at which pair formation occurs increases toward the edge of the polar cap, and the pair formation front asymptotically approaches the last open field line. Particles accelerated in this slot gap can reach very high Lorentz factors, limited to be $\lesssim 10^7$ by curvature-radiation reaction.

It is also possible that the gamma-rays are not tied to any gap-like region of unscreened electric field. Dissipation is expected to be significant in the strong current layers at the magnetic separatrix bounding the closed zone and extending beyond the light cylinder. In particular, magnetic reconnection in the outer current sheet will cause heating, pair production, and copious high-energy emission (Lyubarskii 1996). Estimates for the energy thus released by the Crab and Vela pulsars are consistent with observations. Energy will also be dissipated in the current sheets via Joule heating caused by any anomalous resistivity (Gruzinov 2007). The magnetic field goes through zero in the "strong" outer current sheets, and so force-free electrodynamics—and the $\boldsymbol{E} \cdot \boldsymbol{J} = 0$ restriction—does not locally apply.

Most of the work on fitting emission model predictions to observed light curves has used the rotating vacuum dipole magnetic field configuration. Polar cap models are disfavored because the observed spectrum does not display the expected super-exponential cutoff due to absorption in the strong field near the star, indicating that the emission is coming from farther out in the magnetosphere. Bai & Spitkovsky (2010) investigated the relationship between emission regions and gamma-ray light curves in the more realistic force-free geometry, by tracing gamma-rays intro-

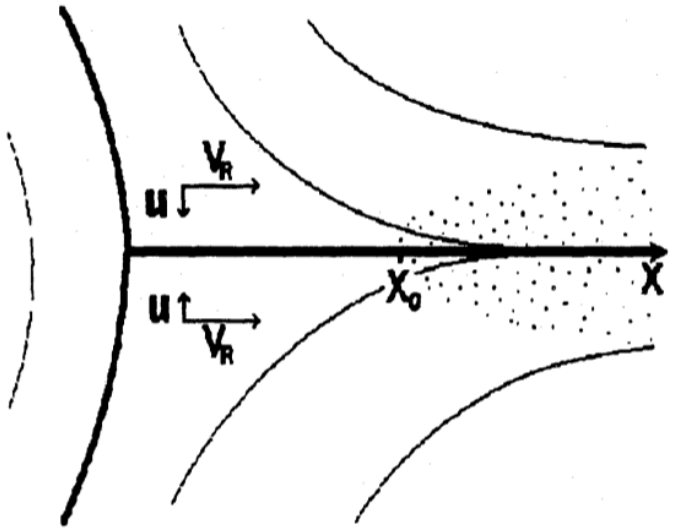

Figure 1.6: Plasma flow V_R close to the inner edge of the outer current sheet, showing magnetic flux surfaces (thin lines) and strong current layers (thick lines). The dotted area is the pair production region. Field lines are closed in the current sheet due to reconnection. The macroscopic Poynting flux is directed along magnetic field lines into the current sheet, where it is dissipated. From Lyubarskii (1996).

duced tangentially to magnetic field lines, which were generated by three-dimensional force-free oblique rotator simulations. They found that the slot and outer gaps do not generate the observed double-peak structure when embedded in a force-free magnetosphere, and proposed a separatrix layer model in which emission originates in the high-current regions as described above.

One can use force-free simulations to find the charge and current densities required at every point by the global problem, allowing the identification of regions which may allow particle acceleration. Simulations also reveal how finite dissipation, which should be present in the current sheets and possibly elsewhere, modifies the large-scale structure of the magnetosphere; this may provide observational signatures with which to compare emission models.

1.2 Magnetars

1.2.1 Soft gamma repeaters

On 5 March 1979, the Venera 11 and 12 spacecraft detected an extremely intense burst of hard X-rays from an "X-ray pulsar," later known as SGR 05266-66, in the Large Magellanic Cloud (Mazets et al. 1979). This unusual event comprised an initial hard spike of duration ~ 0.25 s and peak luminosity 3.6×10^{44} erg s^{-1}, and a softer tail, lasting roughly 200 s (Figure 1.7a). The rise time of the spike, $\sim$ 15 ms, is suggestive of relativistic motion over distances comparable to the neutron star radius, and the tail pulsations are at the neutron star spin period, 8.1 s. The total energy released in the burst was $\sim 5 \times 10^{44}$ erg, assuming isotropic emission. A day later, another burst, about an order of magntitude less energetic, was detected from the same source. Before these events no source of repeated gamma-ray bursts (GRBs) had even been observed. Shortly thereafter—on 24, 25, and 27 March 1979—three short soft gamma-ray bursts were detected from the source SGR 1900+14. These objects were soon known as soft gamma repeaters, or SGRs.

Two further giant flares have been recorded. The 27 August 1998 flare from SGR 1900+14 was detected by Konus-Wind (Mazets et al. 1999), Ulysses (Hurley et al. 1999), and BeppoSAX (Feroci et al. 1999). It was similar in may respects to the 1979 flare, with a hard ($kT \sim 240$ keV) spike of width ~ 0.35 s peaking at luminosity $> 8 \times 10^{44}$ erg s^{-1}, and long soft ($kT \sim 20$ keV) pulsating tail (Figure 1.7b). Just as in 1979, there was a subsequent smaller outburst, this time 47 hours later on 29 August (Ibrahim et al. 2001).

SGRs have been found to have variable spin-down rates which correlate with burst activity. A particularly dramatic example was coincident with the 27 August 1998 giant flare, during which the star experienced a "braking glitch," increasing its spin period by $\Delta P/P \sim 10^{-4}$ (Woods et al. 1999); see Figure 1.8. The spin-down rate appeared to be approximately constant both before and

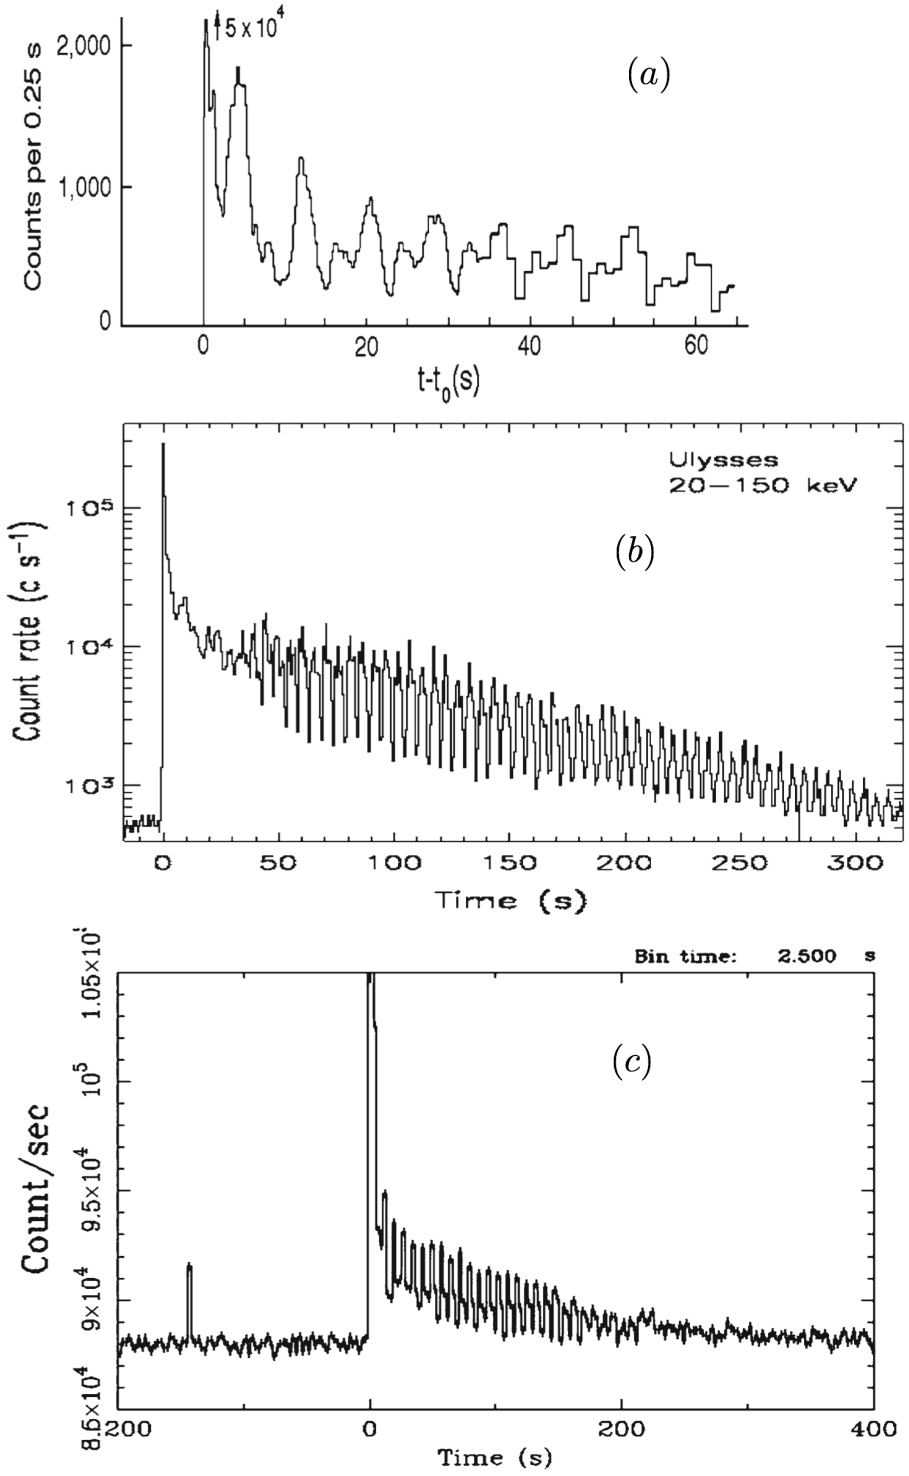

Figure 1.7: Giant flare light curves. (a) SGR 0526-66, 5 March 1979, Venera data at 50–150 keV; (b) SGR 1900+14, 27 August 1998, Ulysses data at 20–150 keV; (c) SGR 1806-20, 27 December 2004, INTEGRAL data at > 80 keV. The initial peaks in (a) and (c) are off the vertical scale. From Mereghetti (2008).

after the giant flare. Due to an 80-day gap in observations it is not clear by how large a factor $\dot{P}$ was enhanced—the average spin-down rate was 2.3 times larger over the observational gap, but may have been much larger if the torque increase was concentrated in a smaller time window (see e.g. Palmer 2002)—or if the accelerated spin-down occurred before or after the flare. Changes in spin-down rate are strongly indicative of large-scale reconfigurations of the magnetosphere, since magnetic braking is the dominant contribution to the spin-down torque. Also suggestive is the simplification of the X-ray pulse profile during the 1998 giant flare (e.g. Woods et al. 2001). As illustrated in Figure 1.9, the persistent X-ray emission changed, over a period of approximately 200 s, from a complicated pulse profile (observed for several years prior to the flare) to a nearly sinusoidal profile. The system of magnetospheric currents which generates the persistent X-ray flux appears to have relaxed, during the flare, to a smoother, less structured configuration.

The third, and to date final, giant flare was detected on 27 December 2004 by INTEGRAL (Mereghetti et al. 2005), Swift (Palmer et al. 2005), and other satellite observatories. The source, SGR 1806-20, has the distinction of being the first observed SGR, with a small burst on 7 January 1979. This event was about a hundred times more energetic than even the previous two giant flares, with a total flare energy of 2×10^{46} erg; most of the difference came from the more intense spike, with luminosity $\sim 2 \times 10^{47}$ erg s^{-1}, as the tail emission was similar. The rise time to the main spike was measured to be approximately 1.5 ms (with exponential time constant 0.3 ms). Additionally, a hard X-ray (> 80 keV) afterglow was detected, peaking about 700 s after the flare and lasting roughly one hour. The interpretation of this emission in terms of a relativistically expanding fireball is supported by VLA observations, which found, from 6 to 19 days after the flare, a resolved linearly polarized radio nebula expanding at a quarter the speed of light (Gaensler et al. 2005).

This SGR showed an increase in spin-down rate before the flare, from $\dot{\nu} = -1.48 \times 10^{12}$

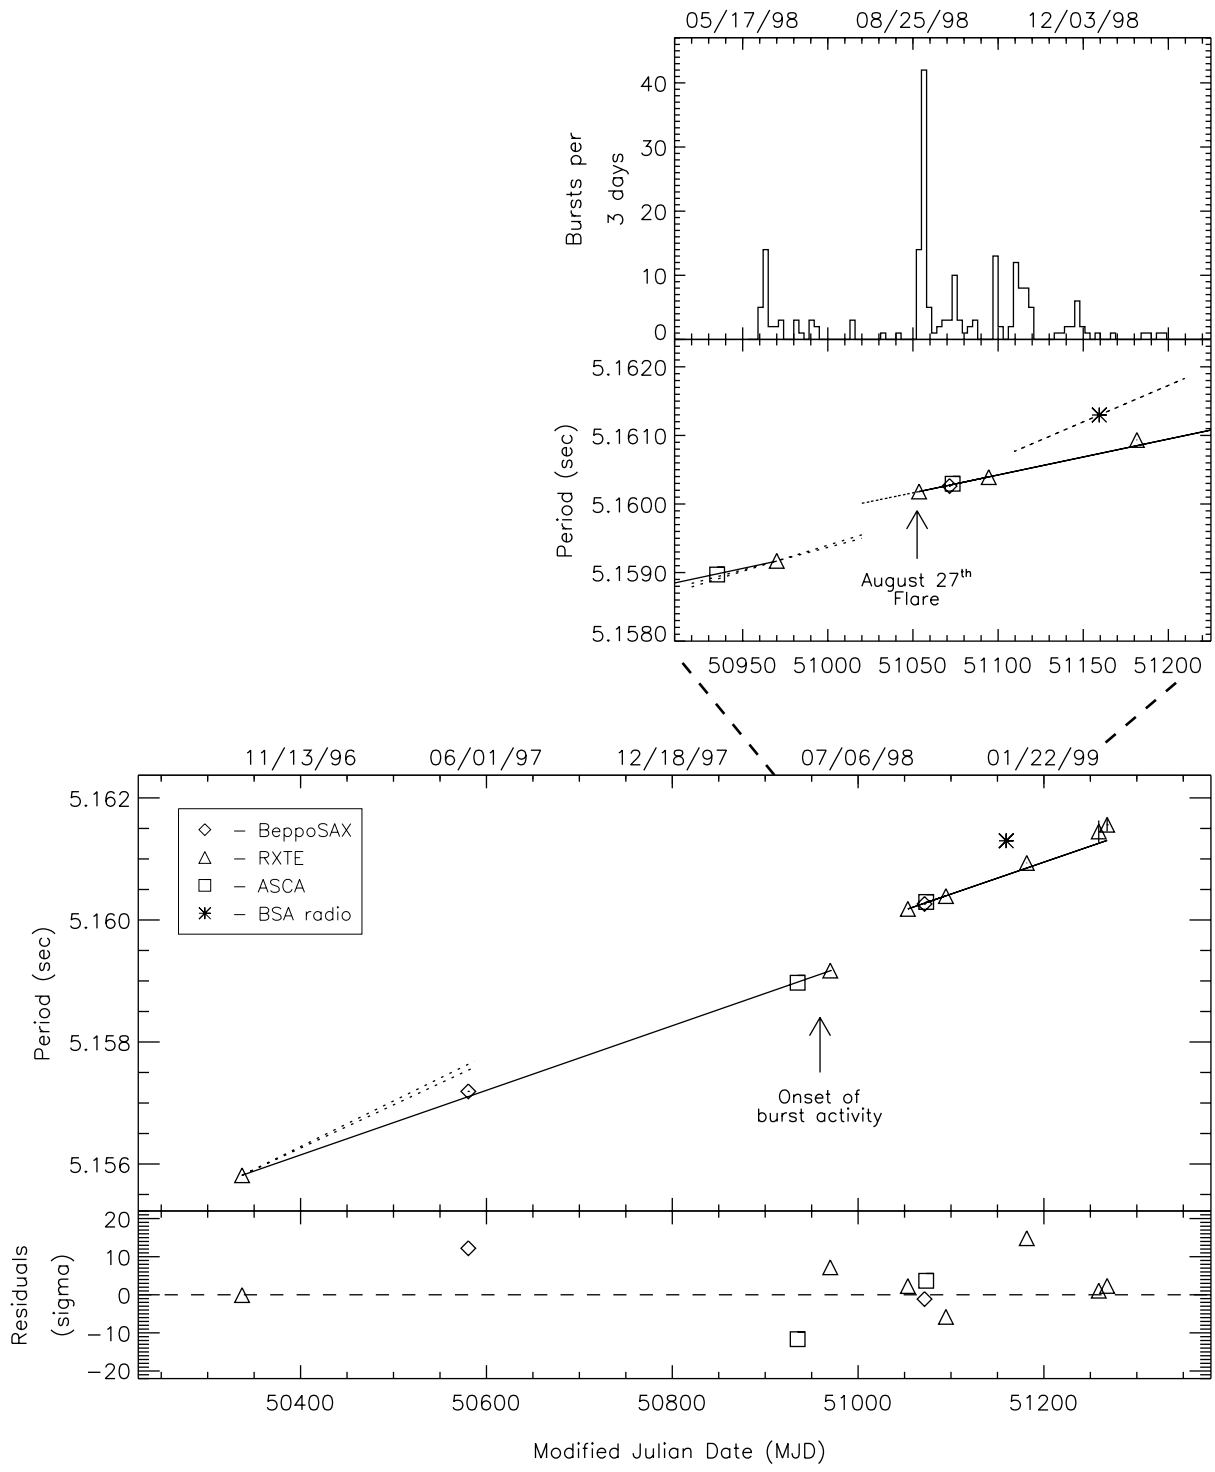

Figure 1.8: Spin period evolution and burst activity of SGR 1900+14, including the "braking glitch" coincident with the 27 August 1998 giant flare. From Woods et al. (1999).

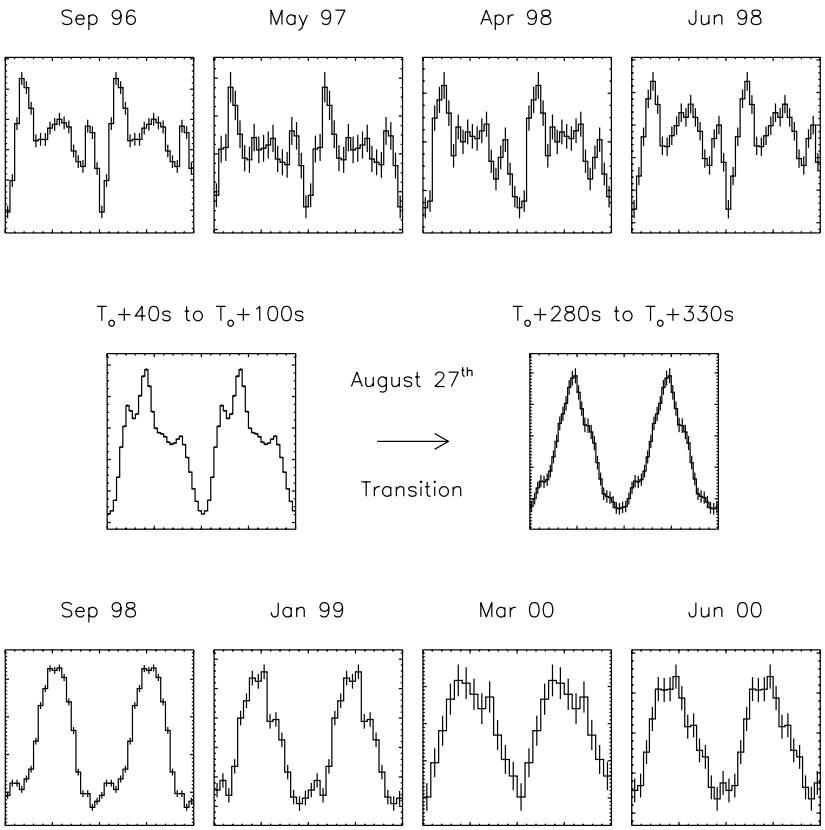

Figure 1.9: Simplification of the SGR 1900+14 X-ray pulse profile, during the 27 August 1998 giant flare. From Woods et al. (2001).

Hz s^{-1} before January 2000 to 5.87 times higher between January 2001 and April 2004 (Woods et al. 2007). The spin-down torque then decreased again just before the giant flare, as shown in Figure 1.10. The upper limit for $\Delta P/P$ over the flare was either 1.5×10^{-4} (Palmer et al. 2005) or 5×10^{-6} (Woods et al. 2007), similar to or less than the braking glitch seen in the 1998 giant flare despite the much greater flare energy. Unlike the 1998 event, the X-ray pulse profile was found to become less sinusoidal after the giant flare.

The 2004 giant flare was the most luminous event ever detected in our galaxy (the spike energy was about 0.5 $L_\odot$ Myr). Observed at cosmological separation, this event would have been very similar to a short-duration hard-spectrum GRB, meaning that at least some of the population

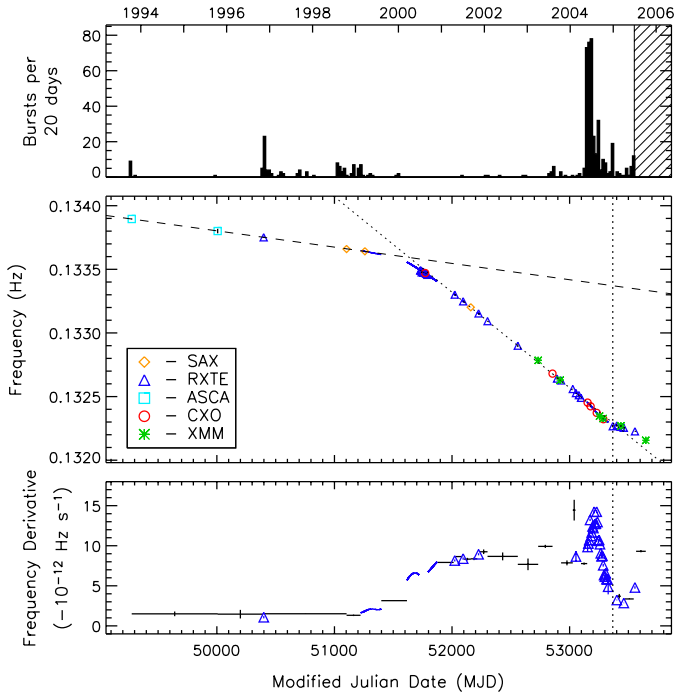

Figure 1.10: Spin-down rate evolution and burst activity of SGR 1806-20 leading up to the 27 December 2004 giant flare. From Woods et al. (2007).

of cosmological GRBs must come from SGR-like sources (Hurley et al. 2005).

The giant flares are not the only, or in some sense even the primary, manifestation of SGRs—the name refers to the much more common, and less energetic, soft repeated bursts. These have peak luminosities up to $\sim 10^{42}$ erg s^{-1}, and durations in the range 0.01–1 s, with a log-normal distribution peaking at 0.1 s. Many statistical properties of the repeated bursts are reminiscent of earthquakes: both have power-law energy distributions, positive correlations between waiting times of successive events, log-normal waiting time distributions, and weak or no correlations between waiting times and intensities (Cheng et al. 1996), hence the popular "starquake" model. The bursts are randomly distributed in rotational phase. SGRs go through episodes of high activity in which many bursts are observed (sometimes hundreds over periods of weeks), followed by quiescent phases which may last for years. No bursts have been observed from SGR 0526-66 since

1983. The incidence of soft bursts increases significantly around the times of giant flares (see the top panels of Figures 1.8 and 1.10). To date, SGR 1627-41 has only been detected via one six-week episode of more than 100 bursts.

1.2.2 Anomalous X-ray pulsars

Anomalous X-ray pulsars (AXPs) are so named because their "anomalous" X-ray luminosity is greater than the rotational energy loss rate inferred from spin-down measurements, Equation 1.1. The first such object was detected by Fahlman & Gregory (1981) in the supernova remnant G109.1–1.0. The X-ray luminosity of the central pulsating source, $\sim 2 \times 10^{35}$ erg s^{-1} in the 2–4 keV band, was in the typical range of accreting X-ray binaries, and the source, 1E 2259+586, was initially taken as such. However neither orbital modulation nor optical or IR counterpart was detected (Davies et al. 1989). The period and period derivative were measured to be 6.98 s and $5{\times}10^{-13}$ s s^{-1}, giving rotational spin-down power about three orders of magnitude less than the X-ray luminosity (Koyama et al. 1987). Soon several other similar objects were discovered, with a narrow 5.4–8.7 s range of pulse periods, and larger spin-down rates 7.32×10^{-13}–4.45×10^{-11} s s^{-1} (Mereghetti & Stella 1995).

The classic AXP characteristic is strong persistent soft X-ray emission. However, SGR-like transient behavior has also been observed from these sources. Gavriil et al. (2002) reported two X-ray bursts from AXP 1E 1048.1-5937, on 29 October and 14 November 2001, with fast rises (21 and 5.9 ms) and slow decays. The peak fluxes, and burst profiles and durations, were similar to those of soft SGR bursts, and the observation of two within about two weeks suggests episodic behavior. This AXP also generated two months-long events, in which the pulsed X-ray emission increased to 2.21 and 3 times the quiescent value over a period of weeks and then decayed on a longer timescale (Figure 1.11, lower panel). Roughly coincidentally, the star experienced a spin-

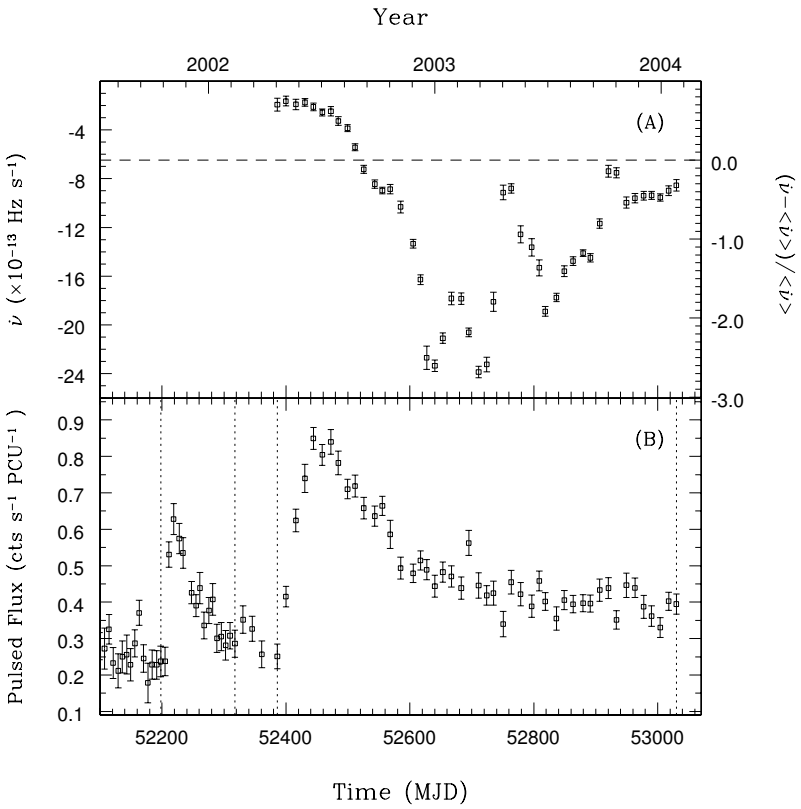

Figure 1.11: Variable spin-down and X-ray flare of AXP 1E 1048.1-5937 over several months. From Gavriil & Kaspi (2004).

down torque enhancement of up to a factor of 12, variable on timescales of weeks to months (Figure 1.11, upper panel).

An unmistakable major episode occurred in AXP 1E 2259+586, during which 80 X-ray bursts, each lasting from 2 ms to 3 s, were detected in four hours (Kaspi et al. 2003). The star also experienced a spin-up glitch of $\Delta\nu/\nu = 4 \times 10^{-6}$, followed by a factor of two increase in spin-down rate lasting for more than 18 days.

XTE J1810-197 ushered in the category of transient AXPs. It was discovered in outburst on 23 January 2003 (Ibrahim et al. 2004), at a flux $\sim$ 100 times greater than its quiescent level (later

recovered in archival data). Since this time its X-ray flux has been declining steadily back to its pre-outburst value (Figure 1.12). Shorter bursts were observed during the larger outburst phase (Woods et al. 2005), and the persistent X-ray emission was similar to that of other AXPs. On 17 March 2006, radio pulsations were detected at the 5.54 s X-ray pulse frequency, making this source the first radio-loud magnetar candidate (Camilo et al. 2006). This emission was not present before the outburst began. Radio timing allowed for more accurate measurements of the spin-down rate, which was found to be decreasing steadily even several years after the onset of the outburst phase (Camilo et al. 2007); see Figure 1.13. Gotthelf & Halpern (2007) found that the spectrum can be modeled by two blackbody components. The cooler component has been decreasing in temperature as its emitting area has expanded to almost the entire stellar surface, possibly due to the spreading of deep crustal heating (Lyubarsky et al. 2002). The effective emitting area of the hotter component has been reduced over time, by at least a factor of 8; Beloborodov (2009b) has attributed this to the shrinking of a hot spot at the base of a high-current field line bundle.

AXPs have soft persistent X-ray spectra in the 0.5–10 keV band which can be described by a combined blackbody-plus-power-law model. The softness of the spectra led to predictions that these sources would not be detected above 10 keV. Contrary to expectations, hard X-ray tails have been detected from 7 magnetars up to at least 150 keV, including several AXPs (Kuiper et al. 2008; Enoto et al. 2010). As shown in Figure 1.14, the measured spectrum has a double-peaked shape, with evidence for a spectral break at high energies. This very high-energy non-thermal emission points toward a magnetospheric origin, as will be described below.

1.2.3 Unification

Despite their many differences, SGRs and AXPs share many general characteristics. The initial separation into bursting and persistent categories has weakened over time, as persistent X-ray

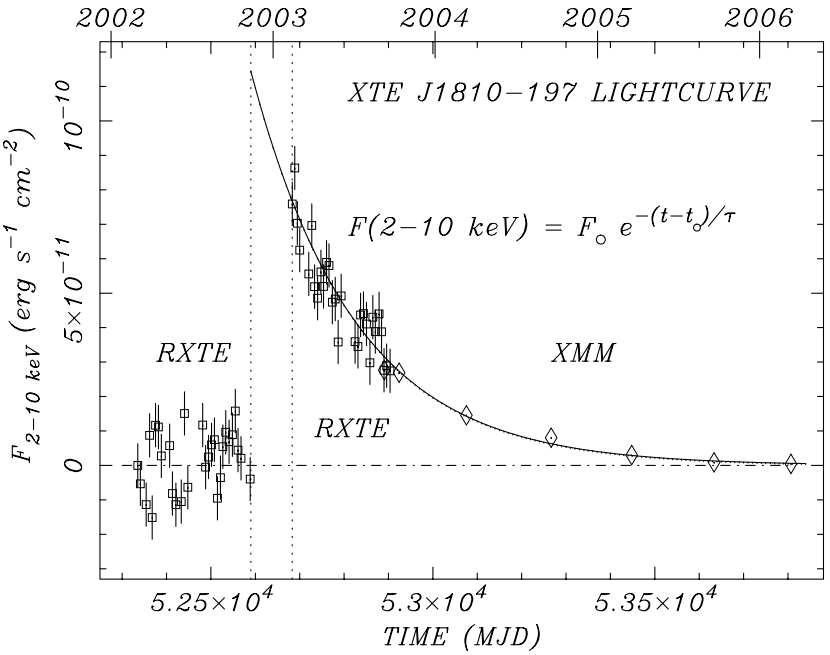

Figure 1.12: XTE 1810-197 light curve following the January 2003 outburst. From Gotthelf & Halpern (2007).

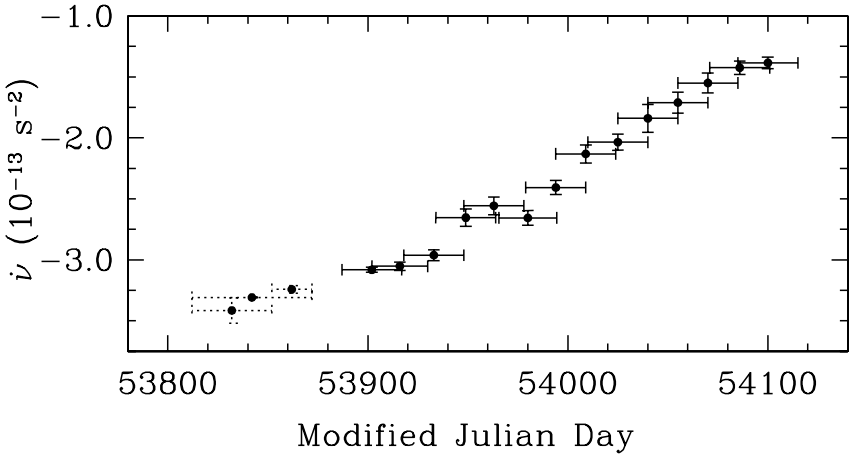

Figure 1.13: XTE 1810-197 spin-down rate during the slow return to the quiescent state, from radio observations. Note that MJD 53800 is roughly March 2006. From Camilo et al. (2007).

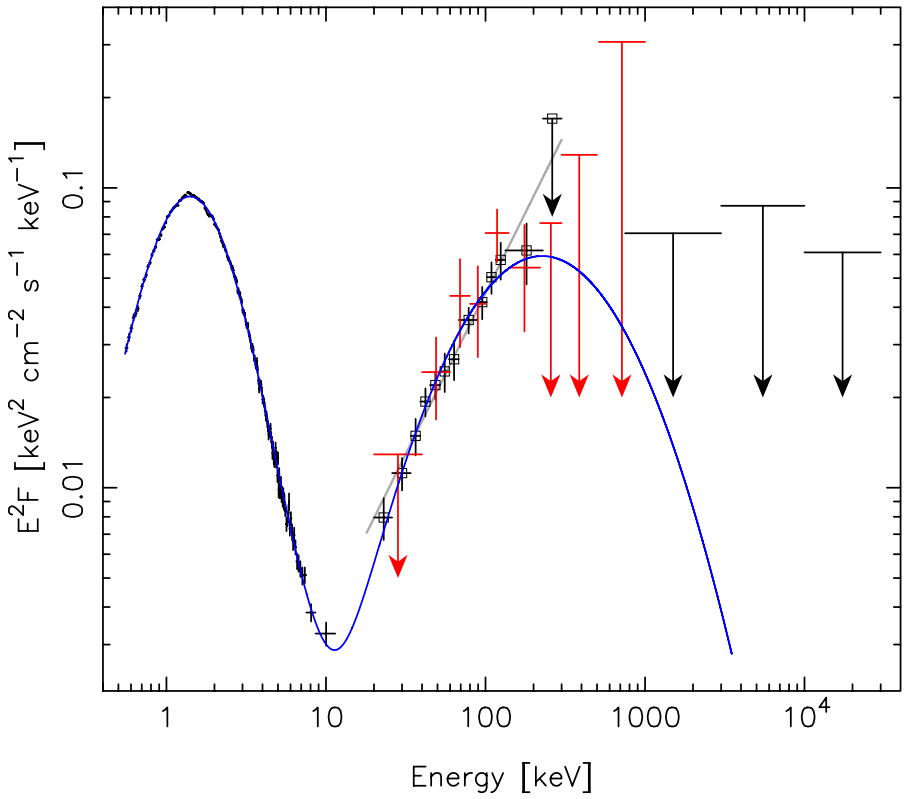

Figure 1.14: The X-ray spectrum of 4U 0142+61 from 0.5 keV to 30 MeV, combining INTEGRAL data (red) and XMM Newton data (black). The highest-energy upper limits are from the Compton Observatory's COMPTEL. The blue line is a best-fit double log-parabola model. From Kuiper et al. (2008).

emission from SGRs, and the AXP outbursts described above, were detected. Taken together, they form a class of objects with periods $P \sim 2$–12 s and spin-down rates $\dot{P} \sim 10^{-13}$–10^{-10} s s^{-1}, giving inferred magnetic fields of $B \sim 10^{14}$–10^{15} G, much higher than the $\sim 10^{12}$ G of normal radio pulsars. They are concentrated in the galactic plane and are sometimes associated with supernova remnants. For these reasons the most convincing current model for these objects is that they are "magnetars"—young neutron stars whose activity is powered by the decay of ultra-strong magnetic fields.

Newly formed neutron stars should undergo vigorous convection during their first 10–30 s

(Burrows & Lattimer 1986; Burrows 1987; Thompson & Duncan 1993), which when coupled to rotation leads to amplification of magnetic fields. Proto-neutron stars with initial spin periods $\sim$ 1 ms and significant differential rotation can support an efficient α–Ω dynamo, in principle generating magnetic fields $\gtrsim 10^{15}$ G. This process becomes much less efficient at longer spin periods, possibly explaining the approximate dichotomy between radio pulsar and magnetar field strengths. Strong magnetic fields cause rapid magnetic braking—spin-down timescales on the order of hours,

$$\tau_{\rm SD} \approx 0.6 \left(\frac{10^{15}\ {\rm G}}{B_{\rm dipole}} \right)^2 \left(\frac{P}{1\ {\rm ms}} \right)^2\ {\rm hr}, \tag{1.12}$$

imply that much of the rotational energy will be dumped into the magnetosphere before supernova shock breakout, giving brighter supernovae and slower remnant stars, which would rapidly spin down to the observed second-scale periods. The magnetar idea was first applied to SGRs, when Duncan & Thompson (1992) proposed that the 1979 giant flare was powered by the $\sim 6 \times 10^{14}$ G field of SGR 0526-66 (estimated using the 8 s rotation period and equating the spin-down time to the age $\sim 10^4$ yr of the surrounding supernova remnant N49).

A 10^{15} G magnetic field has enough total energy to power many giant flares, but explaining how that energy is released from an approximately perfectly conducting star frozen into MHD equilibrium is not straightforward. Goldreich & Reisenegger (1992) explored three mechanisms of magnetic field evolution and decay in neutron stars: Ohmic decay, ambipolar diffusion, and Hall drift. Ohmic dissipation is, by itself, too slow to affect the global magnetic field on timescales of interest. Hall drift is not dissipative in itself, but it can enhance the Ohmic loss rate by promoting a turbulent cascade, driving magnetic energy to smaller scales where resistive dissipation is stronger. This process operates mostly in the outer crust. Ambipolar diffusion may aid the transport of magnetic energy to the dissipation region from deep in the star. Magnetic field evolution due to these non-ideal effects will cause $\boldsymbol{J} \times \boldsymbol{B}$ stresses to build up in the crust, leading either to gradual

plastic deformation of the surface or sudden fracturing starquakes, depending on the crust's tensile strength and other properties like the melt temperature.

The interiors of rotationally powered radio pulsars are filled with superconducting protons and superfluid neutrons, forming a web of interacting quantized flux and vortex tubes, and leading to distinctive long-term magnetic field evolution (e.g. Ruderman et al. 1998). The flux tubes have magnetic fields $\sim 10^{15}$ G; the protons will not become superconducting if the locally averaged magnetic field in the star is much larger than this, and so magnetars are not expected to have superconducting cores (since the total core flux density will be much stronger than the dipole component estimated from spin-down). Reconfiguration of the magnetic field may occur when parts of the star transition into the superconducting phase.

The magnetic flux densities in magnetars' magnetospheres will usually exceed the quantum electrodynamic (QED) flux density,

$$B_{\rm QED} = \frac{m_e^2 c^3}{\hbar e} = 4.4 \times 10^{13} \text{ G}. \tag{1.13}$$

At this strength the energy of the first electron Landau level is comparable to the electron rest mass energy. Exotic QED effects occur when $B \gtrsim B_{\rm QED}$, such as photon splitting and merging, vacuum birefringence, and vacuum polarization, which complicate detailed models of radiative transport in magnetar magnetospheres.

Thompson & Duncan (1995a) proposed just such a detailed radiative model of SGR activity, in which the giant flares are due to a large-scale magnetic interchange instability whose growth time is comparable to the ~ 0.2 s duration of the hard spike. The fast rise time is due to magnetospheric reconnection triggered by the hydromagnetic rearrangement of the star, and the hard spike emission comes from a relativistic outflow. They argue that the timescale on which pairs and photons

are dumped into the emission region must be similar to the observed spike duration; baryon contamination increases the scattering depth and hence prolongs the observed burst, but it also greatly increases the required initial energy input to offset adiabatic cooling losses. This would rule out purely magnetospheric reconnection models, in which the characteristic timescale would be expected to be roughly the light-crossing timescale of the inner magnetosphere $\sim 10\, r_\star/c \sim 3 \times 10^{-4}$ s. The strong magnetic field holds a "trapped fireball" of radiating pair plasma which is responsible for the flare's soft tail. In this model the more common soft SGR bursts are caused by smaller local starquake-like crustal cracking motions.

In a similar vein, SGR 1900+14's "anti-glitch" can be explained by a sudden unpinning of the superfluid neutron vortex tubes, triggered by a sudden crustal fracture (Thompson et al. 2000). This explanation requires the superfluid to spin slower than the crust, opposite to what is inferred from radio pulsars, which experience spin-up glitches. A spin-up glitch was also observed in, for example, AXP 1E 2259+586 as mentioned above.

Magnetic field evolution, besides possibly triggering starquakes, also deforms the magnetosphere. Such a "twisted" magnetosphere is supported by currents of relativistic charges. Magnetars' persistent hard X-ray emission may be explained by resonant cyclotron scattering, of softer thermal photons from the stellar surface, off these charges, and by amplified crustal heating at the bases of the twisted flux bundles where the relativistic charge carriers hit the solid surface.

The $e^{\pm}$ pairs which constitute the magnetospheric currents are created with Lorentz factors $\gamma \sim 10^2$–10^3, by vacuum discharge near the star. These can resonantly up-scatter soft thermal photons, with energy $E_{\rm therm}$, when the photon energy in the electron rest frame matches the electron Landau energy,

$$\gamma\,(1 - \beta\cos\theta)\,E_{\rm therm} = \frac{B}{B_{\rm QED}} m_{\rm e} c^2, \tag{1.14}$$

which occurs when

$$\gamma \approx 10^3 B_{15} \left(\frac{10 \text{ keV}}{E_{\text{therm}}} \right), \tag{1.15}$$

where $B_{15} = B/(10^{15}\,\text{G})$; see Beloborodov (2011) for a review. The scattering boosts the photon energy by a factor of roughly γ^2, $E_{\text{scatt}} \sim \gamma^2 E_{\text{therm}}$, producing hard X-rays.

Pairs are created near the star with large Lorentz factors, and decelerate as they move outward due to resonant-scattering losses. At small radial distances the scattered photons are at very high energies, $E_{\text{scatt}} > 1$ MeV, and are immediately converted into pairs by the strong magnetic field (see Beloborodov & Thompson 2007). At larger radii the electrons have been decelerated to lower γ, and so the scattered photons have energies below the pair-production theshold and can escape, explaining the peak in the hard X-ray spectrum $\lesssim 1$ MeV (Beloborodov 2012).

The twists create a quasi-static hot corona above the star and decay on a timescale $\sim$ 1–10 years (Beloborodov & Thompson 2007). Twisting increases the magnetic energy stored in the magnetosphere, and causes it to be weighted more heavily to larger heights above the surface, which increases the spin-down torque (Thompson et al. 2002). The current configuration partly determines the resonant optical depth (e.g. Viganò et al. 2011) and X-ray spectrum (Fernández & Thompson 2007; Beloborodov 2012; Viganò et al. 2012).

An alternative model for SGR giant flares involves only slow plastic deformation of the neutron star's surface, rather than a sudden global starquake (e.g. Lyutikov 2006). In this model, similar to the standard picture for solar flares and coronal mass ejections, all of the flare's energy is stored in the twisted magnetosphere and released when the magnetic configuration reaches a point of dynamic instability, beyond which no equilibrium can be maintained. The unstable flux bundle expands outward relativistically while still attached to the star, producing a partially collimated outflow (Figure 1.15). As before, the shortest timescales are associated with relativistic wave-crossing times and with the onset of reconnection in high-current regions, but now the ~ 0.25

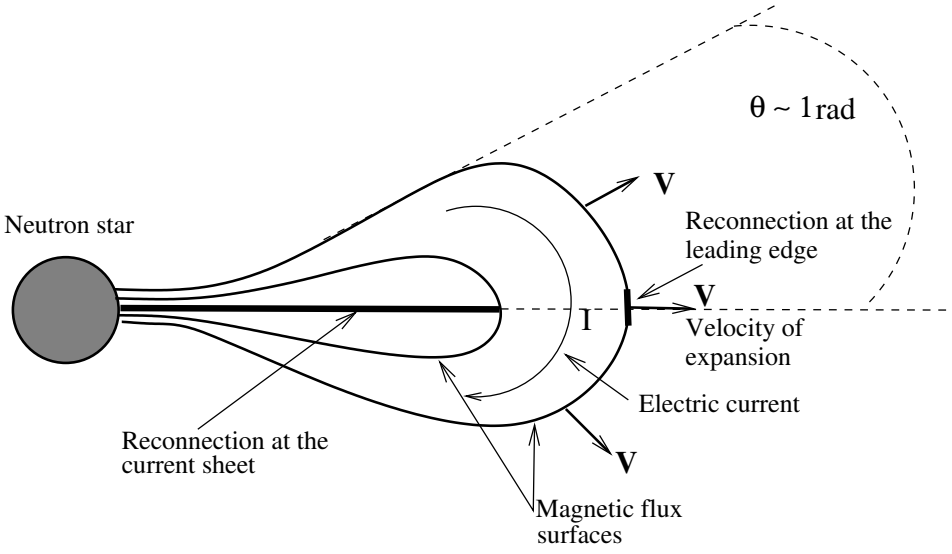

Figure 1.15: Expanding, collimated flux rope in the magnetospheric instability model of SGR giant flares. From Lyutikov (2006).

s duration of the main spike is related to the dynamical timescale of the expanding flux rope, over which the magnetic cloud relaxes to a minimum energy state. The slow distortion of the magnetosphere leading up to the flare naturally explains the increased bursting activity and large spin-down variations observed before the two more recent events, Figures 1.8 and 1.10.

Many details of the two families of giant flare models—global starquake and catastrophic magnetospheric instability—remain to be clarified. In both models most of the energy is released by reconnection in the magnetosphere—they differ principally in where the energy is stored just before the flare, and what kind of instability precipitates the magnetic reconfiguration which drives the reconnection. Simulations allow several aspects of the models to be tested quantitatively, in particular the MHD instability which triggers the giant flare in the magnetospheric-storage model, and the resulting reconnection phase.

1.3 This dissertation

In Chapter 2 we briefly describe force-free electrodynamics, the formulation of MHD we use in our simulations. We present the equations written in covariant (four-dimensional) notation, and in the three-dimensional form used in our code. The force-free current density, which self-consistently closes Maxwell's equations, is described, together with the two "force-free conditions" which must be satisfied by the combined magnetic and electric field configuration. We outline the two classes of waves which are found in the vanishing-inertia limit, the fast magnetosonic and Alfvén modes, and their phase and group velocities.

Our new pseudospectral code for force-free electrodynamic simulations, PHAEDRA, is described in detail in Chapter 3. We present the spatial discretization, describing how it automatically satisfies the regularity conditions at the coordinate poles. Spectral filtering, which maintains stability in the presence of strong discontinuities and is the primary source of numerical dissipation, is discussed at length. Gibbs oscillation introduces large pointwise errors into spectral solutions near discontinuities—we describe techniques of removing these controlled numerical oscillations and demonstrate their application to a current sheet simulated with our code. We give a combined treatment for the non-reflecting outer boundary conditions, involving an approximate removal of incoming information along characteristics and a simple frictional absorbing layer. Unlike many numerical schemes, our code does not exactly maintain the solenoidal condition on the magnetic field—we discuss this issue in detail and describe how we verify that the magnetic field remains well behaved, and the $\nabla \cdot \boldsymbol{B}$ error convergent, even in the neighborhood of current sheets. We outline the implementation of the code and its parallelization. Finally we describe the planned extension of the code to three-dimensional spherical and Cartesian geometries, and the inclusion of explicit physical resistivity prescriptions.

Chapter 4 contains numerical test problems, with which we verify that the code converges to

known analytic solutions, has low numerical diffusivity, and is well-behaved when discontinuities spontaneously form in the solution. We demonstrate exponential convergence to a smooth two-dimensional analytic solution, the Michel rotating monopole. We find that a solution with a strong current sheet exhibits a loss of strict pointwise convergence near the discontinuity, but that the solution remains convergent elsewhere. We describe a twisted dipole test problem, and use it to show that evolving with the full force-free current density term allows for lower numerical diffusion than can be obtained when only the drift current is included.

Our aligned rotator ("axisymmetric pulsar") solution is described in Chapter 5. We start with a star at rest, and show the dynamic evolution to the steady state. This solution is almost ideal everywhere except in the discontinuous current sheet beyond the light cylinder, leading to nearly exact energy conservation inside the light cylinder and the rapid dissipation of outflowing electromagnetic energy in the equatorial current sheet, concentrated near the Y-point. We discuss the viability of the force-free model, by considering the problem of adequate charge supply, necessary for the maintenance of the force-free plasma.

Chapter 6 contains the first part of our study of strongly twisted magnetar magnetospheres. In our numerical models, this twisting is due to gradual shearing of the stellar surface. In this chapter we describe the sequence of force-free equilibria, through which a magnetosphere smoothly moves when the amount of twisting is small and it is applied slowly. We then discuss the MHD instability which has been predicted to occur at large twist angles, and which may be the trigger for large magnetar flares. We find that this instability does take place, at relatively large twist angles $\psi_{\rm crit} \gtrsim 3$ radians when the surface shearing is axisymmetric. We find also that continual twisting has a stabilizing effect, allowing the magnetosphere temporarily to have twist amplitudes significantly larger than the critical value. Despite this stabilization, the formation of a current sheet, large-scale magnetic reconnection, and the release of magnetic energy are inevitable once the critical

twist amplitude is exceeded. We show what happens when the magnetosphere is sheared through multiple reconnection events, and give the fractions of the free twist energy which are retained, dissipated, and expelled as an electromagnetic outflow in our simulations.

In Chapter 7 we extend our study of twisted magnetar magnetospheres to include the effects of solid-body rotation of the neutron star. We again find a sequence of quasi-steady states of increasing twist amplitude, and a critical value dividing gentle relaxation and sudden catastrophic reconnection. The twisting leads to enhancement of the stellar spin-down rate, by factors which, for an axisymmetric twisted polar cap, scale with the square of the twisted flux. As we find in the non-rotating simulations, when subjected to continual twisting through large angles the magnetosphere goes through a sequence of expansion and reconnection events, whose duration and intensity depends on the previous evolution of the system. These events can be gradual or explosive, producing periods of increased torque lasting anywhere from the twisting timescale, which may be as long as months, to the spin period on the order of seconds. We show how an approximate model for an explosive event provides a sudden increase in the rotation period consistent with that observed during the giant flare from SGR 1900+14.

Finally, in Chapter 8 we describe the extension of the simulation code to curved spacetimes. We describe the 3 + 1 approach to foliating spacetime with a progression of spatial hypersurfaces, and show how the equations of general-relativistic force-free electrodynamics can be written in a form closely mirroring Maxwell's equations in a macroscopic medium. We devise a simple dynamical test problem, and compare the evolutions of the magnetospheres of neutron stars and black holes, both in vacuum and in force-free plasma.

The work presented here can be found in three publications:

Introducing PHAEDRA*: a new spectral code for simulations of relativistic magnetospheres*
Kyle Parfrey, Andrei M. Beloborodov, and Lam Hui, 2012, MNRAS, Vol. 423, 1416–1436

Twisting, reconnecting magnetospheres and magnetar spindown
Kyle Parfrey, Andrei M. Beloborodov, and Lam Hui, 2012, arXiv:1201.3635,
submitted to ApJ Letters

Dynamics of strongly twisted relativistic magnetospheres
Kyle Parfrey, Andrei M. Beloborodov, and Lam Hui, 2012, submitted to ApJ

Chapter 2

Force-free electrodynamics

In this chapter we present a concise description of force-free electrodynamics, the system of equations solved by the simulation code described in Chapter 3. This is the most useful formulation of magnetohydrodynamics (MHD) for the study of relativistic magnetospheres, in which

$$\frac{B^2}{8\pi} \gg \rho_{\rm m} c^2, \tag{2.1}$$

where $\rho_{\rm m}$ is the mass density. Standard MHD codes can become inaccurate, and even crash, when the plasma's magnetization is large, because the numerical error in the electromagnetic energy density becomes comparable to the matter energy density (Gammie, McKinney, & Tóth 2003; Komissarov 2004b). They also require the (usually poorly constrained) matter distribution to be set at the beginning of the simulation, and maintained throughout, sometimes by ad-hoc matter injection. Force-free codes do not experience these difficulties.

2.1 Covariant formulation

The system of force-free electrodynamics is the vanishing-inertia, or, equivalently, ultra-relativistic, limit of MHD. The latter can be written as

$$\nabla_\mu {}^*F^{\mu\nu} = 0, \tag{2.2}$$

$$\nabla_\mu T^{\mu\nu} = \nabla_\mu \left(T^{\mu\nu}_{(\mathrm{m})} + T^{\mu\nu}_{(\mathrm{f})} \right) = 0, \tag{2.3}$$

where ${}^*F^{\mu\nu}$ is the Maxwell tensor (Hodge dual of the Faraday tensor $F^{\mu\nu}$), and $T^{\mu\nu}_{(\mathrm{m})}$ and $T^{\mu\nu}_{(\mathrm{f})}$ are the energy-momentum tensors of the matter and electromagnetic fields, respectively. If the matter contribution to $T^{\mu\nu}$ can be neglected, Equation (2.3) simplifies to

$$\nabla_\mu T^{\mu\nu}_{(\mathrm{f})} = 0. \tag{2.4}$$

Combining

$$T^{\mu\nu}_{(\mathrm{f})} = F^\mu{}_\alpha F^{\nu\alpha} - \frac{1}{4} \left(F_{\alpha\beta} F^{\alpha\beta} \right) g^{\mu\nu}, \tag{2.5}$$

where $g_{\mu\nu}$ is the metric tensor of spacetime, with the inhomogeneous Maxwell equations,

$$\nabla_\nu F^{\mu\nu} = 4\pi J^\mu, \tag{2.6}$$

one finds that Equation (2.4) becomes

$$F_{\mu\nu} J^\nu = 0, \tag{2.7}$$

which states that the Lorentz force density vanishes (e.g. Komissarov 2002). Equation (2.7) can alternatively be derived by postulating, in a frame in which the electric and magnetic fields are

parallel, that the electric field vanishes and the current flows along the magnetic field.

2.2 3 + 1 formulation

We move now to a 3 + 1 spacetime point of view. Equation (2.7) becomes

$$\rho_{\mathrm{e}}\boldsymbol{E} + \boldsymbol{J} \times \boldsymbol{B} = 0, \tag{2.8}$$

where $\boldsymbol{E}$ and $\boldsymbol{B}$ are the electric and magnetic fields, and $\boldsymbol{J}$ the current density. We can see that

$$\boldsymbol{E} \cdot \boldsymbol{B} = 0; \tag{2.9}$$

the electric field is "degenerate." This condition, together with $\nabla \cdot \boldsymbol{B} = 0$, implies that the system of electromagnetic fields has only four independent components. Likewise,

$$\boldsymbol{E} \cdot \boldsymbol{J} = 0; \tag{2.10}$$

there is no Joule heating, the system is dissipationless, and formally (conditionally) hyperbolic.

The evolutionary Maxwell equations, Equations (2.2) and (2.6), are the familiar

$$\begin{aligned} \partial_t \boldsymbol{B} &= -\nabla \times \boldsymbol{E}, \\ \partial_t \boldsymbol{E} &= \nabla \times \boldsymbol{B} - 4\pi \boldsymbol{J}, \end{aligned} \tag{2.11}$$

using Gaussian units with $c = 1$. In MHD, an additional relation must be given for the current, closing the equations. In force-free electrodynamics, the current is uniquely determined by Equa-

tions (2.8) and (2.11), together with the condition $\partial_t (\boldsymbol{E} \cdot \boldsymbol{B}) = 0$, to be (Gruzinov 1999)

$$\boldsymbol{J} = \frac{\boldsymbol{B} \cdot \nabla \times \boldsymbol{B} - \boldsymbol{E} \cdot \nabla \times \boldsymbol{E}}{4\pi B^2} \boldsymbol{B} + \left(\frac{\nabla \cdot \boldsymbol{E}}{4\pi} \right) \frac{\boldsymbol{E} \times \boldsymbol{B}}{B^2}. \tag{2.12}$$

The charge density is implicitly defined to be

$$\rho_{\mathrm{e}} = \frac{\nabla \cdot \boldsymbol{E}}{4\pi}. \tag{2.13}$$

Ohm's law in force-free electrodynamics is therefore essentially geometrical. The first term in Equation (2.12), is the conduction current parallel to $\boldsymbol{B}$, which maintains the degeneracy condition, Equation (2.9). The second term is the drift current, being in the form $\rho_{\mathrm{e}} \boldsymbol{v}_{\mathrm{drift}}$, where

$$\boldsymbol{v}_{\mathrm{drift}} = \frac{\boldsymbol{E} \times \boldsymbol{B}}{B^2} \tag{2.14}$$

is the velocity of the magnetic field lines. It is apparent that there is a second condition,

$$B^2 - E^2 > 0, \tag{2.15}$$

equivalent to requiring the drift velocity to be less than the speed of light; since charged particles cannot cross field lines, this is a requirement if we assume that a macroscopic matter velocity field exists. Equations (2.9) and (2.15) are commonly referred to as the "force-free conditions." The second condition implies that the electromagnetic field is intrinsically magnetic, in that a frame exists in which the electric field vanishes.

It is possible for the fields to evolve from a configuration in which this second force-free condition is satisfied everywhere to one in which it is violated at some point, line, surface, or

volume. This local breakdown of the force-free approximation is necessarily accompanied by dissipation, as degeneracy is broken and $\boldsymbol{E} \cdot \boldsymbol{B} \neq 0$. In general, any configuration having field lines of different topology, such as the open and closed lines of the pulsar magnetosphere, will have points, lines, or surfaces at which $|\boldsymbol{B}| = 0$, violating the second condition (Uchida 1997). These sites of force-free breakdown are especially interesting, as the localised dissipation may be responsible for observed radiation.

Force-free electrodynamics supports two classes of waves, which we describe in a frame in which $\boldsymbol{E} = 0$ (Punsly 2003). Fast waves, equivalent to vacuum electromagnetic waves, propagate isotropically with both phase and group velocities equal to the speed of light. They do not carry any charges or currents. Alfvén waves can carry charges and currents, have phase velocity $v_{\text{phase}} = \pm c \cos\theta$, where θ is the angle between $\boldsymbol{B}$ and the wave vector, and have group velocity equal to c and directed along $\boldsymbol{B}$, $\boldsymbol{v}_{\text{group}} = \pm c\boldsymbol{B}/B$.

The "field-evolution" approach we take is not the only way to write the evolutionary equations of force-free electrodynamics. One could equivalently evolve the drift velocity or the Poynting flux vector, $\boldsymbol{S} = \boldsymbol{E} \times \boldsymbol{B}/4\pi$, instead of the electric field. The equations can also be written, using Euler potentials, as a Hamiltonian system (Uchida 1997), or in an axionic formulation (Thompson & Blaes 1998).

Chapter 3

Pseudospectral simulation code

3.1 Introduction

We describe a code for simulations of force-free electrodynamics: PHAEDRA (Pseudospectral High-Accuracy ElectroDynamics for Relativistic Astrophysics). The systems we study are "force-free" in the sense that the Lorentz force density vanishes everywhere, because the electromagnetic fields are strong enough that hydrodynamic forces can be neglected, resulting in self-balancing electromagnetic fields. Relativistic force-free electrodynamics has long been recognised as the appropriate limit for describing the magnetospheres of neutron stars and black holes, yet only recently has a concerted effort begun to study it with direct numerical simulation.

As described in Chapter 1, there exists a wealth of challenging problems for the field of electrodynamic numerical simulation, with applications to pulsars, magnetars, black holes, and gamma-ray bursts. Direct time-dependent simulation is valuable because it permits the study of general realistic initial-value problems, without the restrictions, like self-similarity or stationarity,

that are often necessary in analytical models, and because it naturally tests the stability of field configurations, a question often unanswered by steady-state numerical work.

Several time-dependent force-free electrodynamics codes exist, both those using finite differences (Spitkovsky 2006; Kalapotharakos & Contopoulos 2009; Palenzuela et al. 2010), and those that take a finite volume, or Godunov, approach (Komissarov 2004a; Cho 2005; Asano, Uchida, & Matsumoto 2005; McKinney 2006a; Yu 2011). Our numerical scheme is entirely different, and complementary, being based on orthogonal basis function expansions.

Previous codes have large numerical dissipation or diffusion, introduced either because they do not maintain $\boldsymbol{E} \cdot \boldsymbol{B} = 0$ self-consistently or through the intrinsic diffusivity of the method, while force-free problems often demand long simulations, as the fields may evolve over many wave-crossing times. It is desirable to have a method which can be run for long times without intrinsic dissipation, captures discontinuities, and accurately describes fast dynamics.

The crucial question one asks of a force-free configuration is that of its stability, the onset of instability commonly leading to a dramatic rearrangement of a magnetosphere, sometimes involving explosive reconnection. Spectral calculations tend to have less numerical noise than those of comparable finite-difference or finite-volume ("local") schemes; this noise can erroneously trigger instability. In a study of Sweet-Parker reconnection, the spectral magnetohydrodynamics (MHD) code is found to be largely immune to the secondary island formation, caused by a tearing-mode instability, that is found using local methods for the same problem (Ng & Ragunathan 2011).

In this chapter, we describe a code for axisymmetric simulations, in flat spacetime. It has been designed in such a way as to be extensible with minimal restructuring to a fully three-dimensional setting, in curved spacetime. The extension to curved spacetime is discussed, and some early results presented, in Chapter 8.

Note that in this chapter we distinguish between the contravariant, F^i, and covariant, F_i,

components of a vector field, while when we discuss results in Chapters 4 to 7 we refer only to the components in an orthonormal basis, also written F_i.

3.2 The pseudospectral method

A function $u(x,t)$, the solution of a time-dependent partial differential equation, can be expanded in terms of a set of orthogonal spatial basis functions $\phi_k(x)$,

$$u_N(x,t) = \sum_{k=0}^{N-1} a_k(t)\phi_k(x), \tag{3.1}$$

where $a_k(t)$ are time-dependent expansion coefficients; u_N is an approximation to the function u for some choice of basis functions and truncation N. Spatial derivatives can be taken by analytically differentiating u_N, since the exact derivatives of the basis functions are known. Considering first an equation in x only,

$$Du(x) = f(x),$$

where D is a general differential operator and f is a forcing function, we can think of solving this equation by minimising the residual R: $R = Du_N - f$. In what sense we choose to minimise R will determine the kind of spectral method we construct. In the Galerkin (sometimes just called "the spectral") method, the residual is made orthogonal to the basis functions: $(\phi_k, R) = 0,\ k = 0,\ldots,N-1$, where the brackets indicate an inner product, $(f,g) \equiv \int \omega(x) f(x) g(x) dx$, over a weight function $\omega(x)$. Since the first N spectral coefficients are exact, u_N can be considered to be a *truncation* of the infinite series expansion.

In the pseudospectral method, the residual is made zero at a set of 'collocation' points, $\{x_i\}$: $(\delta[x-x_i], R) = 0,\ i = 0,\ldots,N-1$, where $\delta(x)$ is the Dirac delta-function. The resulting u_N is then an

interpolant of the true u, at the chosen grid points. It can be shown that, if the collocation points are chosen as the abscissas of a Gaussian quadrature associated with the basis set and this quadrature rule is used to calculate the inner products, then the Galerkin and pseudospectral methods are equivalent for linear problems; the error penalty for choosing interpolation over truncation is at worst a factor of two, for trigonometric functions and Chebyshev polynomials (Boyd 2001).

Gaussian quadrature of a function f over a weight function ω,

$$\int_a^b f(x)\omega(x)dx = \sum_{i=0}^{N-1} w_i f(x_i), \tag{3.2}$$

is accomplished by finding the corresponding set of N weights $\{w_i\}$ and N interpolation points $\{x_i\}$; the pay off for being restricted to these interpolation points is that the resulting formula is exact for all $f(x)$ which are polynomials of degree $2N-1$ or less. The weight function determines the basis functions; for example, the Chebyshev polynomials are those polynomials which are orthogonal with respect to the weight $\omega(x) = 1/\sqrt{1-x^2}$ on the interval $[-1, 1]$. For periodic $f(x)$, the composite-trapezoidal and midpoint rules are Gaussian quadratures with an equispaced grid, the corresponding basis being trigonometric functions.

The pseudospectral method can also be thought of as the limiting case of increasing-order finite-difference methods, where the derivative stencil now extends over all grid points. In particular, the Fourier pseudospectral method on a periodic uniform grid is recovered by a finite-difference formalism as stencil width (and formal order of accuracy) goes to infinity (Fornberg 1996). This approach gives a dense differentiation matrix, whose application requires $O(N^2)$ operations. However, identical derivatives can be calculated for the Fourier and Chebyshev basis sets by using the fast Fourier transform (FFT), which requires only $O(N \ln N)$ operations; this is sometimes referred to as the "transform method" (Orszag 1970). In this case, a forward FFT gives the expansion

coefficients, $\{a_k\}$, from which can be constructed the coefficients for the derivative series $\{a'_k\}$: $u'_N(x) = \sum_{k=0}^{N} a'_k \phi_k(x)$. In the Fourier basis, $\phi_k(x) = \mathrm{e}^{\mathrm{i}kx}$ and so $a'_k = \mathrm{i}ka_k$; with Chebyshev polynomials a three-term recurrence relation is used. Finally the derivative at the grid points, u'_N, can be found with an inverse FFT. Given a function to be differentiated, this procedure can be thought of as finding an interpolating function of order N, at a set of N points, and taking the exact analytic derivative of this interpolating function.

The great benefit of these methods is that spectral approximation is exponentially convergent for sufficiently smooth functions: the error decreases faster than any power of the truncation N. It has generally been found that this carries over into exponential convergence of spectral solutions of PDEs, even those with fixed-order time marching. This accuracy has made spectral methods popular in many areas of physics, including meteorology, seismology, shock waves, and reactive flows. Astrophysical applications include accretion disc magnetohydrodynamics (Chan, Psaltis, & Özel 2005, 2009) and general relativity (e.g. Gourgoulhon 1991; Bonazzola et al. 1999; Kidder et al. 2000; Dimmelmeier et al. 2005; Grandclément & Novak 2009). In engineering electrodynamics, the "pseudospectral time-domain" method was introduced by Liu (1997), where it was shown to have much lower diffusion and dispersion error than finite-difference methods, and to require either two (Fourier) or π (Chebyshev) points per wavelength for adequate resolution, in comparison to eight to sixteen points for finite differences. These lower required grid densities make a spectral method more efficient for achieving a given accuracy, despite the higher number of operations per grid point.

3.3 Spatial discretisation

In order to simplify eventual extension to curved spacetime, we adopt a scheme which allows the use of an arbitrary spatial metric. We store, and advance in time, the contravariant vector components of $\boldsymbol{B}$ and $\boldsymbol{E}$; curls are taken with $(\nabla \times \boldsymbol{F})^i = (1/\sqrt{\gamma}) e^{ijk} \partial_j F_k$ and divergences with $\nabla \cdot \boldsymbol{F} = (1/\sqrt{\gamma}) \partial_i (\sqrt{\gamma} F^i)$, where γ is the spatial metric determinant, e^{ijk} the Levi-Civita symbol, and $\boldsymbol{F}$ stands for $\boldsymbol{B}$ or $\boldsymbol{E}$. The quantity to be differentiated, either F_k or $\sqrt{\gamma} F^k$, is first calculated at each point from the contravariant components F^k and the metric, then expanded in orthogonal basis functions. This method requires more forward transforms than one where the derivatives are simplified by the chain rule, since for example both F_r and $\sqrt{\gamma} F^r$ must be transformed into spectral space, but we find it to be more accurate.

Our grid is defined in axisymmetric spherical coordinates (r, θ), with N collocation points in r and L points in θ. We will use i and j to index grid points, and n and l to index wavenumbers, along each direction; $i, n = 0, \ldots, N-1$, and $j, l = 0, \ldots, L-1$.

In the radial direction we use Chebyshev polynomials, T_n, and the collocation points are chosen to be the Chebyshev-Gauss-Lobatto nodes,

$$x_i = -\cos\left(\frac{\pi i}{N-1}\right), \qquad i = 0, \ldots, N-1, \qquad x \in [-1, 1] \,. \tag{3.3}$$

These can be mapped directly onto the physical coordinate, $r_i = r_{\text{in}} + (r_{\text{out}} - r_{\text{in}})(1 + x_i)/2$, $r \in [r_{\text{in}}, r_{\text{out}}]$, or via an additional coordinate mapping (section 3.4). Since the Chebyshev polynomials are mapped cosine functions,

$$T_n(\cos[q]) = \cos(nq) \,, \tag{3.4}$$

with this choice of grid the Chebyshev transform of a function f can be performed with a fast

cosine transform:

$$f(x_i) = \sum_{n=0}^{N-1} f_n T_n(x_i) \,, \\ = \sum_{n=0}^{N-1} f_n \cos\left(\frac{\pi n i}{N-1}\right). \tag{3.5}$$

In the meridional direction we expand in sine or cosine functions, depending on whether the vector component in question is an even or odd function across the pole (its parity). This is related to the "double Fourier" method of expanding functions on a sphere (Merilees 1974; Orszag 1974), which is attractive because it avoids the slow Legendre transform required by spherical harmonics. The following are even, and can be expanded in cosines: $F_r, F^r, F_\phi, F^\phi, \sqrt{\gamma}F^\theta$, whereas odd functions that can be expanded in sines are $F_\theta, F^\theta, \sqrt{\gamma}F^r, \sqrt{\gamma}F^\phi$. Here we are only concerned with axisymmetric modes; in general the parity will depend on whether the azimuthal wavenumber is even or odd. To avoid solving the equations directly on the poles we use a shifted grid:

$$\theta_j = \frac{j + 1/2}{L}\pi \,, \qquad j = 0, \ldots, L-1 \,. \tag{3.6}$$

For instance, the covariant radial component of the magnetic field, once formed by direct index lowering with the metric, can be expanded as

$$B_r = \sum_{n=0}^{N-1} \sum_{l=0}^{L-1} a_{nl} T_n(r) \cos(l\theta) \,, \tag{3.7}$$

and the combination $\sqrt{\gamma}E^r$, required to calculate $\nabla \cdot \boldsymbol{E}$, as

$$\sqrt{\gamma}E^r = \sum_{n=0}^{N-1} \sum_{l=0}^{L-1} a_{nl} T_n(r) \sin(l\theta) \,. \tag{3.8}$$

Once the coefficients a_{nl} of a function have been found, by taking the forward transforms in both directions, the coefficients of the derivative series a'_{nl} can be calculated. For differentiation by θ this is simple: $a'_{nl} = -la_{nl}$ for functions expanded in cosines, and $a'_{nl} = la_{nl}$ for those expanded in sines. Radial differentiation requires the three-term recurrence relation, relating the Chebyshev coefficients of a function to those of its derivative,

$$\begin{aligned} a'_{N-1,l} &= 0 \\ a'_{N-2,l} &= 2(N-1)a_{N-1,l} \\ c_{n-1}a'_{n-1,l} &= a'_{n+1,l} + 2na_{n,l} \,, \end{aligned} \tag{3.9}$$

where $c_0 = 2$, and all other $c_n = 1$.

3.4 Coordinate mappings

The standard Chebyshev nodes, x_n, are not a suitable radial grid for our problems. They are very strongly clustered near the endpoints, leading to a time step restriction for hyperbolic problems that goes like $\Delta t \sim O(1/N^2)$, which makes obtaining high spatial resolution in the rest of the domain unnecessarily expensive. A coordinate map can be used to relieve the endpoint clustering, and also put more nodes in the part of the domain where higher resolution is required. Consider a map from x to a new coordinate y:

$$\begin{aligned} y &= g(x) \\ \frac{\partial}{\partial y} &= \frac{1}{g'}\frac{\partial}{\partial x}, \end{aligned}$$

where $g' = dg/dx$. The derivative coefficients can be calculated as before, and the mapping of the derivative achieved by a multiplication in real space. The arcsine map (Kosloff & Tal-Ezer 1993),

$$x_{\text{asin}} \equiv g_{\text{asin}}(x) = \frac{\arcsin(\alpha x)}{\arcsin(\alpha)}$$
$$g'_{\text{asin}} = \frac{\alpha}{\arcsin(\alpha)} \frac{1}{\sqrt{1-\alpha^2 x^2}} , \tag{3.10}$$

where α is a constant between zero and one and $x_{\text{asin}} \in [-1, 1]$, will create an almost equispaced grid with less clustering at the endpoints; the grid stretching becomes more pronounced as α is increased towards unity. Care must be taken in choosing α: since the map is singular at the endpoints, choosing α too large would impair the convergence of the series. The maximum discrepancy between a function[1] $f(z)$ and its truncated series representation $P_N(z)$ is

$$\max |f(z) - P_N(z)| = c\epsilon , \tag{3.11}$$

where c is a constant which depends on f but is independent of N, and ϵ is related to α by

$$\epsilon = \left(\frac{1 - \sqrt{1-\alpha^2}}{\alpha} \right)^N \longrightarrow \alpha = \operatorname{sech}\left(\frac{|\ln \epsilon|}{N} \right) . \tag{3.12}$$

Therefore choosing ϵ to be sufficiently small will usually make the singularity harmless (Don & Solomonoff 1995; Mead & Renaut 2003). Equation (3.11) determines the *largest* pointwise error, which usually appears at, or close to, a boundary; the error can be smaller than ϵ in the interior of the domain, which we confirm in our tests in Chapter 4. For fixed ϵ the minimum grid spacing only decreases as $O(1/N)$.

The nearly equispaced grid resulting from the arcsine map is still not perfect for computations

[1]The complex function $f(z)$ is the analytic continuation of the real function $f(x)$.

in spherical coordinates when $r_{\rm out} \gg r_{\rm in}$, because the lines of constant θ converge towards the centre of the domain, and so too much radial resolution is used where the angular resolution is low, and not enough where it is high. We would prefer a grid where $\Delta r \leq r\,\Delta\theta$ over most of the domain. To construct this we use a combination of the arcsine map and a smooth algebraic stretching:

$$x_{\rm alg} \equiv g_{\rm alg}(x) = Q\frac{1 + x_{\rm asin}}{Q + 1 - x_{\rm asin}} - 1 \ , \tag{3.13}$$

where Q is a constant and $x_{\rm alg} \in [-1, 1]$. The grid is then linearly transformed to the desired physical coordinates

$$r = R_{\rm in} + \frac{1}{2}(R_{\rm out} - R_{\rm in})(1 + x_{\rm alg}), \quad r \in [R_{\rm in}, R_{\rm out}] \ . \tag{3.14}$$

We generally use $Q \sim 0.1 - 1$, and set the map-induced error to $\epsilon \sim 10^{-9} - 10^{-15}$. It appears to be preferable to have the radial grid spacing somewhat smaller than the meridional spacing throughout most of the domain. Figure 3.1 shows the inter-nodal spacings of an example grid.

Some of our models also include a coordinate transformation in the meridional direction, in order to increase the resolution around a reconnecting current sheet in the equatorial plane. To this end we employ the "Kepler-Burgers" mapping (Boyd 1992),

$$\widetilde{\theta} = \theta + \frac{\gamma}{2}\sin(2\theta), \tag{3.15}$$

where θ is evenly spaced on $[0, \pi]$ and $\widetilde{\theta}$ is the new stretched coordinate. The constant γ controls the degree of stretching; when using this map we set $\gamma \sim 0.3 - 0.5$.

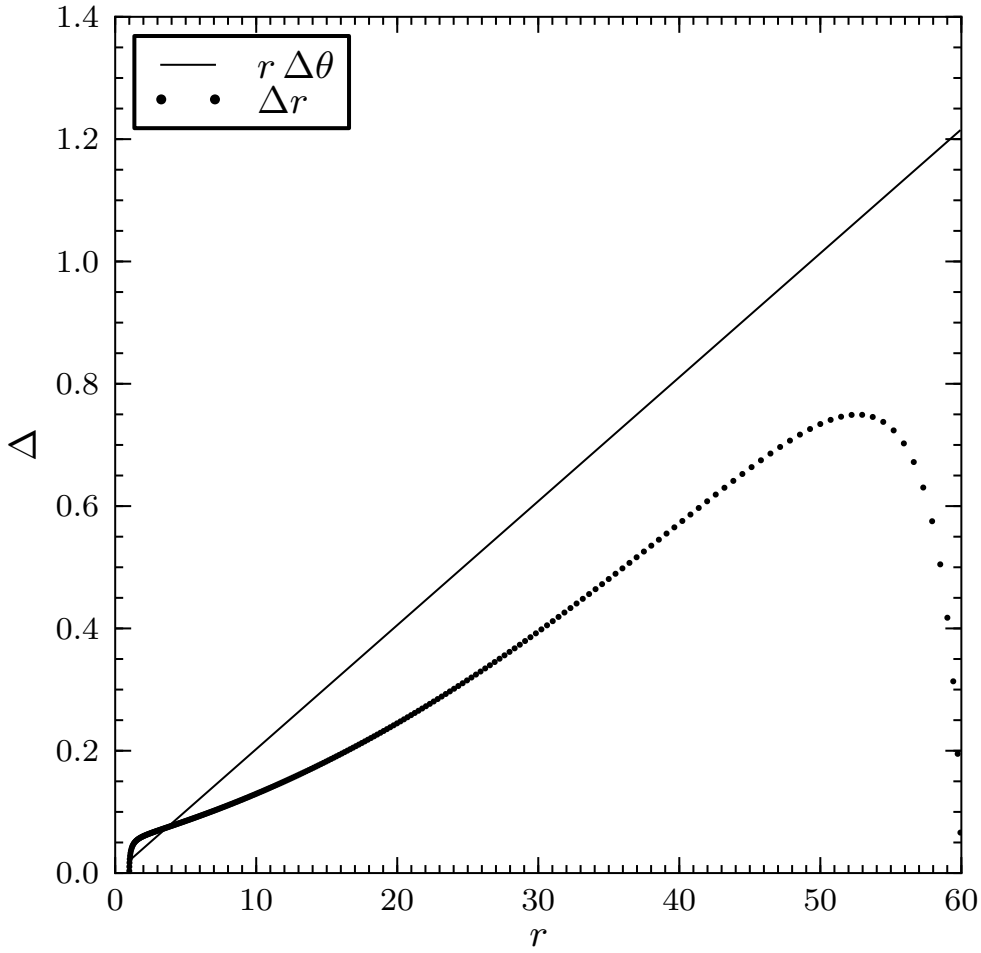

Figure 3.1: Grid spacing versus radius: Chebyshev nodes with a combination arcsine + algebraic mapping; $N = 256$, $L = 155$, $Q = 0.6$, $\epsilon = 10^{-11}$.

3.5 Time evolution

When the spatial derivatives have been found the system of equations becomes a set of ordinary differential equations in time, one for each vector component, which we solve with an explicit Runge-Kutta integrator. We mainly use a third-order, three stage, method (Fornberg 1996). If the ODE to be solved is

$$\frac{du}{dt} = f(u, t),$$

$u \in \left\{B^i, E^i\right\}$, then we can find the solution at time step $n+1$, u^{n+1}, from that at time n by

$$\begin{aligned} k^{(1)} &= \Delta t f(u^n, t^n) \\ k^{(2)} &= \Delta t f(u^n + \frac{1}{3}k^{(1)}, t^n + \frac{1}{3}\Delta t) \\ k^{(3)} &= \Delta t f(u^n + \frac{2}{3}k^{(2)}, t^n + \frac{2}{3}\Delta t) \\ u^{n+1} &= u^n + \frac{1}{4}(k^{(1)} + 3k^{(3)})\,. \end{aligned} \tag{3.16}$$

Since $k^{(2)}$ is not required once $k^{(3)}$ has been calculated, the latter can overwrite the former in memory, and so this method only requires storage for two intermediate arrays. Explicit time advancement is subject to the Courant-Friendrich-Lewey stability constraint on the time step, $\Delta t \leq C_{\rm CFL}\delta_{\rm min}$, where $\delta_{\rm min}$ is the smallest grid spacing, in our case always found at the inner surface. The factor $C_{\rm CFL}$ is, in general, problem and time-integrator dependent; we find that Equation (3.16) is stable for $C_{\rm CFL} \leq 1$, and hence always set $\Delta t = \delta_{\rm min}$.

We compared the above with a fifth-order, six stage, Runge-Kutta integrator. The improvement in accuracy is negligible, even for our stringent test with the Michel monopole solution, indicating that time-stepping errors are subdominant. This is unsurprising given that a fine grid in the radial direction is required near the stellar surface to produce highly-accurate solutions, and so the stability-limited time step produces small errors.

3.6 Spectral filtering

The use of spectral methods for nonlinear hyperbolic problems comes with two principal difficulties. The first is the build-up of power at high wavenumbers due to nonlinear couplings between lower wavenumbers. Hyperbolic problems have no explicit dissipation in the equations of motion,

and spectral methods have very low intrinsic dissipation, so these high modes do not decay, and can lead to the breakdown of the scheme by the aliasing instability. Aliasing occurs because a discrete transform of point values mistakes high-frequency components for low; all modes $e^{i(\omega+qN)x}$, for $q = 0, \pm\pi, \pm 2\pi \ldots$, will be identical when represented on a grid of N points in the interval $[-\pi, \pi]$. nonlinear terms produce high-frequency modes, which are then mistaken as low-frequency modes by the discrete transforms; these phantom low-frequency components are then similarly combined to generate more power at high frequencies, a clearly unstable cycle.

The second difficulty is the ability of nonlinear couplings to create discontinuities in a solution which was previously smooth. Spurious oscillations appear in the spectral interpolant when a function is not resolved by the grid[2]; this is the Gibbs phenomenon. A jump discontinuity causes $O(1)$ errors in its immediate vicinity, and reduces the convergence rate to first order elsewhere.

Both of these difficulties can be largely overcome with spectral filters. If a function u is expanded in a basis ϕ_n, its filtered approximation, $\mathcal{F}u$, is given by

$$\mathcal{F}u = \sum_{n=0}^{N-1} \sigma\left(\frac{n}{N-1}\right)\widetilde{u}_n\phi_n \ , \tag{3.17}$$

where $\widetilde{u}_n$ are the discrete expansion coefficients. Vandeven (1991) showed that if the filter function, $\sigma(\eta)$, is unity at $\eta = 0$, zero for all $|\eta| \geq 1$, and has at least $2p - 1$ continuous derivatives, then $\mathcal{F}u$ will converge with N at $2p$-th order even in the presence of a jump discontinuity, except very close to the jump. In addition, since it strongly damps the high modes, such a filter can prevent the onset of the aliasing instability if regularly applied to each field in a time-dependent simulation. We tried

[2]Any set of points on a grid will be faithfully recovered following transforms into, and back out of, Fourier space; however, the interpolating function, constructed from the Fourier coefficients, can show spurious oscillations *between* the grid points, causing oscillations in the derivative *at* the grid points.

two variable-order examples, the erfc-log filter (Boyd 1996),

$$\bar{\eta} \equiv |\eta| - \frac{1}{2}$$
$$\sigma_{\text{erfc}-\log}(\eta) = \frac{1}{2}\text{erfc}\left\{2(2p)^{1/2}\bar{\eta}\sqrt{\frac{-\ln(1-4\bar{\eta}^2)}{4\bar{\eta}^2}}\right\}, \tag{3.18}$$

and the exponential filter (Majda et al. 1978),

$$\sigma_{\exp}(\eta) = \mathrm{e}^{-\alpha\eta^{2p}}. \tag{3.19}$$

The latter can be made to fulfil approximately the requirement of being zero for all $|\eta| \geq 1$ by setting $\alpha = \alpha_M = -\ln(\epsilon_M)$, where ϵ_M is machine precision; we use $\alpha_M = 35$. We find the exponential filter to give more accurate results and to allow weaker filtering, and so use it exclusively.

The use of the exponential filter to control aliasing was studied in detail by Hou & Li (2007), where they found that a high-order ($2p = 36$) filter can prevent instability in marginally-resolved fluid dynamics simulations while producing more accurate solutions than standard dealiasing methods. We also find this to be the case for our equation system, even though the nonlinear coupling is much stronger than in Euler's equations. For all science runs we use a filter with $\alpha = \alpha_M$ and $2p = 36$, which appears to balance well the conflicting demands of assuring stability while minimising unphysical dissipation. When very low numbers of modes are used (roughly $N < 48$ in any direction) the filtering order needs to be reduced somewhat. This high-order filter is applied to the coefficients of every derivative series[3], and directly to the coefficients of the solution itself at the end of each full Runge-Kutta time step.

[3]When using Chebyshev polynomials it is important to filter the coefficients of the derivative series, a'_{nl}, rather than those of the function before the recurrence relation is used, the a_{nl} (Godon & Shaviv 1993). For this reason our filtering operations are implemented inside the coefficients-to-grid inverse transform.

We turn now to the second issue, concerning stability and convergence in the presence of jump discontinuities. Current sheets, regions of formally infinite current density implying discontinuities in the magnetic field, are a generic feature of force-free electrodynamics. The danger is that the strong production of high-wavenumber power, or the nonlinear interaction of the Gibbs oscillations with the solution, will lead to instability, or at least loss of convergence. Dissipation is required to prevent this, and ensure that the correct entropy solution is selected; see Gottlieb & Hesthaven (2001) for an extensive review of stability and convergence theory for nonlinear hyperbolic problems.

This dissipation can be effectively, and efficiently, provided by a spectral filter (Tadmor 1993; Don 1994). For a time-dependent problem,

$$\frac{du}{dt} = f(u) \ ,$$

we can advance the solution one time step, and then apply a filter to the solution, $u(t + \Delta t)$. Note that in this case we are using the term "filter" more broadly than above; in particular, we will not require that $\sigma(1) = 0$. If we choose to use the exponential filter, Equation (3.19), then the modified equation, taking into account the action of the filter, is

$$\frac{\partial u}{\partial t} = f(u) - \alpha \frac{(-1)^p}{\Delta t \, N^{2p}} \frac{\partial^{2p} u}{\partial x^{2p}} + O(\Delta t^2) \ , \tag{3.20}$$

if u is expanded in Fourier series, and

$$\frac{\partial u}{\partial t} = f(u) - \alpha \frac{(-1)^p}{\Delta t \, N^{2p}} \left[\sqrt{1 - x^2} \frac{\partial}{\partial x} \right]^{2p} u + O(\Delta t^2) \ , \tag{3.21}$$

if it is expanded in Chebyshev polynomials. The second result follows from the relation

$$\left[\sqrt{1-x^2}\frac{\partial}{\partial x}\right]^2 T_n(x) + n^2 T_n(x) = 0 \ . \tag{3.22}$$

In this case the dissipation decreases to zero at the boundaries (recall $x \in [-1, 1]$, and here $u(x)$ is non-periodic), which is useful since then no additional boundary conditions are required (Boyd 1998).

Applying an exponential filter is therefore similar to adding a hyperviscous term to the equation, where the magnitude of the hyperviscosity is

$$\epsilon_N = \frac{\alpha}{\Delta t\, N^{2p}} \ . \tag{3.23}$$

The action of the filter is equivalent to an implicit time integration of the hyperviscous term, and so no additional stiffness is added to the equations.

The similarity of the effect of the exponential filter to an explicit hyperviscous term allows us to import stability and convergence theory derived using such terms. The spectral viscosity (SV) method (Tadmor 1989, 1990) uses second-order viscous regularisation ($2p = 2$), and convergence is obtained by excluding an increasing fraction of low-wavenumber modes from the viscous term as N is increased. Extensions to higher order, $2p \geq 4$, followed for schemes based on Fourier (Tadmor 1993) and Chebyshev (Ma 1998) expansions, known as the super spectral viscosity (SSV) methods. These can be proven to converge to the correct solution of a scalar conservation law for

$$\begin{aligned} \epsilon_N &= \frac{C}{N^{2p-1}} \ , \\ p &\leq O(\ln N) \ . \end{aligned} \tag{3.24}$$

Convergence can't be guaranteed for a system of equations, although it has been proven that if the scheme converges, it converges to the correct entropy solution (Carpenter et al. 2003). Despite the lack of solid theoretical results for nonlinear systems, experience has shown that the method can be stable and convergent; multidimensional examples include shock-vortex interaction (Don 1994; Sun et al. 2006), and problems involving both shocks and combustion (Don & Gottlieb 1998; Gottlieb & Gottlieb 2005), where spectral methods were found to perform well in comparison with high-order shock-capturing schemes.

Using Equation (3.24) as a guide, we find from Equation (3.23) that the filter amplitude should scale as $\alpha = \alpha_{\mathrm{SSV}} = CN\Delta t$. The time step scales as $\Delta t \propto 1/N$ (Section 3.4), and so α_{SSV} should be roughly constant. The filter order, $2p$, should only increase slowly with N, and so we set it to a constant as well. Numerical experiments confirm that fixed α_{SSV} and p, determined by low-resolution simulations, lead to stable and convergent results as resolution is increased. We find best results are obtained with $2p = 8$, $\alpha_{\mathrm{SSV}} \sim 0.01 - 0.1$, corresponding to a weak hyperdiffusion which decreases in strength with resolution like N^{-7}. The SSV filters are applied to every component of the electric and magnetic fields at the end of each full Runge-Kutta time step.

To summarise the filtering procedure, we apply a very high order filter, with $\alpha = 35$ and $2p \sim 36$, to the inverse transform of every derivative series. At the end of each full time step, we apply both the previous filter and one with $\alpha \sim 0.01 - 0.1$ and $2p = 8$ to the field variables themselves.

3.7 Post-processing

The SSV method is exponentially convergent in any error norm, if p increases linearly with N; however stability will often not permit this, and in any case $O(1)$ errors will remain near any

discontinuities. It has long been argued that high-order methods retain enough information to reconstruct a highly accurate (and sometimes spectrally convergent) solution, even if oscillations are present and the pointwise errors are large. The Gibbs oscillations are not noise, and they can be safely removed, without destroying the accuracy of the underlying solution, with a post-processing step after the simulation has finished. There are many ways to do this: for example, real-space filtering or mollification (Gottlieb & Tadmor 1985), one-sided filters (Cai, Gottlieb, & Shu 1992), spatially-varying spectral filters (Boyd 1996; Tadmor & Tanner 2005), and reprojection into a Gibbs-complementary set of basis functions (Gottlieb et al. 1992; Gelb & Tanner 2006). The reprojection method is particularly popular, and many examples exist of successful one-dimensional reconstructions of oscillation-free solutions (e.g. Shu & Wong 1995; Sarra 2003; Ma & Li 2006).

The problem with all of the above methods is that they require the locations of the discontinuities to be known accurately, which is particularly difficult in more than one dimension. We were unable to find a sufficiently robust means of determining the number of unresolved features and their locations, given that the physics generates, and our scheme is capable of resolving, oscillations with wavelengths close to the grid scale.

We show here an example of post-processing applied to a solution with a known discontinuity—the equatorial current sheet in the aligned rotator problem (Chapter 5). Figure 3.2 shows (a) the original SSV solution before post-processing, (b) after applying the optimal spatially-varying spectral filter (Tanner 2006):

$$\sigma_{\rm opt}(k, N, x) = \mathrm{e}^{-z} \sum_{n=0}^{\lfloor \kappa N d(x) \rfloor} \frac{1}{n!} z^n, \qquad (3.25)$$
$$z = \frac{\alpha k^2 d(x)}{2N}$$

where $d(x)$ is the distance to the nearest discontinuity and κ and α are constants, and finally (c)

using the digital total variation (DTV) spatial filter (Rudin et al. 1992; Bürgel & Sonar 2002). The DTV filtered solution u to a noisy variable u^0 is found by minimising the fitted total variation energy

$$W_{\mathrm{DTV}} = \sum_{\beta} |\nabla u_\beta| + \frac{\lambda}{2}\left(u_\beta - u_\beta^0\right)^2, \tag{3.26}$$
$$|\nabla u_\beta| = \sqrt{\sum_{\gamma} (u_\beta - u_\gamma)^2},$$

where β ranges over all points in the dataset, γ denotes each point's neighbours, and λ is related to the expected noise level. The minimisation can be done by linearised Jacobi iteration. This method does not require the locations of the discontinuities, has a natural multi-dimensional form, requires only an estimate of the size of the oscillations to be removed, and has been applied to Chebyshev-based spectral methods with good results (Sarra 2006). However, we find it to perform poorly if resolved physical high-frequency oscillations are present, even with an adaptive local noise estimate.

We do not use any post-processing technique for our results shown in this dissertation, with the exception of Figure 5.9 where we apply a two-dimensional DTV filter. Our intention is simply to highlight that any Gibbs oscillations which are present do not destroy the accuracy of the scheme. The pointwise errors may be large near a jump, and we do not claim spectral convergence of the field values on the grid, but the overall evolution of the system is still consistent with high-order accuracy.

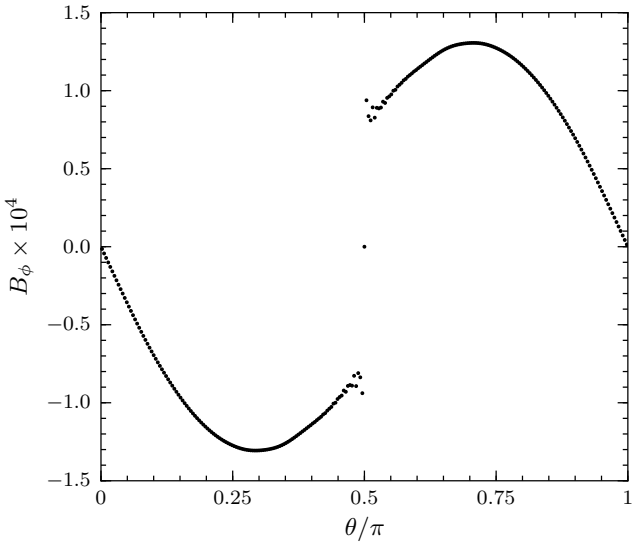

(a) No post-processing

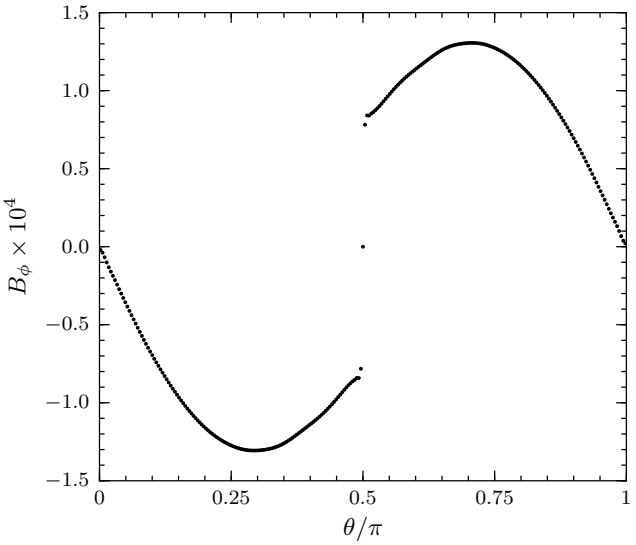

(b) Optimal filter

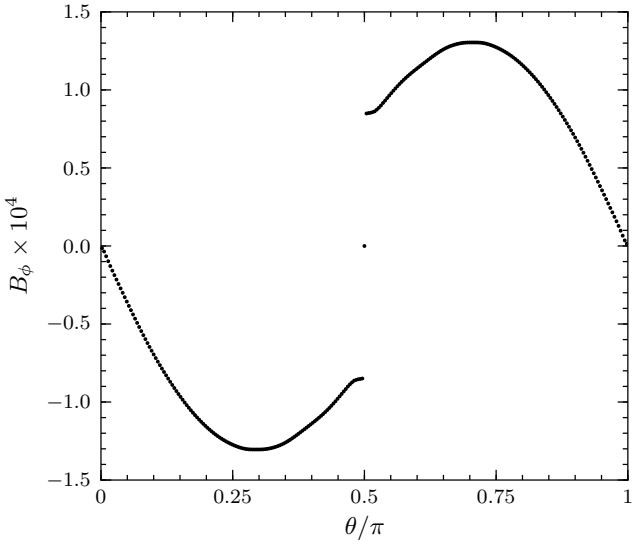

(c) Digital total variation

Figure 3.2: Toroidal magnetic field with discontinuity (current sheet), $L = 255$: (a) no post-processing, SSV only; (b) with optimal filter, $\kappa = 0.05$, $\alpha = 1$; (c) with DTV reconstruction.

3.8 Force-free current

The current density function required to close Maxwell's equations with force-free dynamics can be written as

$$\begin{aligned} \boldsymbol{J} &= \boldsymbol{J}_{\parallel} + \boldsymbol{J}_{\text{drift}}, \\ \boldsymbol{J}_{\parallel} &= \frac{\boldsymbol{B} \cdot \nabla \times \boldsymbol{B} - \boldsymbol{E} \cdot \nabla \times \boldsymbol{E}}{B^2} \boldsymbol{B}, \qquad \boldsymbol{J}_{\text{drift}} = \nabla \cdot \boldsymbol{E} \, \frac{\boldsymbol{E} \times \boldsymbol{B}}{B^2}. \end{aligned} \tag{3.27}$$

The effect of $\boldsymbol{J}_{\parallel}$ is to maintain the force-free condition $\boldsymbol{E} \cdot \boldsymbol{B} = 0$. To date, schemes that rely on staggered grids do not explicitly include this term, because it demands that all components of the fields and their curls be evaluated, by interpolation if necessary, at the locations where the electric field components are defined. Instead, the effect of the parallel current is mimicked by resetting the values of $\boldsymbol{E}$ on the grid such that $\boldsymbol{E}_{\parallel} = 0$. Since the inherent accuracy of spectral derivatives frees us from the need for staggered grids, we can evaluate the full current function without interpolation, and so include both $\boldsymbol{J}_{\parallel}$ and $\boldsymbol{J}_{\text{drift}}$ in the equations of motion. We also manually set $\boldsymbol{E}_{\parallel}$ to zero at the end of each full time step; this has essentially no discernible effect in most of the domain since $\boldsymbol{J}_{\parallel}$ keeps $\boldsymbol{E} \cdot \boldsymbol{B}$ very close to zero anyway, but gives slightly cleaner evolution close to the inner boundary at lower grid resolution. Specifically, $\boldsymbol{E}$ is projected, parallel to $\boldsymbol{B}$, into a plane perpendicular to $\boldsymbol{B}$:

$$\boldsymbol{E} \to \boldsymbol{E} - (\boldsymbol{E} \cdot \boldsymbol{B}) \frac{\boldsymbol{B}}{B^2}. \tag{3.28}$$

Retention of the parallel current results in much lower diffusion error; see the twisted dipole test in Section 4.2.2.

Force-free evolution can lead to configurations in which the second force-free condition, $B^2 - E^2 > 0$, is violated. For example, in the equatorial current sheet of the axisymmetric rotating

dipole (Chapter 5) all components of the magnetic field are zero, and any electric field results in the violation of the condition. An ideal force-free configuration requires $\boldsymbol{E} = 0$ in the current sheet. This can be simulated by immediately reducing the magnitude of the electric field if it is larger than the magnitude of the magnetic field, such that $B^2 - E^2 = 0$ at the end of the operation. At each Runge-Kutta substep, if the condition is violated, we shrink the electric field vector, leaving its direction unchanged:

$$\boldsymbol{E} \rightarrow \sqrt{\frac{B^2}{E^2}}\,\boldsymbol{E}. \tag{3.29}$$

This removal of electric field acts like a small, highly localised, source of dissipation in current sheets. Physically, the removed electromagnetic energy would be converted to thermal energy and radiation; in these simulations it is simply lost from the system.

3.9 Boundary conditions

A system of hyperbolic equations generally requires suitable conditions to be provided at the boundaries of the computational domain. These boundary conditions will depend on the spatial geometry and the physical problem under investigation. We often wish to simulate a volume of force-free plasma surrounding a neutron star, and so describe the boundary treatment for this case.

In these simulations, the inner boundary, $r = r_{\mathrm{in}}$, corresponds to the stellar surface, while the outer boundary at r_{out} should as closely as possible behave as a membrane which perfectly transmits outgoing waves without generating any incoming waves. The behavioural boundary conditions at the poles are satisfied automatically by the choice of either sines or cosines as the basis functions in the meridional direction, as described in Section 3.3.

3.9.1 Inner boundary

For our purposes the star is a perfect rigid conductor, all or part of which can be rotated. Since the star is a perfect conductor, the electric field will be zero in the rotating frame, $\boldsymbol{E}' = 0$. The fields in the lab frame are given by

$$\begin{aligned} \boldsymbol{B} &= \boldsymbol{B}', \\ \boldsymbol{E} &= -(\boldsymbol{\Omega} \times \boldsymbol{r}) \times \boldsymbol{B}', \end{aligned} \tag{3.30}$$

where $\boldsymbol{\Omega}$ is the local angular velocity vector, and therefore

$$\boldsymbol{E} = -(\boldsymbol{\Omega} \times \boldsymbol{r}) \times \boldsymbol{B}. \tag{3.31}$$

Rotation about the z-axis corresponds to the application of an induced poloidal electric field.

Equation (3.31) provides the relationship between $\boldsymbol{E}$ and $\boldsymbol{B}$ infinitesimally below the surface, but our first grid points are in the force-free plasma infinitesimally above. The normal component of the magnetic field, B^r, and the tangential components of the electric field are continuous across the surface, and therefore are known. The required boundary values are $B^r = B^r(\theta)$, $E^\theta = -\Omega B^r \sin\theta$, and $E^\phi = 0$; these are strongly enforced at every Runge-Kutta substep. The other components must be allowed to evolve undisturbed, since they depend on unprescribed surface currents and charges on the star.

A complication introduced by the combination of the above boundary conditions and the SSV filtering of the fields is the anomalous leakage of energy into the domain through the inner boundary. Consider the B^r field: at the end of each time step it is filtered, or smoothed, leading to a slight broadening of the $B^r(r)$ profile, and hence diffusion of field from the boundary into the

domain. According to the above boundary condition, stating that B^r is constant in time, this field at the boundary is immediately replenished, and so over time this cycle can increase the volume-integrated energy on the grid.

A solution is to subtract the initial vacuum fields $(\boldsymbol{B}_0, \boldsymbol{E}_0)$ from the dynamical fields $(\boldsymbol{B}, \boldsymbol{E})$, and evolve $\widetilde{\boldsymbol{B}} = \boldsymbol{B} - \boldsymbol{B}_0$ and $\widetilde{\boldsymbol{E}} = \boldsymbol{E} - \boldsymbol{E}_0$. The current density and electric field in the initial configuration are zero, and so, writing $\boldsymbol{B} = \boldsymbol{B}_0 + \widetilde{\boldsymbol{B}}$, $\boldsymbol{E} = \widetilde{\boldsymbol{E}}$, we get the equations

$$\begin{aligned} \partial_t \widetilde{\boldsymbol{B}} &= -\nabla \times \widetilde{\boldsymbol{E}} \\ \partial_t \widetilde{\boldsymbol{E}} &= \nabla \times \widetilde{\boldsymbol{B}} - \boldsymbol{J}(\boldsymbol{B}_0 + \widetilde{\boldsymbol{B}}, \widetilde{\boldsymbol{E}}). \end{aligned} \tag{3.32}$$

Note also that $\nabla \times \boldsymbol{B}_0 = 0$, and so $\nabla \times \boldsymbol{B} = \nabla \times \widetilde{\boldsymbol{B}}$ in the current density function—no additional derivatives are required. There is no requirement that these new variables be small, and in much of the magnetosphere they will be larger than the corresponding initial field. The initial field is never differentiated or filtered, there is no field leakage by the above mechanism since $\widetilde{B}^r = 0$ is the new magnetic boundary condition, and the energy conservation of the code is improved dramatically (see Section 4.2.2).

To reduce notational clutter, these variables will be neither distinguished nor discussed elsewhere in this paper; the only addition they require is that the evolved magnetic field be temporarily added to a stored initial field before being passed to the current density function (which takes as arguments the electric and magnetic fields, and their curls).

Finally, a small, slowly growing, anomalous toroidal magnetic field was found in certain simulations using a dipole base field, in the equatorial region immediately next to the star. This field appeared only when that region was "dead" ($\nabla \times \boldsymbol{B} = 0$). We have attributed this feature to Alfvén waves on under-resolved field lines, propagated with a numerical scheme with very low

diffusivity.

This phenomenon is completely eradicated by setting a small circular region of the poloidal plane, centred on $(r, \theta) = (r_{\rm in}, \pi/2)$ and with radius $r_{\rm drift}$, to only use the drift current contribution to $\boldsymbol{J}$ in the equations of motion. Neglecting $J_\parallel$ makes the scheme in the "drift" region sufficiently diffusive that the anomalous feature does not develop. Inside this region the electric field is projected to be perpendicular to the magnetic field, using Equation (3.28), at the end of each full time step, exactly as in the rest of the domain. The radius of the drift region can be decreased with increasing resolution; we use $r_{\rm drift} \sim 0.25\, r_{\rm in}$. We have found that modifying the scheme in this small fraction of the domain does not affect the solution in the rest of the domain.

3.9.2 Outer boundary

We implement the non-reflecting boundary condition at the outer boundary using an approximate characteristic decomposition. Spectral methods are very sensitive to "incorrect" boundary conditions, and the simple zero-gradient condition that works well in low-order methods leads to instability. Characteristic-based boundary treatments (e.g. Abarbanel et al. 1991; Godon 1996) specify the outgoing characteristic variables using the calculated data on the grid, and set the incoming characteristic variables to zero.

If we construct a six-component vector of the fields like $\boldsymbol{q} = \{\boldsymbol{B}, \boldsymbol{E}\}$, then the one-dimensional equations of motion in Cartesian coordinates can be written in the general form

$$\partial_t \boldsymbol{q} + \mathsf{A}\partial_x \boldsymbol{q} = 0, \tag{3.33}$$

where the matrix A (the flux Jacobian) should include contributions from the linear and nonlinear

terms. This matrix can be decomposed into its eigenvalues and eigenvectors,

$$\mathsf{A} = \mathsf{S}\Lambda\mathsf{S}^{-1}, \tag{3.34}$$

Λ being a diagonal matrix of the eigenvalues, S containing the right eigenvectors, and S^{-1} the left eigenvectors. The characteristic variables, $\boldsymbol{w}$, are found using the left eigenvectors,

$$\boldsymbol{w} = \mathsf{S}^{-1}\boldsymbol{q}, \tag{3.35}$$

using which the equations of motion decouple, since Λ is diagonal:

$$\partial_t \boldsymbol{w} + \Lambda\partial_x \boldsymbol{w} = 0. \tag{3.36}$$

In this one-dimensional case each component of $\boldsymbol{w}$ will move in either the positive or negative x direction, depending on the sign of the corresponding eigenvalue. The incoming variables can be identified, set to zero, and the primitive variables recovered using $\boldsymbol{q} = \mathsf{S}\boldsymbol{w}$.

Rather than using the exact characteristic variables, we have implemented an approximate boundary condition using the characteristics of the vacuum Maxwell's equations.

At each point on the outer boundary, construct a local Cartesian vector basis, with $\hat{\mathbf{x}}$ in the radial direction, $\hat{\mathbf{y}}$ along $\hat{\boldsymbol{\theta}}$, and $\hat{\mathbf{z}}$ along $\hat{\boldsymbol{\phi}}$. We will then identify $B_y = rB^\theta$, $B_z = r\sin\theta B^\phi$ etc., where the Cartesian components are in a normalised basis, and the spherical components are contravariant

as usual. The four propagating characteristic variables are

$$\begin{pmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} E_z + B_y \\ E_y - B_z \\ E_z - B_y \\ E_y + B_z \end{pmatrix}. \tag{3.37}$$

The variables w_1 and w_2 propagate in the negative r direction, these are incoming variables and will be set to zero; w_3 and w_4 propagate in the positive r direction and will be calculated from the field values on the grid. Inverting Equation (3.37), zeroing the incoming characteristics, and replacing Cartesian with spherical components gives

$$\begin{pmatrix} B^\theta \\ B^\phi \\ E^\theta \\ E^\phi \end{pmatrix} = \frac{1}{r} \begin{pmatrix} -w_3 \\ w_4/\sin\theta \\ w_4 \\ w_3/\sin\theta \end{pmatrix}; \tag{3.38}$$

the indicated fields at the boundary are replaced with these values at every Runge-Kutta substep, the radial components being left unchanged.

This approximate boundary condition is stable and works well for predominantly radial waves, but can generate artefacts when sizeable tangential waves are present. To prevent these from appearing we use a thin sponge layer next to the outer boundary, which absorbs outgoing waves (If et al. 1987). Any waves reaching the boundary are much attenuated, and so the vacuum boundary condition is a better approximation. Likewise, the approximate non-reflecting condition allows the use of a thinner and weaker absorbing layer than would be sufficient with a reflecting boundary.

The sponge layer is introduced by adding a frictional term to Maxwell's equations, which

become

$$\begin{aligned}\partial_t \boldsymbol{B} &= -\nabla \times \boldsymbol{E} - \sigma_s(r)\boldsymbol{B}_{\text{ang}} \\ \partial_t \boldsymbol{E} &= \nabla \times \boldsymbol{B} - \boldsymbol{J} - \sigma_s(r)\boldsymbol{E},\end{aligned} \tag{3.39}$$

where $\boldsymbol{B}_{\text{ang}} = \{0, B^\theta, B^\phi\}$, not including B^r—we found that damping the radial magnetic field led to unphysical currents leaking back into the normal domain. This approach is similar to the matched-layer method of Yang et al. (1997). The frictional coefficient, σ_s, is chosen to be zero in most of the domain, and to rise smoothly near the boundary; the functional form we use is

$$\sigma_s(r) = \begin{cases} 0 & \text{if } r < r_s \\ \sigma_0 \left(1 - \exp\left[-\gamma \left(\dfrac{r - r_s}{r_{\text{out}} - r_s}\right)^\beta\right]\right) & \text{if } r > r_s \end{cases} \tag{3.40}$$

with the values $\sigma_0 \sim 0.5 - 1$, $\gamma \sim 6$, $\beta \sim 4$. This boundary treatment is robust, effective, and insensitive to the values of the sponge layer coefficients.

3.10 Magnetic field divergence

Maxwell's equations comprise two evolution equations, for the electric and magnetic field vectors, and two constraint equations. One of these constraints, Gauss's law, is automatically satisfied, since the charge density has been replaced by $\nabla \cdot \boldsymbol{E}$ in the drift current term in Equation (3.27). The other constraint is the solenoidal condition on the magnetic field: $\nabla \cdot \boldsymbol{B} = 0$.

In theory this should not be a concern, since the evolution equation for $\boldsymbol{B}$ implies $\partial_t(\nabla \cdot \boldsymbol{B}) = -\nabla \cdot (\nabla \times \boldsymbol{E}) = 0$. However, in numerical schemes the operators for calculating divergences and curls usually do not satisfy $\nabla \cdot (\nabla \times \boldsymbol{V}) = 0$ exactly for any vector $\boldsymbol{V}$, raising the worry that this truncation-

level magnetic field divergence might build up over time. The presence of such magnetic charges can lead to unphysical forces and even instability.

To maintain stability many MHD codes use constrained transport, which can ensure that some representation of $\nabla \cdot \boldsymbol{B}$ is kept at machine zero (see Tóth (2000) for a review). These methods are incompatible with our global spatial derivatives.

Another option is to evolve the magnetic field using a vector potential: writing $\boldsymbol{B} = \nabla \times \boldsymbol{A}$, the first evolution equation becomes $\partial_t \boldsymbol{A} = -\boldsymbol{E}$; see Chan et al. (2009) for a spectral implementation. Although $\nabla \cdot (\nabla \times \boldsymbol{A})$ is only zero to truncation error, this error would not grow over time. The disadvantage of a vector potential is the introduction of second-order spatial derivatives, and, in our case, an increase in the number of derivatives that must be taken. More problematically, we found this method to be less stable than the direct magnetic field-evolving method, especially at the boundary.

It appears that the best approach in the context of a spectral method may be to do nothing, and rely on the accuracy of the derivatives, and hence the smallness of the truncation error, to maintain the solenoidal condition to high precision. Let us define a normalised magnetic divergence,

$$(\nabla \cdot \boldsymbol{B})_{\text{norm}} \equiv \frac{\nabla \cdot \boldsymbol{B}}{|\boldsymbol{B}| / \sqrt{A}}, \tag{3.41}$$

where A is the cell area[4]. Our results are highly divergence-free, as we illustrate with steady-state solutions to three problems. For the Michel rotating monopole (Section 4.2.1), the normalised magnetic divergence is $\sim 10^{-14}$ for an 84×56 grid, and $\sim 10^{-18}$ for a grid of 192×128. The twisted magnetosphere (Section 4.2.2) has normalised $\nabla \cdot \boldsymbol{B}$ mostly around $\sim 10^{-12}$, rising to 10^{-9} in the region with largest current. Our fiducial aligned rotator solution (Chapter 5) has normalised

[4]Our discretisation is based on nodes rather than cells, but here it is unimportant which adjacent nodes are chosen to form a fictitious cell.

$\nabla \cdot \boldsymbol{B}$ of $\sim 10^{-6}$ in and near current sheets and $\sim 10^{-10}$ elsewhere; the highest-resolution solution has values roughly ten times lower everywhere.

The higher $\nabla \cdot \boldsymbol{B}$ near current sheets is at least partly due to the presence of the aforementioned Gibbs oscillations, and would be expected to be much smaller in the recoverable high-accuracy solution (see Section 3.7). This Gibbs-generated divergence would also be present if a vector potential had been used. Finally, $\nabla \cdot \boldsymbol{B}$ is higher than normal right at the inner boundary, presumably because of the filtering applied to the Chebyshev derivatives, which become increasingly dependent on the highest frequencies as the boundaries are approached (e.g. Godon & Shaviv 1993).

Most importantly, no unusual or troubling behaviour has been observed to be correlated to an increase of the magnetic field divergence on the grid, the evolution appears to be stable for very long run times, and the measured divergence decreases with increasing resolution.

3.11 Code infrastructure

PHAEDRA is written in C for speed and portability. The spectral transforms are performed with the FFTW3 library (Frigo & Johnson 2005), which allows transforms of arbitrary size with $O(N \log N)$ complexity. The code is fully MPI-parallel, with a simple automatic domain-decomposition function which does not require the grid dimensions to be a multiple of the number of processors.

The parallelisation works similarly to the method of (Pelz 1991). The domain, in real space, is slab-decomposed in the radial direction. In the forward transform to spectral space, $f(r, \theta) \rightarrow a_{nl}$, the θ-transforms are first performed using data local to each processor, after which the mixed $f_l(r)$ data undergo a parallel matrix transpose, done with a single `MPI_Alltoallv()` call. The r-transforms can then be performed, and the coefficients a_{nl}, now decomposed in the l-direction, used to calculate the coefficients of the desired derivative. The real-space derivative values are

found by repeating the above steps in the reverse order. The filter is applied just before the inverse transform. In the SSV filtering step, performed at the end of each full time step, the coefficients are calculated, filtered, and immediately transformed back to real space. To date the code has been run on between one and 64 processors.

The data output is done collectively, in parallel, using the Parallel HDF5 library. The data are accessed via the XDMF standard, in which an auxiliary XML file, also written by the code, describes the contents of the HDF file containing the data arrays. This format allows the data to be easily opened by the VisIt visualisation software, among others, without the need of a custom plugin.

3.12 Discussion

3.12.1 PHAEDRA

We have found that PHAEDRA converges exponentially quickly to smooth solutions, is stable and accurate when discontinuities are present, and exhibits very low numerical diffusion outside of current sheets. A critical ingredient is the full force-free current density, in particular the parallel conduction current, which maintains $\boldsymbol{E} \cdot \boldsymbol{B} = 0$ without introducing any stiffness to the equations of motion. We are able to use the full current because our mesh is unstaggered, and so its evaluation doesn't require any interpolation of fields or their derivatives. The unstaggered grid is itself enabled by the inherent accuracy of the pseudospectral spatial derivatives.

The flexibility of the mapped-coordinate method allows efficient calculation in large domains, while retaining accuracy near the stellar surface at the inner boundary. It also permits resolution to be concentrated where large gradients are expected to form, a kind of fixed mesh refinement. With the adjustable parameter in the arcsine mapping of the radial Chebyshev series, a deliberate

tradeoff can be made between accuracy close to the boundary and stable time step. The ability to model interactions of force-free waves with a solid boundary to high accuracy will be useful when studying, for example, turbulence driven by neutron star vibrations. The treatment of non-periodic boundaries is one of the advantages of spectral expansions over high-order finite difference methods, which can experience difficulties at surfaces.

The code is parallelised with the standard MPI library, and can be run in both shared- and distributed-memory environments. It is efficient enough to run on dense grids, using around 16–32 processors, for many millions of time steps. Simulations on coarser grids can be run on one or a few processors, on consumer laptops and workstations. One concern is that the $O(N^2)$ communication time required for the global MPI call will become a problem when scaling to a very large number of processors, for example for a three-dimensional calculation. We expect that the code should scale well up to several hundred processors on existing hardware, and that this will be sufficient.

The principal issue we have encountered is the resistivity imparted to current sheets by the spectral filters (see Chapter 5). There are numerous tricks that can be used to reduce or eliminate this in simple cases, like the aligned rotator we describe later. However, we prefer not to use any method insufficiently robust, flexible, or efficient to also be applicable to general three-dimensional current sheets propagating across the grid. The ideal aligned rotator may represent the worst-case comparison of PHAEDRA to codes employing finite-difference or finite-volume methods, which can evolve the solution to a steady-state with very little flux closed outside the light cylinder, either with use of a staggered mesh (Spitkovsky 2006), or manual nulling of the inflow of flux into the current sheet (McKinney 2006b). We are investigating adaptive spectral- and real-space filtering, which may be able to stabilise current sheets with less dissipation than global filtering.

In any case, physical current sheets are believed to possess finite resistivity (e.g. Zenitani & Hoshino 2007; Zenitani & Hesse 2008). The effective hyper-resistivity of the filters confines

the dissipation to grid-scale features like current sheets where it is physically expected, leaving resolved features to evolve ideally. In this sense, our aligned rotator solution may be realistic.

3.12.2 Planned Extensions

We are planning to extend the scheme in several ways. The derivatives are already being calculated in a metric-independent fashion, and so it should be straightforward to adapt the code to the Schwarzschild or Kerr metrics, using the formalism of Komissarov (2004a). The first stages of this extension are described in Chapter 8. Eventually it should even be possible to include effects due to a time-dependent metric in this framework (Komissarov 2011), possibly using data from a code which evolves the Einstein equations.

The method currently assumes axisymmetry; we intend to make it fully three-dimensional, using either the complete "double Fourier" method on spheres or expansions in spherical harmonics. We expect the overall scheme to be adaptable to different choices of basis functions with only minor changes. This would allow us to investigate the oblique pulsar magnetosphere in spherical coordinates, and so alleviate some of the difficulties, such as stair-stepping on the inner surface and an inflexible equispaced Cartesian grid, encountered by some existing 3D force-free codes. Removing the axisymmetric restriction will also permit us to evaluate the stability of field configurations produced by general surface footpoint motions. Considering the lower resolution that is possible with 3D calculations, it is encouraging that we see close agreement between simulations on coarse and fine grids, even for the pulsar solution.

Another promising modification is to a 3D Cartesian geometry, using Fourier series in all directions. This geometry would be suitable for studying ultra-relativistic turbulent cascades (Thompson & Blaes 1998; Cho 2005) and instabilities surrounding nearly force-free current sheets. Aside from being simpler than expansions in spherical coordinates, the Cartesian-plus-Fourier

combination has the benefit of allowing the solenoidal constraint on the magnetic field to be easily enforced in spectral space.

The low intrinsic numerical diffusivity of our code makes it ideal for studying the effect of physical plasma resistivity. Although force-free electrodynamics lacks a well-defined fluid frame, and hence a preferred form for any dissipative terms, several formulations of resistive, nearly force-free, electrodynamics have been proposed (Lyutikov 2003; Gruzinov 2008; Li et al. 2012; Kalapotharakos et al. 2011). Being able to resolve waves with many fewer points per wavelength, spectral codes require less diffusivity than lower-order schemes to accurately transport, without unphysical oscillations, a given profile on the same grid. Brandenburg (2003) found that spectral derivatives permitted a viscosity five times lower than that needed for even a sixth-order finite-difference method. If enough physical resistivity is added to resolve otherwise sharp current layers, we will probably be able to dispense with the "super spectral viscosity" filtering, and the code should again achieve exponential convergence.

Chapter 4

Test problems

This chapter presents the results of several numerical test problems, with which we have verified that the code is convergent to known analytic solutions. The code converges exponentially quickly with increasing resolution when the solution is smooth, and has very low numerical diffusivity on resolved scales. We demonstrate that the evolution remains stable when current sheets—discontinuities in the magnetic field—form spontaneously on the grid, and that the oscillations which appear are restricted to a region near the discontinuity whose width decreases as the number of spectral modes retained increases. We describe a test problem employing a twisted dipole field, and use it to demonstrate that the effective numerical diffusivity is lowest when the full force-free current density term is included in the equations of motion.

4.1 1D tests

For the following two test cases we use a one-dimensional simplification of the code described in Chapter 3, with Chebyshev polynomials as the basis functions, an arcsine coordinate map, and a Cartesian vector basis (which simply requires using the metric $\gamma_{ij} = \delta_{ij}$). The boundary conditions are enforced strongly at $n = 0, N-1$; no sponge layers or non-reflecting boundary conditions are used.

4.1.1 Stationary Alfvén wave

Komissarov (2004a) describes an analytical solution for a stationary Alfvén wave: $B_x = B_y = E_z = 1$, $E_y = 0$,

$$B_z(x) = \begin{cases} 1 & \text{for } x < 0, \\ 1 + 0.15\left\{1 + \sin\left[5\pi\left(x - 0.1\right)\right]\right\} & \text{for } 0 < x < 0.2, \\ 1.3 & \text{for } x > 0.2, \end{cases} \tag{4.1}$$

and $E_x = -B_z$. For ease of comparison with this paper, we also use $N = 200$ and a domain $x \in [-1.5, 1.5]$. The SSV filter strength is set to $\alpha_{\mathrm{SSV}} = 0.1$ (even though no such filtering is required for this smooth problem) so that the effect of a discontinuity-capturing level of diffusion can be observed.

Fig. 4.1 shows the numerical solution for B_z at $t = 2$ and $t = 4000$. The profile does not broaden noticeably, even after more than a million time steps. The small wiggles, on either side of the jump in the $t = 4000$ solution, are imprinted by the action of the SSV filter on a function which is not infinitely smooth (B_z has discontinuous first derivatives at $x = 0$ and $x = 0.2$).

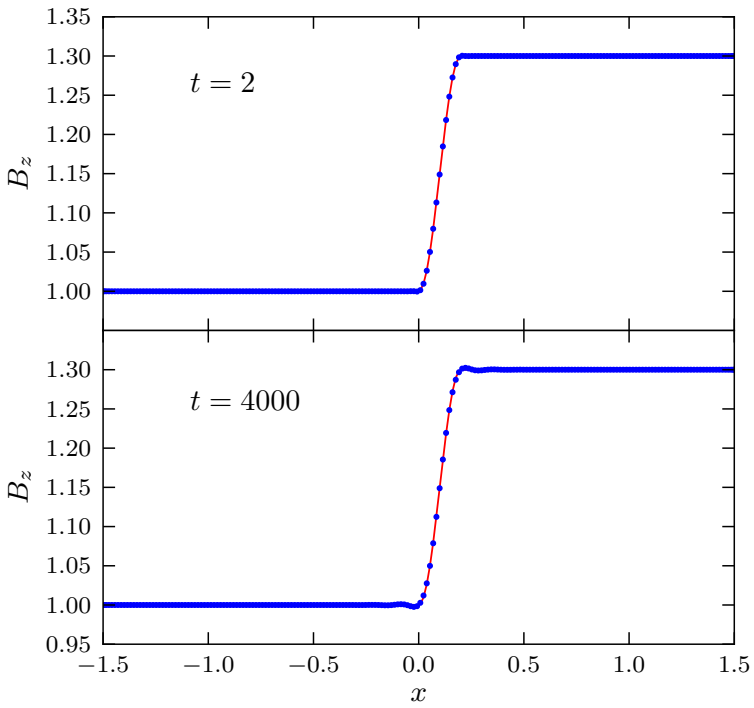

Figure 4.1: Stationary Alfvén wave test problem. The line shows the exact solution, the dots show the numerical solution.

4.1.2 Riemann problem

A Riemann problem which results in a current sheet is described by Komissarov (2004a). The initial conditions are: $\boldsymbol{E} = \boldsymbol{0}$, $B_z = 0$, $B_x = 1$, and

$$B_y(x) = \begin{cases} B_0 & \text{for } x < 0, \\ -B_0 & \text{for } x > 0. \end{cases} \tag{4.2}$$

The current sheet forms spontaneously at $x = 0$, and two fast step waves are emitted, one in either direction. The numerical solution for B_y at $t = 1$ is shown in Fig. 4.2; $\alpha_{\mathrm{SSV}} = 0.1$, and $2p_{\mathrm{SSV}} = 8$ as usual. The three jump discontinuities are clearly unresolved on the grid, and the fast waves remain

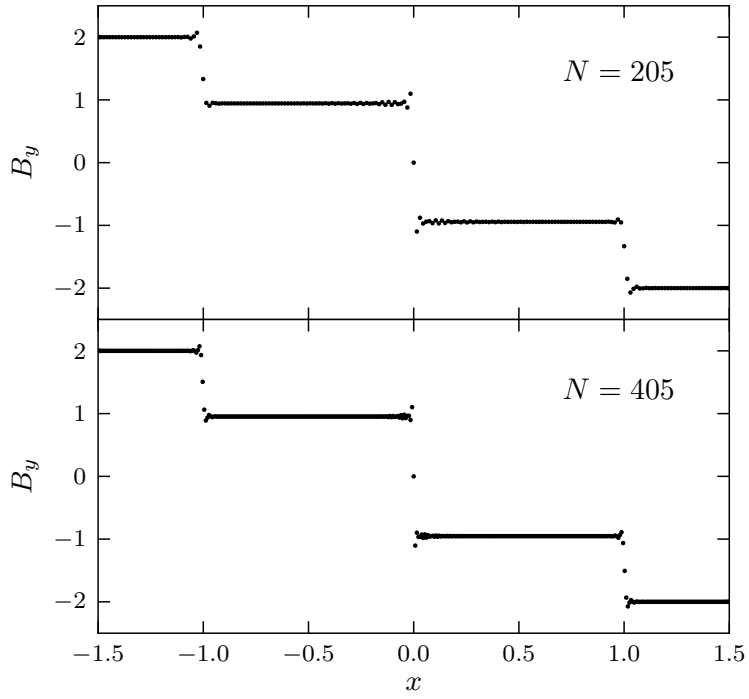

Figure 4.2: Riemann problem at $t = 1$, with current sheet at $x = 0$, for two grid sizes N.

sharp. The Gibbs oscillations are confined to the immediate vicinity of a discontinuity, and the affected region shrinks as resolution is increased.

4.2 2D tests

In these problems we used the full 2D code, in spherical coordinates. In all cases the star has a radius of unity, $r_{\rm in} = 1$; its radial light-crossing time-scale is then $\Delta t = 1$.

4.2.1 Michel monopole & split monopole

The exact three-dimensional solution for the field configuration, in force-free electrodynamics, surrounding a rotating magnetic monopole was derived by Michel (1973). Given initial conditions $B^r = f_0/r^2$, $B^\theta = B^\phi = 0$, $\boldsymbol{E} = 0$, and subject to uniform rotation with angular velocity Ω, the steady-state solution has a toroidal magnetic field

$$B_\phi = f_0\,\Omega\frac{\sin\theta}{r}. \tag{4.3}$$

This analytic solution is well suited to a multi-dimensional convergence test. We set the domain to be $1 \le r \le 30$ and vary the grid size $N \times L$, holding $N \approx 1.5L$. No sponge layer is used, since monopole field lines imply radial outgoing waves, for which the characteristic boundary treatment is very effective (Section 3.9.2). The order of the aliasing-controlling filter is $2p = 36$ for $L \ge 32$, and slightly lower for smaller grid sizes; no SSV filtering is applied because the solution is smooth. The angular velocity is smoothly increased from zero to $\Omega = 0.1$ in twenty radial light-crossing times (i.e. between $t = 0$ and $t = 20$), and the solution is sampled at $t = 100$, on a surface of constant radius at $r = 5$. The fractional errors in B_ϕ, defined as

$$\text{fractional error} = \left|\frac{B_\phi^{\text{numerical}}}{B_\phi^{\text{analytic}}} - 1\right|, \tag{4.4}$$

are plotted in Fig. 4.3. The errors decrease approximately exponentially with increasing resolution, reaching a level of roughly 10^{-12} with a grid size of 84×56. No effort was made to optimise the domain size, or the radius or time at which the numerical and analytic solutions are compared.

A similar test can be performed for the split monopole, which simply involves reversing the sign of f_0, and therefore of B^r and B^ϕ in the solution, across the equator. The discontinuous

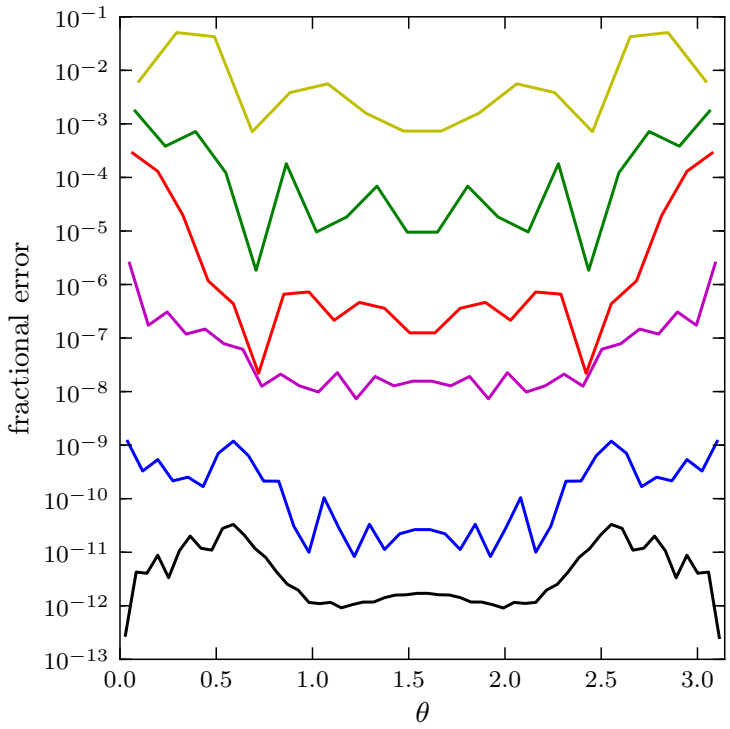

Figure 4.3: Fractional errors in B_ϕ for a rotating monopole. Grid sizes are, from top, $L = 16, 20, 24, 32, 40, 56$; $N \approx 1.5L$. N and L are the number of nodes in the radial and angular directions, respectively.

magnetic field implies the existence of an equatorial current sheet. This configuration allows us to test the behaviour of the code in the presence of a realistic discontinuity; we use SSV filtering with $\alpha_{\text{SSV}} = 0.05$. Fig. 4.4 shows the errors for the split monopole, for a problem otherwise identical to that described previously.

In this case the convergence is not uniform, being faster farther from the discontinuity at the equator. Order unity errors remain immediately beside the current sheet, as described in Section 3.6, but the solution converges strongly everywhere else, reaching a fractional error floor of 10^{-7} for the 224×155 grid. The large pointwise errors near the equator are not worrying, because they are due to the superposition of well-behaved, understood, and controlled Gibbs oscillations

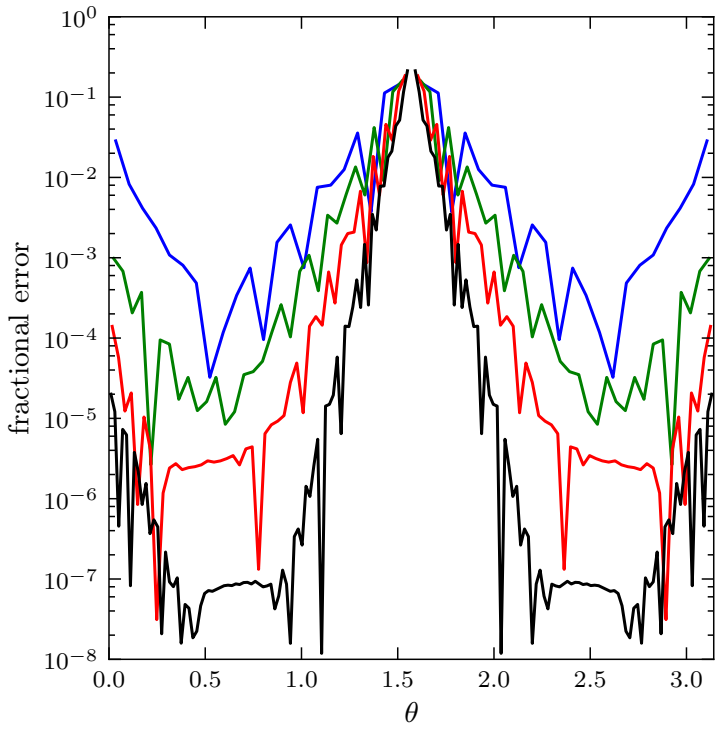

Figure 4.4: Fractional errors in B_ϕ for a rotating split monopole. Grid sizes are, from top, $L = 45, 65, 95, 155$; $N \approx 1.5L$.

on top of an otherwise accurate solution, rather than uncontrolled numerical noise or error; see Section 3.7 for a discussion.

4.2.2 Twisted dipole

As another 2D test, consider a dipole magnetosphere that is twisted by a differential rotation of the star's surface. We assume a latitude-dependent surface motion that is symmetric about the dipole axis; then the magnetosphere remains axisymmetric. The twisted field lines become extended in the azimuthal direction between their footpoints in the northern and southern hemispheres.

We start from a normal dipole configuration at $t = 0$ and gradually impart a twist by moving

the surface; at $t = T$ the surface motion stops. If the magnetosphere is an ideal conductor, it must remain static at $t > T$ — the twist will remain 'frozen' and persist forever. In the presence of resistivity, the magnetosphere must untwist with time (Beloborodov 2009a). Our ideal force-free model has no physical dissipation, and the rate of untwisting at $t > T$ will measure the numerical dissipation for our scheme.

The speed of the surface motion is much smaller than the speed of light, and so the twist will be implanted slowly, relative to the relevant wave-crossing time-scales. In our simulation the stellar rotation is applied only in an annular region of one hemisphere:

$$\Omega(\theta) = \frac{\Omega_{\text{centre}}}{1 + \exp\{\kappa\,(|\theta - \theta_{\text{ctr}}| - \Delta)\}}, \tag{4.5}$$

where θ_{ctr} is the colatitude of the centre of the annulus, Δ is its angular half-width, and κ determines the sharpness of its edges.

The rotation rate is increased smoothly from zero to Ω_{max} in time $T/2$, and returned symmetrically to zero,

$$\Omega_{\text{centre}} = \begin{cases} (\Omega_{\text{max}}/2)\,[1 - \cos(2\pi t/T)] & \text{if } t \le T, \\ 0 & \text{if } t > T, \end{cases} \tag{4.6}$$

twisting the affected region of the star, $\theta_{\text{c}} - \theta_{\text{HW}} \le \theta \le \theta_{\text{c}} + \theta_{\text{HW}}$, through an angle of $\psi = \Omega_{\text{max}} T/2$.

Here we use a grid of $N \times L = 320 \times 255$ on a domain $r \in [1, 40]$, and twist an annular region of the northern hemisphere, given by $\theta_{\text{ctr}} = 0.14\pi$, $\Delta = 0.04\pi$, and $\kappa = 60$. For this test we want only a small perturbation on top of a dipole field (small total twist amplitude), and for the magnetosphere to move slowly through a sequence of quasi-equilibrium states (requiring small Ω_{max} and large T), and so choose $\Omega_{\text{max}} = 10^{-3}$ and $T = 200$. This implants a twist of $\psi = 0.1$ radians by time $t = 200$, after which the stellar surface is at rest.

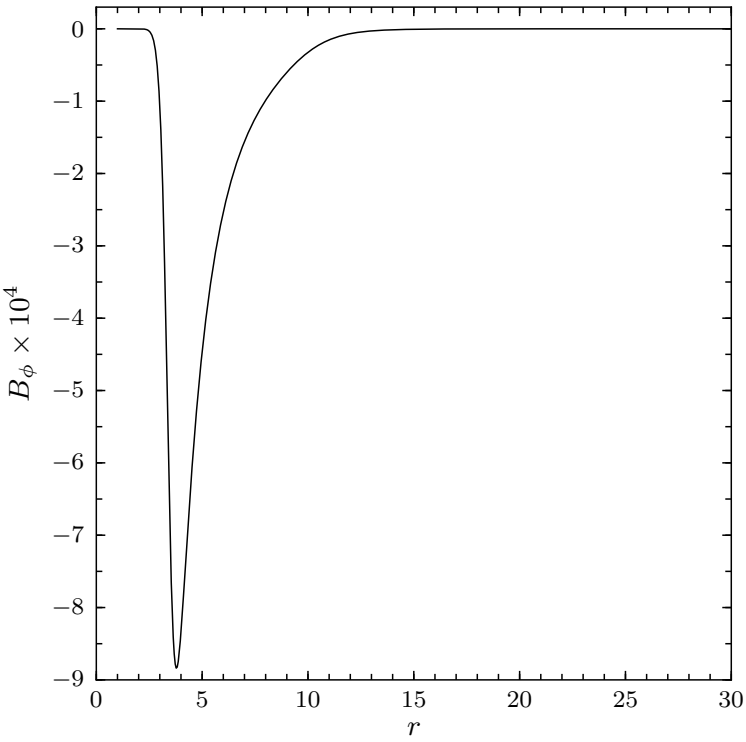

Figure 4.5: B_ϕ versus r at $\theta = \pi/2$, at $t = 200$, once the twist profile has been implanted; simulation (a) (see text).

To highlight the importance of the parallel current (see Section 3.8) we perform three simulations, all using the usual eighth-order SSV filtering with $\alpha_{\rm SSV} = 0.05$. In run (a) we use the full force-free current in the equations of motion, $\boldsymbol{J} = \boldsymbol{J}_\parallel + \boldsymbol{J}_{\rm drift}$ from Equation (3.27); in (b) we keep only the drift current, $\boldsymbol{J} = \boldsymbol{J}_{\rm drift}$, and remove accumulated $\boldsymbol{E}_\parallel$ at every Runge-Kutta substep; in (c) we use only the drift current, and remove parallel electric field at the end of every full time step. Run (a) therefore uses our standard numerical scheme; for this simulation the profile of B_ϕ along the equator at $t = 200$ is shown in Fig. 4.5.

The total twist along any field line can be calculated by numerically integrating its path in space: the twist is $\psi = |\phi_{\rm end} - \phi_{\rm start}|$. We label the field lines with the fractional flux function, u. It measures the fraction of the star's magnetic flux passing through a circular contour, centred on

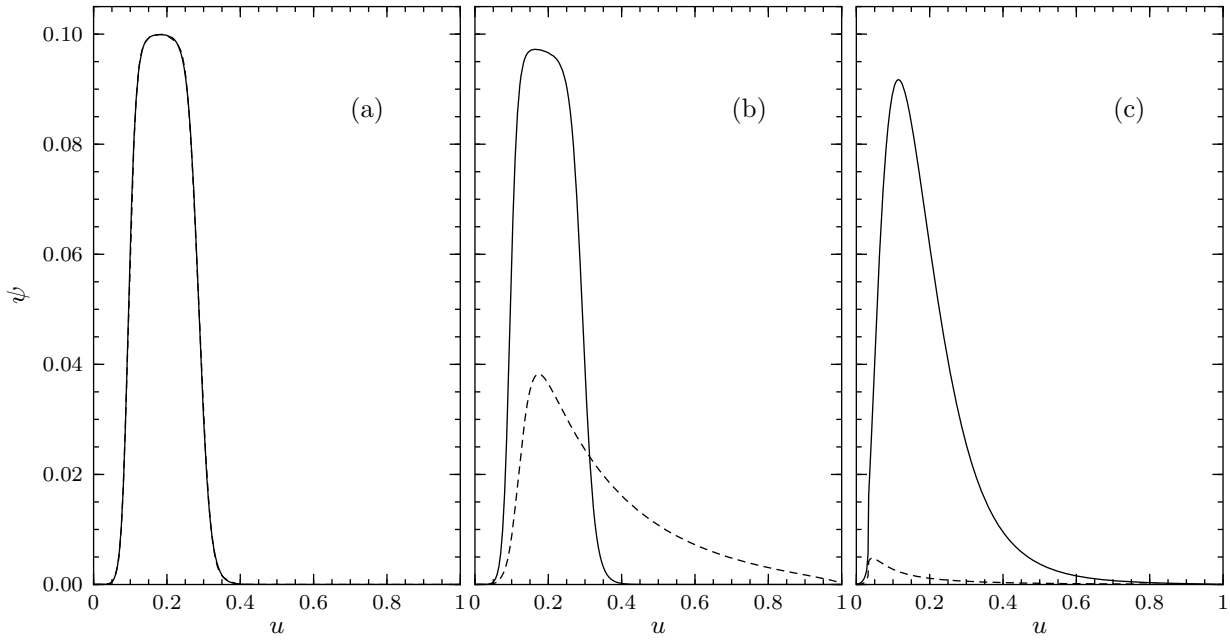

Figure 4.6: Twist versus flux function at $t = 200$ (solid line) and $t = 4000$ (dashed line), using (a) full current density, (b) drift current only, enforcing the $\boldsymbol{E} \cdot \boldsymbol{B} = 0$ constraint at each Runge-Kutta substep, (c) drift current only, enforcing the constraint at the end of a full time step.

the magnetic axis, which goes through the field line's footpoint. This function is $u = \sin^2 \theta_1$ for a dipole field, where θ_1 is the co-latitude of the northern footpoint, a relationship which is unchanged by any axisymmetric motion of the stellar surface.

Fig. 4.6 shows the measured twists at two times, $t = 200$ (solid lines) and $t = 4000$ (dashed lines), for each of the three runs. In run (a) the curves lie on top of one another; in fact, sharp twist profiles are preserved, at this resolution, even for tens of thousands of light-crossing times. In run (b) the profile at $t = 200$ is fairly good, but it has diffused significantly by $t = 4000$, while in run (c) the profile is already inaccurate by the earlier time and has almost completely disappeared by the later. These three cases demonstrate how important the parallel current is for self-consistently maintaining $\boldsymbol{E} \cdot \boldsymbol{B} = 0$ throughout each time step; omitting it implies making errors which demand

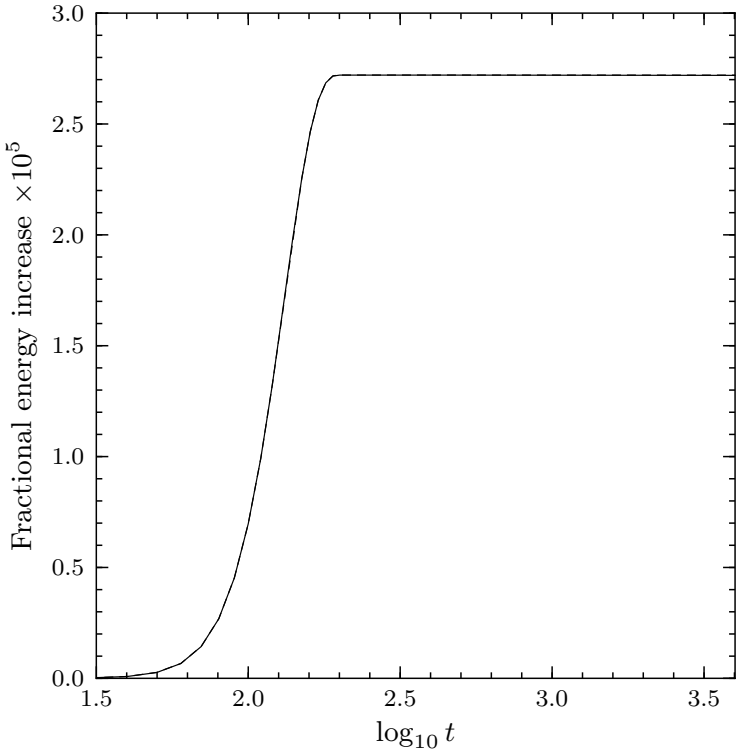

Figure 4.7: Energy injected into magnetosphere by twisting, in units of initial field energy: $W_1 - W_0$ (solid line), W_2 (dashed line). The lines lie on top of one another.

the removal of electric field from the system, introducing an artificial source of dissipation.

We focus now on run (a). The permanence of the implanted twist seen in Fig. 4.6 can also be interpreted in terms of energy conservation. Define W_1 as the total electromagnetic energy of the twisted configuration, instantaneously measured on the grid,

$$W_1 = \frac{1}{8\pi} \int_V \left(B^2 + E^2 \right) \mathrm{d}V, \tag{4.7}$$

where V is the three-dimensional volume of the computational domain. Let W_0 be the energy for the initial pure dipole field. Define W_2 as the energy calculated by numerically integrating the net

energy flux into the grid over time,

$$W_2 = \int_0^t \left[L_{\rm in}(t') - L_{\rm out}(t')\right] \mathrm{d}t', \tag{4.8}$$

where $L = (r^2/2) \int_0^\pi (\boldsymbol{E} \times \boldsymbol{B})_r \sin\theta \, \mathrm{d}\theta$, and the subscripts denote the luminosity measured at the inner and outer boundary surfaces respectively. The part of the grid containing the absorbing layer is not included in calculating W_1 or W_2. If there is no dissipation, we expect $W_1 - W_0 = W_2$.

The measured twist energy $W_1 - W_0$ and the imparted energy W_2 are shown in Fig. 4.7 as a function of time. The two curves lie on top of one another to within their widths. The energy on the grid is stable once the twist-up phase is complete. Between $t = 200$ and $t = 4000$ the twist energy decreases by a factor of 4×10^{-4}. At $t = 4000$, the fractional difference between accumulated twist energy and integrated net Poynting flux is $[(W_1 - W_0) - W_2] \,/\, (W_1 - W_0) = -3.9 \times 10^{-4}$.

Chapter 5

The aligned rotator

The objects which we observe as radio pulsars are generally accepted to be rotating magnetised neutron stars, with magnetic fields of 10^{12} G or more. If the star's basic magnetic field is modeled as a dipole, the electric field induced by its rotation, Equation (3.31), has a large radial component, which is able to rip charged particles from the stellar surface (Goldreich & Julian 1969). These particles, and $e^{\pm}$ pairs created in the magnetosphere, surround the star with force-free plasma.

To simplify the problem, one can consider the axisymmetric configuration with the magnetic moment parallel to the rotation axis: the aligned rotator. A steady-state solution was found by Contopoulos et al. (1999), which included a region of closed field lines extending to the light cylinder (defined as the cylindrical radius $R_{\rm LC}$ at which the co-rotation speed is the speed of light), and open, asymptotically monopolar, field lines extending to infinity. A current sheet is present at the equator beyond the light cylinder, which splits at the 'Y-point' to follow the last closed flux surface. It was later found that equilibrium solutions exist with the Y-point at any distance from the rotation axis, within the light cylinder (Goodwin et al. 2004; Contopoulos 2005; Timokhin 2006).

Absent a unique solution, one turns to time-dependent studies (Spitkovsky 2006; McKinney 2006b; Komissarov 2006; Kalapotharakos & Contopoulos 2009; Yu 2011), which have all found that the Y-point moves toward the light cylinder, where it subsequently remains.

5.1 Numerical setup

In our simulations, spatial and temporal scales are fixed by setting $r_* = 1$, where r_* is the stellar radius; $c = 1$ as usual. The star is smoothly spun up from rest, over a few light-crossing times, to the rotation frequency Ω, implying a light cylinder at $R_{\rm LC} = 1/\Omega$. We have investigated cases with $R_{\rm LC} = 5$, 10, 20, 30, and 40; there were only minor differences between the solutions, and here we concentrate on those with $R_{\rm LC} = 20$. In these we set $r_{\rm out} = 70$, with an absorbing layer beginning at $r = 50$. A radial arcsine-plus-algebraic coordinate mapping is used, with parameters $\epsilon = 10^{-11}$ and $Q = 0.7$; the grid is equally spaced in the meridional direction. The current sheets are captured with super spectral viscosity (SSV): $2p_{\rm SSV} = 8$, $\alpha_{\rm SSV} = 0.1$.

We have performed simulations with a range of grid sizes, from $N \times L = 81 \times 49$ to 768×507. In order to demonstrate that a very fine grid is not needed for accurate results, we will concentrate on our run with a grid of 384×255—all the following results are from this simulation unless explicitly noted otherwise. The behaviour is similar for all resolutions.

The initial magnetic field is set to a dipole with unit magnetic moment: $B_r = 2\cos\theta/r^3$, $B_\theta = \sin\theta/r^3$, $B_\phi = 0$. Rotation is introduced by applying an electric field, Equation (3.31), bringing the star to its final angular velocity by $t = 10$.

5.2 Evolution to the steady state

During the spin-up phase, Alfvén waves are injected along field lines into the magnetosphere, filling it with charges and currents. The magnetic energy increases and the poloidal field lines inflate, appearing to be pulled outward along the equator, where waves from opposite hemispheres meet. By $t = 40$ a clear distinction can be seen between those lines which are too close to the star to be strongly inflated, and those which are on the path to opening. Waves can be seen to propagate backwards and forwards on the former, while the latter are expanding at nearly the speed of light. Poloidal field line projections are drawn in Fig. 5.1.

Let us define the closed zone as that region with oscillating fields and currents, and the Y-point as the point at which the currents around the closed zone meet at the equator. All the fields are smooth before $t \approx 35$, at which time the radial and azimuthal currents collapse to an unresolved equatorial current sheet. As the configuration evolves, all components of the current around the closed zone, and in the equatorial current sheet, change direction, as one can see in Fig. 5.2. The Y-point is at $r = 0.5R_{\rm LC}$ before the reconfiguration, and at $r = 0.6R_{\rm LC}$ afterwards. By $t = 90$ Alfvén waves have again filled the closed zone up to the Y-point, and the quasi-steady march of the Y-point to the light cylinder begins.

The equatorial current sheet is unresolved on the grid, having discontinuities in both B_r and B_ϕ; Fig. 3.2(a) shows B_ϕ against θ at $r = 22.25$. The total magnetic field drops close to zero directly on the equator, and so the magnitudes of any electric fields must be reduced so that the second force-free condition, Equation (2.15), is not violated (see Section 3.8). When choosing the grid, it is advantageous to have a line of nodes on the equator, otherwise the current sheet must choose which of the two equidistant nodes to collapse onto—since neither choice obeys the symmetry of the problem, the sheet will periodically move from one line to the other.

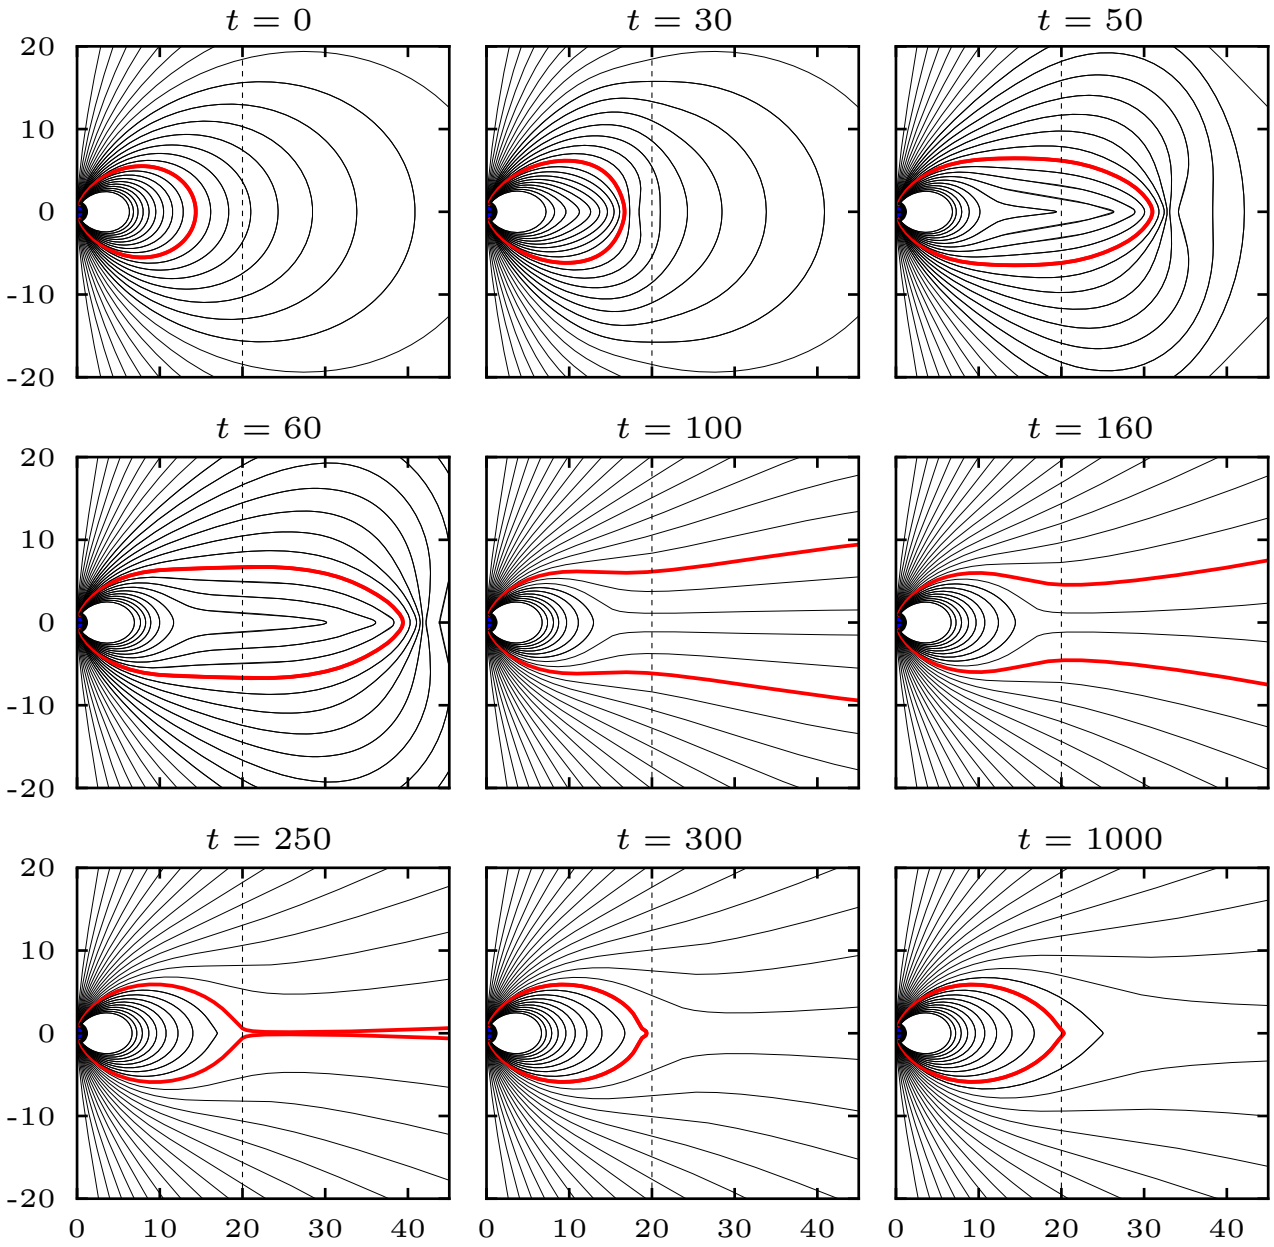

Figure 5.1: Poloidal field lines for the aligned rotator, from non-rotating at $t = 0$ to the equilibrium solution at $t = 1000$. The thicker, red field line closes at the Y-point in the equilibrium solution; the dashed line at $R = 20$ indicates the light cylinder. There are 25 lines drawn from each pole, evenly spaced in colatitude between 0.01π and 0.13π.

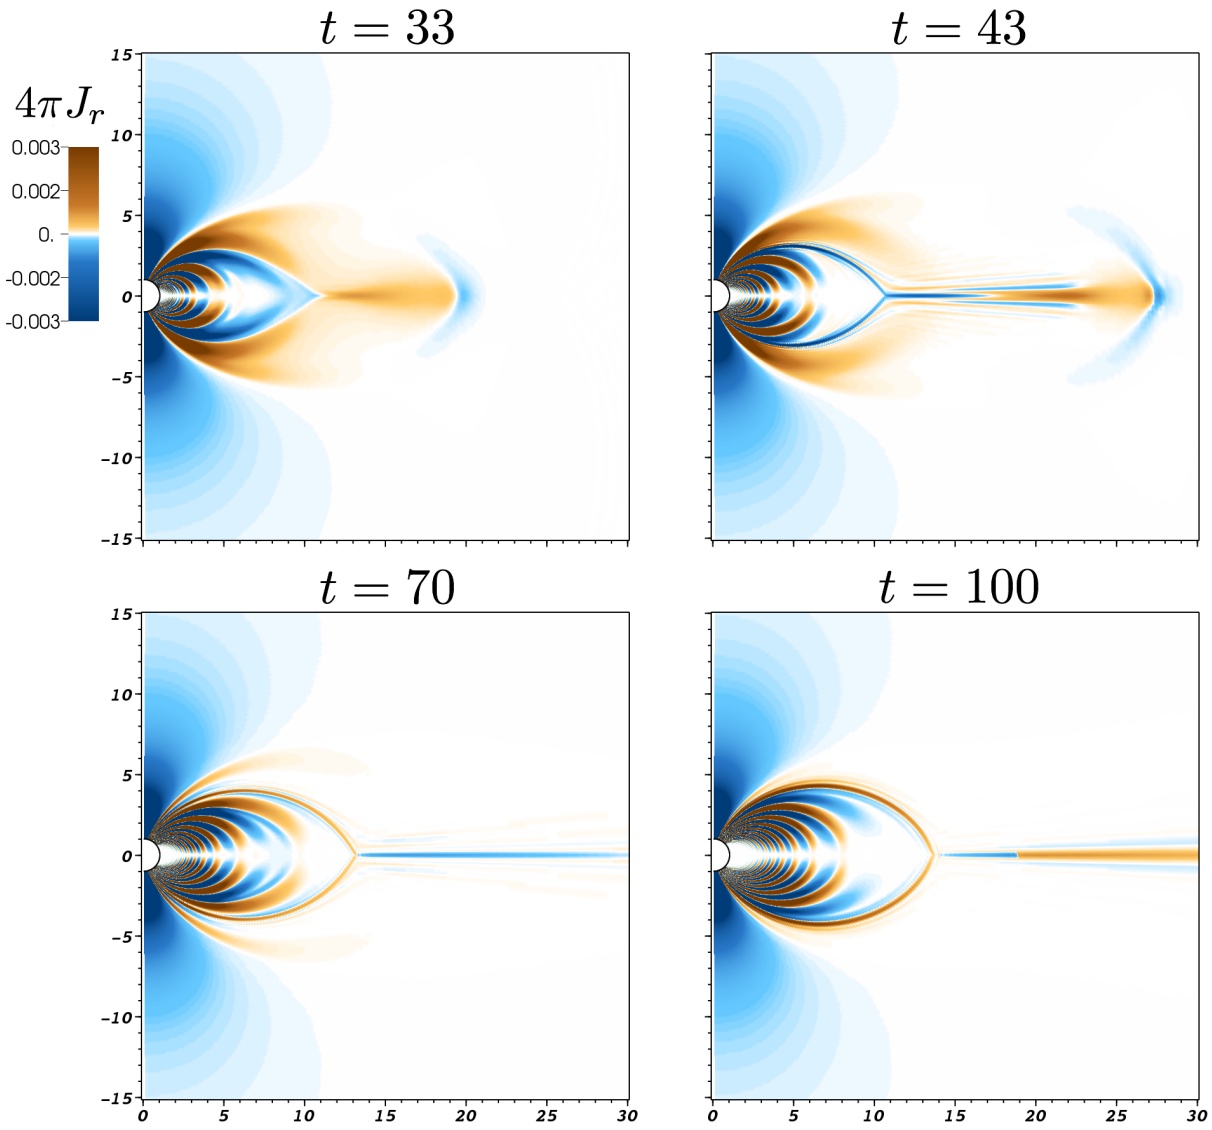

Figure 5.2: Formation of the Y-point. The color shows radial current density; gold is positive, blue is negative. Charge flows into, and accumulates at, the Y-point, until the Y-point reaches the light cylinder. The light cylinder is at $R = 20$.

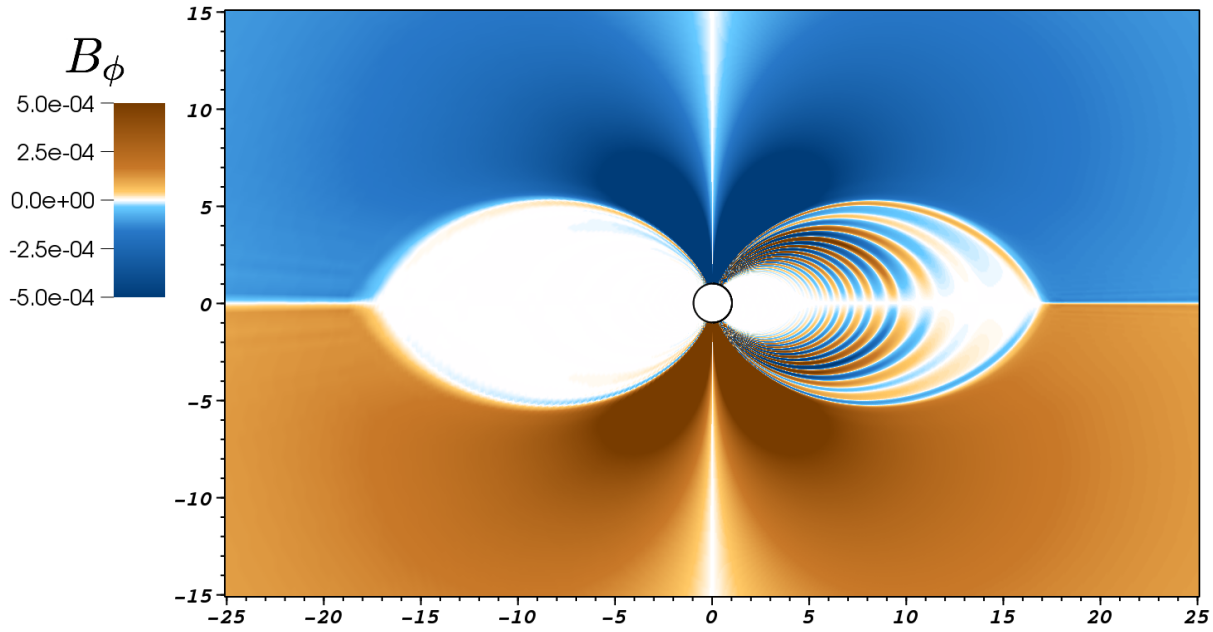

Figure 5.3: Comparing low resolution (256×155), left, and high resolution (768×507), right, simulations at $t = 250$. The color is toroidal magnetic field. The Y-point has not yet reached the light cylinder at $R = 20$.

At $t \sim 100$, more open flux has been created than can be supported by the energy being pumped into the magnetosphere by the star's rotation. The Y-point moves very slowly towards the light cylinder, as some of the open field lines reconnect in the current sheet. The Alfvén waves in the closed zone are eventually damped by numerical dissipation[1]; Fig. 5.3 shows the toroidal magnetic field for low (256×155) and high (768×507) resolution runs at $t = 250$. The oscillations survive longer with higher resolution, but the Y-point moves only slightly slower; here, it is at $r = 16.8$, versus 17.5 for the low-resolution run.

The Y-point reaches the light cylinder at $t \sim 500$, but open flux continues to reconnect slowly until $t \sim 700$, about 5.5 rotation periods. After this time the solution is stationary. The time taken to reach equilibrium is practically independent of grid size. The equilibrium solution has some

[1]The waves are sheared by field line curvature, becoming longer and thinner. Numerical dissipation becomes significant when they approach the grid scale, and are attenuated by the filters.

field lines closed outside the light cylinder due to the effective resistivity of the current sheet; for example see the last panel of Fig. 5.1 or Fig. 5.6 below. Note however that the zero-toroidal-field closed zone is strictly within the light cylinder.

5.3 Steady state

We now concentrate on the obtained equilibrium solution. Unlike numerical solutions in previous works, ours does not exhibit plasmoid emission by the Y-point once it has approached the light cylinder. The Y-point and current sheet are steady for long times; we ran a 256×155 simulation until $t = 10,000$, or 80 rotational periods, without seeing any indications of Y-point instability. We found plasmoid emission from the Y-point in only two circumstances. Firstly, if the level of filtering was too low; no plasmoids were seen if the filtering was strong enough to prevent the Gibbs oscillations on either side of the current sheets from growing over time. Secondly, if the radial grid spacing was too large near the light cylinder; we found the Y-point was stable if $\Delta r \lesssim 0.75\, r\Delta\theta$. This is probably due to the action of the spectral filters, which make the current sheet mildly resistive, slightly diffusing the Y-point.

The current sheet resistivity largely comes from filtering B^θ in the meridional direction, across the sheet. As the sheet forms, $B^\theta(\theta)$ goes to zero at the equator, with a cusp-like profile on each side, implying significant high-wavenumber content. The filters damp these high wavenumbers, causing the smoothed B^θ to be non-zero on the equator, which closes field lines. Eventually an equilibrium is reached between filtering and the electromagnetic forces trying to compress the current sheet. We verified this picture by using unfiltered values of B^θ whenever the $B^2 - E^2 > 0$ condition is violated; this resulted in near-zero magnetic field in the current sheet, and a Y-point that moved outwards much more slowly. However, without filtering the evolution eventually became unstable.

To investigate the dissipation in the current sheet, we performed a simulation using a 576×255 grid, with an outer boundary at $r = 1000$, an absorbing layer beginning at $r = 700$, and a more severe coordinate mapping; again, $R_{\rm LC} = 20$. In Fig. 5.4 we plot integrated Poynting flux through concentric spheres, as a function of radius. The outgoing flux is constant within the light cylinder, with maximum fractional variation of about 3×10^{-4} near the star. The flux inside the light cylinder agrees with the value found in previous works, $\mu^2\Omega^4/c^3$, where μ is the star's magnetic moment, to fractional accuracy of 6×10^{-3} for our fiducial simulation and 2×10^{-3} for our high-resolution one.

Outside the light cylinder, some of the outgoing flux is lost in the current sheet due to its effective resistivity due to the filters; this deficit decreases with increased resolution. The energy loss is relatively large, because the resistivity is confined to the current sheet, which is kept sharp by the fully ideal surrounding magnetosphere. Solutions with global resistivity dissipate a smaller fraction of their luminosity in the current sheet (Kalapotharakos et al. 2011). As shown in Fig. 5.5, the open field lines are asymptotically radial.

The magnetic field lines closing inside the Y-point form the closed zone, with no poloidal currents and zero toroidal magnetic field. The density of field lines is higher around the boundary of the closed zone. This can be seen in Fig. 5.6, where we plot field lines and contours of toroidal magnetic field. The closed field lines outside the light cylinder form cusps at the equatorial current sheet (we do not consider these lines to be part of the 'closed zone' because particles are not trapped on them, being able to flow out through the current sheet). The total magnetic flux through the light cylinder, in one hemisphere, is $1.39 \pm 0.01\ \mu\Omega/c$. This is larger than the value, 1.23, obtained by Contopoulos (2005) and Timokhin (2006) for ideal steady-state force-free models, because some of the closed flux in our solution has diffused through the light cylinder due to finite resistivity.

The magnetosphere has current leaving the polar cap and, outside the light cylinder, returning

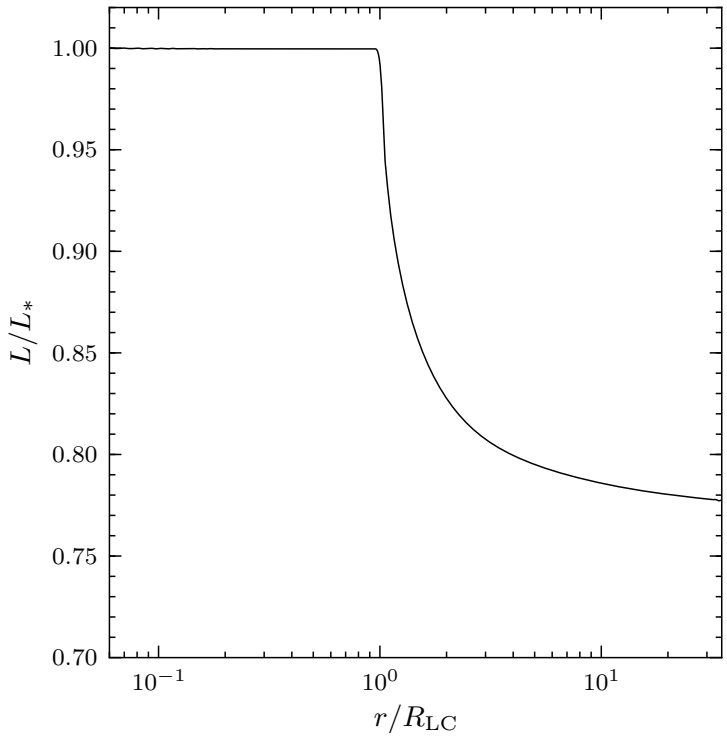

Figure 5.4: Luminosity (integrated Poynting flux) through a sphere of radius r, as a function of r. L_* is the luminosity measured at the stellar surface; it equals $\mu^2\Omega^4/c^3$ (to a fractional accuracy of 6×10^{-3}).

to the star mostly in the current sheet. At the Y-point the current sheet splits in two, and follows the boundary of the closed zone. In Fig. 5.8 we plot the poloidal current density measured on the star, normalised to the Goldreich-Julian current density (the speed of light times the equilibrium charge density); a similar plot was obtained for the steady-state solution by Timokhin (2006). To aid comparison with Fig. 5 of that paper, we have scaled the axis to the width of the polar cap. The current sheet is seen just inside $\theta/\theta_{\rm pc} = 1$.

The three components of the current density, and the charge density, of the equilibrium solution are shown to a common scale in Fig. 5.7. Some Gibbs oscillations are apparent near the Y-point and beside the outer current sheet, but they are controlled by the filtering and do not appear to cause any problems. The closed zone has no poloidal current, but there is toroidal current

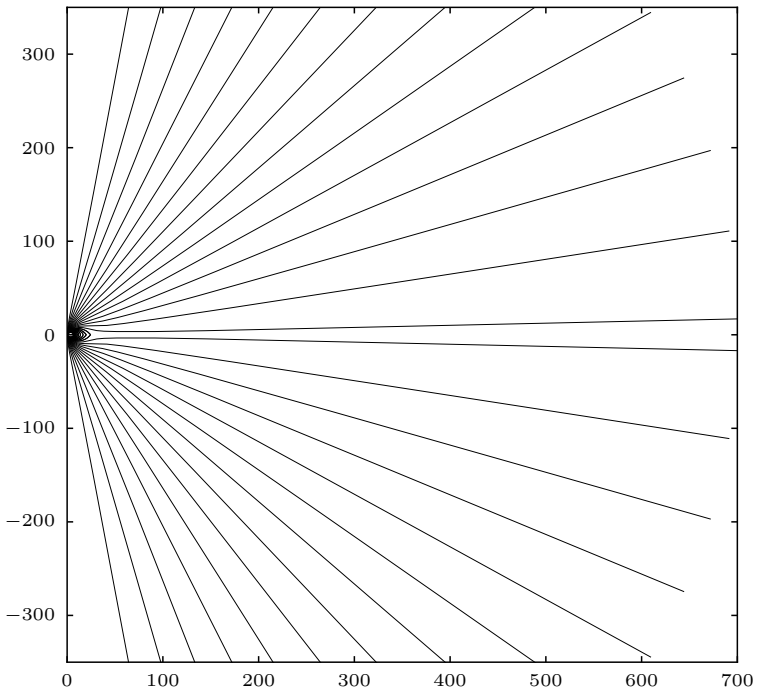

Figure 5.5: Poloidal field lines, out to $r = 35R_{\rm LC}$. The field lines have the same footpoints as in Fig. 5.1.

from the corotating charge density.

In a steady, ideal, force-free magnetosphere the electric field has a potential that is constant along the magnetic field lines. Our solution is everywhere close to this behaviour (except in the equatorial current sheet where numerical resistivity is significant). The equipotentials follow the shape of magnetic field lines, and are smooth and accurate. The charge density corresponds to the Laplacian of the potential, and the second derivatives have more numerical noise. Nevertheless, the obtained charge density reproduces all the expected features, including steep gradients near the current sheet and the singular behaviour at the Y-point. The current sheet is positively charged outside the Y-point and negatively charged around the closed zone, in agreement with the previously obtained steady-state solutions (see Timokhin 2006). The negatively-charged current sheet around the closed zone appears to be resolved, and the thickness of the negatively-charged region decreases slowly with increased resolution. This suggests a thickening of the current sheet due to

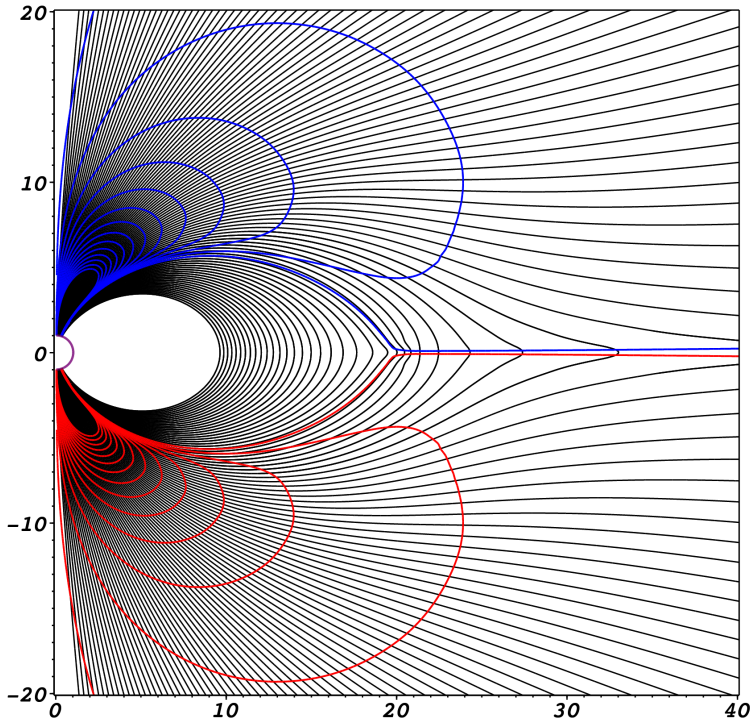

Figure 5.6: Poloidal field lines and contours of constant B_ϕ. There are 100 lines drawn from each pole, equally spaced in colatitude between $\theta = 0.005\pi$ and 0.11π. Contours, 14 in each hemisphere, are drawn in blue for negative values and red for positive values. They show $|B_\phi|$ equally spaced between 3×10^{-5} and 10^{-3}.

finite resistivity, as argued by Gruzinov (2011). The thickening must occur due to resistivity near and outside the Y-point, as dissipation inside the light cylinder is negligible (see Fig. 5.4).

5.4 Viability of the force-free model

The force-free model relies on the availability of charges that sustain the required charge density and electric currents. Charges can be pulled out from the star or supplied by $e^{\pm}$ pair creation. Both mechanisms require a longitudinal voltage, i.e. $\boldsymbol{E} \cdot \boldsymbol{B} \neq 0$; pair creation, in particular, requires a significant voltage. In most observed pulsars, this voltage is not so large as to make the force-free approximation unreasonable. A real danger for the force-free model appears if the required charge

density or current cannot be created.

The electron-positron discharge operates at the polar cap if $\alpha^{-1} < 1$, where $\alpha \equiv J/c\rho_e$ (Beloborodov 2008). This condition is satisfied in the zone of return current ($J_r > 0$), near the edge of the polar cap where $\alpha < 0$ (Fig. 5.8). Pairs are created with a high multiplicity and outflow along the magnetic field lines, screening $\boldsymbol{E} \cdot \boldsymbol{B}$. The presence of dense pair plasma makes the force-free approximation safe in the return-current zone (except in the current sheet). The boundary of this zone ($J_r = 0$) is shown by the blue curve in Fig. 5.9.

The discharge does not occur in the central parts of the polar cap where $J_r < 0$ and $0 < \alpha < 1$. Instead, the required ρ_e and J are supplied by the charge-separated flow pulled out from the star (Beloborodov 2008; Chen & Beloborodov, in preparation). The force-free approximation remains accurate along the field lines extending from this region as long as $0 < \alpha < 1$ remains satisfied. We find that $\alpha < 1$ everywhere in the zone $J_r < 0$, inside and outside of the light cylinder. However, $\alpha > 0$ is *not* satisfied. There is a small region in the zone $J_r > 0$ outside the light cylinder where α becomes negative (because ρ_e changes sign, see Fig. 5.9). The charge-separated outflow passing through this region fails to supply the charge density of the required sign, and a large $\boldsymbol{E} \cdot \boldsymbol{B}$ must develop. $\boldsymbol{E} \cdot \boldsymbol{B}$ may be limited by locally initiated pair creation in young, fast-rotating pulsars; however, in most pulsars pair creation is inefficient so far from the neutron star.

We conclude that, for the aligned rotator, the force-free approximation is expected to fail in the region where $J_r < 0$ and $\rho_e > 0$. This region is, however, small, and this problem may not impact the obtained global solution. We note that a similar, but larger, region is seen in Fig. 4 of Timokhin (2006) and in Fig. 3 of Contopoulos et al. (1999). The difference between their and our solutions is due to the fact that their model is strictly ideal everywhere, while our model is (nearly) ideal only outside the equatorial current sheet. The dissipation in the current sheet affects the magnetosphere as discussed above and shrinks the region that is dangerous for the force-free

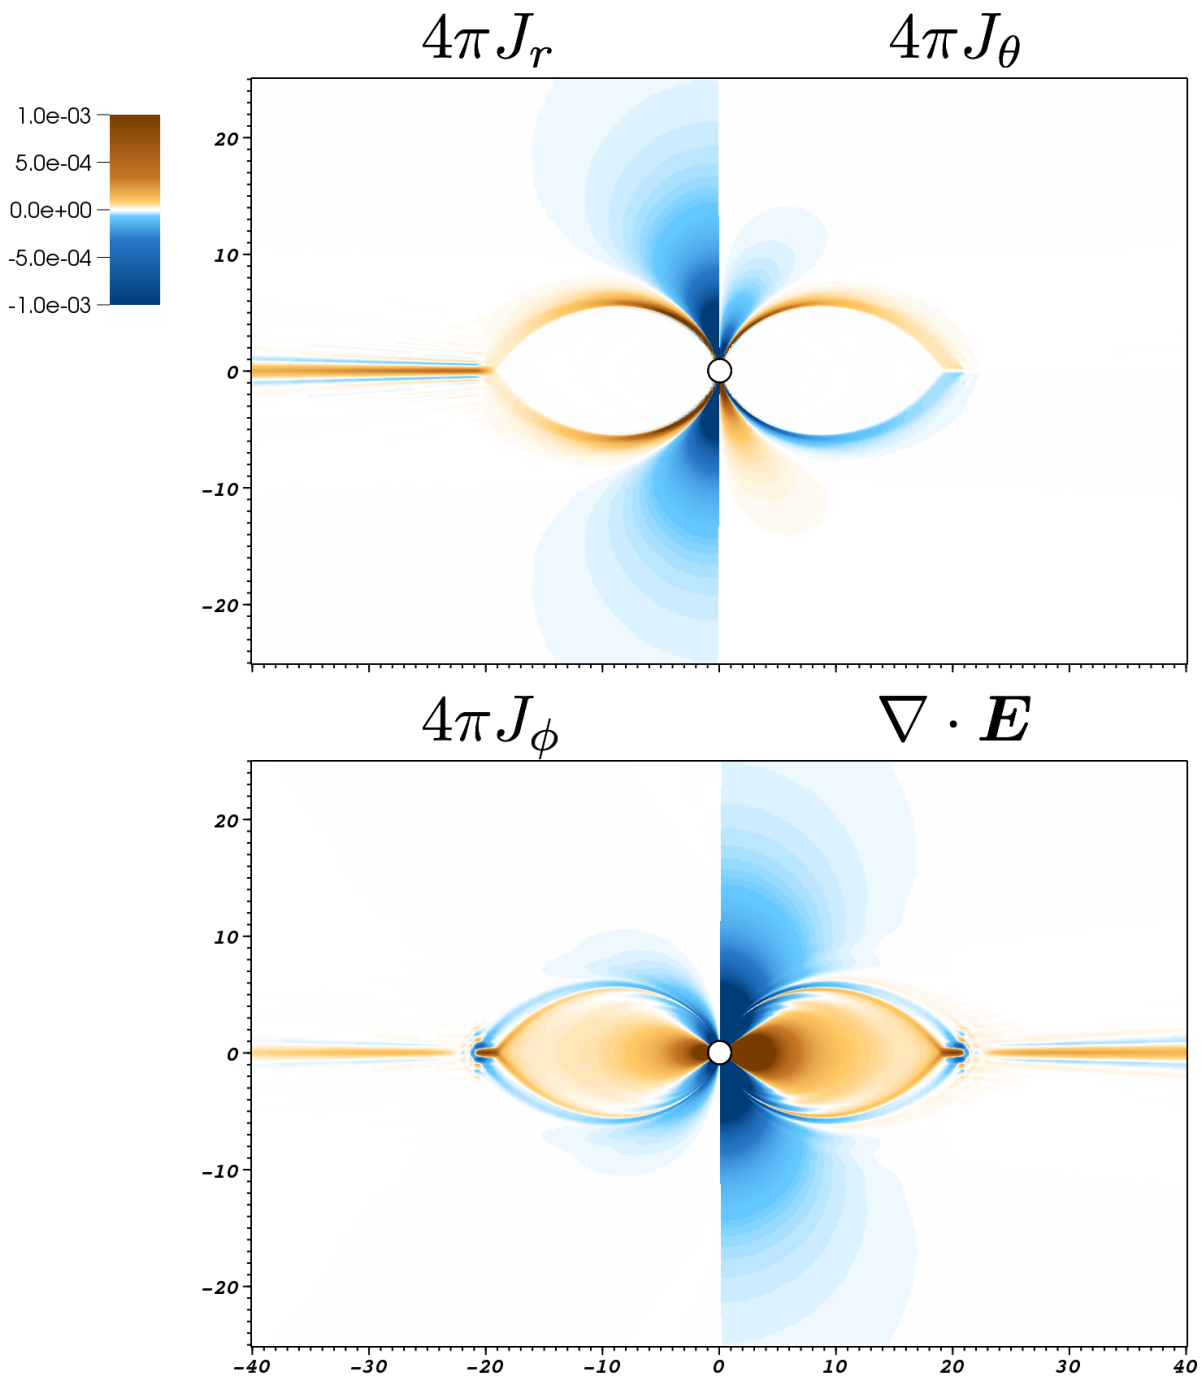

Figure 5.7: Steady-state current, $\boldsymbol{J} \equiv \nabla \times \boldsymbol{B}$, and electric charge distributions of the equilibrium solution.

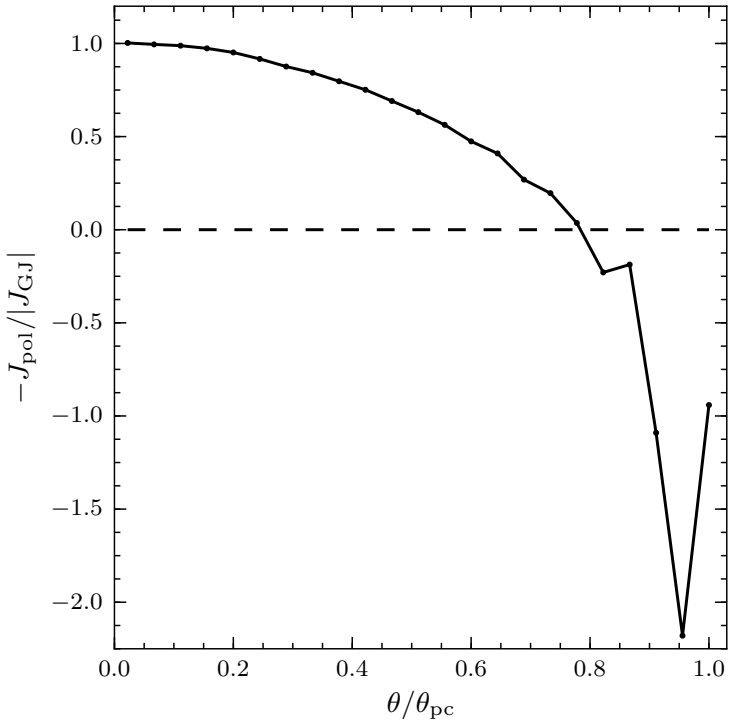

Figure 5.8: Normalised poloidal current density across the polar cap. $J_{\rm GJ} = \rho_{\rm GJ}$, $\rho_{\rm GJ} = -\boldsymbol{\Omega} \cdot \boldsymbol{B}/(2\pi c)$. $\theta_{\rm pc} = 0.088\pi$ is the half-width of the polar cap, defined by the footpoint of the last field line in the closed zone.

model outside the current sheet.

We note also that the polar-cap accelerator in the zone of return current $J_r > 0$ can supply some $e^{\pm}$ pairs to the neighbouring field lines with $J_r < 0$ (near the blue curve in Fig. 5.9). Pair creation is not exactly local to the acceleration region because it involves an intermediate step– the emission of a high-energy photon which must propagate a finite distance *across* the field lines before converting to $e^{\pm}$. Pairs created on, and outflowing along, the field lines slightly outside the return-current zone may supply enough positive charges to the problematic small region ($J_r < 0$, $\rho_e > 0$) and validate the force-free condition there.

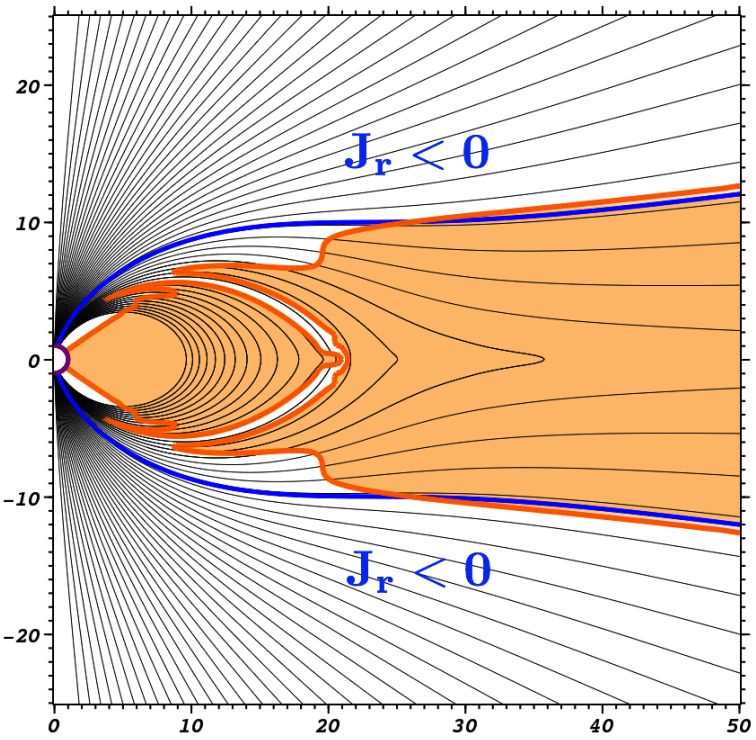

Figure 5.9: The blue line separates the outflowing current region, with $J_r < 0$, near the poles, and the return current region nearer the equator. The orange line is the contour of zero charge density ρ_e, and the shaded region has $\rho_e > 0$. Poloidal field lines are drawn in black. The charge density has been cleaned with the DTV filter (Section 3.7) to remove Gibbs oscillations near the equatorial current sheet.

Chapter 6

Dynamics of strongly twisted magnetospheres

6.1 Introduction

Force-free magnetic fields, embedded in low-pressure perfectly conducting plasma, can store energy when the field is twisted into a non-potential state. This deformation can be the result of smooth shearing of a boundary surface, such as a heavy plasma or a conducting solid, in which the field's footpoints are frozen. The sudden release of the stored energy is a compelling model for solar flares and coronal mass ejections: energy is accumulated in the coronal force-free field due to differential rotation of the heavy photospheric plasma, and liberated following a loss of magnetic equilibrium. The dynamic behavior of the solar corona is appropriately studied using non-relativistic magnetohydrodynamics (MHD), as there the Alfvén speed is much smaller than the speed of light (e.g. Tomczyk et al. 2007).

Magnetars, neutron stars with ultra-strong magnetic fields $B \gtrsim 10^{14}$ G, also release magnetic energy in sudden bursts and flares (e.g. Woods & Thompson 2006; Mereghetti 2008). An evaluation of the ability of the magnetospheres of these objects to store magnetic energy, and the stability of the resulting field configurations to energy release on dynamic timescales, requires study of the twisting problem in the relativistic plasma limit,

$$\frac{B^2}{8\pi} \gg \rho_{\rm m} c^2 \,. \tag{6.1}$$

In the magnetospheres of compact objects the Alfvén speed is very nearly the speed of light, and the Lorentz force density vanishes even during rapid dynamic motion. Here, the evolution of the magnetosphere is described by force-free electrodynamics (also known as relativistic force-free MHD or magnetodynamics).

When the magnetosphere is axisymmetric, the "twist" ψ can be measured as the azimuthal angular displacement between a field line's two footpoints due to their relative shearing motion; $\psi = 0$ on all lines in the potential state. The study of current-carrying magnetar field configurations has largely been restricted to either the linear, weakly sheared limit $\psi \lesssim 1$ (Beloborodov 2009b), or the case of self-similar fields (Thompson et al. 2002; Pavan et al. 2009). While a self-similar model allows one to generate formal magnetospheric solutions with different amounts of shear, they cannot be connected as a realistic evolutionary sequence because this sequence would require compressive motions of the footpoints of magnetic field lines, and thus compression of the crust into which the field lines are frozen; this is forbidden as a neutron star's crust is nearly incompressible.

Additionally, magnetar magnetospheres are not globally twisted—the shear is expected to be confined to a fraction of the stellar surface (Beloborodov 2009b). Further progress in magnetar

theory requires the development of a fully nonlinear model for twists in realistic geometries, their stored energy, and the dynamical release of this energy when the field lines are overtwisted and force-free equilibrium is lost. One would also like to know the effect of the twists on the spin-down rate of the star, as will be discussed in Chapter 7.

6.1.1 Sequence of equilibria

As long as stellar rotation is neglected and purely static (equilibrium) configurations are considered, relativistic and non-relativistic force-free magnetospheres are described by the same equations. These equilibria have been extensively studied in the context of the solar corona. An equilibrium force-free configuration satisfies

$$(\nabla \times \boldsymbol{B}) \times \boldsymbol{B} = 0, \tag{6.2}$$

and so the magnetic field and steady-state current are parallel,

$$\nabla \times \boldsymbol{B} = \alpha(\mathcal{L})\boldsymbol{B}, \tag{6.3}$$

where α is constant on each field line, labeled by $\mathcal{L}$. The boundary value problem posed by Equation (6.2) can be investigated by setting $\alpha = \lambda\, g(\mathcal{L})$, where λ is a constant and $g(\mathcal{L})$ is some prescribed function that is constant along field lines[1]. A sequence of equilibrium states can be constructed by increasing λ from zero (the potential field). Under quite general conditions, it can be shown using a virial theorem approach that there is a limiting value λ^*, such that there are no

[1]In axial symmetry, the poloidal field can be written in terms of one vector potential component A_ϕ as $\boldsymbol{B}_\mathrm{p} = \nabla A_\phi \times \hat{\boldsymbol{\phi}}/r\sin\theta$. (This $A_\phi = f/2\pi$, where f is the poloidal flux function defined in Equation (6.5)) Equation (6.2) is satisfied if the toroidal field has the form $B_\phi = F(A_\phi)/r\sin\theta$; then $\alpha = \mathrm{dF/dA}_\phi$. A popular method of constructing a sequence of equilibria is by imposing a power-law relation between F and A_ϕ, $F = \lambda A_\phi^\gamma$ (Raadu 1971; Milsom & Wright 1976).

equilibrium solutions for $\lambda > \lambda^*$ (Aly 1984). It was argued that reaching this limit implied the formation of current sheets (Milsom & Wright 1976) and the onset of dynamic behavior, which could trigger a solar flare (Low 1977).

Using virial relations, it can also be proven that there is a maximum energy associated with the sequence of force-free fields with a fixed normal magnetic field distribution on the stellar surface (Aly 1984), and that the field lines become open to infinity for this maximum energy configuration (Aly 1991; Sturrock 1991). For a twisted dipole configuration, the related asymptotic open field is a split-monopole-like state having the same surface flux. Evaluating the energy of twisted configurations is useful as it shows how much energy can be stored in the non-potential magnetosphere; it also helps identify possible spontaneous transitions between twisted states. All field lines cannot spontaneously open to infinity, since the fully open field has the maximum energy. However, it is possible for a closed magnetic field to have higher free energy than a partially open field which is accessible by ideal MHD plasma motions (Wolfson & Low 1992). Therefore spontaneous partial opening is energetically allowed—the closed field can be unstable in an absolute sense.

A more physical way of constructing twisted equilibria is by specifying the shear, or angular separation between footpoints, of each field line; then a sequence of solutions can be generated by monotonically increasing the shear (Low 1978). Shearing the field line footpoints increases the magnetic pressure and energy of the field, causing the field lines to expand outward. As the field lines expand, a current layer forms, which becomes thinner and whose current density grows as the open-field configuration is approached. Following this sequence, one finds that the field lines open to infinity and that the layer becomes a current sheet, infinitely thin and with infinite current density, implying a tangential discontinuity in the magnetic field (Wolfson 1989, 1995). In the presence of any non-zero resistivity, these current layers would eventually be subject to reconnection, triggering dynamic motion and the release of stored energy in such an "overtwisted"

magnetosphere.

The critical-point behavior of overtwisted magnetospheres has been investigated using several techniques. The α-specified equilibrium sequence ends at a maximum λ^*, as described above. Using a perturbative expansion around the dipole potential field outside a magnetized sphere, Low (1986) showed that some smooth shearing profiles result in unbounded plasma displacements far from the sphere even for infinitesimally small shear. The field line expansion accelerates dramatically above a certain shear, as seen in studies using the magnetofrictional method (Roumeliotis et al. 1994) and a self-similar model (Wolfson 1995). Uzdensky (2002) argues from principles of magnetostatic balance that field configurations, with footpoints frozen into spheres, will open to infinity at a finite shear (i.e. that no equilibrium solutions with closed field lines exist, where the field lines' footpoints have angular separation larger than a critical value). This is equivalent to saying that a finite shear suffices to take the energy of the field configuration to its theoretical maximum.

6.1.2 Dynamics of overtwisted configurations

The shearing problem has also been studied using time-dependent numerical MHD simulations, with parameters suitable for the solar corona. Simulations have the advantage of naturally testing the stability of each equilibrium state. Resistive simulations in a Cartesian geometry show that strongly sheared fields eventually form current sheets, undergo reconnection, and eject large plasmoids (Mikic et al. 1988; Biskamp & Welter 1989; Finn et al. 1992). Differential azimuthal footpoint motion of a dipole field, in opposite directions above and below the equator in a spherical geometry, eventually leads to accelerating field line expansion and unstable eruption (Steinolfson 1991).

The question of critical shear was investigated in detail by Mikic & Linker (1994, henceforth

ML94). Using ideal MHD simulations, they showed that smoothly sheared field configurations first move, in a quasi-steady manner, through a series of quasi-equilibrium states. They proposed that at a certain value of the applied shear the field enters a state of magnetic non-equilibrium and the field lines open to infinity, eventually forming a tangential discontinuity in the magnetic field (current sheet). The critical shear also separates distinct responses to resistivity: before the critical state is reached, the field slowly relaxes if resistivity is introduced, while after this point the introduction of a resistive term leads to reconnection in the current layer and the expulsion of a plasmoid. In their simulations, no eruptive behavior occurs during ideal evolution—after entering the non-equilibrium state the field is conjectured to transition smoothly to the (partially) open configuration (although this behavior is not observed, possibly due to the long equilibration timescales of inflated field lines).

6.1.3 This chapter

In this chapter, we study the dynamics of twisted relativistic magnetospheres (Equation (6.1)) using time-dependent numerical simulations with the new simulation code PHAEDRA (described in Chapter 3). The formulation of the problem is given in Section 6.2, where we write down the equations of force-free electrodynamics, and describe the initial and boundary conditions. The numerical method is briefly described in Section 6.3. In Section 6.4, we concentrate on non-rotating stars, and consider the sequence of nonlinear equilibria that is obtained by gradual shearing of the magnetic footpoints, the resulting inflation of poloidal field lines, and the free energy stored in the twisted magnetosphere. In Section 6.5 we describe the dynamic phase, entered when the magnetosphere becomes "overtwisted" and loses equilibrium. We address the critical twist amplitude at which equilibrium is lost and its dependence on the shearing rate and profile, the reconnection rate following the formation of the current sheet, and the relative amounts of energy expelled and

dissipated.

6.2 Problem formulation

6.2.1 Model setup

The initial state in each simulation is a dipole magnetic field at rest, in a computational domain between the perfectly conducting stellar surface, at $r = r_\star$, and the outer boundary at $r_{\rm out}$. This state has total energy W_0 in the computational domain,

$$W_0\left(r_\star, r_{\rm out}\right) = \frac{\mu^2}{3}\left(\frac{1}{r_\star^3} - \frac{1}{r_{\rm out}^3}\right). \tag{6.4}$$

We use spherical coordinates (r, θ, ϕ) with the polar axis aligned with the magnetic dipole axis. Field lines are labeled using the colatitude of the field-line footpoint θ_f. We will also label field lines with the flux function f, which is the total magnetic flux through the stellar surface at colatitudes $\theta < \theta_{\rm f}$. In the northern hemisphere ($\theta < \pi/2$) f is given by

$$f(\theta_{\rm f}) = 2\pi r_\star^2 \int_0^{\theta_{\rm f}} B_r\left(r_\star, \theta\right) \sin\theta \, {\rm d}\theta. \tag{6.5}$$

For the surface normal field distribution of a dipole (which is unchanged by our axisymmetric azimuthal shearing) the flux function is

$$f(\theta_{\rm f}) = 2\pi\mu \frac{\sin^2\theta_{\rm f}}{r_\star} = 2\pi\mu \frac{u(\theta_{\rm f})}{r_\star}, \tag{6.6}$$

where $u(\theta_{\rm f}) = \sin^2\theta_{\rm f}$ is the fractional flux function.

6.2.2 Sources of twisting

In a real magnetar system there are two sources of field line twisting. The first is shearing of the stellar crust ω_c, driven by evolution of the magnetar's strong magnetic field caused by non-ideal MHD effects (such as Hall drift, ambipolar diffusion, and resistivity). The surface may move slowly and plastically, or be subject to sudden cracking motions or "starquakes." In the latter case, a propagating fracture would move at the Alfvén speed in the stellar interior (Thompson & Duncan 1995b),

$$v_{\rm A} = 9 \times 10^6 B_{15} \rho_{15}^{-1/2} \,{\rm cm\,s^{-1}}, \tag{6.7}$$

which is much less than the speed of light, and so both plastic and cracking motions are slow in the sense that they develop over a timescale much longer than the light-crossing timescale of the inner magnetosphere.

The second source of twisting is resistivity in the magnetosphere itself, which is predominantly due to the existence of a threshold voltage for pair production, $\Phi_\pm \sim 10^9$ V (Beloborodov & Thompson 2007; Beloborodov 2009b). Voltages below the threshold can be sustained along field lines, resulting in Ohmic dissipation $\boldsymbol{E} \cdot \boldsymbol{J}$ and twisting of the field line at rate (Beloborodov 2011)

$$\dot{\psi} = 2\pi c \, \frac{\partial \Phi}{\partial f}, \tag{6.8}$$

where Φ is the voltage between the field line's footpoints, expected to be close to $\Phi_\pm$. This process leads to characteristic timescales on the order of years. The dissipation tends to quickly erase electric currents on field lines closing near the star, forming a potential cavity ($\nabla \times \boldsymbol{B} = 0$), but the mechanism can actually increase the twist on more extended field lines where $\partial\Phi/\partial f > 0$. Reported shrinking hot spots on transient magnetars support the localization of twist on extended field lines (Beloborodov 2011).

These two twisting mechanisms operate concurrently, and their combined effect can be represented by an effective shearing rate ω,

$$\dot{\psi}(f) = \omega_c(f) + 2\pi c \frac{\partial \Phi}{\partial f} \equiv \omega(f). \tag{6.9}$$

In our simulations we use the effective rate ω, attributing all of the twisting to surface shearing. This allows us to model the magnetosphere as being entirely ideal and force-free, except in thin current sheets where resistivity is expected.

6.2.3 Boundary conditions & shearing profiles

Our simulations concern stars whose magnetospheres are twisted by differential rotation of their surfaces. As the conducting stellar surface drags the frozen-in magnetic field, an electric field is induced in the static laboratory frame (see Section 3.9.1)

$$\boldsymbol{E} = -\left[(\boldsymbol{\Omega} + \boldsymbol{\omega}) \times \boldsymbol{r}\right] \times \boldsymbol{B}\,, \tag{6.10}$$

where we have decomposed the surface motion into solid-body rotation, $\boldsymbol{\Omega}$, and shearing, $\boldsymbol{\omega}$. The simulations are axisymmetric: $\boldsymbol{\Omega}$ and $\boldsymbol{\omega}$ are both always parallel to the magnetic axis.

In this paper we present results with three shearing profiles $\omega(\theta)$, which are shown in Figure 6.1. The first is identical to that used by ML94,

$$\omega_{\rm ML}(\theta) = \omega_0 \frac{\Theta}{\sin\theta} \exp\left[\left(1 - \Theta^4\right)/4\right] \tag{6.11}$$

where $\Theta = (\theta - \pi/2)/\Delta\theta_m$ and $\Delta\theta_m = \pi/9$. This profile concentrates the shear near the equator, peaking at $\theta_\pm = \pi/2 \pm \Delta\theta_m$. We label the applied shear by the azimuthal displacement ψ between two

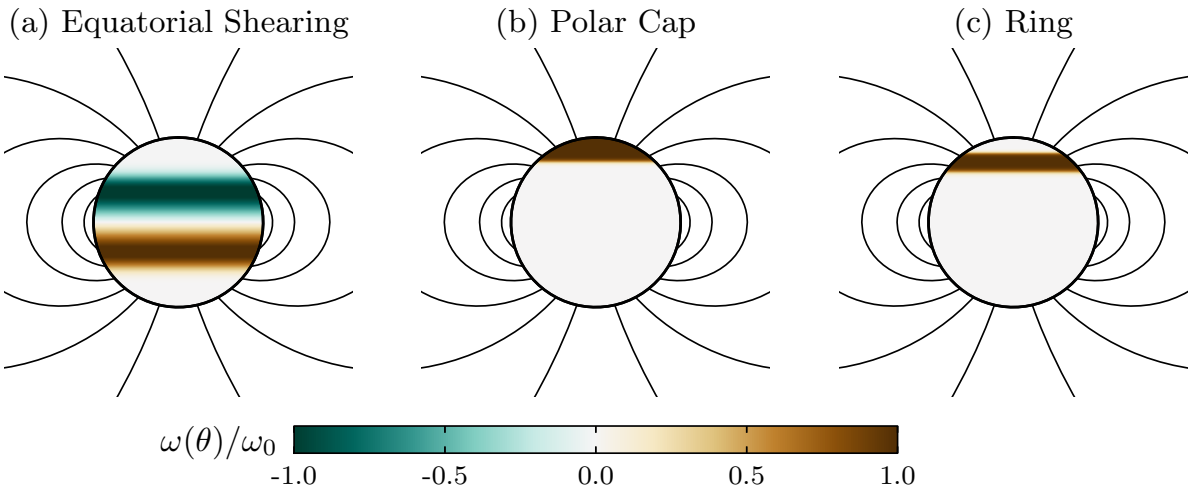

Figure 6.1: The three shearing profiles used, normalized to the maximum shearing rate ω_0. (a) Equatorial shearing, Equation (6.11), peaking at $\theta = \pi/2 \pm \pi/9$; (b) polar cap, Equation (6.12), extending down to $\theta_{\rm pc} = 0.25\pi$; (c) ring profile, Equation (6.13), centered at $\theta_{\rm ctr} = 0.25\pi$ and with width $\Delta = 0.05\pi$. The dipole field lines are equally spaced in u, between $u = 0.05$ and 0.9, and $\kappa = 50$ in (b) and (c).

footpoints located at $\theta_\pm$; all other field lines have twist amplitude smaller than ψ. Field lines above and below the equator are dragged in opposite directions, and so equatorial reflection symmetry is preserved.

In order to survey the dependence of solutions on shear profile, we also study two other models: the "polar cap" and "ring." In the polar cap model, the twisting rate is nearly constant from the magnetic axis down to some colatitude, and then decreases exponentially quickly to zero,

$$\omega_{\rm cap}(\theta) = \frac{\omega_0}{1 + \exp\left[\kappa\left(\theta - \theta_{\rm pc}\right)\right]}, \tag{6.12}$$

where κ determines the sharpness of the transition region whose center is at $\theta_{\rm pc}$ (i.e. $\omega_{\rm cap}(\theta_{\rm pc}) =$

$\omega_0/2$). The ring model twists a band of latitudes, centered at $\theta_{\rm ctr}$ and with angular half-width Δ,

$$\omega_{\rm ring}(\theta) = \frac{\omega_0}{1 + \exp\left\{\kappa\left[|\theta - \theta_{\rm ctr}| - \Delta\right]\right\}}. \tag{6.13}$$

The ring extends in colatitude from $\theta_1 \equiv \theta_{\rm ctr} - \Delta$ to $\theta_2 \equiv \theta_{\rm ctr} + \Delta$. We generally set $\kappa \approx 50$. In both the polar cap and ring twisting profiles, only the northern ($\theta < \pi/2$) footpoints are moved, and the fields do not remain symmetric about the equator. We choose these models as simple one- and two-parameter families of shearing profiles, approximating step-function selection of twisted field lines.

In each simulation, the shearing rate $\omega_0(t)$ is smoothly increased from zero to its maximum value using a cosine bell. Depending on the problem, the shearing is either maintained at this constant rate until the end of the simulation, or, if a specific twist is to be implanted, the shearing rate is eventually smoothly returned to zero.

A non-reflecting outer boundary condition is applied at $r_{\rm out}$; its numerical implementation is described briefly in Section 6.3.2, and in more detail in Section 3.9.2.

6.2.4 Units

When describing the results of our simulations we will use the following units. Distance will be measured in units of $r_\star$, time in units of $r_\star/c$, angular velocity in units of $c/r_\star$, magnetic field in units of $\mu/r_\star^3$, and current density in units of $c\mu/r_\star^4$. All angles are measured in radians.

6.3 Numerical method

We solve the equations of force-free electrodynamics, Equations (2.11) and (2.12), using the parallel pseudospectral simulation code PHAEDRA; see Chapter 3 for a detailed description of the code.

The code calculates spatial derivatives by expanding the fields along each coordinate direction in global orthogonal basis functions (Chebyshev polynomials in the radial direction, cosine and sine functions in the meridional), calculating the coefficients of the derivative series in spectral space, and transforming back to real space. The variables are temporally integrated at each point in real space using Runge-Kutta time stepping. This method converges quickly to smooth solutions, and has low numerical diffusion and dissipation.

6.3.1 Filtering & resistivity

Spectral filtering is applied at each time step to prevent aliasing instability and the accumulation of high-wavenumber noise from sharp features in the solution; see Section 3.6 for a detailed description. If a function $u(x)$ is expanded in a set of basis functions $T_n(x)$, the filtered function, $\mathcal{F}u(x)$, is given by

$$\mathcal{F}u(x) = \sum_{n=0}^{N-1} \sigma\left(\frac{n}{N-1}\right)\widetilde{u}_n T_n(x)\,, \tag{6.14}$$

where $\widetilde{u}_n$ are the discrete expansion coefficients and $\sigma(\eta)$ is the filter function,

$$\sigma(\eta) = \mathrm{e}^{-\alpha\eta^{2\mathrm{p}}}. \tag{6.15}$$

We use two filters: one of very high order to prevent aliasing instability ($2p = 36$, $\alpha = -\ln\epsilon_M$, where ϵ_M is machine precision) and one of eighth order to maintain stability in the presence of

discontinuous current sheets ($2p = 8$, $\alpha = \alpha_{\rm SSV} = 0.01 - 0.1$). The effect of this second "super spectral viscosity" filter can be interpreted as an eighth-order hyper-resistivity, and is the dominant source of dissipation in our simulations. Because it is of high order, this resistivity is negligible on all resolved scales, only acting when the field gradients approach the grid scale. This restricts the dissipation to regions with high current density, where it is physically expected, either from collective plasma effects or charge starvation. Equivalently, the filter sets the reconnection length scale to be very close to the grid scale. Our solutions can be viewed as being ideal up to the point of current sheet formation, at which point resistivity is introduced where it is required. Increasing resolution decreases the reconnection length scale, allowing current sheet formation to proceed further before the current layer becomes resistive.

6.3.2 Coordinate maps

In simulations of non-rotating stars, where the initial field is a standard dipole, the closed field lines extend to arbitrary distances, and outgoing waves on all field lines should be able to return to the star. This can be problematic when performing calculations in a finite domain. (This problem does not arise in studies of rotating stars, because outside the light cylinder the field lines are open to infinity, and can be safely truncated.) In some cases in Section6.4 and Section6.5, we use an exponential coordinate mapping to place the computational boundary at sufficient distance that no waves reach it from the star over the length of the simulation; in longer simulations we include an absorbing layer near the outer boundary, placed far enough away that the removal of outgoing waves has little effect. In all simulations, a non-reflecting boundary treatment, based on an approximate characteristic decomposition of the equations of motion, is implemented at $r = r_{\rm out}$.

The exponential coordinate map we use is

$$r(x) = \exp\left\{Q\left[g_{\rm asin}(x) + 1\right]\right\} r_{\rm in} \ , \tag{6.16}$$
$$Q = \frac{1}{2}\ln\left(\frac{2+\zeta\Delta\theta}{2-\zeta\Delta\theta}\right)(N-1) \ ,$$

giving $r \in [r_{\rm in}, r_{\rm out}]$ where $r_{\rm out} = e^{2Q} r_{\rm in}$, and $g_{\rm asin}(x)$ is the arcsine map which reduces the clustering of the Chebyshev nodes, $x \in [-1, 1]$, near the end points. If $g_{\rm asin}(x)$ were exactly equispaced the resulting grid would have spacings Δr that lay on a line through the origin with slope $\zeta\Delta\theta$.

In simulations of non-rotating stars we use a transformation in the meridional direction,

$$\widetilde{\theta} = \theta + \frac{\gamma}{2}\sin(2\theta) \tag{6.17}$$

with $\gamma = 0.4$, to smoothly increase the grid resolution near the equator. With rotation, it is important to resolve well the narrow open flux bundles near the poles, and so we do not use this mapping.

6.4 Equilibria

6.4.1 Reaching equilibrium

In this section we discuss the equilibrium solutions of a dipole magnetosphere, part of whose flux has been twisted through an angle ψ. Our simulations begin with a dipole field with no twist, the field being everywhere poloidal and potential. Shearing of the stellar surface begins at $t = 0$, launching current- and charge-carrying Alfvén waves into the magnetosphere and causing the twisted field lines to acquire toroidal magnetic field B_ϕ. An equilibrium solution at a specific

ψ can be found by shearing the surface up to that angle, halting the shear motion, and allowing the magnetosphere to relax to a static configuration over several wave-crossing times of the twisted field lines[2]. Alternatively, if the shearing is slow enough the magnetosphere will smoothly move through a sequence of quasi-equilibrium states. These are close to the true equilibria, which are approached as the twisting rate is decreased; the deviation can be measured with a scalar virial equation, as described in Section 6.4.3. In this paper, most "equilibrium" solutions described, including all where a quantity is shown smoothly varying with ψ, are these quasi-equilibria.

The evolution of an initially untwisted configuration to a twisted quasi-equilibrium state is shown in Figure 6.2. A ring in the northern hemisphere is brought from rest to a constant twisting rate ω by $t = 5$, injecting an Alfvén wave into the magnetosphere (Figure 6.2a). The wavefront is sheared by the dipole field geometry, as the wave moves along diverging field lines. The wave reaches the field lines' southern footpoints and is reflected (Figure 6.2b); subsequently, waves bounce backward and forward on the closed field lines, establishing quasi-equilibrium over the majority of the twisted flux by $t \approx 35$ (Figure 6.2d), even as the ring continues to slowly rotate. Note that the quasi-equilibrium field is very close to being reflection-symmetric about the equator, even though only the northern footpoints are in motion, because the shearing timescale is much longer than the wave-crossing time scale on these field lines. This final state is very similar to that which is found if the twisting is immediately turned off and the configuration allowed to relax to equilibrium (the wiggles in the B_ϕ contours are erased as the corresponding field lines equilibrate).

One can think of the ideal (dissipationless) twisted configuration as a superposition of Alfvén waves created by the polar-cap motion, which are trapped in the closed magnetosphere where they continue to bounce off the perfectly conducting surface. The waves are more easily seen in the charge density distribution (lower panels of Figure 6.2). They are stretched to smaller perpendic-

[2]The simulated magnetosphere is nearly ideal (as long as no current sheets form); therefore the resulting twisted configuration will have the same twist amplitude ψ as imparted by the surface motion.

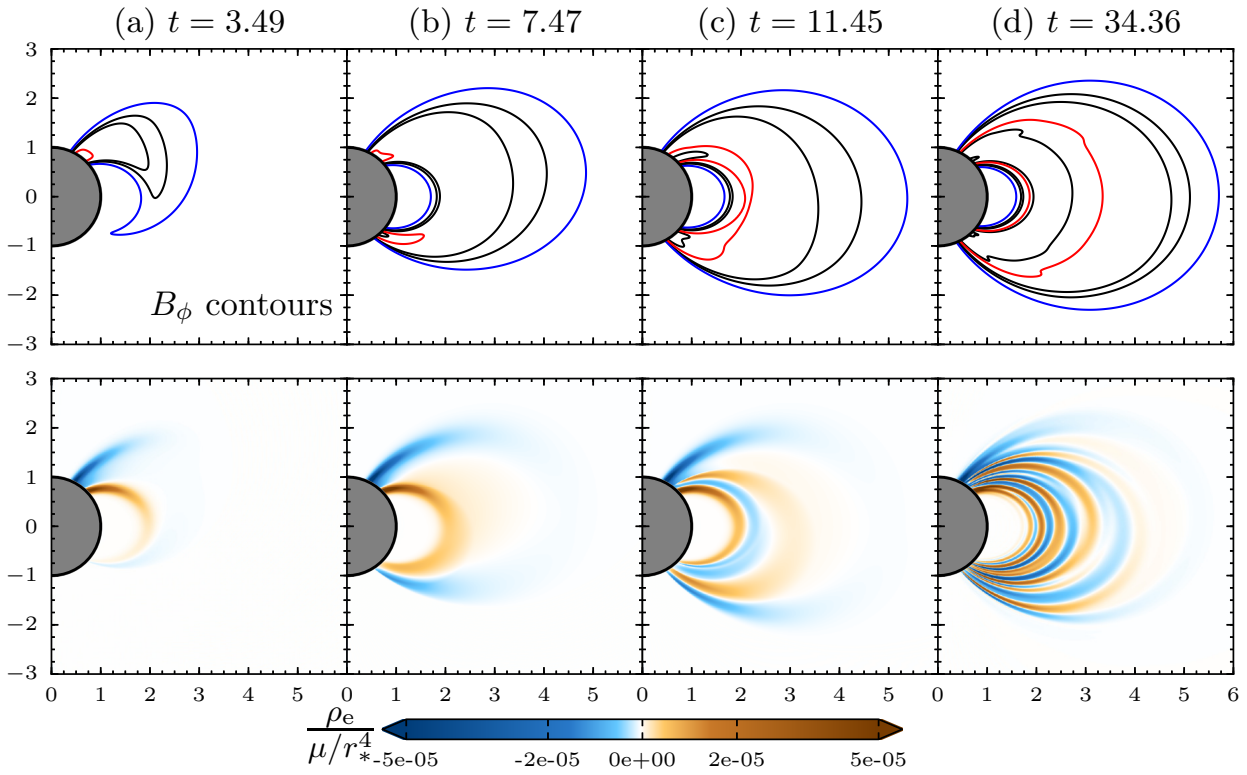

Figure 6.2: Continuously twisted magnetosphere reaching quasi-equilibrium (initial transient phase). A ring extending from $\theta/\pi = 0.15$ to 0.25 is smoothly brought from rest to $\omega_0 = 10^{-4}\, c/r_\star$ by $t = 5\, r_\star/c$. Upper panels: contours of toroidal field at $-B_\phi = 10^{-6}$ (blue), 5×10^{-6}, 10^{-5}, 5×10^{-5} (red), 10^{-4}, and $5\times10^{-4}\ \mu/r_\star^3$. Lower panels: charge density $\rho_e = \nabla\cdot\boldsymbol{E}/4\pi$ in units of $\mu/r_\star^4$. Time is indicated in units of $r_\star/c$.

ular scales over time, as the background geometry increasingly shears each succeeding reflected wave. In a perfectly ideal plasma the waves would bounce backward and forward indefinitely, and be sheared ever thinner; in our numerical solutions, their thickness reaches the grid scale, and they are gradually smoothed to larger scales by the spectral filters. This has a negligible effect on the overall field distribution.

Eventually, the waves excited by the initial acceleration of the ring are damped, and the subsequent evolution is a slow progression through a sequence of configurations with increasing ψ, which are almost identical to exact equilibria.

6.4.2 Sequence of equilibrium solutions

Once the twisted field lines have reached quasi-equilibrium, the sequence of force-free equilibria can be investigated, by slowly shearing the stellar surface at a constant rate. The numerical solution moves quasi-statically through the sequence as the accumulated twist ψ on the field lines increases.

For a perturbative expansion in small $\psi \ll 1$, the first-order departure from the potential field is the addition of a toroidal component; the second-order effect is a slow expansion of the poloidal field lines, as the additional magnetic pressure due to the toroidal field modifies the pressure–tension force balance, pushing the field lines outward (Low 1986).

Our numerical procedure allows us to generate equilibria for general surface shearing profiles and twist angles. Solutions at $\psi = 0$, 1.5, and 3 for three shearing profiles are shown in Figure 6.3 and Figure 6.4. These profiles were chosen to illustrate shearing confined to regions near the equator, at the pole, and in mid-latitudes. The solutions shown are stable; if the shearing is arrested they remain static indefinitely, which we have explicitly verified over many magnetospheric light-crossing times.

In all three cases it can be seen that there is noticeably more poloidal expansion between $\psi =$ 1.5 and 3 than between $\psi = 0$ and 1.5. The toroidal current density, J_ϕ, is everywhere positive for the equatorial shearing model. In contrast, for the polar cap and ring models, J_ϕ changes sign as one moves away from the pole toward the region where the shearing rate is zero. The early stages of the formation of the equatorial current sheet are evident in Figs. 6.3 and 6.4b, including the concentration of toroidal current near the equator; B_ϕ likewise increases near the equator (and decreases elsewhere) as ψ grows. Equatorial shearing causes all the field lines to expand significantly; eventually (in the absence of resistivity) every field line will open to infinity given sufficient twisting, and the magnetosphere will be brought to the maximum-energy configuration for a dipole surface flux distribution. The ring shearing profile, on the other hand, displays a clear distinction

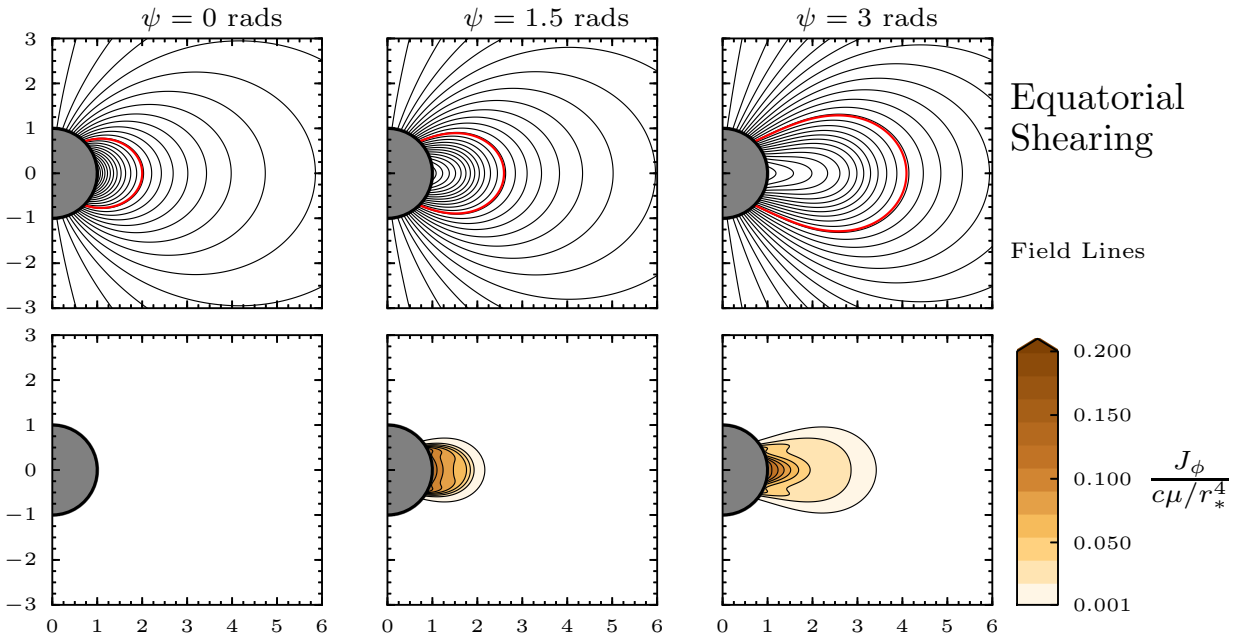

Figure 6.3: Equilibrium solutions for the equatorial (ML94) shearing profile, Equation (6.11). Upper panels: poloidal field lines, equally spaced in flux function in the range $u = 0.01$–0.975 with spacing $\Delta u \approx 0.04$, and a red field line at $u = 0.5$. Lower panels: filled contours of J_ϕ, equally spaced between $J_\phi = 0.001$ and $0.2\ c\mu/r_\star^4$ with step $\Delta J_\phi \approx 0.018\, c\mu/r_\star^4$. Axes are labeled in units of $r_\star$.

between twisted, inflating, field lines and the untwisted field below; the final state in this sequence will be a partially open field. The polar cap solution (Figure 6.4a) shows less progress toward forming the current sheet by $\psi = 3$; as with ring shearing, the final state at large ψ will also be partially open.

The equatorial shearing solutions are similar to those found by ML94, using the same shearing profile but evolving the time-dependent equations of non-relativistic MHD in a restricted "zero (plasma) beta" formulation. This similarity is expected, because the force-free equilibrium states, as determined by Equation 6.2, are independent of the plasma motions by which they are produced.

Field lines, drawn in three-dimensional space, are shown in Figure 6.5 for the polar cap and ring shearing models used in Figure 6.4, again at $\psi = 3$. At this large twist angle the toroidal

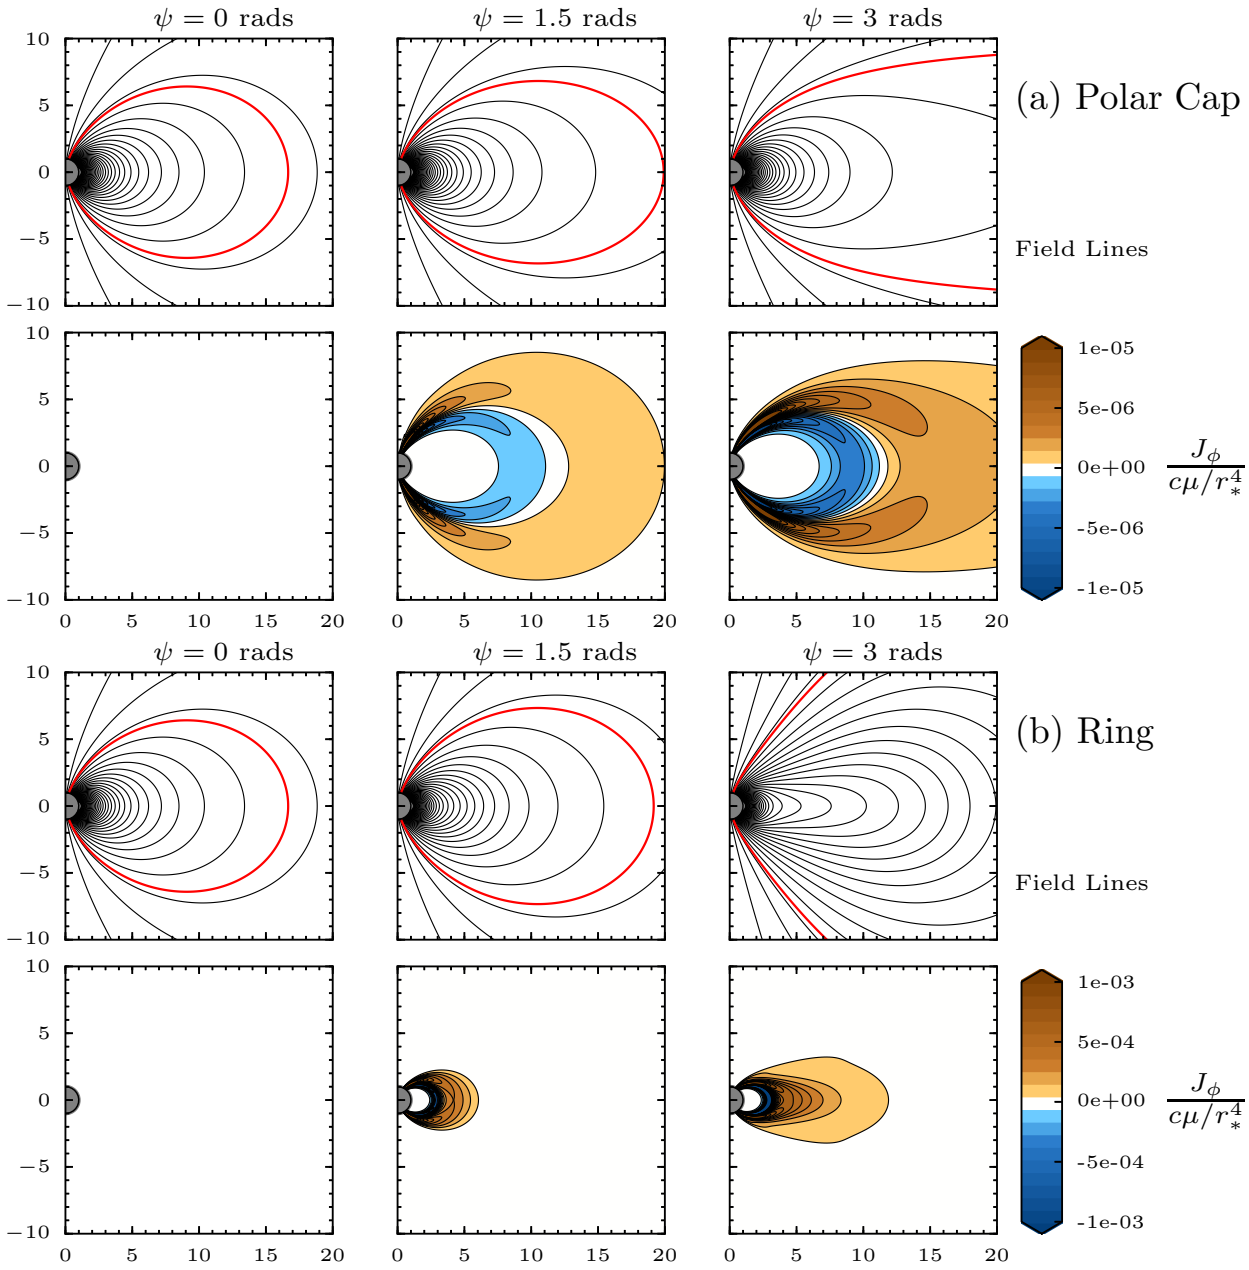

Figure 6.4: Equilibrium solutions for two shearing profiles: (a) polar cap $\left(\theta_{\rm pc} = 0.1\pi\right)$ and (b) ring (bounded by $\theta_1 = 0.15\pi$ and $\theta_2 = 0.25\pi$). Upper panels: poloidal field lines, equally spaced in flux function in the range $u = 0.01$–0.8 with spacing $\Delta u \approx 0.02$, and a red field line at $u = 0.06$. Lower panels: 20 filled contours of J_ϕ, equally spaced in the ranges shown. Axes are labeled in units of $r_\star$.

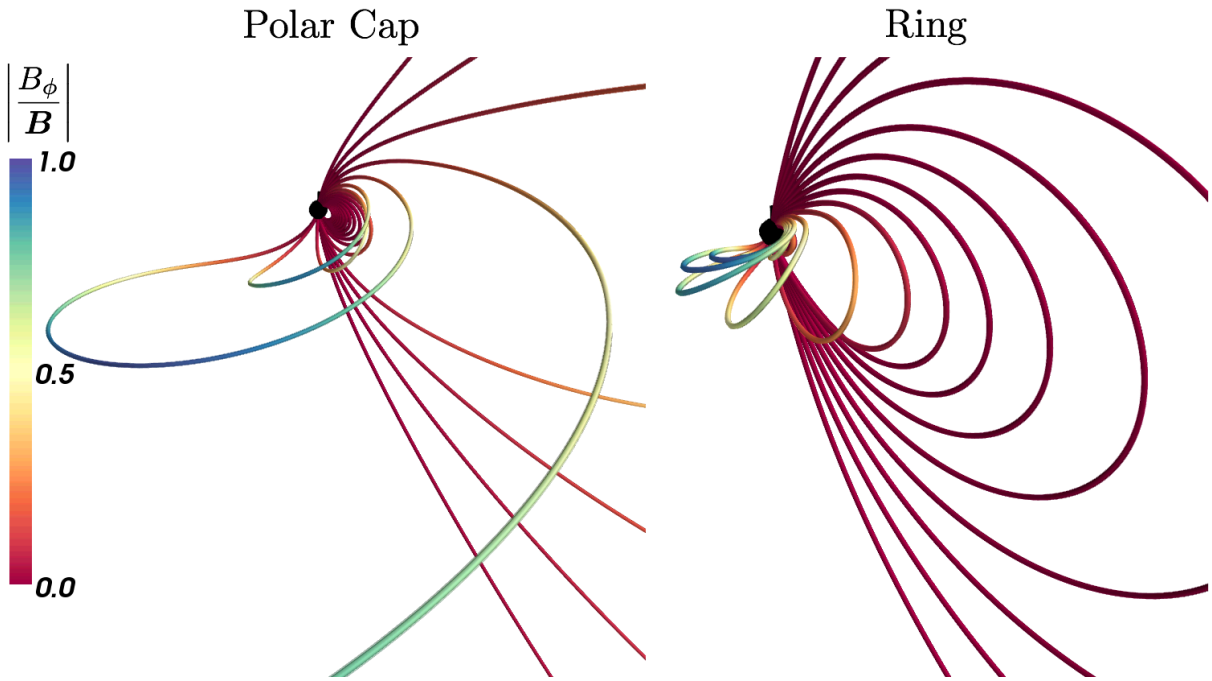

Figure 6.5: Three-dimensional field lines, for the same polar cap and ring shearing models as in Figure 6.4, at $\psi = 3$. The lines are colored by the fractional contribution of the toroidal field, $\left|B_\phi/\boldsymbol{B}\right|$, at each point along their lengths. Twenty lines are drawn from each hemisphere, equally spaced in colatitude between $\theta/\pi = 0.04$ and 0.25.

magnetic field becomes dominant near the equator and small away from it. This means that field lines which are pushed away from the equatorial region by the expansion, including those attached to the twisted polar cap, become predominantly radial near the star (and out to increasing distances from it as ψ grows).

The expansion of the field lines occurs in two phases. At small twist amplitude, the poloidal field is only weakly affected by shearing, and the maximum height of each field line above the stellar surface, $R_{\rm max}$, increases slowly with ψ. As one steps through the sequence of quasi-equilibria, the magnetosphere eventually becomes much more sensitive to ψ, and $R_{\rm max}$ then increases exponentially with twist angle.

When the magnetosphere enters the second regime depends on the profile of the applied

shear, and in particular on whether the twisting is applied all the way to the pole (polar cap) or if there are untwisted field lines surrounding the sheared flux (ring profile); in the latter case the untwisted lines help to contain the twisted flux, and entry into the rapidly inflating regime is delayed. In Figure 6.6 the two modes are illustrated for a polar cap shearing model. In this case, the magnetosphere inflates exponentially above a twist of $\psi \sim 1.5$–1.75. The lowest-lying field lines are in the untwisted zone inside the sheared flux, and do not participate in the expansion. We stress that even in the rapidly inflating regime the magnetosphere progresses smoothly through stationary quasi-equilibria, because ψ is increased very slowly in our simulations. There is no critical behavior in the solutions described in this section; the deviation from equilibrium becomes significant only at $\psi \gtrsim 3$ and will be discussed in Section 6.5.

6.4.3 Energy of equilibria

The minimum energy state for a magnetosphere with a given distribution of B_r on the stellar surface is the potential ($\nabla \times \boldsymbol{B} = 0$) field that has only closed field lines. Shearing does work against the field lines' tension, and transfers energy from the star to the magnetosphere. At small twist angle $\psi \ll 1$, this increase is due to the B_ϕ^2 contribution added to the magnetic energy density, while the poloidal field is hardly changed. As ψ increases and the field lines expand, the additional energy is increasingly stored in the poloidal field components. The limiting maximum energy configuration of sheared fields is the fully open field, in which all field lines extend to infinity and the toroidal component is everywhere zero (Aly 1991; Sturrock 1991). For a dipole potential field in an infinite domain ($r_{\rm out} \to \infty$) having energy W_0, this fully open state has energy $W_{\rm open} = 1.662\, W_0$ (Barnes & Sturrock 1972; ML94).

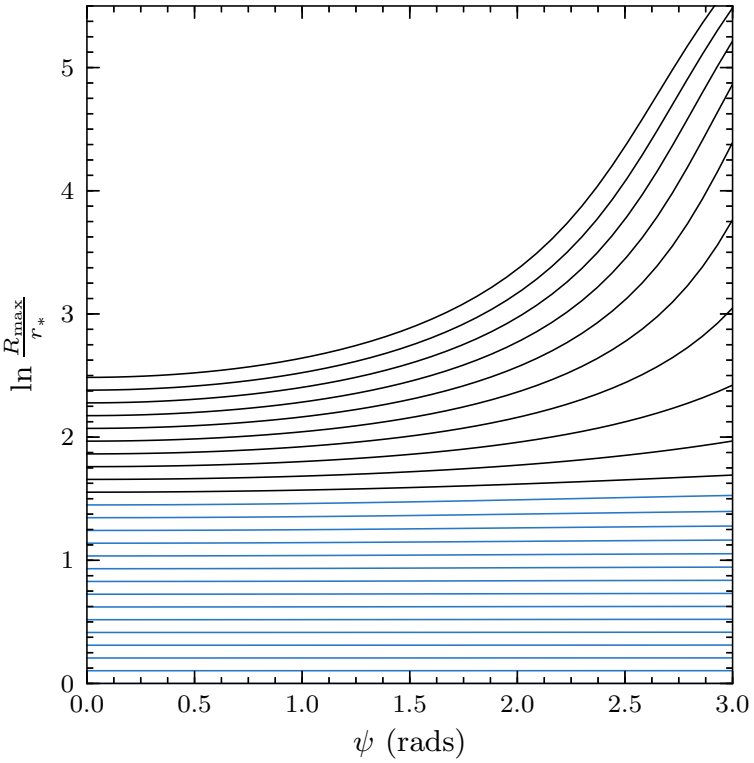

Figure 6.6: Natural log of $R_{\max}$ versus twist angle, where $R_{\max}$ is the maximum height of each field line, for a twisted polar cap with $\theta_{\rm pc} = 0.15\,\pi$. The blue curves trace untwisted field lines.

The total magnetic energy, W, can be found by integrating B^2 over volume,

$$W \equiv \frac{1}{8\pi} \int_V B^2 \mathrm{d}V, \tag{6.18}$$

where V is the volume of the computational domain (excluding the outer absorbing layer if present). The energy of any equilibrium force-free configuration outside a surface is related to the distribution of $\boldsymbol{B}$ on that surface by a scalar virial theorem (Aly 1984). The energy expected in the computational domain from the virial theorem, $W_{\rm vir}$, can then be found by subtracting the energy that,

in equilibrium, should lie beyond its outer boundary:

$$W_{\rm vir} \equiv W_\infty(r_\star) - W_\infty(r_{\rm out}), \tag{6.19}$$

where

$$W_\infty(r) \equiv \frac{r^3}{4}\int_0^\pi \left(B_r^2 - B_\theta^2 - B_\phi^2\right)\sin\theta\,{\rm d}\theta \tag{6.20}$$

is the energy of an equilibrium state, integrated from r to infinity. Since an equilibrium solution has $W = W_{\rm vir}$, these quantities can be used to test how close our quasi-equilibrium solutions are to equilibrium.

The twist free energy is defined by $W_{\rm tw} = W - W_0$ where W_0 is the untwisted dipole energy given by Equation (6.4). At small twist amplitude $\psi \ll 1$, the free energy is just the energy of the toroidal magnetic field, which is given by Beloborodov (2009b),

$$W_{\rm tw} \approx \int_{r>r_\star} \frac{B_\phi^2}{8\pi}\,{\rm d}V = \frac{\mu}{2cr_\star}\int_0^1 I(u)\,\psi(u)\,{\rm d}u, \tag{6.21}$$

where $I(u)$ is the poloidal current function (defined in the same way as the poloidal flux function f, except that one replaces B_r with J_r in Equation (6.5)). By Stokes' theorem, the toroidal magnetic field and poloidal current function are related by $B_\phi = 2I/cr\sin\theta$, and integrating B_ϕ along a field line gives a relation between I and ψ,

$$\psi = \frac{4Ir_\star^2}{u^2c\mu}\sqrt{1-u}. \tag{6.22}$$

In our polar cap models, $\psi(u) \approx$ constant for $0 \le u \le u_{\rm pc}$ and $\psi(u) = 0$ elsewhere; combining this profile, Equations (6.21) and (6.22), and $W_0 = \mu^2/3r_\star^3$ one estimates the total energy of the twisted

configuration to be

$$W \approx \left(1 + \frac{\psi^2 u_{\rm pc}^3}{8}\right) W_0, \qquad \psi \lesssim 1. \tag{6.23}$$

(Note that $u_{\rm pc}$ is not identical to the u_* used in Beloborodov (2009b), hence the difference in numerical coefficients.)

The measured energy W as a function of ψ is shown for several models in Figure 6.7, where it is also compared to $W_{\rm vir}$ (dashed lines) and the analytical estimate in Equation (6.23) (dashed-dotted line). For the equatorial shearing model (Figure 6.7a) the energy approaches, but does not exceed, the energy of the completely open field (dotted line), and $\mathrm{d}^2W/\mathrm{d}\psi^2 < 0$ near the open configuration. The solution is always very close to equilibrium ($W = W_{\rm vir}$) for the twist angles displayed.

In panels (b) and (c) of Figure 6.7 the energy is given for three polar cap models, and three ring models, of thickness $2\Delta = 0.05\pi$. Each polar cap model, of extent $\theta_{\rm pc}$, corresponds to one ring model whose lower boundary θ_2 is at the same colatitude $\theta_{\rm pc}$. In each case, the energy of the twisted polar cap solution is greater than that of the corresponding ring, and the difference increases as the ring is moved toward the equator. The analytical estimate, Equation (6.23), is a good approximation to the energy of the polar cap solution at twist angles $\psi \lesssim 1.25$, and overestimates the energy at larger twists.

The top polar cap solution begins to depart slightly from equilibrium above $\psi \sim 3$ indicating that the magnetosphere leaves the sequence of equilibria and enters a dynamical state. The ring solutions are still in equilibrium even at $\psi = 3.5$, because of the confining effect of unsheared field lines blanketing the twisted flux. We next turn our attention to what happens when the magnetosphere is subjected to large twist angles, and the quasi-equilibrium approximation ceases to be valid.

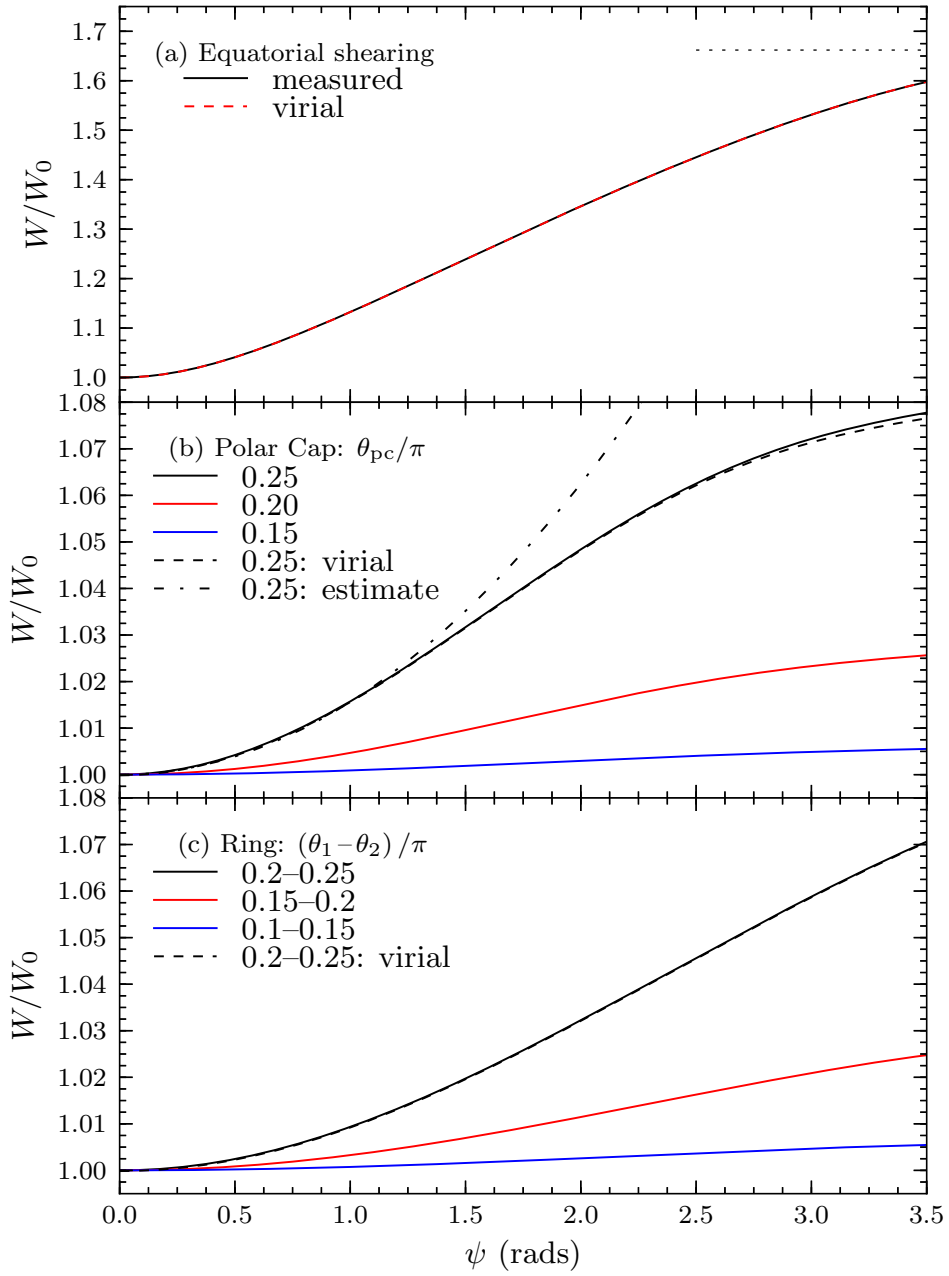

Figure 6.7: Total energy on the grid, in units of the energy of the untwisted dipole W_0, for (a) equatorial (ML94), (b) polar cap, and (c) ring twisting profiles. The dashed lines show the value expected from the equilibrium force-free virial theorem, Equation (6.19), for the top curve in each panel, the dashed-dotted curve gives the analytic estimate, Equation (6.23), corresponding to the top polar cap model, and the dotted line in (a) indicates the energy of the completely open field.

6.5 Dynamics of overtwisted magnetospheres

When a sufficiently large twist ψ is applied, the magnetosphere must necessarily leave the sequence of force-free equilibria and enter a fully dynamical state. This has been argued in two ways in previous work. Firstly, that as the magnetic field is sheared, regions of increasingly high current density are created, evolving into thin current sheets separating regions of potential open field (Barnes & Sturrock 1972; Yang et al. 1986; Sturrock 1991). In the presence of any resistivity, which must be present even if only at very high current densities, such strong current layers will eventually suffer reconnection and dynamic reconfiguration, involving the dissipation and ejection of magnetic energy (e.g. Steinolfson 1991).

The second argument posits the existence of a critical twist angle, $\psi_{\rm crit}$, at which the magnetosphere experiences ideal magnetic non-equilibrium and transitions from a closed to a (partially) open configuration (ML94; Uzdensky 2002). This dynamic opening would be a purely ideal process; however, it results in the formation of a current sheet between oppositely directed open field lines, and so eventually must also be accompanied by reconnection-powered dynamics due to non-ideal physics in the sheet.

It is of interest to identify the critical twist amplitude at which the magnetosphere is reconfigured. In numerical simulations, it can be difficult to distinguish dynamic motion due to the ideal loss of equilibrium from the rapid progression through quasi-equilibrium states, because in a simulation twisting is applied on a finite timescale and at large ψ the magnetosphere becomes very sensitive to further increasing the twist angle. Furthermore, it can be difficult to distinguish the collapse of current layers to discontinuous current sheets due to loss of equilibrium from collapse caused by effective numerical resistivity when the layers' thickness approaches the grid scale.

We choose to measure $\psi_{\rm rec}$, the twist amplitude at which the fast reconnection phase begins,

because this moment can be clearly defined: it is the first instant at which $E^2 > B^2$ anywhere in the domain (which necessitates the removal of electric field to mimic dissipation). We find that this inequality is satisfied only in the discontinuous current sheets that arise after the collapse of thicker current layers. A small amount of flux can reconnect before $\psi_{\rm rec}$ due to effective resistivity in the current layer before it has fully collapsed; however this reconnection is slow and only involves a small fraction of the total reconnecting magnetic flux. If the shearing rate is slow then $\psi_{\rm rec}$ should be approximately equal to $\psi_{\rm crit}$; this is investigated and confirmed in Section 6.5.3.

The critical twist may be measured by other means. For instance, ML94 defined $\psi_{\rm crit}$ as follows. After evolving to an ideal twisted configuration (which was still stable), they changed the equations of motion by introducing a resistive term $\eta\nabla^2\boldsymbol{B}$. If $\psi < \psi_{\rm crit}$, the twisted field lines slowly relaxed toward the initial potential state, while if $\psi > \psi_{\rm crit}$ the thick current layer at the equator collapsed, forming an X-point at which there was rapid reconnection of twisted field lines, and severing a large plasmoid of twisted flux which was ejected from the system.

6.5.1 Equatorial shear

Here we repeat the experiment of ML94, using the same surface shearing profile, Equation (6.11), but with relativistic force-free MHD. Rather than manually switching on magnetic diffusivity in the whole domain, we use spectral filters to consistently introduce resistivity in regions with very sharp field gradients (and hence high current density), as described in Section 6.3.

These simulations have grid size $N \times N_\theta = 384 \times 255$, and take place in a domain $1 < r < 60$, of which the last ten stellar radii comprise an absorbing zone. In order to keep the magnetosphere close to equilibrium, even at large shear, we evolve in three stages: from $\psi = 0$ to 2.8, then from 2.8 to 3.6, and finally from 3.6 on, at shearing rates $\omega_0 = 1.4 \times 10^{-3}$, 8×10^{-4}, and 2.5×10^{-4} respectively. (We quote shears by the total azimuthal angular separation of the footpoints of the

most twisted field line: $\psi = 2\int \omega_0(t)\,\mathrm{d}t$.) The magnetosphere is allowed to equilibrate between stages (although when twisting this slowly it anyway remains close to equilibrium). Increasing ψ in three stages with decreasing twisting rate allows us to approach the critical point very gently.

The magnetosphere is still stable at $\psi = 3.6$. As additional shear is applied, the field lines expand outwards and the current becomes more concentrated near the equator. The rates of expansion and concentration accelerate as ψ increases, and eventually the current layer becomes sufficiently strong that non-negligible resistivity is introduced there by the spectral filters, and the layer's thickness suddenly decreases at $r \approx 3.4$, initiating reconnection. This occurs at a twist angle of $\psi_{\mathrm{rec}} = 3.678$. Contours of toroidal current density at $\psi = 3.6$ and 3.678 are shown in Figure 6.8.

To test the sensitivity of this result to grid resolution, and hence to resistive (or reconnection) scale, we repeat the above procedure with a coarser 256×155 grid. Again the solution is indefinitely stationary at $\psi = 3.6$, and reconnection begins at $\psi_{\mathrm{rec}} = 3.654$. Increasing the strength of the eighth-order filter (from $\alpha_{\mathrm{SSV}} = 0.025$ to 0.05) and reducing the order of the high-order filter (from $2p = 36$ to 26) each only reduces ψ_{rec} by about 0.05%. We conclude that the magnetosphere is extremely sensitive to shear above $\psi \approx 3.65$, and that the point at which a discontinuous current sheet forms is insensitive to numerical parameters and the resistive length scale. This is consistent with the magnetosphere entering a state of magnetic non-equilibrium at $\psi_{\mathrm{crit}} \sim 3.65$.

These results are in approximate agreement with the findings of ML94; they estimated the critical shear for this profile to be roughly 4 radians, and found stable ideal configurations at this shear level. The difference is due to their inclusion of gas pressure and the gravitational field of the sun, which resists the expansion of plasma, and hence field lines, away from the star.

Let us now consider the change in energy of the magnetosphere during the dynamical phase; the numbers given below are for the higher-resolution simulation (using the 384×255 grid). At the onset of reconnection, the energy on the grid (in the domain $1 < r < 50$) is $W = 1.619\,W_0$,

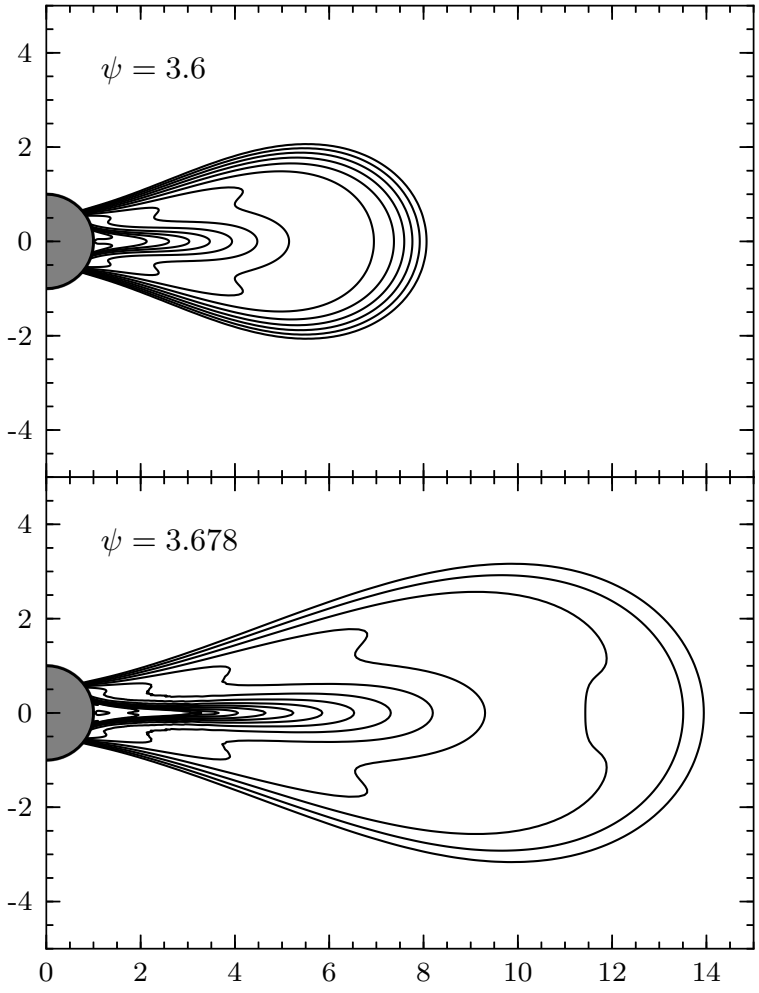

Figure 6.8: Contours of toroidal current density, J_ϕ, for the equilibrium state at $\psi = 3.6$ and just before the onset of reconnection at $\psi = 3.678$. The contours are equally spaced in $\ln\left(J_\phi\right)$ between $J_\phi = 10^{-4}$ and $0.25\ c\mu/r_\star^4$, with spacing $\Delta \ln\left(J_\phi\right) \approx 0.56$.

$2p_{\rm high}$	36	36	26
$\alpha_{\rm SSV}$	0.025	0.05	0.025
retained	31.65%	31.01%	31.65%
expelled	53.65%	51.84%	53.64%
dissipated	14.67%	17.15%	14.71%

Table 6.1: Fraction of the magnetic free energy retained, expelled, and dissipated during the dynamic phase, for three simulations with different spectral filtering parameters.

where W_0 is the energy of the dipole potential field in the same volume. The totally open field has $W \approx 1.662\, W_0$, and so the magnetosphere is, energetically speaking, about 94% of the way from the dipole toward the limiting configuration (an underestimate, since more of the energy is stored at large radii beyond $r = 50$ in the twisted state).

The dynamic evolution triggered at $\psi = \psi_{\rm rec}$ involves the expulsion of magnetic energy from the system, in the form of a plasmoid of helical magnetic field disconnected by reconnection from the stellar surface, and the dissipation of magnetic energy in the current sheet.

During the dynamic phase the magnetosphere expels 53.68% of the twist energy (of magnitude $0.619\, W_0$), dissipates 14.67%, and retains 31.65% in the form of static twisted flux. The result of doubling the strength of the eighth-order filter (to $\alpha_{\rm SSV} = 0.05$) and reducing the order of the high-order filter (from $2p_{\rm high} = 36$ to 26) is given in Table 6.1. We can safely say that more energy is expelled than dissipated, but the precise contribution of each is dependent on numerical parameters; the amount of energy retained by the system is less sensitive to numerics. The dissipation fraction is predominantly controlled by the strength of the low-order filter, which removes energy from grid-scale features like current sheets. A precise measurement of the energy dissipated into heat by these reconnection events must await a resistivity prescription more closely modeling the relevant microphysics.

6.5.2 Polar cap shear

In the preceding section we argued, using one shearing profile, that at a certain twist amplitude the magnetosphere becomes extremely sensitive to further twisting, and that the critical point is largely insensitive to numerical parameters (such as grid resolution). In the following sections we investigate the critical point's dependence on the shearing rate, ω_0, and on the surface shearing profile, $\omega(\theta)$. We use two grids, stretched in the radial direction using the exponential map, Equation (6.16): the smaller grid has $N \times N_\theta = 640 \times 255$ and an outer boundary at $r_{\rm out} = 2155\, r_\star$, the larger grid has a size of 1024×375 and $r_{\rm out} = 2773\, r_\star$.

We will primarily study polar caps twisted at a constant rate until reconnection. Here, we illustrate the general behavior with reference to a single model, a cap extending to $\theta_{\rm pc} = 0.15\pi$ twisted at $\omega_0 = 2.5 \times 10^{-3}$. The calculation was performed using the larger grid.

The evolution of the magnetosphere through the reconnection event is shown in Figure 6.9. In Figure 6.9a, the magnetosphere is significantly inflated ($\psi \approx 3.1$) but a strong equatorial current layer has not yet formed. The magnetosphere is very sensitive to additional shearing at this point, and the current layer appears soon thereafter. By $\psi \approx 4.4$ (Figure 6.9b) almost all the flux which will eventually reconnect has been opened, yet the configuration is stabilized by the continual gradual expansion of closed field lines (see Section 6.5.3).

The collapse of the current layer to a discontinuous sheet does not begin until $\psi \approx 5.5$ (Figure 6.9c). During the collapse, the Y-point-like cusp separating closed and open field lines retreats toward the star at speed $\sim c/4$, and the reconnection criterion is first satisfied at $\psi = \psi_{\rm rec} = 5.59$, $t = t_{\rm rec} = 2246$ (Figure 6.9d, in which the pinch point is clearly visible). The withdrawal of the closed flux further removes magnetic pressure support from the reconnection region, increasing the rates of collapse and reconnection.

Reconnection proceeds across the equatorial current sheet, producing plasmoids of a range

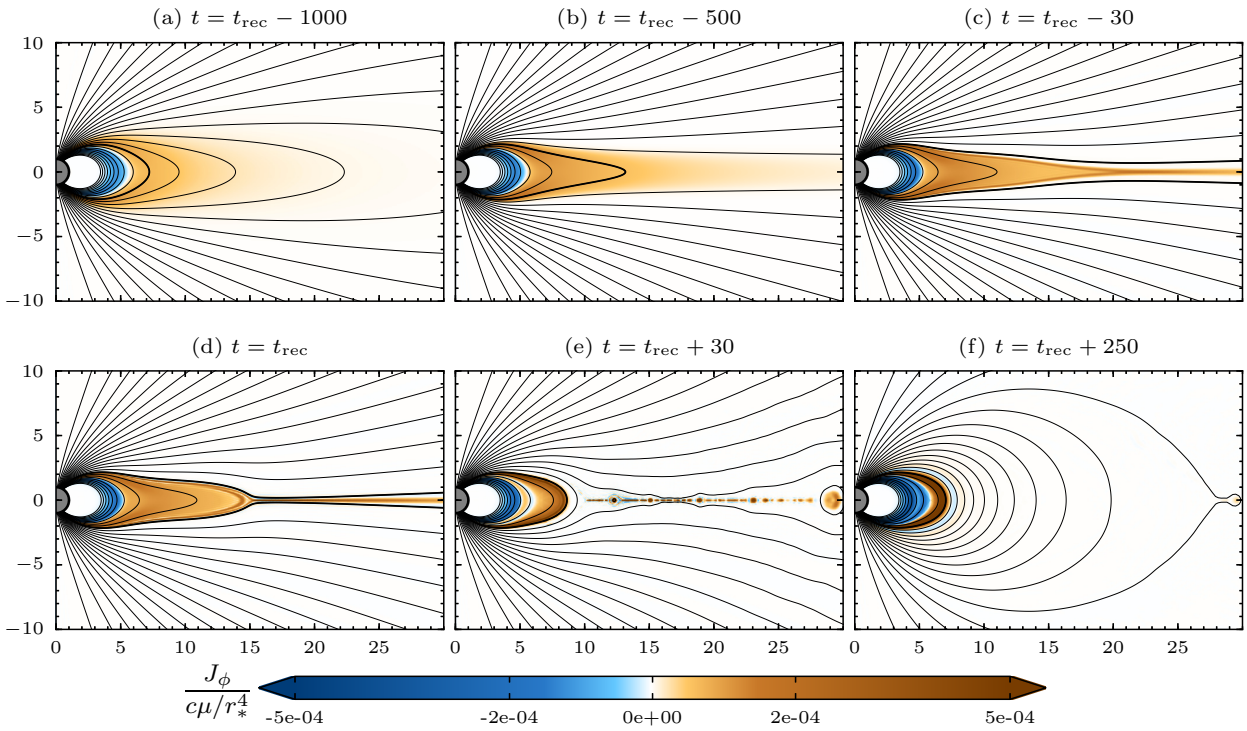

Figure 6.9: Formation of the current sheet: states leading up to and following reconnection, for a polar cap of $\theta_{\rm pc} = 0.15\pi$ and shearing rate $\omega_0 = 2.5 \times 10^{-3}\ c/r_\star$. The reconnection begins at $t_{\rm rec}$, panel (d). Color shows toroidal current density, lines are poloidal field lines projections, equally spaced in flux function in the range $u = 0.01$–0.3 with spacing $\Delta u \approx 0.012$. One field line is highlighted by a thicker line; this field line first opens and then closes again when it reconnects. Time is in units of $r_\star/c$.

of sizes, which are either ejected from the system or bounce backward and forward on closed field lines (Figure 6.9e). This process reaches an approximate steady state, with magnetic flux moving toward the current sheet at $v_\theta \sim 0.1\,c$, and the newly reconnected field lines rushing away from the reconnection point at $v_r \sim c$ in the equatorial plane.

Numerical dissipation eventually removes the trapped bouncing Alfvén waves, and the magnetosphere relaxes to a steady state, retaining some twisted flux (Figure 6.9f). In this simulation, the configuration reaches this new equilibrium before any waves from the star interact with the outer boundary.

The stellar shearing has continued throughout this phase, although it is too slow to have much effect on the dynamical evolution described. After the dynamical phase, it begins to slowly re-twist the field lines which reconnection left purely poloidal. More importantly, the twist amplitude continues to grow on those field lines that did not reconnect and therefore did not lose any twist, which we will refer to as the "twisted reservoir." Eventually the magnetic pressure in this strongly twisted region becomes so great that it explodes outward unstably, initiating a second reconnection event. The sudden explosive behavior is similar to the "magnetic detonation" described by Cowley & Artun (1997).

The two rounds of overtwisting (inflation) and loss of twist (reconnection) are illustrated in Figure 6.10, which shows field line heights above the surface versus applied shear. For the second event, the twist amplitude in the flux which had previously reconnected is much lower than the total twist accumulated in the twisted reservoir, and so the evolution is more like that of a twisted ring solution than a polar cap model. The weakly twisted overlying flux acts as a kind of nozzle, keeping the expanding flux near the equator, which causes it to reconnect earlier (see Section 6.5.5). For example, the field line indicated by the blue line in Figure 6.10 reconnects when its apex is at $R_{\rm max} \approx 140$ in the first event, but it only reaches $R_{\rm max} \approx 75$ in the second event. The red line in Figure 6.10 represents a field line which expands but remains closed during the first event (the drop in $R_{\rm max}$ is due to the rapid retraction of the closed bundle) and experiences reconnection in the second event, during which more field lines expand and reconnect because the expansion is driven by deeper flux.

In both events shown in Figure 6.10, reconnection of the field lines shown takes $\Delta t \approx 100\, r_\star/c$. Reconnection continues beyond this time on field lines having $R_{\rm max}(t=0) > 10$ (not shown); however, most of the energy will be released by the lowest-lying reconnecting field lines. The entire reconnection phase takes about $\Delta t \approx 200$–$300\, r_\star/c$.

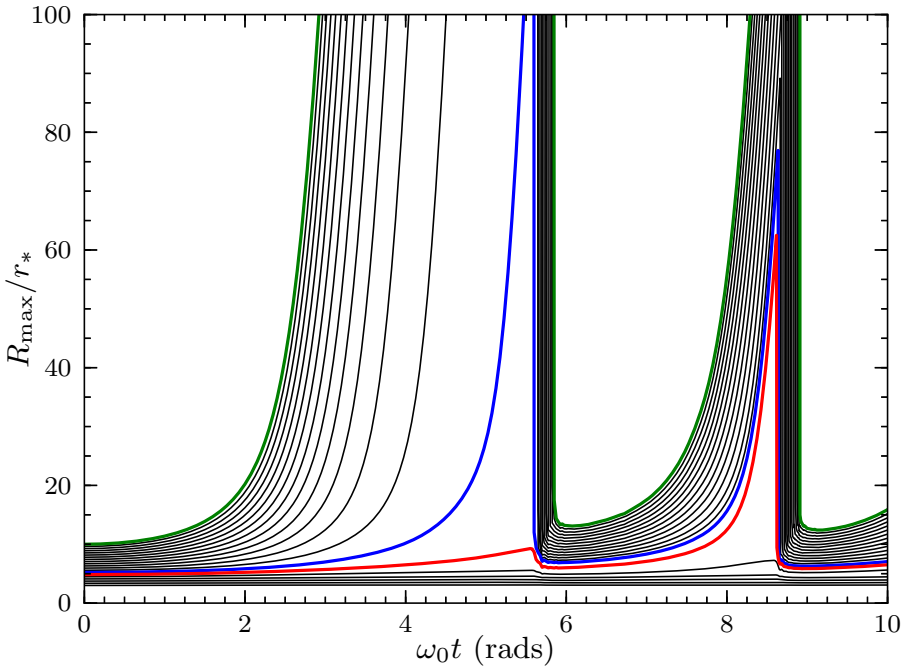

Figure 6.10: Maximum field line heights versus applied shear, for the polar cap model twisted through two reconnection events. At $t = 0$, the twenty field lines shown are equally spaced in $R_{\rm max}$, from $R_{\rm max} = 3$ to 10 $r_\star$. Three curves are highlighted with thicker colored lines.

At the onset of the first reconnection event the magnetosphere has free energy $W_{\rm tw} = 0.0069\, W_0$ in the volume $1 < r < 500$. During the dynamic phase 9.9% of the free energy is expelled from the system, 43.8% is dissipated, and 46.3% is retained in the new quasi-equilibrium state. This dissipation fraction is approximately three times as large as was found in the fiducial equatorial shearing model in Section 6.5.1.

6.5.3 Dependence of $\psi_{\rm rec}$ on shearing rate

We now turn to the effect of the shearing rate, which we study with the shearing profile employed in the preceding section, a polar cap extending from the northern axis to $\theta_{\rm pc} = 0.15\pi$. In these simulations the shearing rate is smoothly increased from zero to ω_0 and then held constant. In

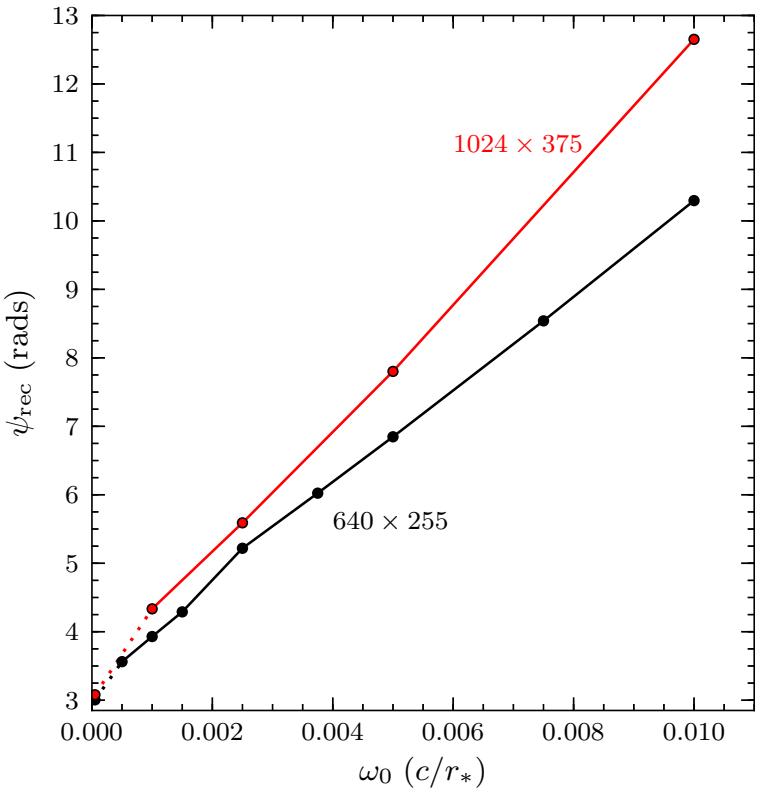

Figure 6.11: Twist angle at the onset of reconnection versus shearing rate, for a polar cap with $\theta_{\rm pc} = 0.15\pi$. The points at smallest ω_0, connected to the others with dotted lines, were found by shearing in stages with progressively lower ω_0. Results for the smaller grid ($N \times L = 640 \times 255$) are shown in black, those from the larger grid (1024×375) are in red.

Figure 6.11, the twist angle at which reconnection begins, $\psi_{\rm rec}$, is plotted against ω_0 for shearing rates from $\omega_0 = 5 \times 10^{-4}$ to 10^{-2}.

The twist angle at reconnection is a strong, approximately linear, function of the shearing rate,

$$\psi_{\rm rec} \approx \psi^0_{\rm rec} + \text{const} \cdot \omega_0. \tag{6.24}$$

This effect appears to be the result of a competition between the formation of a thin current layer following the onset of non-equilibrium at $\psi_{\rm crit}$ and field inflation driven by twisting (see Figure 6.9). The thin current layer appears first at a pinch point, just outside those closed field

lines which do not inflate rapidly and will not undergo reconnection, because the expansion of the opening field lines forms a region of lower magnetic pressure there. The formation of the current layer is slow at first and becomes faster as it proceeds. In its early stages, the narrowing of the high-current region may be counteracted by continual twisting, which causes the last closed field lines to expand, pushing them through the pinch point and preventing the overlying open lines from moving toward the current layer. The rate at which closed flux is pushed through the pinch point decreases as the last closed field line is anchored ever deeper in the magnetosphere, and the thickness of the current layer decreases slowly. Eventually the current layer's thickness approaches the grid scale and it becomes resistive. The resistivity causes the layer to collapse quickly to a discontinuous current sheet, initiating reconnection. The point at which this occurs, for a given shearing profile, will depend on the shearing rate and the resistive length scale set by the numerical grid.

This "dynamical stabilization" explains the deviation of the total magnetic energy from the virial theorem energy at large twist amplitude in Figure 6.7b—the magnetosphere is in the dynamically stabilized state for $\psi \gtrsim 3$. This is confirmed by further simulations, in which we gently implant a certain twist amplitude, the shearing rate being smoothly decreased such that it becomes zero at $t_{\rm tw}$ when the total shear reaches $\psi_{\rm tw}$. The magnetosphere is stable after being twisted to $\psi_{\rm tw} = 3.1$, while current sheet formation and reconnection occur if $\psi_{\rm tw} = 3.25$.

In Figure 6.11, the curves for grids A and B converge as ω_0 is decreased—in the limit of quasi-static twisting, $\psi_{\rm rec} = \psi_{\rm crit}$ is independent of grid resolution and hence resistive length scale. The dynamical stabilization disappears as $\omega_0 \to 0$. The slope of each $\psi_{\rm rec}(\omega_0)$ curve is determined by the resistive length scale of the simulations (here roughly the grid scale). As the resistive length scale is decreased, the dynamical stabilization is effective to larger ψ because higher field gradients must be created before the current layer will resistively collapse, and so the slope $\Delta\psi_{\rm rec}/\Delta\omega_0$

increases with grid resolution.

A finite shearing rate complicates the picture of a magnetosphere evolving quasi-statically toward a well-defined unstable critical point. Any twisting rate, even one which keeps the magnetosphere almost perfectly in equilibrium up to large ψ, as in Figure 6.7, eventually becomes dynamically important, delaying the onset of reconnection. Each model represented in Figure 6.11 undergoes reconnection (on a timescale shorter than the twisting timescale) if the surface shearing is smoothly halted above a twist of $\psi \sim 3$. Beyond this angle they are in dynamically stabilized states—there are no corresponding true equilibrium configurations to which they may gently relax.

The stabilizing effect of twisting can be thought of as being analogous to the effect of pedaling when riding a bicycle—one becomes more stable to spontaneous toppling when pedaling more quickly. Of course, faster pedaling may lead one more quickly to a catastrophic end, such as collision with a tree, whose intensity is amplified by the greater speed. Similarly, a higher twisting rate will drive a magnetosphere to large-scale reconnection in a shorter time (though at a larger twist amplitude), and increase the total amount of magnetic energy expelled and dissipated in the dynamic phase. The application of this "bicycle effect" to the magnetospheres of compact objects with be discussed in Section 6.6.

As $\omega_0 \to 0$, the measured $\psi_{\rm rec}$ should go to the critical angle $\psi_{\rm crit}$. To study the behavior in the very slow twisting limit we evolved the magnetosphere in a series of stages, with decreasing ω_0, in a manner similar to Section 6.5.1; the configuration is allowed time to equilibrate between stages. In the final stage, we evolve from a stable equilibrium state through the point of reconnection, with a shearing rate of $\omega_0 = 5 \times 10^{-5}$. Using the smaller grid, we bring the configuration from a stable state at $\psi = 2.75$ to reconnection at $\psi_{\rm rec} = 3.005$; with the larger grid, we find a stable equilibrium at $\psi = 2.95$ and measure reconnection at $\psi_{\rm rec} = 3.079$. These measurements (shown in Figure 6.11, connected to the constant-twisting values by a dotted line), agree to a fractional discrepancy of

about two percent, confirming that at low twisting rate $\psi_{\rm rec}$ is insensitive to numerical resolution and the details of the spectral filtering.

These measurements may in fact slightly underestimate $\psi_{\rm crit}$. Because of the slow twisting rates and time allowed for equilibration, the outer boundary is not out of causal contact with the star at all times, and eventually outgoing waves on a fraction of the very inflated field lines are removed by the absorbing layer. This artificially reduces the magnetic pressure near the boundary, drawing out the inflated flux, and generally causing the configuration to be further along the path to current sheet formation than it would be naturally at a given ψ. We have confirmed this effect using domains with $r_{\rm out} \sim 100$–250, finding that decreasing $r_{\rm out}$ leads to reconnection at smaller $\psi_{\rm rec}$. This effect may explain why the very-slow-shearing values lie below the lines formed by the constant-twisting measurements, and may be responsible for the kink in the black curve (for the smaller grid) at $\omega_0 = 2.5 \times 10^{-3}$ in Figure 6.11—below this shearing rate, the outer boundary is no longer out of contact for the whole simulation; the effect is not clear in the curve for the larger grid, possibly because of its larger $r_{\rm out}$.

6.5.4 Formation of the current sheet

When the shear applied to a magnetosphere is slowly brought to the critical twist amplitude the field lines expand outward to large heights above the stellar surface. As the field lines expand the current layer separating expanding flux of opposite directions (or sign of B_r) becomes thinner, tending toward a discontinuous current sheet. This field line expansion is an ideal MHD process.

The critical behavior can also be seen if a magnetosphere is twisted past the critical angle at a finite shearing rate, and the shearing rate is then slowly reduced to zero. We use the toroidal current density J_ϕ to measure the strength of the current layer at the equator; a discontinuous current sheet has infinite J_ϕ. Using the larger 1024×375 grid, a polar cap with $\theta_{\rm pc} = 0.15\pi$ is

brought to a maximum twisting rate of $\omega_0 = 2.5 \times 10^{-3}$, which is held constant for approximately $\Delta t = 10^3\, r_\star/c$ and then reduced slowly to zero by $t_{\rm tw}$, such that a total twist amplitude of $\psi_{\rm tw} = 3.25$ is implanted by $t_{\rm tw} = 1690\, r_\star/c$. For $t < t_{\rm tw}$, J_ϕ grows approximately linearly with time. Following $t_{\rm tw}$, the current density continues to increase, albeit more slowly, even though no further shearing is applied. About 100–200 $r_\star/c$ after $t_{\rm tw}$, the current layer's thickness approaches the grid scale, resistivity becomes dynamically important, and the thickness of the current layer dives increasingly quickly to zero, giving $\mathrm{d}^2 J_\phi/\mathrm{d}t^2 > 0$.

The ideal and resistive phases are illustrated in Figure 6.12, which shows J_ϕ on the equator at $r = 40$, approximately the radius at which the X-point first begins to form, for four values of the filtering strength $\alpha_{\rm SSV}$ (a higher value implies more resistivity, and when $\alpha_{\rm SSV} = 0$ the numerical resistivity is due to the high-order filter; see Section 3.6). For $t < t_{\rm tw}$ the evolution is independent of the resistivity level. After this time the models diverge, the current layers in those with higher resistivity collapsing earlier toward a discontinuous current sheet. It is possible that there is a slow ideal collapse phase, and that in the perfect absence of resistivity the magnetosphere would still gradually evolve toward a state with open field lines and an infinitely thin tangential magnetic field discontinuity. However, at the critical point the magnetosphere becomes extremely sensitive to resistivity, and we were unable to isolate a filtering- or resistivity-independent collapse phase. In all our simulations—both using the equatorial shearing and the polar cap twisting models—the identifiable collapse behavior occurs when the current layer thickness is near the grid scale; increasing grid resolution merely slightly increases the twist angle at which resistivity takes over and initiates collapse.

The $\alpha_{\rm SSV} = 0.005$ model (red curve in Figure 6.12) begins to collapse rapidly at $t \approx t_{\rm tw} + 200$; in this simulation the jump in B_r across the equator becomes unresolved by the grid at $t \approx t_{\rm tw} + 436$ (see Figure 6.13). The X-point moves inward as the current sheet becomes stronger (see

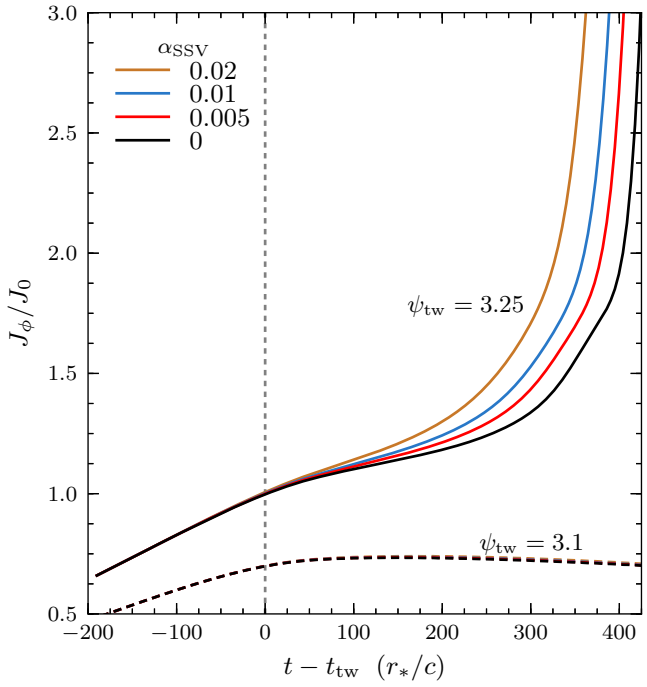

Figure 6.12: Collapse to a current sheet, showing J_ϕ at $(r,\theta) = (40,\pi/2)$, in units of its value at $t = t_{\rm tw}$ ($J_0 = 4.09\times10^{-5}\ c\mu/r_\star^4$), for different values of the filtering strength $\alpha_{\rm SSV}$. The evolution becomes sensitive to the effective resistivity level after the shearing rate reaches zero at $t_{\rm tw}$.

Figure 6.9), and the reconnection criterion $E^2 > B^2$ is first satisfied at $t = t_{\rm tw} + 500.1$ and $r = 24.3$. The collapse and subsequent reconnection occur before any waves from the star have reached the outer boundary.

Figure 6.12 also illustrates how sensitive the magnetosphere becomes to shearing at large twist amplitudes. At $t = t_{\rm tw} - 100$ the shearing rate is only 1.4×10^{-4} (and is slowly decreasing), yet the current density at the equator grows strongly. Once the shearing is withdrawn completely the current density growth slows, and under the influence of resistivity the magnetosphere continues to evolve toward a partially open state. At the filtering levels shown, our code applies very little dissipation on the scales at which the resistive collapse begins. The configuration is very sensitive to resistivity, and even the small amount introduced by the spectral filters is enough to initiate the

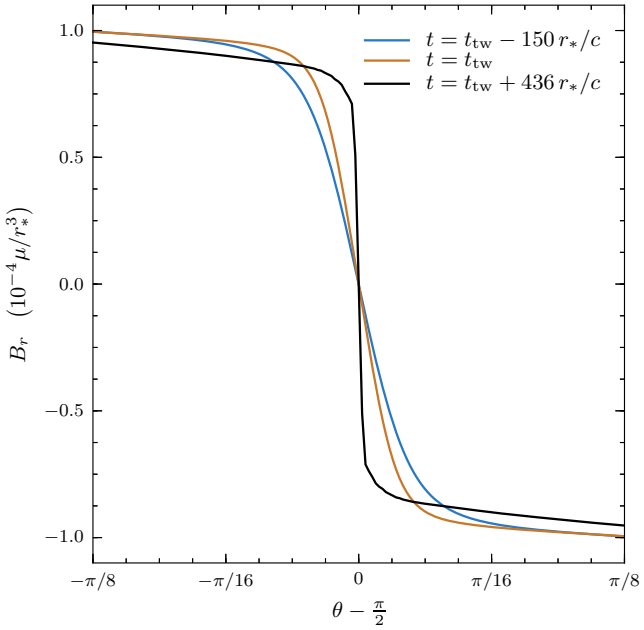

Figure 6.13: Jump in B_r across the equator during the formation of the current sheet, at $r = 40\, r_\star$, from the $\alpha_{\rm SSV} = 0.005$ simulation. The gold curve shows the current layer when the shearing rate reaches zero, $t = t_{\rm tw}$. The jump becomes unresolved by the grid at $t = t_{\rm tw} + 436\, r_\star/c$ (black curve).

resistive instability. As the layer becomes thinner it becomes more resistive, and so the collapse accelerates unstably.

If $\psi_{\rm tw}$ is below the critical amplitude, the magnetosphere instead gently relaxes toward a stable equilibrium. This is shown in Figure 6.12 for $\psi_{\rm tw} = 3.1$, with colored dashed lines representing the same four resistivity levels. These lines nearly lie on top of one another, implying that the relaxation is an ideal process. Eventually, when $t \gtrsim t_{\rm tw} + 600$, the equilibrium begins to resistively diffuse, and the different $J_\phi(t)$ curves begin slowly to diverge depending on the filtering level.

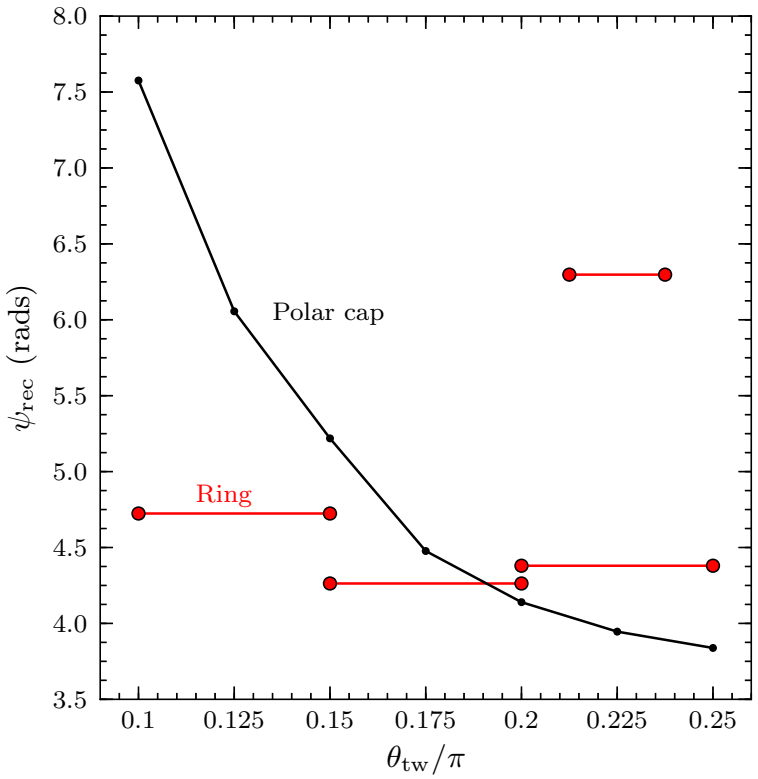

Figure 6.14: Twist angle at onset of reconnection for various shearing profiles, for shearing rate $\omega_0 = 2.5 \times 10^{-3}\ c/r_\star$. The horizontal axis denotes $\theta_{\rm pc}$ for the polar cap models (black curve); for the ring models, the red horizontal lines begin and end at θ_1 and θ_2 respectively.

6.5.5 Dependence on shearing profile

To investigate the effect of changing the profile of the applied surface shear, we performed several simulations (using the smaller grid), at the same shearing rate $\omega_0 = 2.5 \times 10^{-3}$, for polar caps of different sizes and rings of different central latitudes and widths. Figure 6.14 plots $\psi_{\rm rec}$ for these models (for reference, the black point at $\theta_{\rm pc} = 0.15\pi$ refers to the same simulation as the black point at $\omega_0 = 2.5 \times 10^{-3}$ in Figure 6.11).

The polar cap models show a strong dependence of $\psi_{\rm rec}$ on $\theta_{\rm pc}$. Very large twist angles are required to induce reconnection for small polar caps. However, this is a dynamic effect. When

the procedure of evolving increasingly slowly in stages is used, a polar cap model with $\theta_{\rm pc} = 0.1\pi$ undergoes reconnection at $\psi_{\rm rec} = 3.086$, close to the value given in the preceding section ($\psi_{\rm rec} = 3.005$, for a model with $\theta_{\rm pc} = 0.15\pi$). Loosely speaking, the energy required to open a flux bundle centered on f, $W_{\rm f}$, increases more rapidly with flux function near the equator than near the pole: $\mathrm{d}^2 W_{\rm f}/\mathrm{d}^2 f > 0$. This means that continual twisting can dynamically stabilize a flux bundle at smaller f more effectively against collapse and reconnection. In the quasi-static limit, the critical twist angle for a sheared polar cap is roughly independent of its size. This is not surprising given the self-similarity of the dipole field, since changing $\theta_{\rm pc}$ is equivalent to changing the radius of the star.

The ring models are more complicated. They have untwisted field blanketing the twisted flux; this untwisted field has no B_ϕ^2 pressure, and would not naturally undergo poloidal expansion. The untwisted flux acts to confine the twisted field underneath, which must push against its magnetic tension in order to inflate. Consider a ring extending from θ_1 to θ_2, where θ_1 is the colatitude of the footpoint nearer the pole. Equilibrium ring solutions, at a given ψ, have decreasing poloidal expansion with increasing θ_1; in particular, they have inflated less than a polar cap model with $\theta_{\rm pc} = \theta_2$ (a polar cap can be thought of as a ring with $\theta_1 = 0$). Magnetic pressure builds up in the weakly expanding twisted flux bundle; when the configuration loses equilibrium, this magnetic overpressure region can inflate explosively at nearly the speed of light. This explosive mode of field opening will be discussed further in Chapter 7.

There are three principal effects determining $\psi_{\rm rec}$ for a twisted ring.

1. The untwisted flux acts like a nozzle, compressing the expanding flux around the equator. This prevents the dynamic stabilizing mechanism from operating as effectively, because closed field lines need to push overlying open lines out of the way in order to open themselves. Ring solutions therefore tend to experience reconnection at angles closer to their

quasi-static critical angle than polar cap solutions; this is why, in Figure 6.14, the ring extending from 0.1π to 0.15π reconnects earlier than the polar cap with $\theta_{\rm pc} = 0.15\pi$.

2. Decreasing the size of the ring increases the twist required for instability and reconnection, because narrower rings inject less energy into the magnetosphere; this is illustrated by the two rings centered at $\theta_{\rm ctr} = 0.225\pi$.

3. $\psi_{\rm rec}$ increases as a ring is moved closer to the equator, because larger twist angles are needed to inflate high-tension deeply buried field lines, and to push away increasingly "heavy" enveloping untwisted field; this explains why the ring from 0.2π to 0.25π reconnects at a larger twist angle than the ring extending from 0.15π to 0.2π (even though it contains 11% more twisted flux).

6.6 Discussion

Using numerical simulations, we have studied the response of relativistic magnetospheres to slow twisting of the magnetic field lines by shearing of the stellar surface into which the lines are frozen. Although the general force-free twisting problem has relevance to many astrophysical objects, including accretion discs and gamma-ray bursts, here we focus on applications to magnetar theory.

Twisting injects current-carrying Alfvén waves into the magnetosphere, setting up a system of large-scale magnetospheric currents. Resonant Compton scattering of soft thermal photons from the stellar surface, off the particles which constitute these currents, may explain magnetars' non-thermal emission. Magnetic configurations can become unstable when strongly sheared (with twist amplitude $\psi > 1$)—the ensuing dynamic evolution results in reconnection and the release of magnetic energy, a possible mechanism for magnetar bursts and giant flares. The magnetospheric twisting model therefore has the potential to unify these two very different aspects of magnetar

behavior. Our simulations demonstrate what happens when an axisymmetric relativistic magnetosphere is twisted beyond the instability threshold.

6.6.1 Equilibrium solutions

When the shearing is applied slowly and the twist amplitude is small, the magnetosphere moves quasi-statically through a sequence of quasi-equilibrium states (which are arbitrarily close to the true equilibria). As the twist amplitude increases, these equilibria have increasing toroidal magnetic field, current density, and total magnetic energy. The field lines in the poloidal plane expand outwards under the increased magnetic pressure from the toroidal field. The expansion rate, $\mathrm{d}R_{\max}/\mathrm{d}\psi$, is negligible at $\psi = 0$ and increases with ψ (Figure 6.6). Eventually many field lines inflate so rapidly that the toroidal field at each point along them decreases even as the total integrated twist angle continues to grow. In the theoretical final state in the equilibrium sequence, all sheared field lines are open to infinity (i.e. have infinite $R_{\max}$), there is no toroidal magnetic field, and the current density is confined to discontinuous current sheets.

The particular final equilibrium state in which the entire stellar surface is sheared has the maximum energy of any configuration with the same surface normal magnetic field distribution; for the dipole distribution we use, this maximum energy is $W_{\max} = 1.662\, W_0$. This is the greatest magnetic energy that can be stably stored in the magnetosphere. Our numerical solutions clearly evolve toward this state when shearing is applied down to the equator. We produce a configuration that has energy $1.619\, W_0$, energetically 94% of the way from the initial dipole field to the end point of the equilibrium sequence.

The energy stored in the magnetosphere is less if the twisting is not applied globally. For example, a twisted polar cap encompassing half of the star's total magnetic flux can only store approximately $0.07\, W_0$ of free energy and remain in equilibrium (Figure 6.7). This reduces the

size of the energy reservoir which can be released in a subsequent explosion. More energy can be stored if the twisting is confined to a ring of deeply buried field lines, because the tension in the overlying untwisted flux has a confining effect, preserving equilibrium to larger twist amplitudes. One can estimate the magnetic free energy using the approximation that poloidal field expansion is negligible, Equation (6.23); we find that this estimate is accurate for $\psi \lesssim 1.25$, and overpredicts the configuration's energy above this twist angle.

6.6.2 Critical twist amplitude, overtwisting, and reconnection

It has long been argued that at some critical point the twisted magnetosphere would lose equilibrium and enter a fully dynamic state, possibly resulting in magnetic reconnection and the dissipative release of energy. Our simulations support this view. For every shearing profile on the star $\omega(\theta)$, there appears to be a critical twist amplitude $\psi_{\rm crit}$. Configurations with $\psi < \psi_{\rm crit}$ are indefinitely stable, while no stable state having $\psi > \psi_{\rm crit}$ can be created. When the magnetosphere loses equilibrium, field lines inflate at the speed of light toward the open state, forming a thin current layer which collapses into a discontinuous current sheet.

The collapse initiates a phase of unstable reconnection, and all of the flux opened by twisting reconnects to form a fully closed magnetosphere, as illustrated in Figure 6.9. In the process some of the stored magnetic energy is expelled as large-scale Poynting flux and some is dissipated into heat. During reconnection, magnetic flux moves toward the current sheet at $v \sim 0.1\,c$ and newly reconnected field lines leave the reconnection region at nearly the speed of light. This phase lasts $\Delta t_{\rm rec} \approx 50$–$200\, r_\star/c$, which is about 1.5 to 6 ms for $r_\star = 10^6$ cm, depending on where reconnection occurs within $r \lesssim 50\, r_\star$. This is much shorter than the duration of SGR bursts (~ 0.1 s) or the hard spikes of SGR giant flares (~ 0.25 s), but is similar to the rise time to a giant flare's main peak ($\sim$ 1.5–5 ms). The duration of the main peak may be related to the time taken for the relativistic

outflow to become optically thin. The magnetic field is driven into the reconnection region by global magnetic stresses and leaves it at almost the speed of light; therefore the reconnection rate may turn out to be insensitive to the resistive mechanism and length scale.

We find that the fraction of the free energy of the twist that is dissipated in the dynamic phase is about 15% for the equatorial shearing model (Section 6.5.1) and nearly 44% in our fiducial polar cap simulation (Section 6.5.2). While the dissipation fraction appears to be responsive to some numerical parameters, these large fractions are encouraging if one wishes to power energetic gamma-ray flares with liberated twist energy.

Finding the critical twist amplitude is complicated by an additional effect, which is that the continual twisting is itself temporarily stabilizing (Section 6.5.3). Reconnection is inevitable when ψ exceeds $\psi_{\rm crit}$, but its onset can be delayed by continuing to twist beyond $\psi_{\rm crit}$; the faster the twisting, the larger the twist amplitude which can be reached before reconnection begins. A larger twist amplitude translates into more energy stored in the magnetosphere and potentially more energetic bursts or flares when the energy is released. This is analogous to riding a bicycle—one is more stable to falling over when one pedals more quickly, which also stores more kinetic energy which can be released in a collision. At twist amplitude above $\psi_{\rm crit}$, the solution is no longer in a quasi-equilibrium state, because there is no corresponding equilibrium to which the solution would gently relax (without change of magnetic topology) if the shearing were halted; the solution is rather in a "dynamically stabilized" state, entirely dependent for its stability on the maintenance of the shearing of the stellar surface.

The plastic motion of a magnetar's surface is expected to be slow, and so the dynamical stabilization would likely not be a large effect, leading in this case to instability and reconnection close to the critical twist angle. On the other hand, proto-magnetars and other highly energetic objects may have large, even relativistic, shearing rates; these systems may be able to maintain

stability up to larger twist angles, store more energy in their magnetospheres, and so produce more powerful flares during the subsequent reconnection phase. For example, a polar cap model whose critical twist angle is approximately 3 radians can be stabilized up to a twist angle of more than 12 radians by shearing at $\omega = 0.01\ c/r_\star$ (this is likely to be an underestimate, because the realistic resistive length scale may be smaller than that in our simulations).

We find that twisting leads to configurations which are very sensitive to further increases in ψ, and our results are consistent with there being no stable states with twist angle greater than some critical value $\psi_{\rm crit}$. Thin current layers, separating regions of nearly radial magnetic field lines, collapse to discontinuous current sheets when a twist amplitude with $\psi > \psi_{\rm crit}$ is implanted and the surface shearing rate is returned to zero (Section 6.5.4). The collapse dynamics are sensitive to the resistivity level, and we have found no clear evidence for an ideal or resistivity-independent unstable collapse phase. Although the fast current layer collapse appears to be a resistive process, the critical twist angle at which it occurs is only weakly dependent on the effective resistivity (which in our simulations can be changed via either the grid scale or the spectral filtering level). This is because at large twist amplitudes the magnetosphere becomes sensitive to further shearing, and so a large decrease in reconnection length scale would require only a small increase in applied shear to reach the resistive instability threshold. It is in this sense that we speak of an effectively resistivity-independent critical twist amplitude, the end point of the sequence of equilibria constructed by arbitrarily slow twisting, determined only by the shearing profile on the star. These conclusions apply only to the axisymmetric configurations we have studied; in a fully three-dimensional scenario additional instability modes become available, possibly leading to ideal unstable field opening.

For sheared polar caps, $\psi_{\rm crit}$ is approximately independent of the polar cap size, decreasing weakly as the polar cap is enlarged. More extended field lines are more susceptible to dynamic stabilization, and so at a finite shearing rate a small twisted polar cap configuration can be stable

to much larger twist amplitudes than one constructed by shearing a large polar cap, even though the two profiles have similar critical angles. When a ring of flux is twisted, the critical twist angle depends on the size of the ring: narrow rings are stable to large twist amplitudes.

In polar cap simulations, the magnetic field expands gradually and many field lines are inflated to large distances (effectively becoming open) before reconnection. When a ring is twisted the expansion is delayed by the blanketing untwisted flux, and near the instability point the twisted flux bundle suddenly inflates violently. The untwisted flux acts as an elastic nozzle, preventing the outflow from expanding easily in the meridional direction, and many field lines reconnect while their apexes are still relatively close to the star, $R_{\rm max} \lesssim 100 r_\star$. An explosive ring-like event can in fact be produced by a twisting polar cap, if the shearing is continued following the first reconnection event. Some of the strongly twisted low-lying field lines, which did not reconnect in the first event, become unstable as they are twisted further, opening explosively before appreciable twist accumulates on the field lines above them.

Chapter 7

Twisting, rotating magnetospheres

7.1 Introduction

In this chapter, we study the evolution of the differentially twisted magnetosphere of a star that is already in solid body rotation. Stellar rotation adds a new characteristic scale to the problem. Spatially, this scale is the light cylinder radius, $R_{\rm LC}$, the distance from the rotation axis at which corotation with the stellar surface implies motion at the speed of light: $R_{\rm LC} \equiv c/\Omega$. No field lines can close outside the light cylinder in an ideal magnetosphere; rotation therefore puts an upper limit on the possible size of a stable bundle of closed flux. Beyond the light cylinder, field lines are open to infinity. The new characteristic scale can also be viewed as the flux function of the last open field line, $u_{\rm rot} = \sin^2\theta_{\rm rot}$, where $\theta_{\rm rot}$ is the colatitudinal extent of the polar cap of open flux, $\theta_{\rm rot} \approx (r_\star/R_{\rm LC})^{1/2}$.

Observed magnetars rotate with periods $P = 2$–12 s and are gradually spun down, on a timescale comparable to the age of the star, $t \sim 10^4$ yr (e.g. Mereghetti 2008). The spin-down

torque acting on the star is controlled by the open magnetic flux passing through the light cylinder of radius $R_{\rm LC} = cP/2\pi$. As the magnetosphere is twisted, its energy density grows and it tends to inflate (Section 6.4); this opens more magnetic flux, increasing the spin-down torque. Thus, the magnetic twists, besides generating non-thermal emission, can also account for the observed temporal variations in spin-down rate $\dot{P}$.

For example, an outburst from XTE J1810-197 was followed by a change in $\dot{P}$ of more than a factor of three over nine months (Camilo et al. 2007). The spin-down rates of SGRs 1900+14 and 1806-20 have been observed to vary by a factor of four over timescales of months (Woods et al. 2002). The 27 August 1998 giant flare from SGR 1900+14 was coincident with a fractional increase in the spin period of $\Delta P/P = 10^{-4}$ (Woods et al. 1999). Because of an 80 day gap in observations before the flare, the behavior of $\dot{P}$ is unknown. It is unclear if there was a gradual change in spin-down over this period, if ΔP resulted from a brief and dramatic increase in magnetic torque, or if it was a result of a sudden change inside the neutron star, a sort of "anti-glitch" (Thompson et al. 2000).

7.2 Model setup

Twisting of open field lines simply results in the twist energy flowing away from the system. More interesting effects appear if surface shearing twists field lines at larger u, in the closed zone of the magnetosphere. Each twisting, rotating model can be characterized by the dimensionless number a, which relates the amount of twisted flux to the rotationally opened flux, $u_{\rm rot}$. For a polar cap model, it is simply the ratio of the magnetic fluxes emerging through the polar caps $\theta < \theta_{\rm pc}$ and $\theta < \theta_{\rm rot}$,

$$a = \frac{u_{\rm pc}}{u_{\rm rot}} = \frac{\sin^2\theta_{\rm pc}}{\sin^2\theta_{\rm rot}}; \tag{7.1}$$

all models considered below have $a > 1$. Similarly, a ring model is described by two numbers, a_1 and a_2, which label the two flux surfaces that bound the sheared flux bundle in units of $u_{\rm rot}$; the model will be referred to as "a_1–a_2." A ring extending from θ_1 to θ_2 is labeled by $a_1 = \sin^2\theta_1/\sin^2\theta_{\rm rot}$ and $a_2 = \sin^2\theta_2/\sin^2\theta_{\rm rot}$; the fractional magnetic flux through the twisted ring is $(a_2 - a_1)u_{\rm rot}$.

Two numerical grids are employed in the simulations described in this section, depending on the light cylinder radius. The first grid has a resolution of $N \times L = 384 \times 255$ and is used when $R_{\rm LC} = 20\,r_\star$, the second has 768×507 grid points and is used when the rotation rate is half as fast, $R_{\rm LC} = 40\,r_\star$. Increasing the light cylinder radius allows one to twist a larger multiple of the rotationally-opened flux, which shrinks as $R_{\rm LC}$ increases. In particular, all simulations with a or a_2 greater than 7 are performed with the $R_{\rm LC} = 40\,r_\star$ grid.

At each instant of time, there are two further numbers which quantify the changing configuration. The first is the rate at which energy is transferred from the star to the magnetosphere; this is the Poynting flux integrated over the stellar surface,

$$L = \frac{r_\star^2}{2}\int_0^\pi (\boldsymbol{E}\times\boldsymbol{B})_r \sin\theta\,\mathrm{d}\theta. \tag{7.2}$$

The second number is the torque on the star applied by the magnetosphere (the magnetic braking rate), which is the integrated angular momentum flux,

$$T = -\frac{r_\star^3}{2}\int_0^\pi \left(E_r E_\phi + B_r B_\phi\right)\sin^2\theta\,\mathrm{d}\theta. \tag{7.3}$$

In a steady state, the torque is applied solely by open magnetic flux[1]. Shearing increases the torque

[1]During twisting, energy and angular momentum are transferred from the star to the closed zone in the magnetosphere, where they are stored.

on the star because it pushes previously closed field lines through the light cylinder. This is a strong effect at large ψ, which we will study with several models in the sections which follow. We label the Poynting luminosity of, and magnetospheric torque on, the untwisted rotating star by L_0 and T_0 respectively.

7.3 Quasi-steady twisted rotating states

We first consider the combined effects of twisting and solid-body rotation on the sequence of twisted force-free equilibria. Before the shearing is begun, we construct the equilibrium untwisted rotating configuration. A star is brought from rest to a constant rotation rate of $\Omega = 0.05$ (giving $R_{\rm LC} = 20\, r_\star$) over several stellar light-crossing times, and the solution is allowed to relax to a steady state (Figure 7.1a) as described in Chapter 5. This rotating configuration has fractional open flux of $u_{\rm rot} \approx 0.071$. Surface shearing is then applied to this equilibrium state on top of the stellar rotation. Starting at $t = 0$, we shear a polar cap, the magnetic flux through which is three times the flux opened by rotation (i.e. $a = 3$). This cap is slowly brought from corotation to a maximum shear angular velocity of $\omega_0 = \Omega/50$. The twisted polar cap extends down to $\theta_{\rm pc} = \sin^{-1}(3u_{\rm rot})^{1/2} \approx 0.153\pi$, which is comparable to the polar cap model studied in Section 6.5.3.

As the surface is sheared, Alfvén waves are injected onto the field lines rooted in the twisted cap, causing them to bend in the azimuthal direction. Those waves which travel along open field lines flow out of the system at the speed of light, and the toroidal field on these lines does not increase significantly beyond its rotationally induced value. On the other hand, waves on closed field lines are trapped, and so the total twist increases on the twisted closed flux, storing energy in the magnetosphere. As in the non-rotating twisted solutions (see Figure 6.4), the additional magnetic pressure from the toroidal field causes these twisted closed lines to expand outward,

pushing some previously closed field lines through the light cylinder (Figure 7.1b).

When a field line expands through the light cylinder, it initially remains closed in the equatorial current sheet due to finite resistivity there; it opens completely to infinity with only a small amount of further poloidal expansion. The newly opened field line has B_ϕ of opposite signs at its two footpoints, and so makes a net contribution to T. Thus, as more field lines are pushed outside $R_{\rm LC}$ the spin-down torque increases. Field lines undergo erratic reconnection as they are pushed through the resistive current sheet, sometimes creating discrete plasmoids which outflow through the sheet. The solution[2] at a given ψ, such as in Figure 7.1b, is quasi-steady, in that most of the flux (and integrated quantities like L and T) are stable, but a small fraction of the field lines (those closing in the sheet) continue to fluctuate on short time scales.

7.4 Overtwisting of rotating magnetospheres

7.4.1 Relaxation versus reconnection

Here we describe three simulations, in which twists of three different amplitudes, $\psi_{\rm tw} = 1$, 1.5, and 1.75, are carefully implanted in a rotating magnetosphere. We begin with the same steady state rotating solution as in the previous section. At $t = 0$, a polar cap with $a = 3$ is brought from corotation to a shear angular velocity of $\omega_0 = \Omega/50$. The shearing rate is held constant for some time (on the order of ten rotational periods), and then smoothly reduced to zero over a similar time scale, giving a final twist amplitude of $\psi_{\rm tw} = \int_0^{t_{\rm tw}} \omega_0(t)\,{\rm d}t$, where $t_{\rm tw}$ is the time at which the shearing rate returns to zero. The rotating configuration is then evolved with no further shearing.

[2]The bulging of the closed lines above, but not below, the equator in Figure 7.1b occurs because only the field lines' northern footpoints are sheared. The twisting causes B_ϕ to increase by a small amount on northern open field lines—they expand slightly and pull the closed flux below along with them. This effect becomes less pronounced as ω_0/Ω is decreased; it is not evident in Figure 7.3, because by the times shown the shearing has been halted and the extra twist energy on open lines has propagated out of the domain.

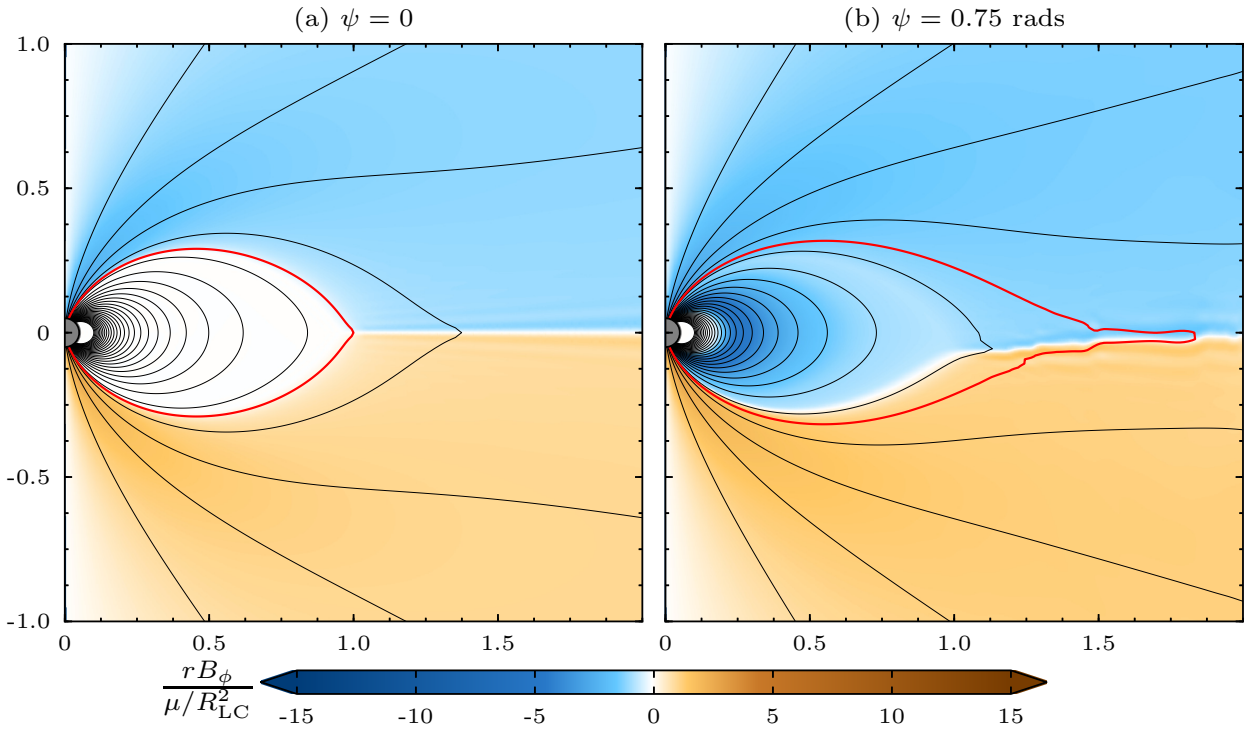

Figure 7.1: (a) Initial solution: untwisted rotating star; (b) quasi-steady twisted magnetosphere. Color shows toroidal magnetic field times the radial coordinate, in units where this quantity is initially ~ 1 at the light cylinder; poloidal field lines, thirty equally spaced in flux function between $u = 0.01$ and 0.5, are shown in black; an additional field line, that initially closes at the light cylinder, is marked in red. Axes are labeled in units of $R_{\rm LC} = 20\, r_\star$.

In the $\psi_{\rm tw} = 1.0$ simulation, $t_{\rm tw}$ is reached without any large-scale dynamic motion taking place. For $t > t_{\rm tw}$, some of the field slowly diffuses back through the light cylinder due to dissipation in the current sheet. As the open flux decreases, so does the spin-down torque, which fluctuates as plasmoids are created and expelled (Figure 7.2a). Over time, a narrow band of approximately untwisted flux is formed in the closed zone, between the light cylinder and the twisted field lines which were never opened (Figure 7.3a).

In the $\psi_{\rm tw} = 1.75$ simulation, the configuration gradually becomes less stable as the twisting rate is reduced to zero, as the dynamical stabilization effect due to the finite shearing rate

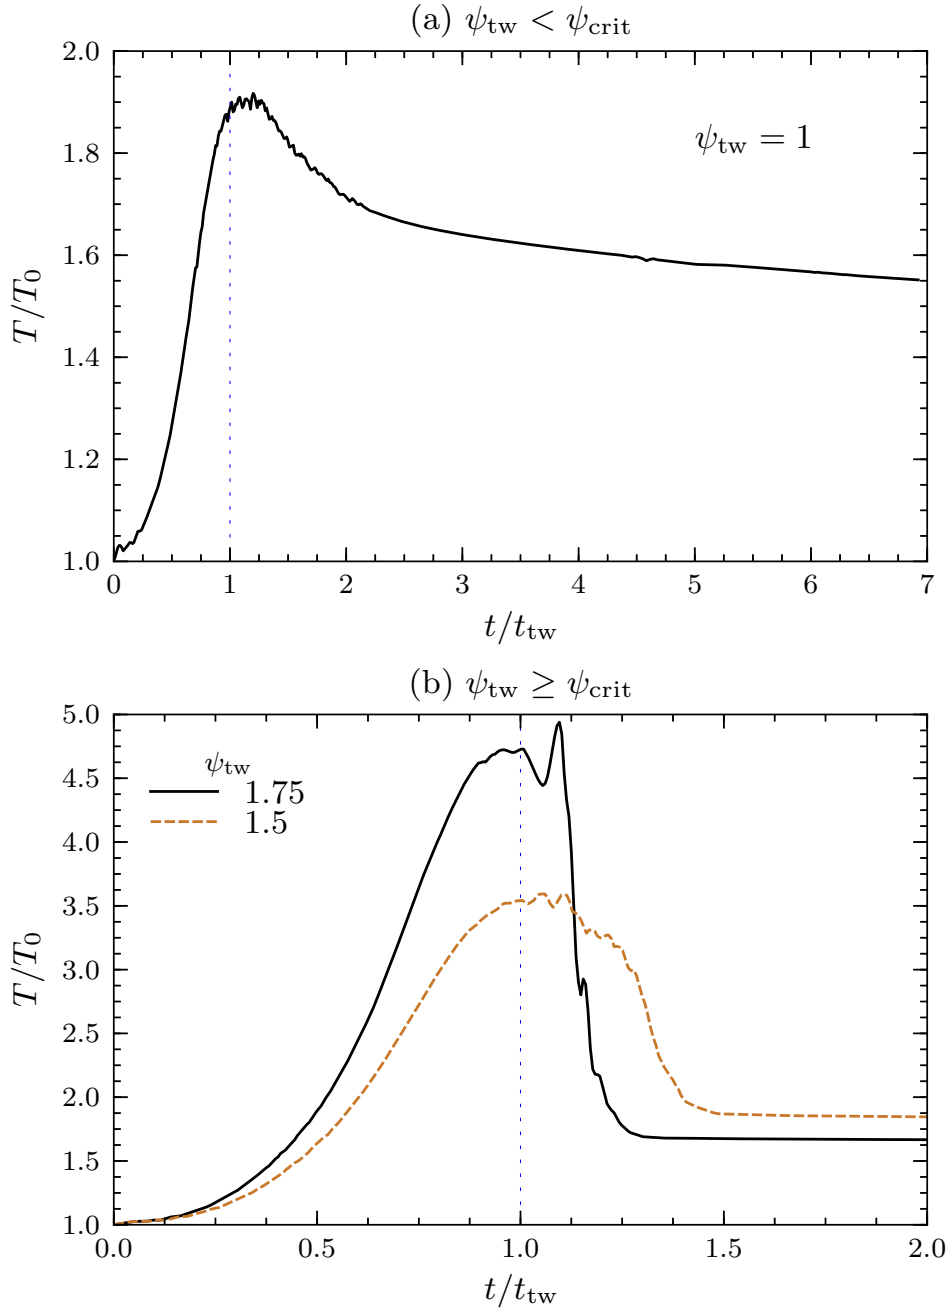

Figure 7.2: Torque on the star, in units of the initial torque due to rotation, for polar caps sheared to $\psi_{\rm tw}$ over time $t_{\rm tw}$. (a) $\psi_{\rm tw} < \psi_{\rm crit}$: slow diffusion of magnetic flux back through the current sheet; (b) $\psi_{\rm tw} \gtrsim \psi_{\rm crit}$: the magnetosphere undergoes large-scale reconnection to a twisted rotating steady state. A dotted line indicates $t = t_{\rm tw}$.

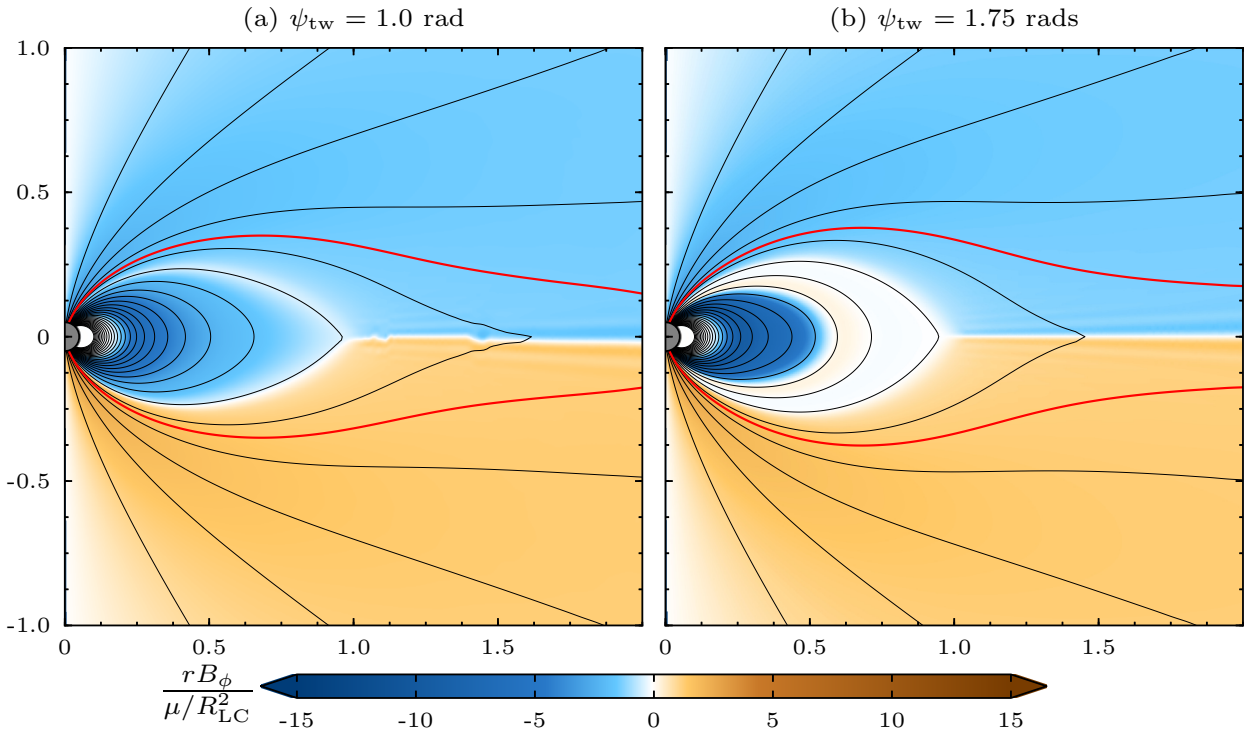

Figure 7.3: Twisted rotating solutions at $t = 2\,t_{\rm tw}$, where $t_{\rm tw}$ is the time taken to implant a twist of amplitude $\psi_{\rm tw}$. (a) $\psi_{\rm tw} = 1 < \psi_{\rm crit}$: a zone of untwisted flux slowly grows just inside the light cylinder, as some open flux diffuses back through the current sheet; (b) $\psi_{\rm tw} = 1.75 > \psi_{\rm crit}$: steady state following sudden reconnection of overtwisted flux. The bundle of twisted flux (with large B_ϕ) inside the light cylinder is the twisted reservoir; the flux with $B_\phi \sim 0$ is the cavity. Color and field lines are as in Figure 7.1.

(Section 6.5.3) becomes weaker. At some point near $t_{\rm tw}$, a bundle of closed flux expands rapidly through the light cylinder, but the field lines beneath it are still in approximate pressure equilibrium and remain closed. The inflation of this last opening field line is shown in Figure 7.4: the field line with $u = 0.1435$ expands rapidly through the light cylinder, while the next field line shown (with $u = 0.1465$) does not. A current layer forms behind the $u = 0.1435$ field line, which becomes thinner as its apex moves outward at approximately the speed of light—this is an ideal (dissipationless) collapse to a current sheet. Note that this collapsing current layer is distinct from the discontinuous current sheet just below it; the former comprises twisted field lines having B_ϕ of

only one sign (here, negative) along their entire lengths, while the latter, created by rotation, is a discontinuity across which both B_ϕ and B_r change sign.

An X-point (or pinch) geometry is created at the cusp separating closed and inflating flux at $r \approx 0.8\,R_{\rm LC}$, and the field gradients (and hence effective numerical resistivity) become large enough to trigger reconnection at the X-point. Reconnection first occurs across the twisted current layer, not the rotationally-induced current sheet. Following reconnection, those parts of field lines which are attached to the star snap back toward it, removing pressure support around the X-point; this causes more flux to dive into the reconnection region, including field lines which previously straddled the current sheet (such as the $u = 0.125$ and 0.135 field lines in Figure 7.4). The O-geometry field line sections are untethered from the stellar surface, and are expelled, carrying most of the twist energy out of the system. The torque on the star plummets as more open flux becomes closed inside the light cylinder (Figure 7.2b).

Reconnection continues until a new equilibrium is reached (Figure 7.3b); since no further shearing is applied, this is a true equilibrium state, in contrast to the quasi-steady state in Figure 7.1b. In the equilibrium solution, there is a bundle of strongly twisted field lines, the "twisted reservoir," which were never open and so did not experience reconnection. There is also a region of untwisted closed flux, the "cavity," between the reservoir and the flux surface closing at the light cylinder; the twist on these field lines was emitted to infinity while the field lines were open. The twist in the reservoir causes more flux to be open than in the initial untwisted state at $t = 0$. In general, the larger the twist amplitude $\psi_{\rm tw}$, the more flux is opened and reconnects, and so the smaller the twisted reservoir (and the closer the new equilibrium state is to the untwisted rotating solution; Figure 7.2b).

We find that a slow mode of relaxation, due to ejection of small plasmoids, operates if $\psi_{\rm tw} = 1$; after $t_{\rm tw}$ a new equilibrium is gradually approached without large-scale reconnection or significant

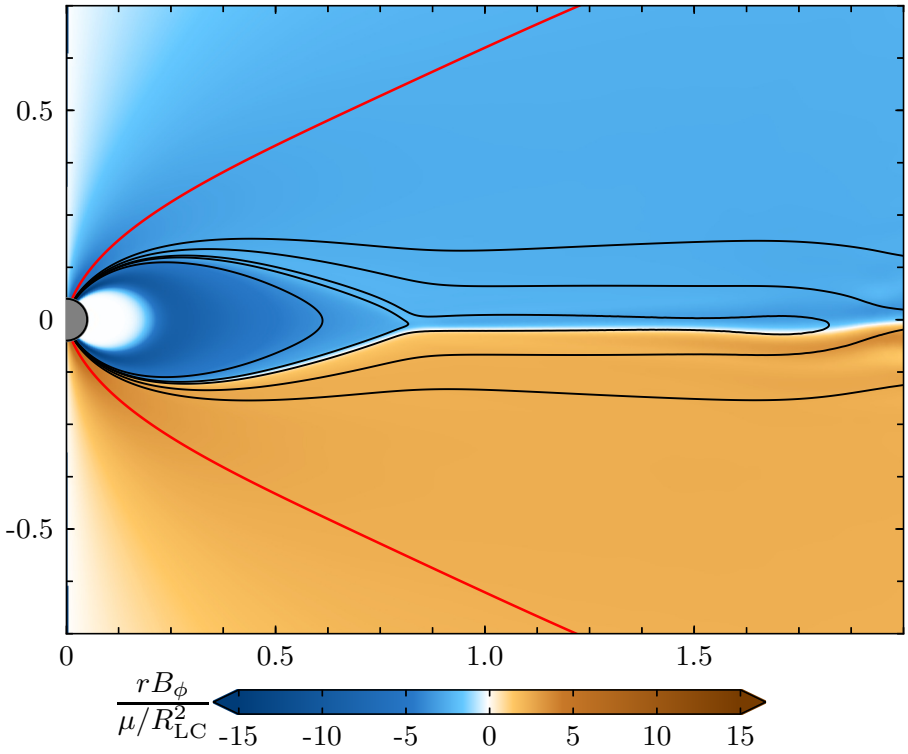

Figure 7.4: Formation of the reconnecting current layer in the $\psi_{\rm tw} = 1.75$ simulation, at $t = 1.09\,t_{\rm tw}$. Black curves indicate field lines at $u = 0.125$ (outermost line), 0.135, 0.1435, 0.1465, and 0.155 (innermost line); the red curve shows the field line that closes at the light cylinder at $t = 0$. Reconnection will begin in the current layer between the $u = 0.1435$ and 0.1465 field lines, which lies above the discontinuous current sheet (indicated by the color discontinuity). Color is toroidal magnetic field multiplied by the radial coordinate. Axes are labeled in units of $R_{\rm LC}$.

dynamics. If $\psi_{\rm tw} = 1.75$ a catastrophic reconnection event occurs at $t_{\rm tw}$, and the configuration reaches a new equilibrium state on a much shorter timescale. The simulation with $\psi_{\rm tw} = 1.5$ displays intermediate behavior: first several smaller plasmoids are released, but the system soon enters an unstable rapid-reconnection phase (see Figure 7.2b). This division into gradual relaxation and large-scale fast reconnection leads us to conclude that this system displays the same critical behavior as non-rotating twisted magnetospheres: there is a dramatic reconfiguration when the system is overtwisted, which occurs when the twist amplitude exceeds a critical value $\psi_{\rm crit}$. In the

slow shearing limit, magnetic reconnection immediately follows the loss of equilibrium, however reconnection can be delayed by a finite shearing rate. The particular model we describe here has $\psi_{\rm crit} \approx 1.5$, which is smaller than the value for the corresponding non-rotating problem, $\psi_{\rm crit} \approx 3$ (Section 6.5.3).

7.4.2 Shearing through multiple reconnection events

In the preceding section we described simulations in which a twist of a specific amplitude is implanted in the magnetosphere, and no shearing is applied thereafter. Now we consider models in which the shearing rate, rather than being returned to zero, is maintained at a constant value ω_0, thus shearing the magnetic footpoints through large angles $\omega_0 t$ over the duration of the simulations. In the following we will use $\omega_0 t$ to label the applied surface shear instead of ψ, which we reserve for the angular displacement between the footpoints of a given field line (this will not be the same for all twisted field lines once the magnetosphere has experienced a reconnection event).

The evolution of the twisting rotating polar cap magnetosphere is as follows; we describe in particular a simulation with a slow shearing rate of $\omega_0 = \Omega/200$. As the twisting begins, toroidal magnetic field builds up on the closed field lines connected to the sheared cap (on open field lines the twist is emitted to infinity rather than accumulated). As the twist pumping $\dot{\psi}$ is slow compared with the Alfvén crossing time, the magnetosphere first evolves through a sequence of quasi-equilibrium states. The increased magnetic pressure in the closed zone causes these field lines to expand outwards, pushing more flux through the light cylinder and hence increasing the spin-down torque on the star.

The newly opened field lines lose their twist—it is emitted to infinity—and establish the usual spin-down-wind configuration where the field changes sign across the equatorial plane. The magnetic field discontinuity is supported by the equatorial current sheet, terminating at a Y-point

which separates closed and open flux. The magnetic field remains almost reflection-symmetric about the equator (modulo differences in sign between oppositely directed flux), because the twisting timescale is much longer than the wave-crossing timescale on the closed field lines, allowing waves to distribute the twist nearly symmetrically.

The expansion rate increases while the twist ψ grows with constant rate ω, and eventually the field lines expand through the light cylinder faster than they adjust to a new spin-down steady state; this sets up a pattern of magnetospheric "breathing," where the amount of open flux oscillates, via reconnection in the current sheet, with increasing amplitude and a period of $\sim 7\ R_{\rm LC}/c$ (Figure 7.5b).

After a significant fraction of the polar cap flux has been opened, at $\psi \sim 2$, the recently opened field lines become unstable to catastrophic reconnection and a large fraction of the open flux reconnects, bringing the Y-point within the light cylinder. The released magnetic energy is expelled in a plasmoid-fragmented outflow in the equatorial plane. The Y-point quickly moves back to the light cylinder, opening up a cavity of zero toroidal field between the low-lying strongly twisted region and the open flux bundle (Figure 7.5c). We will term the evolution from Figure 7.5a to Figure 7.5c a "gradual" expansion and reconnection event.

The accumulated twist $\psi = \omega_0 t$ remains intact on those lines, lying closer to the star, which never opened; let us call these field lines the "twisted reservoir." As the shear motion of the polar cap continues, toroidal field is added to both the twisted reservoir and the cavity. Before a new significant ψ builds up in the cavity, the over-twisted reservoir becomes unstable and expands outward explosively, and all the overlying field lines are briefly pushed through the light cylinder (Figure 7.5d). This process is reminiscent of magnetic detonation (Cowley & Artun 1997). There is a narrow spike in the spin-down torque, larger than in the previous event because more flux is opened. The enhanced open flux immediately reconnects and closes, leaving a reservoir of lesser

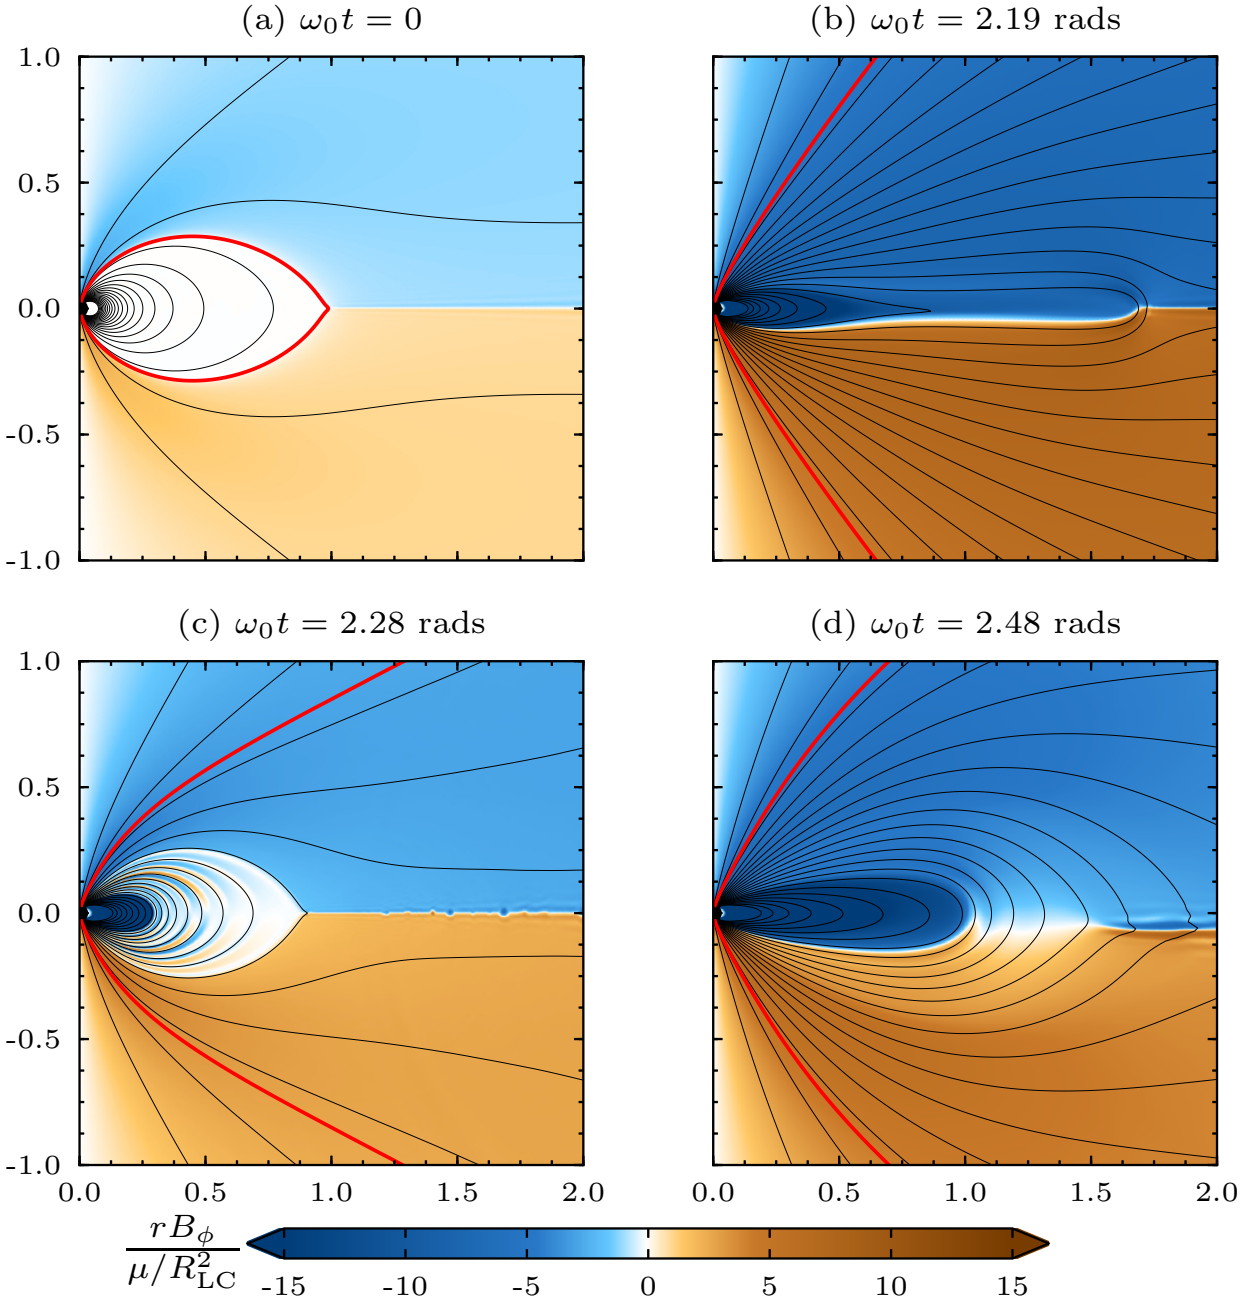

Figure 7.5: The first two expansion and reconnection events from a simulation with $a = 3$ and $\omega_0 = \Omega/200$. Color shows toroidal magnetic field times the radial coordinate, in units where this quantity is initially ~ 1 at the light cylinder. Black curves show poloidal field lines, equally spaced in poloidal flux function between $u = 0.01$ and $1/3$ with separation $\Delta u \approx 0.015$; one additional field line, that initially closes at the light cylinder, is shown in red. Axes are labeled in units of $R_{\rm LC} = 40\, r_\star$. (a) The initial state (untwisted rotating star); (b) the last outward "breathing" motion before the first reconnection; (c) after the first reconnection, the inner dark blue region is the twisted reservoir, the almost white area is the cavity; (d) detonation—the reservoir expands explosively.

volume with $\psi \approx \omega_0 t$ and a bigger cavity with $\psi \approx 0$. This is an "explosive" expansion and reconnection event.

This cycle of a twisted reservoir exploding and reconnecting is repeated several times as the shearing continues. In the first few events successively more flux is opened, leading to higher spikes in spin-down rate, and a deeper trough in torque immediately following reconnection, since the twisted reservoir comprises ever fewer field lines. The spikes in spin-down power become increasingly narrow as they become more energetic, with the most dramatic events having a duration of only a few light crossing times of the light cylinder. The torque evolution for this simulation is given by the black curve in Figure 7.6.

The time between events grows, because the twisted reservoir is shrinking. Eventually, the additional twist needed to detonate the reservoir becomes greater than ~ 2 rads, and the cavity fills with enough toroidal field to initiate a gradual reconnection event before the reservoir explodes. After this point the picture becomes more complicated, as there can be a twist-free cavity on top of a twisted reservoir comprising both flux that did not open in the last event, and flux that has never opened. Different amounts of flux will then expand and open depending on the history of the system.

7.4.3 Varying the shearing rate and polar cap size

In Section 6.5.3 we described how the surface shearing rate strongly affects the twist amplitude at which the magnetosphere undergoes reconnection. We now investigate the effect of varying the rate at which a rotating magnetosphere is sheared. The spin-down torque T, for simulations having the same sheared polar cap ($a = 3$) but varying ω_0/Ω, is plotted against applied surface shear in Figure 7.6.

Up to $\omega_0 t \sim 1.5$, the torque is very similar for the three runs shown in Figure 7.6. At this

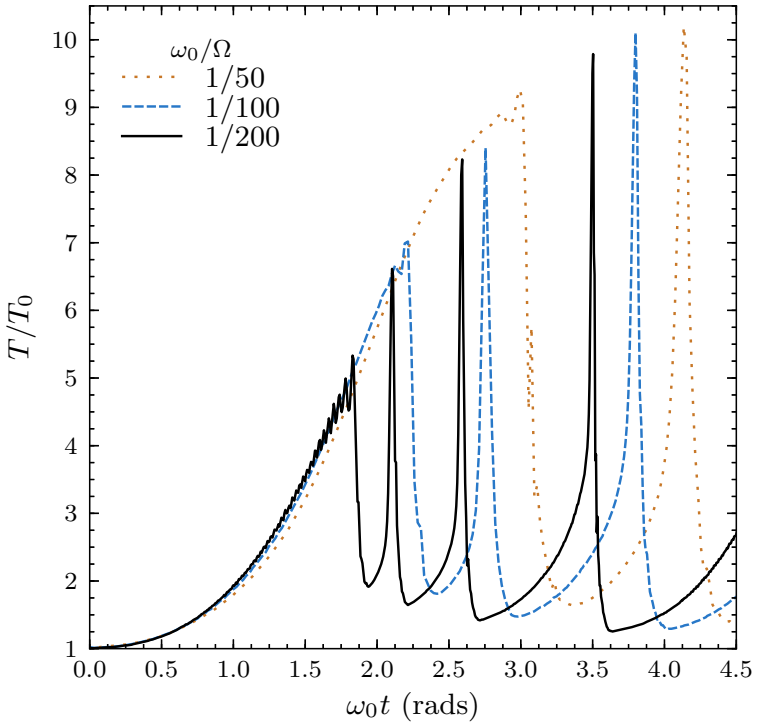

Figure 7.6: Torque on the star, in units of the torque at $t = 0$, for a polar cap with $a = 3$ and varying twisting rate ω_0.

point the run with lowest shearing rate, $\Omega/200$, enters a phase of magnetospheric "breathing," where bundles of closed flux inflate rapidly through the light cylinder, are initially unsupported by expanding underlying flux and so begin to close again, but are prevented from initiating runaway large-scale reconnection by the eventual expansion of the lower-lying field. These increasingly vigorous motions are accompanied by oscillations in T of growing amplitude and period $\sim 7R_{\rm LC}/c$. At $\omega_0 t \approx 1.8$ the newly opened field lines begin to collapse on a timescale shorter than the time needed for the lower-lying flux to push them back through the light cylinder, and a large fraction of the flux opened by twisting reconnects.

Faster shearing can stabilize the magnetosphere against large-scale reconnection, similar to the results of simulations without solid-body rotation $\Omega = 0$; in these simulations with larger ω_0/Ω

the breathing phase is absent. However, the twist amplitude needed to move between subsequent brief ring-like explosive events is less sensitive to the shearing rate, because these are due to sudden unstable dynamic expansion of the twisted reservoir, and the overlying less-twisted flux acts like an elastic nozzle, inducing reconnection soon after inflation (see Section 6.5.2).

For $\omega_0 = \Omega/200$, $\psi_{\rm rec}$ is roughly equal to the critical twist amplitude estimated in the preceding section; this shearing rate is used in the simulations described below.

Increasing the size of the twisted polar cap allows a larger multiple of the star's rotationally opened flux to be pushed through the light cylinder, giving a greater spin-down torque enhancement. We have run simulations with $a = 2$ to 12 with shearing rate $\omega_0 = \Omega/200$, and find that the peak spin-down torque scales as, and for the largest events is roughly equal to, a^2,

$$\frac{T_{\rm peak}}{T_0} \approx a^2. \tag{7.4}$$

The torque is plotted against applied shear in Figure 7.7, where the quadratic scaling of the spike heights is apparent. The evolution to the maximum footpoint displacement shown ($\omega_0 t = 8$) takes roughly 263 stellar rotation periods with shearing rate $\omega_0 = \Omega/200$. Note also that solutions with larger a undergo their first reconnection event at larger twist amplitude, since twisted deep flux pushes the overlying field lines outward; this slightly greater expansion allows the dynamical stabilization mechanism to be effective to larger twist. The average torque enhancement over the length $t_{\rm tot}$ of the simulation,

$$\langle T \rangle = \frac{1}{t_{\rm tot} T_0} \int_0^{t_{\rm tot}} T(t')\, {\rm d}t', \tag{7.5}$$

is 2.64 for $a = 3$, 5.45 for $a = 6$, and 10.75 for $a = 12$ (where $\omega_0 t_{\rm tot} = 8$). We conclude that for twisted polar caps, $\langle T \rangle \sim a$ when averaged over several flux expansion and reconnection events.

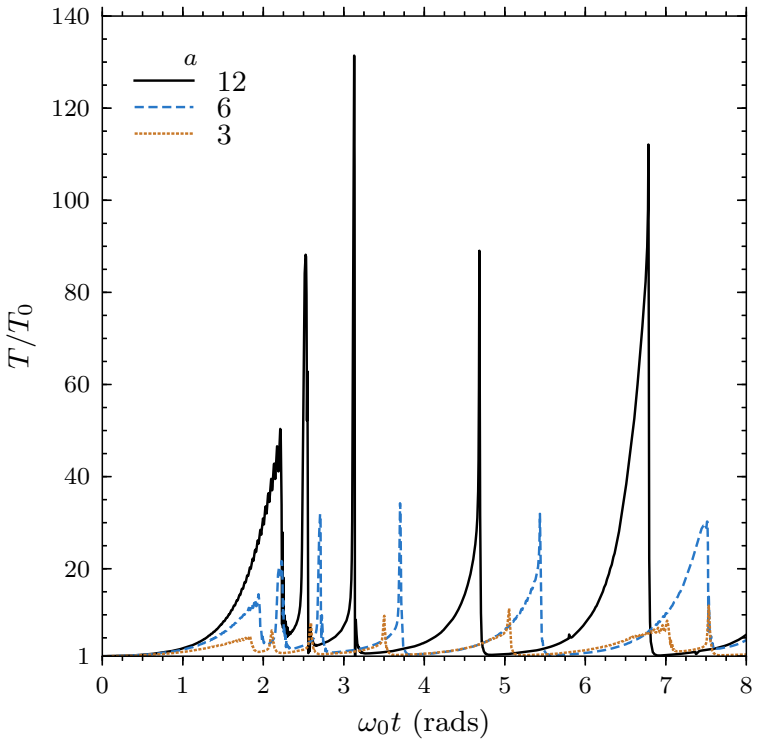

Figure 7.7: Torque on the star for polar caps of varying sizes a, with shearing rate $\omega_0 = \Omega/200$. The red dotted curve in this figure represents the same simulation as the solid black curve in Figure 7.6.

7.4.4 Ring shearing in the closed magnetosphere

The behavior seen in Figure 7.7 is complicated—expansion and reconnection events can be gradual (on the twisting timescale) or explosive (on the light-cylinder-crossing timescale), or lie somewhere between the two, depending on the how the magnetosphere has evolved up to the time of the event.

The evolution of a narrow ring of deeply buried flux is simpler: there is little poloidal expansion (and hence increase in torque) until a large twist amplitude has accumulated, at which point the twisted flux bundle inflates explosively under high magnetic pressure. Figure 7.8 shows the torque enhancement for two ring models having the same ratio of twisted to rotationally opened flux, but located at different latitudes. Both simulations have, on average, narrower torque spikes than the

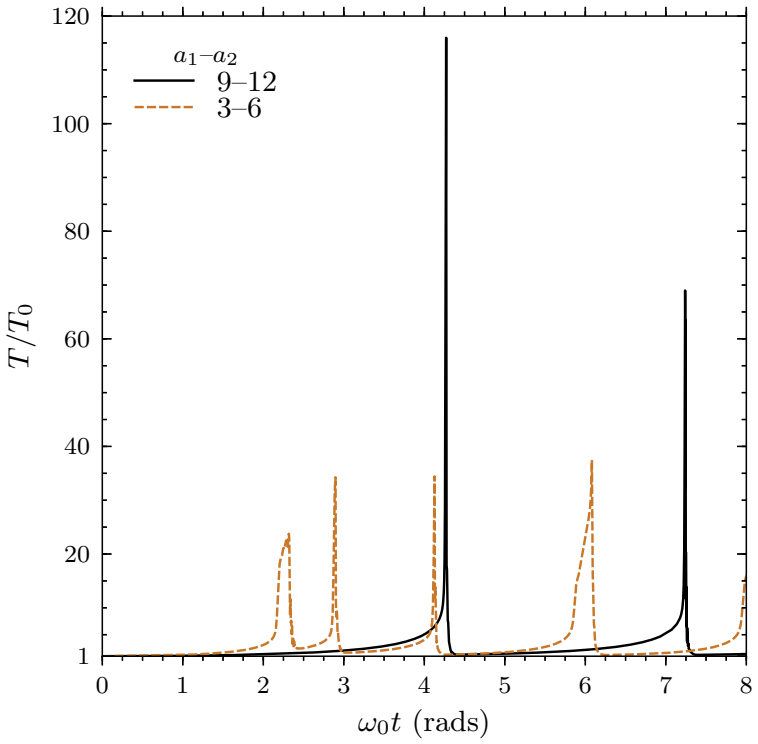

Figure 7.8: Torque enhancement due to slow twisting of two rings, having different a_1–a_2 but containing the same multiple of the rationally opened flux.

polar cap models (Figure 7.7), with the more deeply buried twisted ring powering distinctly more explosive events (only two in eight radians of twisting, and both on the light-cylinder-crossing timescale). The lower flux surface of the ring, a_2, determines the torque peak height, because during an eruption the twisted field line bundle pushes some of its own flux and all of the overlying flux (at smaller u) through the light cylinder,

$$\frac{T_{\rm peak}}{T_0} \approx a_2^2. \tag{7.6}$$

The more deeply buried (far from the open flux bundle) the ring, the larger the twist amplitude required for detonation, and the briefer and more powerful the expansion and reconnection event.

The first detonation of the $a = 9$–12 simulation is shown in Figure 7.9. At small to moderate applied shear, the deeply buried twisted flux bundle aquired a large toroidal field component, but there is little poloidal expansion—even by $\omega_0 t = 2$, only a small amount of additional flux has been pushed through the light cylinder (Figure 7.9a). When the shear is close to the critical twist angle for the twisted flux, the twisted bundle begins to expand poloidally, causing more of the overlying field lines to open; in this phase the magnetosphere becomes more sensitive to the twist amplitude, but the field expands stably on the shearing timescale (Figure 7.9b). At the critical twist amplitude, the twisted flux bundle loses equilibrium, and inflates rapidly under its high magnetic pressure (Figure 7.9c). This unstable inflation occurs on a dynamical timescale, on the order of $R_{\rm LC}/c$. Near the base of the expanding flux, a current sheet forms behind the last expanding field line; in Figure 7.9c, this current sheet will be just below the field line drawn at $u = 0.4$. Catastrophic unstable reconnection then takes place in the current sheet, as described in Section 7.4, leaving a smaller twisted reservoir. The Y-point at the boundary of the newly formed cavity retreats back to the light cylinder during the reconnection process (Figure 7.9d). The magnetosphere returns to an equilibrium spin-down steady state after the Y-point reaches the light cylinder.

The widths of the various torque spikes for each of the five simulations are given in Table 7.1, where their full widths at half maximum (FWHM) are scaled to the light-crossing time of the light cylinder, $R_{\rm LC}/c$. The polar cap models ($a = 3, 6, 12$) have several kinds of events: rapid explosions on roughly the light-crossing time, gradual inflations on the twisting timescale $\gg R_{\rm LC}/c$, and those which are somewhere between the two. The ring simulations generally have briefer, more dramatic, events—the $a = 9$–12 run in particular has only two torque spikes, both of duration $\sim 2R_{\rm LC}/c$.

Magnetic energy is dissipated into thermal energy during the reconnection phase, which corresponds to the sharp downstrokes in the torque curves in Figures 7.6 to 7.8. The reconnection timescale $\Delta t_{\rm rec}$ mostly depends on the twist distribution: following gradual "polar-cap-like" expan-

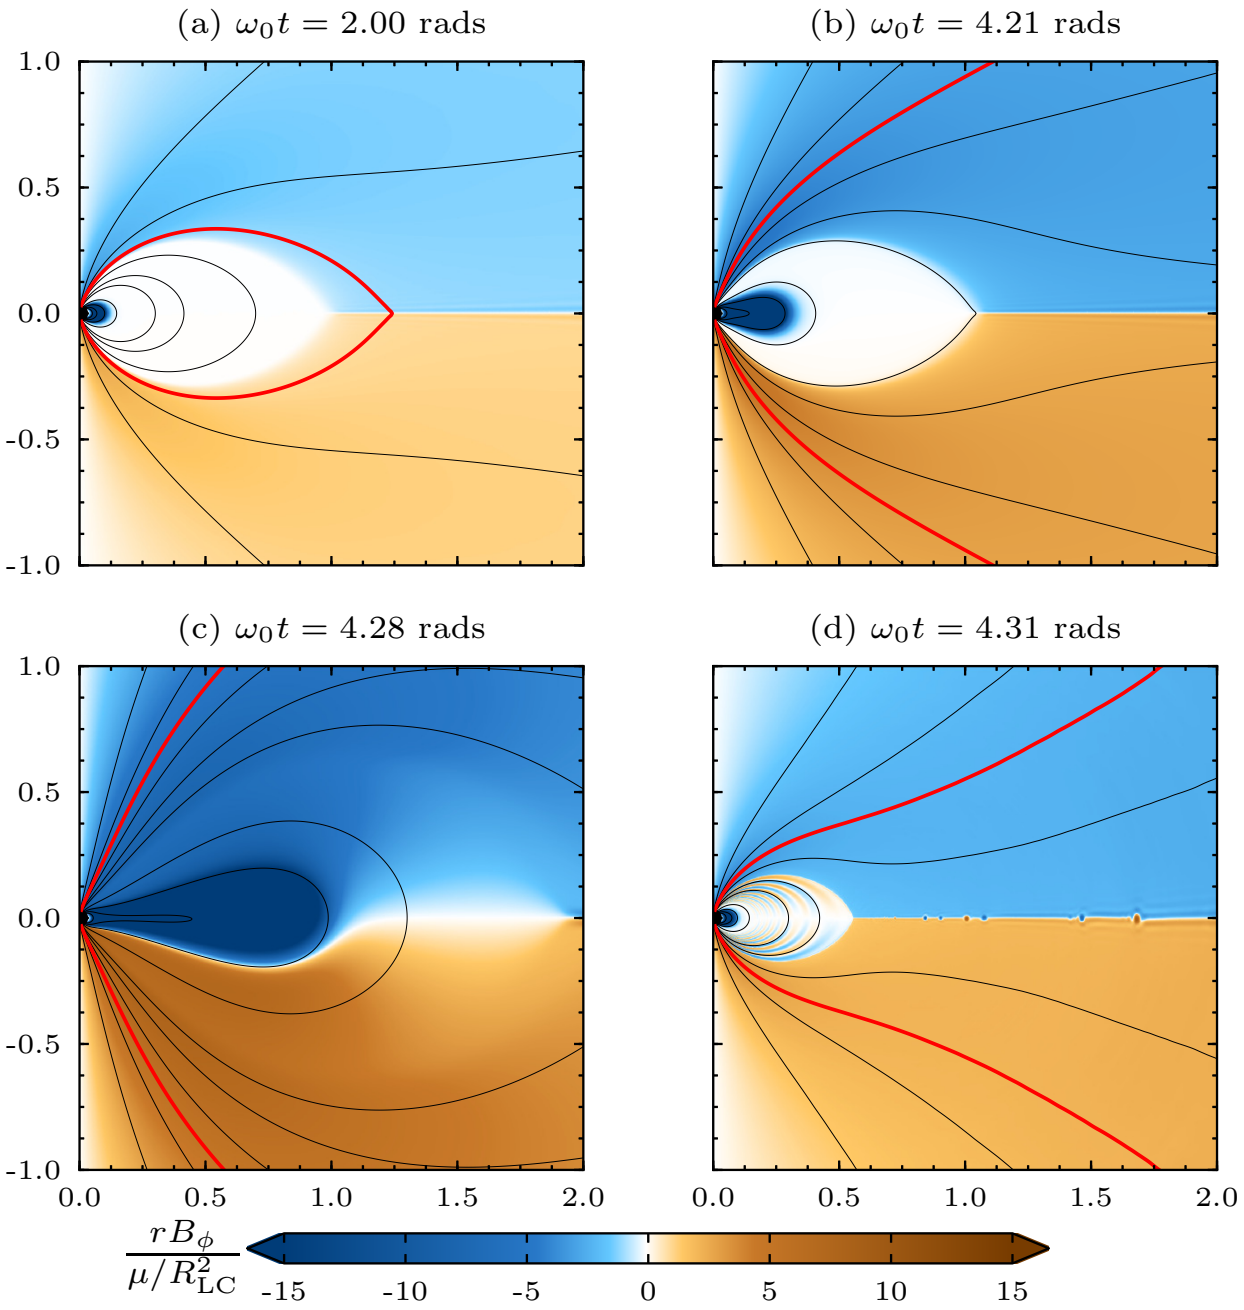

Figure 7.9: The first explosive reconnection event for a twisted ring with $a = 9$–12. Color is as in Figure 7.1; the black curves show poloidal field lines, one drawn at $u = 0.01$, four equally spaced between $u = 0.025$ and 0.1, and five between $u = 0.1$ and 0.5. The red curve indicates the field line initially closing at the light cylinder. Axes are labeled in units of $R_{\rm LC} = 40\, r_\star$.

a	$R_{\rm LC}$	Peak 1	Peak 2	Peak 3	Peak 4	Peak 5
3	20	79.5	8.4	7.5	6.7	12.3
6	20	65.3	15.2	6.4	4.4	17.9
12	40	55.7	11.9	4.5	2.8	30.6
3–6	20	29.9	5.9	3.1	25.9	56.9
9–12	40	2.3	2.1			

Table 7.1: Torque peak widths: $\Delta t_{\rm FWHM}/(R_{\rm LC}/c)$

sion, $\Delta t_{\rm rec} \approx 100$–$200\ r_\star/c$, while after explosive "ring-like" flux breakout the timescale is shorter, $\Delta t_{\rm rec} \approx 50$–$70\ r_\star/c$, because of the nozzle formed by the confining untwisted field lines (whether the actual shearing profile on the stellar surface is a ring or polar cap is unimportant). These results are independent of $R_{\rm LC}$ and grid resolution; scaled to the radius of a neutron star, they predict reconnection lasting roughly 1.5 to 6 ms.

7.5 Asymmetry, linear momentum transfer, and trapped Alfvén waves

We have discussed how twisting of the magnetosphere increases the spin-down torque on the star. The rotating simulations are slightly asymmetric about the equator, since only northern latitudes are sheared, raising the possibility of a net transfer of linear momentum along the star's rotation axis, which we shall call the z-axis. A net force could arise from asymmetric spin-down as the magnetosphere is inflated, or from asymmetry in the final dynamic phase in which the overtwisted flux is ejected.

An upper bound can be placed on the net force by noting that the momentum density carried by the electromagnetic field, p, does not exceed its energy density divided by c. Therefore the greatest momentum that can be transferred from the magnetosphere, if it is entirely emitted to

infinity from one side, is $p_{\rm max} = W_{\rm max}/c$, where $W_{\rm max} \approx W_{\rm pot}$ is the maximum energy that can be stored, and so

$$p_{\rm max} \approx \frac{\mu^2}{3r_\star^3 c}. \tag{7.7}$$

The maximum estimated force, if this momentum is transfered over a time $r_\star/c$, is then

$$F_{\rm max} \approx \frac{\mu^2}{3r_\star^4} = 3\times 10^{41} \left(\frac{\mu}{10^{33}\,{\rm G\,cm^3}}\right)^2 \left(\frac{10^6\,{\rm cm}}{r_\star}\right)^4 {\rm dynes}. \tag{7.8}$$

For a neutron star of mass $M_* = 1.4\,M_\odot$, momentum conservation $M_* v_{\rm kick} = p_{\rm max}$ implies a maximum kick velocity of only $v_{\rm kick} \sim 0.03$ km s^{-1}, for an ultra-strong magnetic field $B \sim 10^{15}$ G that corresponds to $\mu \sim 10^{33}$ G cm^3.

Even though a dramatic rocket effect is impossible, it is interesting to calculate the applied force along the z-axis, as it provides both a measure of asymmetry and a window into the inner magnetosphere during a flare-like eruption. The flux of z-momentum onto the $r = r_\star$ surface is given by the T_{rz} component of the electromagnetic stress-energy tensor, which can be integrated over the stellar surface to give the total instantaneous force on the star.

The net force during the first explosive event of the $a = 9$–12 ring simulation (as illustrated in Figure 7.9) is shown in Figure 7.10, where it is scaled to $10^{-4}\,F_{\rm max}$ and $t_{\rm rec}$ is the time at which reconnection begins. The torque spike for this event has a full width at half maximum of $\Delta t = 92\,r_\star/c$ (Table 1). Before reconnection begins, there is a small force pushing the star in the negative z-direction, which increases on the same timescale as the torque, forming a peak of similar FWHM and with minimum value $-6\times 10^{-6}\,F_{\rm max}$. This small force implies that the magnetosphere is very nearly symmetrical about the equator even during the explosive inflation phase.

The reconnection process injects Alfvén waves onto the newly closed field lines. These waves, clearly visible in our simulations, travel from the reconnection region to the stellar sur-

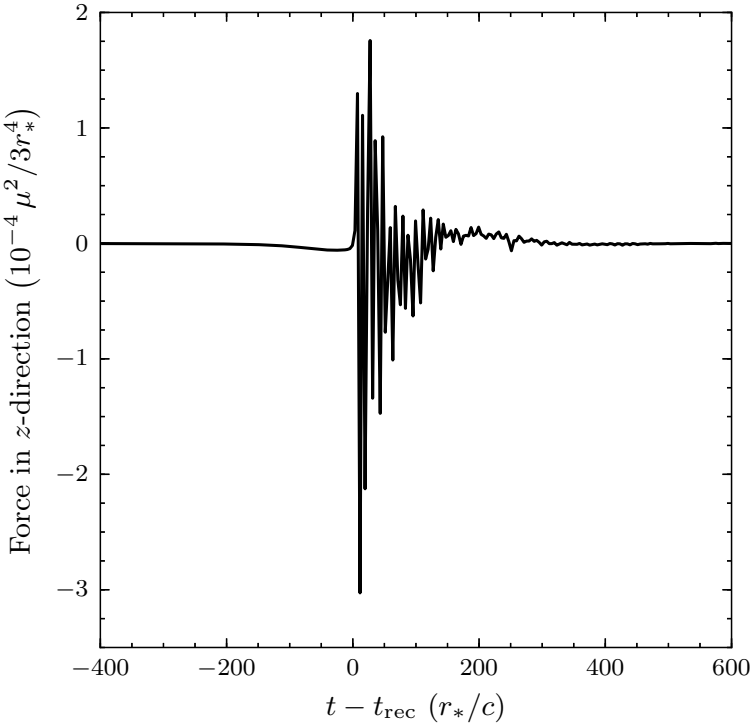

Figure 7.10: Force along the rotation axis during a flare-like event, where $t_{\rm rec}$ indicates the onset of reconnection and the force is scaled to 10^{-4} times the estimated maximum force.

face, are reflected, and continue to bounce backward and forward on the closed lines; at the same time, reconnection continues to inject new waves, and the closed flux bundle oscillates as it adjusts to a new pressure equilibrium state. The field lines are not exactly symmetrical, and so waves from the reconnection region (and subsequent reflected waves) are incident on the surface at slightly different times, giving the star small net kicks of alternating sign along the axis. These impulses are responsible for the rapidly oscillating net force after $t_{\rm rec}$ seen in Figure 7.10, which can be significantly larger than the maximum force before reconnection. The trends on longer timescales may be due to the equilibration of the closed flux. The small-scale Alfvén waves are gradually removed by numerical dissipation, and are gone by roughly $t = t_{\rm rec} + 500\, r_\star/c$.

7.6 Discussion

Many features of the non-rotating twisting problem described in Chapter 6 are also present when stellar solid-body rotation is added. There is still a critical twist amplitude beyond which the magnetosphere is unstable to catastrophic large-scale reconnection, and an approximate dichotomy between gradual, or polar-cap-like, events and those which are explosive, or ring-like. Rotation changes the evolution in detail; for example, the relaxation of a rotating quasi-equilibrium configuration to a true stationary equilibrium (when the stellar shearing is withdrawn) involves some diffusion of open flux back through the light cylinder, and the critical twist amplitude is smaller in the rotating than in the non-rotating solutions.

Most importantly, rotation adds a new spatial scale, the light cylinder radius $R_{\rm LC}$, and a new physical quantity, the spin-down torque T applied to the star by the magnetosphere. As in the non-rotating solutions, twisting causes expansion of the magnetic field in the poloidal plane; this pushes more flux through the light cylinder, increasing the spin-down torque: $T/T_0 > 1$, where T_0 is the torque applied by the rotating untwisted magnetosphere. When twisting a polar cap, initially extending over both open and closed flux, the torque first increases on the twisting timescale ω^{-1}; eventually, field lines open faster than the underlying flux can take their place, and the magnetosphere becomes unstable to large-scale reconnection of most of the opened flux, as illustrated in Figure 7.4.

Reconnection leaves a ring-like "reservoir" of strongly twisted field lines (e.g. Figure 7.3b); as these continue to be twisted, they soon become unstable to explosive opening, resulting in narrow torque spikes of duration only a few times the light-crossing time of the light cylinder (Figure 7.7). The heights of these spikes are approximately equal to a^2, where a is the ratio of the twisted flux to the flux opened by rotation alone. The additional shear required to go from one reconnection event to the next varies in our simulations from as little as 0.25 rads to as much

as 3 rads. After a few opening and reconnection events, the additional twist required to open the remaining most strongly twisted flux becomes similar to the total twist which must be applied to open the less twisted overlying field lines, and there follows a complicated mixture of gradual and fast evolution. Only the brief, explosive events occur if, instead of a polar cap, a ring containing only closed flux is twisted (Figures 7.8 and 7.9).

Significant enhancement of the stellar spin-down rate during a flare is a prediction of this model. In particular, the spin-down rate should increase rapidly before the flare is observed. It is difficult to make specific predictions, because the timescale of the torque enhancement could be anywhere from a few times $R_{\rm LC}/c$ (approximately the spin period, which is on the order of 2–12 seconds for most magnetars) to the possibly very long twisting timescale, and depends on the shearing profile and the history of the system (i.e. whether the magnetosphere has recently produced a flare). The most luminous flares require the participation of a large fraction of the magnetosphere; since magnetars spin relatively slowly, this translates into large values of the twisted-to-rotationally-opened flux ratio $a \gtrsim 10^3$, implying enormously accelerated spin-down, on the order of $T/T_0 \gtrsim 10^6$, during the field expansion phase.

The energy stored by twisting a polar cap $\theta_{\rm pc}$ of a star with magnetic moment μ and radius $r_\star$ is given by Equation (6.23),

$$W_{\rm tw} \approx \frac{\mu^2 \psi^2 \sin^6\theta_{\rm pc}}{24 r_\star^3} \sim 4 \times 10^{46} \mu_{33}^2 \psi^2 \sin^6\theta_{\rm pc}\ {\rm erg}, \tag{7.9}$$

where $\mu_{33} = \mu/10^{33}$ G cm^3 and $r_\star \approx 10^6$ cm. During an explosive event, this energy is partially released in a magnetic outflow and partially radiated in a powerful flare. Our simulations suggest that such events are accompanied by a brief increase in spin-down torque by a factor as large as a^2 which leads to an abrupt increase in spin period. Then the fractional change in period is

approximately given by

$$\frac{\Delta P}{P} \sim a^2 \frac{\Delta t}{t_0} \sim a^2 \left(\frac{c\,\Delta t}{2\pi R_{\rm LC}} \right) \dot{P}_0, \tag{7.10}$$

where $\dot{P}_0 = P/t_0$ is the spin-down rate for the untwisted magnetosphere and Δt is the duration of the torque increase by a^2. A typical $\theta_{\rm pc} \sim 0.3$–0.6, which is sufficient for observed giant flares (Equation 7.9), corresponds to $a \sim 3 \times 10^3$ (Equation 7.1). SGR 1900+14 has $\dot{P}_0 \approx 10^{-10}\,{\rm s\ s^{-1}}$, so a short-duration spike in spin-down, with Δt comparable to $R_{\rm LC}/c$, can give $\Delta P/P \sim 10^{-4}$ as observed for the August 1998 flare. This suggests that huge anti-glitches may be explained without recourse to sudden changes in the stellar interior. For the December 2004 flare in SGR 1806-20, which was two orders of magnitude more energetic, one could expect even larger ΔP, which is not observed—the measured upper limit for $\Delta P/P$ is 5×10^{-6} (Woods et al. 2007). Note that $\Delta P \propto a^2$ while $W_{\rm tw} \propto a^3$. Variations in a, μ, and twist geometry may lead to significant variations in ΔP.

The average $\dot{P}$ in the twisted-cap model is increased by a factor of a (Section 7.4.3), suggesting that the use of the standard dipole estimate for $\dot{P}$ may overestimate the magnetic field. This effect is weaker if a ring is twisted instead of a polar cap—then the explosive mode of field opening is dominant and the average $\dot{P}$ is not increased as much.

The spin-down rate is predicted to increase most dramatically just before flares or bursts. However, twisting can greatly increase the spin-down rate even when the magnetosphere is in a stable steady state. For example, Figures 7.1b and 7.3b show stable configurations in which the spin-down rate is 1.64 and 1.67 times the spin-down rate of the untwisted rotating solution respectively (Figure 7.2). Care must be taken when estimating magnetars' magnetic fields from the standard dipole spin-down relation, because doing so may significantly overestimate the strength of the dipole component if the magnetosphere is twisted.

Chapter 8

General-relativistic simulations

Gravitational effects have not been considered in the preceding chapters, in which we solved the equations of force-free electrodynamics in flat spacetime. This is unlikely to be a bad approximation for the pulsar and magnetar applications described, in which the dominant effects are produced by the rotation and shearing of the footpoints of magnetic field lines frozen into a solid stellar crust. Even so, the spacetime immediately surrounding a real neutron star will have significant Schwarzschild-like curvature, as they have $r_\star \approx 2.4\, r_\mathrm{s}$, where $r_\mathrm{s} = 2GM/c^2$ is the Schwarzschild radius, and rapidly rotating pulsars and proto-neutron stars will produce an appreciable frame-dragging effect.

Furthermore, the relativistic force-free limit is expected to apply in many cases to plasmas surrounding black holes, and force-free electrodynamics has historically been a standard tool for studying the behaviour of conducting plasmas in these systems (e.g. Blandford & Znajek 1977; MacDonald & Thorne 1982). We are not free to make a Newtonian or post-Newtonian gravity approximation because the particle inertia is explicitly taken to zero in the force-free limit, and so

a fully general-relativistic formulation is required.

In this final chapter, we present the formulation of curved spacetime electrodynamics we have implemented in PHAEDRA, and some early test problems in which we compare dynamics in neutron star and black hole magnetospheres, both with and without force-free plasma.

8.1 Hypersurface foliation

There are many ways to extend the force-free equations into a curved spacetime setting, as there is freedom both in the choice of spatial surfaces and timelike observers, and in how the four-dimensional energy conservation equation is expressed in a three-dimensional form which can be evolved forward in the chosen time coordinate. The method we have implemented uses what is often called the "3+1 split" of spacetime into a foliation of spacelike hypersurfaces, linked together by a universal time parameter (see e.g. Misner et al. 1973; Thorne et al. 1986). This is in contrast to the timelike congruence-centred approach of Landau & Lifshitz (1971) (which could be thought of as a "1+3 split"). The hypersurface foliation method is a standard technique in general relativity, and has become a popular means of posing the MHD equations in a form suitable for numerical simulation (e.g. Baumgarte & Shapiro 2003b).

Consider a spacetime metric written in the ADM form (Arnowitt et al. 1962),

$$\begin{aligned} \mathrm{d}s^2 &= -\alpha^2 \mathrm{d}t^2 + \gamma_{ij}\left(\mathrm{d}x^i + \beta^i \mathrm{d}t\right)\left(\mathrm{d}x^j + \beta^j \mathrm{d}t\right) \\ &= -(\alpha^2 - \beta^2)\mathrm{d}t^2 + 2\beta_i \mathrm{d}x^i \mathrm{d}t + \gamma_{ij}\mathrm{d}x^i \mathrm{d}x^j, \end{aligned} \tag{8.1}$$

where Roman indices indicate spatial components, $i \in 1, 2, 3$. The coordinates x^i label positions on the spatial hypersurfaces, which have three-dimensional metric γ_{ij}. The spatial components of this

three-metric coincide with those of the four-metric, $\gamma_{ij} = g_{ij}$; note however that $\gamma^{ij} \neq g^{ij}$. These hypersurfaces can be thought of as the "absolute space" at different values of the universal time parameter t. In our simulations, the computational grid represents a spatial hypersurface, and with each time step we march forward everywhere by the same amount Δt of coordinate time.

A unit normal vector n^μ can be constructed at every point on each hypersurface,

$$n_\mu = (-\alpha, 0, 0, 0). \tag{8.2}$$

This four-velocity defines our fiducial observer (FIDO), which is instantaneously at rest with rest to the absolute space. We will use this observer to measure the electromagnetic fields at individual points in spacetime. The metric field $\alpha(t, x^i)$ is the "lapse function," which relates the passage of coordinate time t to that of the proper time τ experienced by each FIDO,

$$\mathrm{d}\tau = \alpha \mathrm{d}t. \tag{8.3}$$

The lapse function likewise relates velocities measured per unit coordinate time to those measured by a FIDO using its proper time,

$$v_{\text{proper}} = \frac{\mathrm{d}x^i}{\mathrm{d}\tau} = \frac{1}{\alpha}\frac{\mathrm{d}x^i}{\mathrm{d}t} = \frac{1}{\alpha} v_{\text{coord}}. \tag{8.4}$$

The coordinate grid can move with respect to the absolute space defined by the FIDOs; this means that although FIDOs are at rest with respect to a specific hypersurface, their four-velocities can have non-zero spatial components. Raising the index in Equation (8.2) gives the FIDO four-

velocity vector

$$n^\mu = \frac{1}{\alpha}(1, -\beta^i) \tag{8.5}$$

$$= (\mathrm{d}t, -\beta^i \mathrm{d}t), \tag{8.6}$$

since $\mathrm{d}t = 1/\alpha$ along a unit vector ($\mathrm{d}\tau = 1$). Here the "shift vector" $\beta(t, x^i)$ is clearly the velocity of the coordinate grid with respect to the FIDO, per unit coordinate time.

In this chapter we will restrict attention to electrodynamic solutions in the Schwarzschild metric. This spacetime can be foliated in many different ways, each choice corresponding to a different set of coordinates; we will describe two foliations. We use normalized units in which $G = M = c = 1$; in these units the Schwarzschild horizon is at $r = r_\mathrm{s} = 2$.

The first gives the metric in standard static Schwarzschild coordinates,

$$g_{tt} = -(1 - z), \tag{8.7}$$

$$g_{rr} = \frac{1}{1 - z}, \tag{8.8}$$

$$g_{\theta\theta} = r^2, \tag{8.9}$$

$$g_{\phi\phi} = r^2 \sin^2\theta, \tag{8.10}$$

where $z \equiv r_\mathrm{s}/r$. In the ADM language, this coordinate system has

$$\gamma_{ij} = \mathrm{diag}(g_{rr},\, g_{\theta\theta},\, g_{\phi\phi}), \tag{8.11}$$

$$\alpha = \sqrt{1 - z}, \tag{8.12}$$

$$\beta^i = 0. \tag{8.13}$$

The shift vector is zero, and so the fiducial observers are static with respect to the spatial coordinates. The lapse function is unity at $r = \infty$ ($t = \tau$), and decreases with decreasing r; it becomes zero at the horizon, implying that proper time would pass infinitely slowly for an (unphysical) FIDO at rest at the horizon.

The second foliation produces the ingoing Eddington-Finkelstein (IEF) coordinate system. These coordinates are the non-rotating ($a \to 0$) limit of the Kerr metric expressed in Kerr-Schild coordinates. They can be derived by a transformation of the time coordinate of the static Schwarzschild system,

$$t \to \bar{t} = t + 2\ln(r - 2), \tag{8.14}$$

giving a new set of coordinates ($\bar{t}$, r, θ, ϕ); these are regular everywhere in $0 < r < \infty$ even though the transformation becomes singular at $r = 2$. Writing $\bar{t} \to t$, in the new system the metric is given, again with $z \equiv r_s/r$, by

$$g_{tt} = -(1 - z), \tag{8.15}$$

$$g_{tr} = z, \tag{8.16}$$

$$g_{rr} = 1 + z, \tag{8.17}$$

$$g_{\theta\theta} = r^2, \tag{8.18}$$

$$g_{\phi\phi} = r^2 \sin^2\theta. \tag{8.19}$$

Since (r, θ, ϕ) are the same as in the Schwarzschild coordinate system, r is still the areal radius

(i.e. the area of a surface of constant r is $4\pi r^2$). The lapse and shift are now

$$\alpha = \frac{1}{\sqrt{1+z}}, \tag{8.20}$$

$$\beta^i = \left(\frac{z}{1+z},\, 0,\, 0\right). \tag{8.21}$$

The lapse is again unity at infinity and decreases as the horizon is approached, but now it remains finite through the horizon. The radial component of the shift vector is positive, implying that the FIDOs are moving radially inward with respect to the coordinates, and with increasing velocity. In a sense, it is this infalling motion of the observers which removes the Schwarzschild coordinate singularity—in these coordinates there are no unphysical static observers at or inside the horizon. Even though the FIDOs are at rest with respect to a given spatial hypersurface, it is usually more helpful to think of the coordinate grid as being rigid and static and the observers as falling inward.

Both Schwarzschild and IEF coordinates have diagonal spatial metrics $\gamma_{ij} = g_{ij}$. We can therefore express components of a three-dimensional spatial vector using a local orthonormal vector basis (just as one introduces an orthonormal basis $(\hat{\mathbf{e}}_r, \hat{\mathbf{e}}_\theta, \hat{\mathbf{e}}_\phi)$ when working in spherical coordinates in flat spacetime). The orthonormal components $A_{\hat{\mathbf{i}}}$ of a spatial vector $\boldsymbol{A}$ are given by

$$A_{\hat{\mathbf{i}}} = \sqrt{\gamma_{ii}} A^i = \sqrt{A_i A^i}, \tag{8.22}$$

where repeated indices do not imply summation (i.e. $A_{\hat{\mathbf{r}}} = \sqrt{\gamma_{rr}} A^r = \sqrt{A_r A^r}$).

8.2 Equations of electrodynamics in curved spacetime

We use the formulation of curved spacetime electrodynamics developed by Komissarov (2004a, 2011). The covariant Maxwell equations are

$$\nabla_\mu {}^*F^{\mu\nu} = 0, \tag{8.23}$$

$$\nabla_\nu F^{\mu\nu} = 4\pi I^\mu, \tag{8.24}$$

where $F^{\alpha\beta}$ is the Faraday tensor and I^α the current four-vector. Our strategy is to write these equations down in component form, and replace combinations of metric terms and components of $F^{\alpha\beta}$ with new defined quantities. For example, one of the component equations of Equation (8.23) reads

$$\frac{1}{\sqrt{\gamma}}\partial_i\left(\alpha\sqrt{\gamma}\,{}^*F^{ti}\right) = 0, \tag{8.25}$$

where $\sqrt{\gamma} = \det\left(\gamma_{ij}\right)$. Making the substitution $B^i = \alpha\,{}^*F^{it}$, this equation can be written as $\nabla \cdot \boldsymbol{B} = 0$. B^i is a three-dimensional spatial vector existing only on a given spatial hypersurface. Here the divergence and curl operators imply the standard three-dimensional curvilinear coordinate expressions

$$\nabla \cdot \boldsymbol{A} = \frac{1}{\sqrt{\gamma}}\partial_i\left(\sqrt{\gamma}A^i\right), \tag{8.26}$$

$$\nabla \times \boldsymbol{A} = \frac{1}{\sqrt{\gamma}}\epsilon^{ijk}\partial_j A_k, \tag{8.27}$$

where $\epsilon^{ijk} = \epsilon_{ijk}$ is the three-dimensional Levi-Civita symbol.

We introduce the following quantities,

$$B^i = \frac{1}{2} e^{ijk} F_{jk}, \tag{8.28}$$

$$H_i = \frac{\alpha}{2} e_{ijk} F^{jk}, \tag{8.29}$$

$$D^i = \alpha F^{ti}, \tag{8.30}$$

$$E_i = F_{it}, \tag{8.31}$$

$$J^i = \alpha I^i, \tag{8.32}$$

$$\rho_{\mathrm{e}} = \alpha I^t, \tag{8.33}$$

where $e_{ijk} = \sqrt{\gamma}\epsilon_{ijk}$ is the Levi-Civita tensor in the spatial hypersurface. Maxwell's equations in a curved spacetime can now be written as

$$\nabla \cdot \boldsymbol{B} = 0, \tag{8.34}$$

$$\frac{1}{\sqrt{\gamma}} \partial_t \left(\sqrt{\gamma} \boldsymbol{B} \right) = -\nabla \times \boldsymbol{E}, \tag{8.35}$$

$$\nabla \cdot \boldsymbol{D} = 4\pi \rho_{\mathrm{e}}, \tag{8.36}$$

$$\frac{1}{\sqrt{\gamma}} \partial_t \left(\sqrt{\gamma} \boldsymbol{D} \right) = \nabla \times \boldsymbol{H} - \boldsymbol{J}. \tag{8.37}$$

The correspondence between these equations and those of electrodynamics in a macroscopic medium is made complete by restricting consideration to time-independent metrics (in which $\partial_t \sqrt{\gamma} = 0$), as we do in this chapter. These equations can easily be solved in a time-dependent metric, by evolving the vector densities $\sqrt{\gamma}\boldsymbol{B}$ and $\sqrt{\gamma}\boldsymbol{D}$ instead of the vectors $\boldsymbol{B}$ and $\boldsymbol{D}$. One of the advantages of using this formulation is that the equations of motion are expressed solely using three-dimensional divergence and curl operators, and so no Christoffel symbols need be stored or calculated.

In this formulation the effect of spacetime curvature can be thought of as transforming the vacuum into an electromagnetically active medium. The constitutive relations can be found from Equations (8.28–8.33) to be

$$\boldsymbol{E} = \alpha \boldsymbol{D} + \boldsymbol{\beta} \times \boldsymbol{B}, \tag{8.38}$$

$$\boldsymbol{H} = \alpha \boldsymbol{B} - \boldsymbol{\beta} \times \boldsymbol{D}. \tag{8.39}$$

In flat spacetime, $\alpha = 1$ and $\beta^i = 0$, and so $\boldsymbol{B} = \boldsymbol{H}$ and $\boldsymbol{D} = \boldsymbol{E}$.

Each of the spatial vectors we have introduced can be extended into a four-vector, whose time component vanishes and whose spatial components are the same as those of the spatial vector, i.e. $B^\mu = (0, B^i)$. In the cases of $\boldsymbol{B}$ and $\boldsymbol{D}$, these four-vectors can be written as

$$B^\mu = -{}^*F^{\mu\nu} n_\nu, \tag{8.40}$$

$$D^\mu = F^{\mu\nu} n_\nu, \tag{8.41}$$

and so have a particularly simple interpretation—they are the magnetic and electric fields measured by the fiducial observers at every spacetime event. Equation (8.38) is suggestive of removing an inductive term due to the motion of the observer across the coordinates, and transforming to a per-unit-coordinate-time basis (see Equation (8.4)); loosely speaking, $\boldsymbol{E}$ and $\boldsymbol{H}$ are the coordinate-based fields and $\boldsymbol{D}$ and $\boldsymbol{B}$ are the FIDO-based fields. For example,

$$S_\infty = \boldsymbol{E} \times \boldsymbol{H} / 4\pi \tag{8.42}$$

is the Poynting flux of energy-at-infinity (or gravitationally redshifted energy), while

$$S_{\rm loc} = \boldsymbol{D} \times \boldsymbol{B}/4\pi \tag{8.43}$$

is the FIDO-measured local Poynting flux.

The electromagnetic energy-momentum tensor is

$$T^{\mu}{}_{\nu} = \frac{1}{4\pi}\left[F^{\mu\gamma}F_{\nu\gamma} - \frac{1}{4}\left(F_{\alpha\beta}F^{\alpha\beta}\right)\delta^{\mu}_{\nu}\right], \tag{8.44}$$

and the density of energy-at-infinity is given by

$$e_{\infty} = -\alpha T^{t}{}_{t} \tag{8.45}$$

$$= \frac{1}{8\pi}\left(\boldsymbol{E}\cdot\boldsymbol{D} + \boldsymbol{B}\cdot\boldsymbol{H}\right). \tag{8.46}$$

8.3 Force-free current in curved spacetime

It can be shown that $\rho_{\rm e} = -I^{\nu}n_{\nu}$, and so $\rho_{\rm e}$ is the charge density measured by FIDOs. However, $\boldsymbol{J}$ is not the same as the FIDO measured current, which we will call $\boldsymbol{j}$, there being an additional term due to the motion of the observers,

$$\boldsymbol{J} = \alpha\boldsymbol{j} - \rho_{\rm e}\boldsymbol{\beta}. \tag{8.47}$$

The force-free limit is defined by the vanishing of the Lorentz force density,

$$F_{\mu\nu}I^{\mu} = 0. \tag{8.48}$$

In the 3+1 language developed above, this can be split into

$$\boldsymbol{E} \cdot \boldsymbol{J} = 0, \tag{8.49}$$

$$\rho_{\mathrm{e}} \boldsymbol{E} + \boldsymbol{J} \times \boldsymbol{B} = 0, \tag{8.50}$$

using the current $\boldsymbol{J}$ appearing in Equation (8.37), or into

$$\boldsymbol{D} \cdot \boldsymbol{j} = 0, \tag{8.51}$$

$$\rho_{\mathrm{e}} \boldsymbol{D} + \boldsymbol{j} \times \boldsymbol{B} = 0, \tag{8.52}$$

using only the electromagnetic fields and currents measured by FIDOs. Equations (8.50) and (8.52) imply that both $\boldsymbol{E}$ and $\boldsymbol{D}$ are perpendicular to $\boldsymbol{B}$,

$$\boldsymbol{E} \cdot \boldsymbol{B} = 0, \tag{8.53}$$

$$\boldsymbol{D} \cdot \boldsymbol{B} = 0, \tag{8.54}$$

and so $\partial_t \left(\sqrt{\gamma} \boldsymbol{D} \cdot \boldsymbol{B} \right) = 0$; combining this equation with Equations (8.50), (8.35), and (8.37) gives the current density

$$\boldsymbol{J} = \frac{\boldsymbol{B} \cdot (\nabla \times \boldsymbol{H}) - \boldsymbol{D} \cdot (\nabla \times \boldsymbol{E})}{4\pi B^2} \boldsymbol{B} + \left(\frac{\nabla \cdot \boldsymbol{D}}{4\pi} \right) \frac{\boldsymbol{E} \times \boldsymbol{B}}{B^2}. \tag{8.55}$$

The current density measured by FIDOs can be found by going to the FIDO's locally inertial frame and using the standard flat spacetime current expression, Equation (2.12), with the replacement

$\boldsymbol{E} \to \boldsymbol{D}$ since we require the FIDO-measured electric field,

$$\boldsymbol{j} = \frac{\boldsymbol{B} \cdot (\nabla \times \boldsymbol{B}) - \boldsymbol{D} \cdot (\nabla \times \boldsymbol{D})}{4\pi B^2} \boldsymbol{B} + \left(\frac{\nabla \cdot \boldsymbol{D}}{4\pi} \right) \frac{\boldsymbol{D} \times \boldsymbol{B}}{B^2}. \tag{8.56}$$

We evaluate Equation (8.55) in the code, since it requires the same spatial derivatives ($\nabla \times \boldsymbol{H}$, $\nabla \times \boldsymbol{E}$, and $\nabla \cdot \boldsymbol{D}$) as Maxwell's equations, Equations (8.34–8.37). This way we only require the same number of derivatives (and hence spectral transforms) as in flat spacetime.

The force-free conditions,

$$^*F_{\mu\nu} F^{\mu\nu} = 0, \tag{8.57}$$

$$F_{\mu\nu} F^{\mu\nu} > 0, \tag{8.58}$$

become in 3+1 notation

$$\boldsymbol{B} \cdot \boldsymbol{D} = 0, \tag{8.59}$$

$$B^2 - D^2 > 0. \tag{8.60}$$

The first constraint is preserved exactly by the equations of motion, and very closely maintained by free evolution in the code; it can be easily imposed periodically if one wishes, as in the non-GR code. The evolution does not necessarily preserve the second constraint, and it is violated in current sheets in which $\boldsymbol{B} = 0$, requiring electric field to be removed as described in Section 3.8. Again, this is straightforward as it simply requires the rescaling of $\boldsymbol{D}$, the evolved variable in Equation (8.37).

8.4 Boundary conditions

In these simulations we place the outer boundary at sufficiently large distance, $r \gg r_s$, that the spacetime there is approximately flat, and implement the same non-reflecting boundary treatment as when working in flat spacetime (Section 3.9.2).

We exclusively use the infalling IEF coordinates when studying black hole systems; this coordinate system is regular for all $r > 0$, and so the inner domain boundary can be placed inside the event horizon. All of the characteristic curves (see Section 3.9.2) are directed radially inward at all points $r \leq r_s$, and so each point is causally disconnected from points at smaller r. Therefore no boundary conditions can be placed on the inner surface, and the equations of motion are simply solved as normal there.

Inner boundary conditions are required in neutron star problems. In a frame comoving with the stellar surface, the normal magnetic field and tangential electric field components are continuous across the interface between the solid surface and the plasma. We therefore impose the known values of these components in the crust (where they are determined by the prescribed motion of the surface) on the fields at the innermost grid points, infinitesimally above the crust in the force-free plasma, as described in Section 3.9.1.

For simplicity, we will concentrate on the restricted case of a static star, which will form the basis of our test problem in Section 8.6; then the electric field should be zero in the surface's rest frame. This is straightforward to implement in Schwarzschild coordinates, as here the stellar surface has the same rest frame as the FIDOs of the spacetime slicing, and so the boundary conditions

imposed are

$$B^r = f(\theta), \tag{8.61}$$

$$D^\theta = 0, \tag{8.62}$$

$$D^\phi = 0, \tag{8.63}$$

where $f(\theta)$ is constant in time.

When using IEF coordinates, one must Lorentz boost from the static (Schwarzschild) frame of the stellar surface to the infalling frame of the FIDOs, in which the components of $\boldsymbol{B}$ and $\boldsymbol{D}$ are determined. The boost is in the radial direction, and so the radial component of $\boldsymbol{B}$ is unchanged. The tangential components of the electric field obey

$$\boldsymbol{D}_{\rm sch} = \Gamma\left(\boldsymbol{D}_{\rm IEF} + \boldsymbol{v} \times \boldsymbol{B}_{\rm IEF}\right) = 0, \tag{8.64}$$

where $\Gamma = 1/\sqrt{1 - v_i v^i}$ and $\boldsymbol{v}$ is the velocity of the static observer as measured by the infalling FIDO, on a per unit FIDO proper time basis,

$$v^i = \frac{\beta^i}{\alpha}. \tag{8.65}$$

Equation (8.64) gives

$$D^i = -\frac{\epsilon^{ijk}}{\sqrt{\gamma}} v_j B_k, \tag{8.66}$$

where all fields are measured in IEF coordinate; therefore using $v^i = (v^r, 0, 0)$ we find the boundary

conditions to be imposed are

$$B^r = f(\theta), \tag{8.67}$$

$$D^\theta = \frac{v_r}{\sqrt{\gamma}} B_\phi, \tag{8.68}$$

$$D^\phi = -\frac{v_r}{\sqrt{\gamma}} B_\theta, \tag{8.69}$$

where $v_r = \gamma_{rr}\beta^r/\alpha = z\sqrt{1+z}$.

8.5 Static fields in the Schwarzschild spacetime

The Schwarzschild metric is spherically symmetric. In this spacetime, the monopole magnetic field,

$$\boldsymbol{B} = \frac{\text{const}}{\sqrt{\gamma}} \left(\frac{1}{r^2}, 0, 0 \right), \tag{8.70}$$

$$\boldsymbol{D} = 0, \tag{8.71}$$

is a trivial static solution to Maxwell's equations, because $H^i = (\alpha(r)B^r, 0, 0)$ and hence $\nabla \times \boldsymbol{H} = \boldsymbol{J} = \partial_t \boldsymbol{D} = \partial_t \boldsymbol{B} = 0$. The split monopole, similar to the above but where the sign of B^r is reversed below the equator (i.e. for $\theta > \pi/2$), is also a static solution in the absence of dissipation. When resistivity is present, reconnection occurs in the equatorial current sheet. In the absence of rotation or some other energy source the field lines will all eventually reconnect, resulting in an entirely closed magnetosphere.

The unique poloidal static field outside a star having a dipole-like distribution of normal magnetic flux on its surface, $B^r = f(r)\cos\theta$, was found by Wasserman & Shapiro (1983); expressed in

orthonormal components and measured by the static Schwarzschild observer, it is

$$B^{\rm Sch}_{\hat{\mathbf{r}}} = -\frac{6\mu\cos\theta}{r^3\,z^2}\left[\frac{1}{z}\ln(1-z)+1+\frac{z}{2}\right],$$
$$B^{\rm Sch}_{\hat{\theta}} = \frac{6\mu\sin\theta}{r^3\,z^2}\left[\frac{1}{z}\sqrt{1-z}\ln(1-z)+\frac{1-z/2}{\sqrt{1-z}}\right], \tag{8.72}$$
$$B^{\rm Sch}_{\hat{\phi}} = 0, \tag{8.73}$$

where $z \equiv r_{\rm s}/r$ and μ is the magnetic dipole moment. In the IEF coordinate system one must Lorentz boost to the frame of the infalling observers; they measure the static configuration to have magnetic field

$$B^{\rm IEF}_{\hat{\mathbf{r}}} = B^{\rm Sch}_{\hat{\mathbf{r}}}, \tag{8.74}$$
$$B^{\rm IEF}_{\hat{\theta}} = \Gamma\, B^{\rm Sch}_{\hat{\theta}}, \tag{8.75}$$
$$B^{\rm IEF}_{\hat{\phi}} = 0, \tag{8.76}$$

and electric field

$$D^{\rm IEF}_{\hat{\mathbf{r}}} = D^{\rm IEF}_{\hat{\theta}} = 0, \tag{8.77}$$
$$D^{\rm IEF}_{\hat{\phi}} = -\Gamma v B^{\rm Sch}_{\hat{\theta}} = -v B^{\rm IEF}_{\hat{\theta}}, \tag{8.78}$$

where $v = v_{\hat{\mathbf{r}}} = \sqrt{\beta_r\beta^r}/\alpha$ and $\Gamma = 1/\sqrt{1-v^2}$.

As a first test problem, we have confirmed that these two field configurations are static when evolved in time by the code using each coordinate system's appropriate boundary conditions. Because this requires a non-trivial balancing of forces it gives us confidence that the two coordinate systems are implemented correctly.

8.6 Dynamic magnetospheres: a test problem

We present here a simple test problem to illustrate dynamic electromagnetic fields outside neutron stars and black holes, in vacuum and in force-free plasma. We will describe four simulations, one for each combination of astrophysical object and current type ($\boldsymbol{J} = 0$ or $\boldsymbol{J} = \boldsymbol{J}_{\rm FFE}$). For consistency we run all four models using the infalling IEF coordinate system, on a grid of resolution $N_r \times N_\theta = 256 \times 155$ and with $r_{\rm out} = 60\,GM/c^2$.

We devise a simple initial field configuration that is not static in Schwarzschild spacetime. This dipole-like field can be constructed by noting that

$$B^i = \frac{1}{2} e^{ijk} F_{jk} \tag{8.79}$$

$$= \frac{1}{2\sqrt{\gamma}} \epsilon^{ijk} F_{jk} \tag{8.80}$$

$$= \frac{1}{2\sqrt{\gamma}} \epsilon^{ijk} \left(\partial_j A_k - \partial_k A_j \right), \tag{8.81}$$

where A_i is the vector potential. If we use the same vector potential form $\boldsymbol{A}(r, \theta)$ that produces the standard flat spacetime dipole magnetic field $\boldsymbol{B}_{\rm flat}$, the resulting "corresponding" dipole-like magnetic field $\boldsymbol{B}_{\rm gr}$ in a different metric will be given by

$$\boldsymbol{B}_{\rm gr} = \sqrt{\frac{\gamma_{\rm flat}}{\gamma_{\rm gr}}} \boldsymbol{B}_{\rm flat}. \tag{8.82}$$

The $\sqrt{\gamma_{\rm flat}/\gamma_{\rm gr}}$ prefactor is just what is needed to make the new field divergence free, $\nabla \cdot \boldsymbol{B}_{\rm gr} = 0$ (see Equation 8.26). The original dipole field $\boldsymbol{B}_{\rm flat}$ was a static solution in flat spacetime, however the new configuration will in general be non-static, $\nabla \times \boldsymbol{H} \neq 0 \rightarrow \partial_t \boldsymbol{D} \neq 0$. For this initial $\boldsymbol{B}_{\rm gr}$ the resulting evolution will preserve $B_{\hat{\phi}} = D_{\hat{\mathbf{r}}} = D_{\hat{\theta}} = 0$.

The evolution of the magnetic and electric fields for the case of a neutron star in vacuum is

shown in Figure 8.1a; field lines are drawn as contours of constant vector potential A_ϕ,

$$A_\phi(r,\theta) = \int_0^\theta \sqrt{\gamma} B^r(r,\theta') \mathrm{d}\theta'. \tag{8.83}$$

We initialize the solution with $D_{\hat{\phi}} = -vB_{\hat{\theta}}$ everywhere, so that the boundary conditions, Equations (8.67–8.69), do not create a discontinuity in D^ϕ at the stellar surface. As the configuration begins to evolve, the field lines fall toward the star, and a region of strong electric field forms near the stellar surface. The magnetic field then bounces back outward, emitting a transverse electromagnetic wave of loops of field lines unconnected to the star. The splitting of the magnetic field lines, into sections connected to the star and disconnected loops, takes place in regions with large electric field, $E^2 > B^2$. This is allowed in vacuum, because there is no matter field which must have a subluminal velocity (see Section 2.2). The rebounding magnetic field approaches an equilibrium configuration, which is the Wasserman-Shapiro field; this process is mostly complete by $t = 70\,GM/c^3$. In this steady state the magnetic field is more concentrally concentrated than in the initial configuration, as can be seen by comparing the first and last panels of Figure 8.1a.

The behaviour is similar when the magnetosphere is filled with force-free plasma (Figure 8.1b). However, the imposition of the $B^2 - D^2 > 0$ constraint prevents the splitting of field lines into sections connected to and disconnected from the stellar surface, and hence the emission of the transverse fast wave. Small current sheets form and prevent magnetic reconnection on dynamic timescales, and so the field lines retain their closed topology. The magnetosphere approaches the same Wasserman-Shapiro equilibrium configuration, though on a longer timescale than in the vacuum simulation.

In a black hole magnetosphere there is no conducting surface to anchor the magnetic field. In vacuum, the magnetic field is partly swallowed by the black hole and partly emitted to infinity

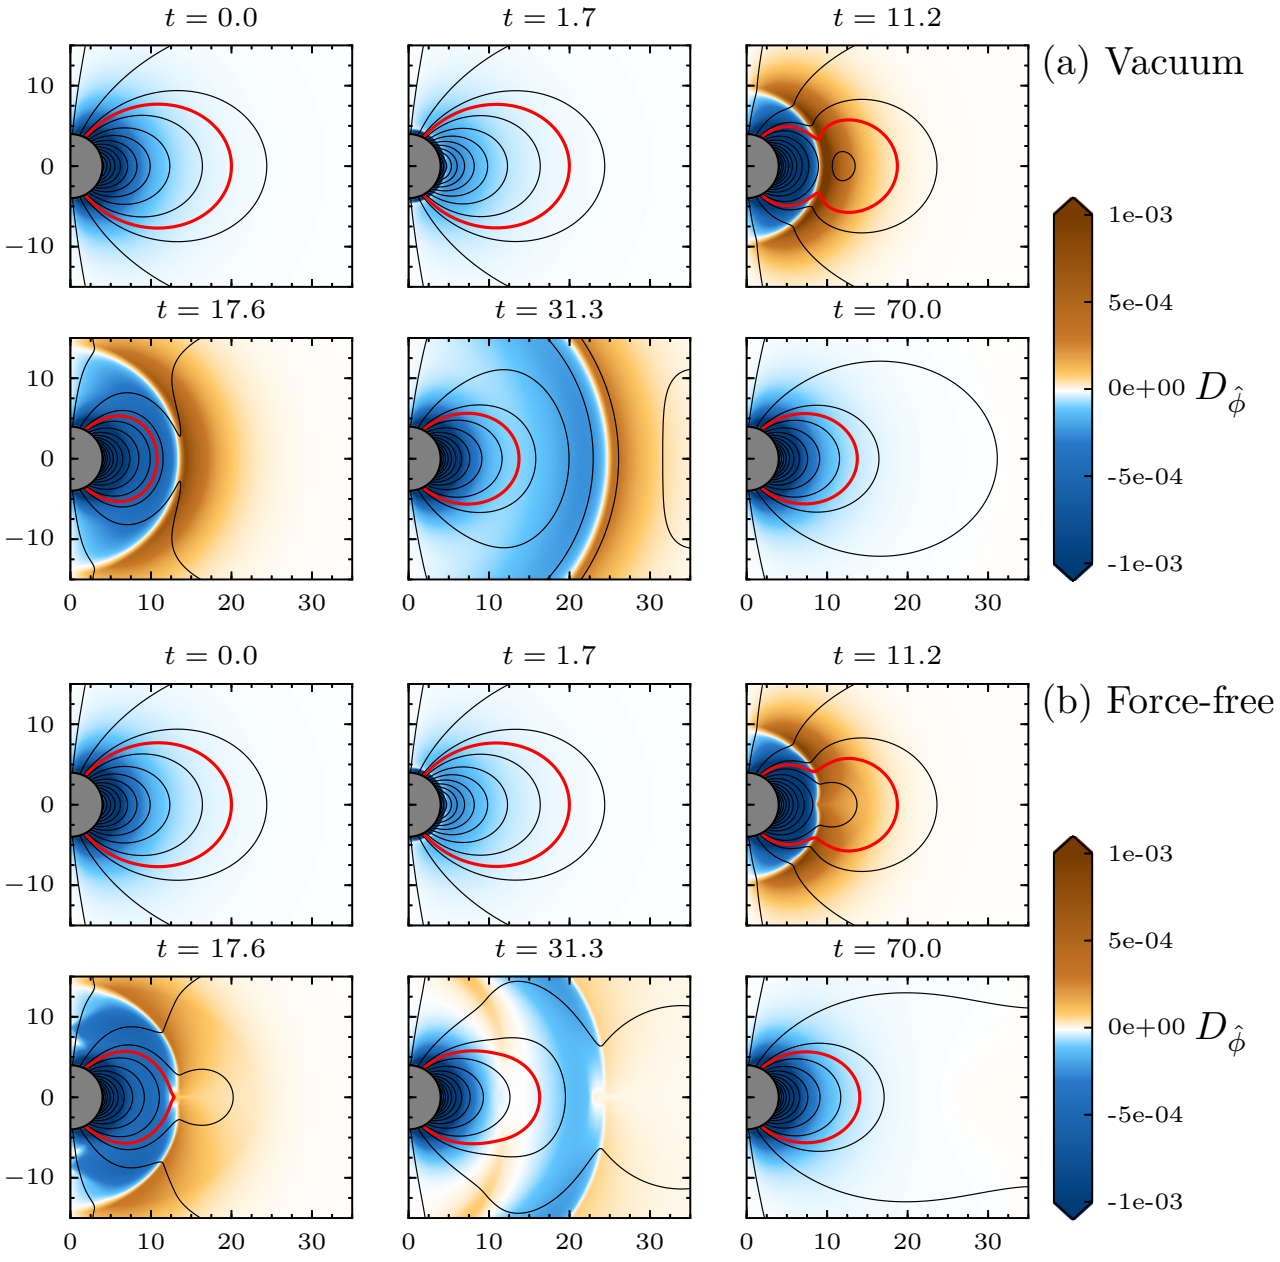

Figure 8.1: Neutron star evolution: relaxation to Wasserman-Shapiro field. Distance is measured in units of GM/c^2, time in units of GM/c^3, and $D_{\hat{\phi}}$ in units of $\mu c^6/G^3M^3$. The stellar radius is $r_\star = 2r_\mathrm{s} = 4GM/c^2$. Poloidal field lines are shown from $A_\phi = 10^{-3}\,\mu c^4/G^2M^2$ onward, equally spaced with $\Delta A_\phi = 0.02\,\mu c^4/G^2M^2$. An additional field line, with $A_\phi = 0.05$, is shown in red.

in the form of a transverse electromagnetic wave, similar to the wave seen in the vacuum neutron star simulation (Figure 8.2a). We verify that the wavelength of the outgoing wave satisfies $\lambda \approx 25\,GM/c^2$, the quasi-normal mode for dipole radiation scattering off a Schwarzschild black hole (Ferrari & Mashhoon 1984); this behavior has also been seen in simulations of the vacuum magnetosphere of a star collapsing to a black hole (Baumgarte & Shapiro 2003a). In this simulation the magnetic field disappears on a dynamical timescale set by the mass of the black hole.

Our final simulation contains a black hole surrounded by force-free plasma (Figure 8.2b). As in the force-free neutron star model, the plasma places a topological constraint on the field lines, preventing them from splitting except via resistive processes. A current sheet forms at the equator which prevents closed field lines from being quickly sucked into the black hole; instead they must diffuse through the current sheet, and the magnetosphere is destroyed on a resistive, rather than dynamical, timescale. This effect has been seen in simulations of rotating black holes (Lyutikov & McKinney 2011). We find that a current sheet is able to forestall the loss of magnetic flux even when the back hole is not rotating. There is significantly more magnetic energy near the black hole, $r \lesssim 10\,GM/c^2$, at $t = 67\,GM/c^3$ in the force-free simulation than in the vacuum case. At late times the force-free magnetosphere evolves towards a split monopole-like configuration.

The force-free current density, as calculated from Equation (8.55), is in fact zero everywhere throughout the simulations in both the neutron star and black hole force-free models. The difference between vacuum and force-free evolution is due entirely to the imposition of the second force-free constraint $B^2 - D^2 > 0$, which involves the removal of electric field $\boldsymbol{D}$ and therefore acts as an effective current. The absence of currents means that there are no Alfvén waves outside current sheets—all of the dynamics are mediated by the fast mode.

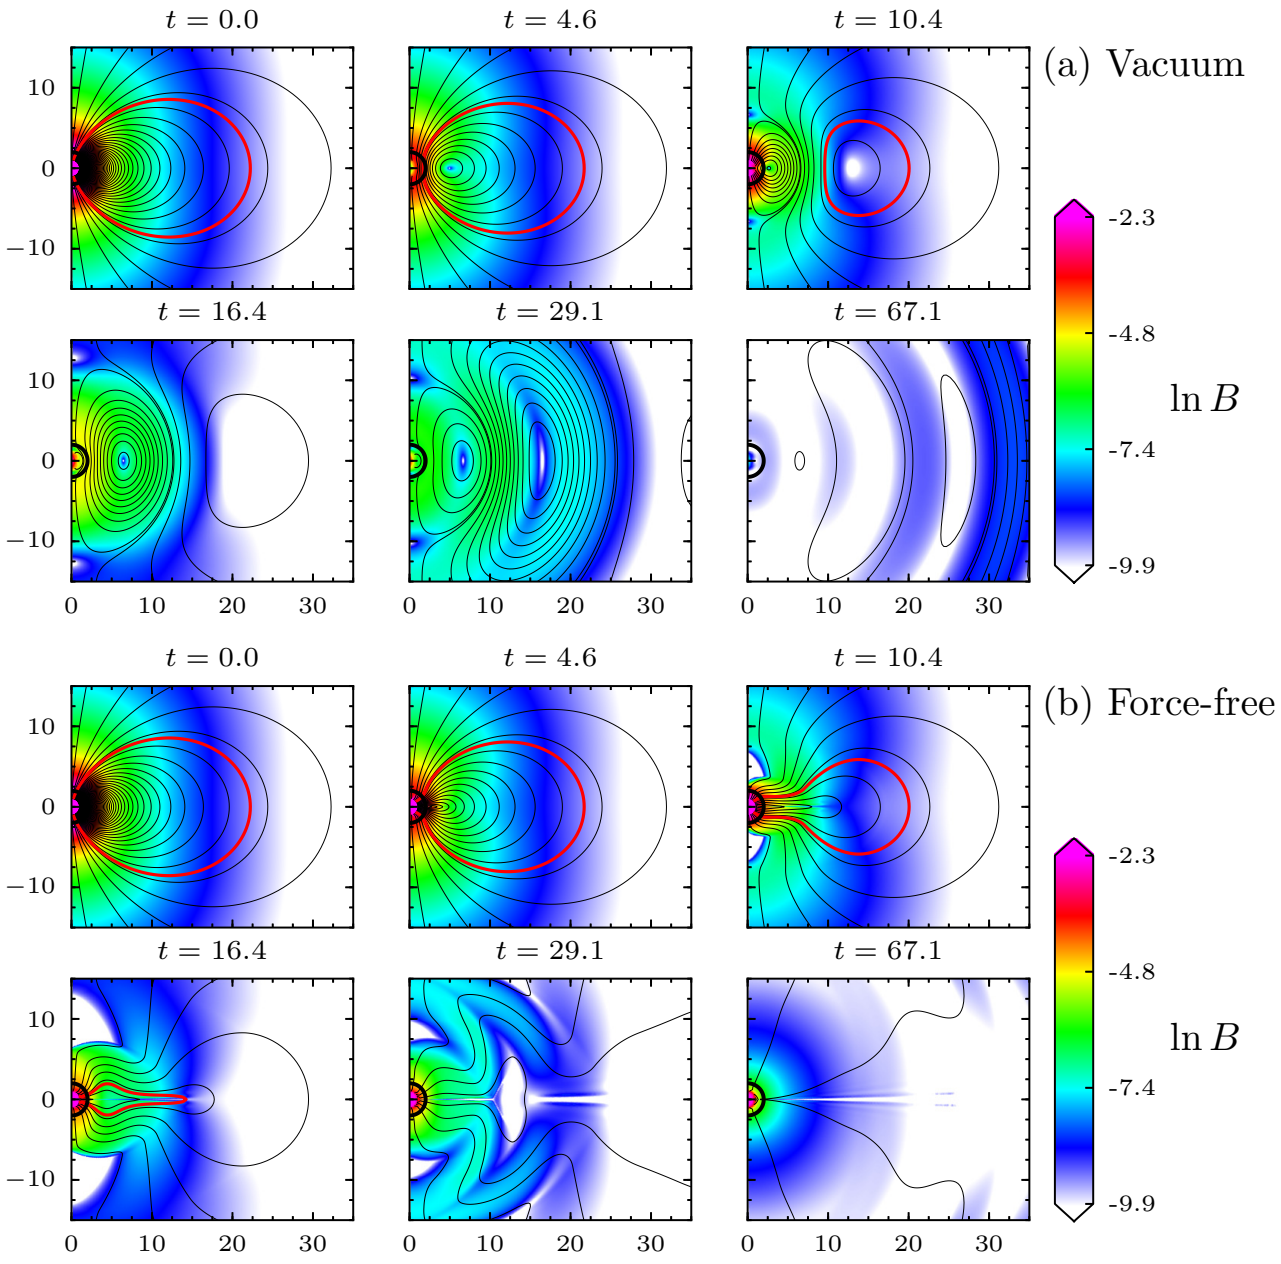

Figure 8.2: Black hole evolution. Distance is measured in units of GM/c^2, time in units of GM/c^3, and the magnitude of the magnetic field, B, in units of $\mu c^6/G^3M^3$. A thick black line is drawn at the horizon $r = 2GM/c^2$. Poloidal field lines are shown from $|A_\phi| = 10^{-3}$ to $0.991\,\mu c^4/G^2M^2$ inclusive, equally spaced with $\Delta A_\phi = 0.01\,\mu c^4/G^2M^2$. An additional field line, with $A_\phi = 0.045$, is shown in red.

Bibliography

Abarbanel, S. S., Don, W. S., Gottlieb, D., Rudy, D. H., & Townsend, J. C. 1991, Journal of Fluid Mechanics, 225, 557

Abdo, A. A., Ackermann, M., Ajello, M., Atwood, W. B., Axelsson, M., Baldini, L., Ballet, J., Barbiellini, G., Baring, M. G., Bastieri, D., & et al. 2010, ApJS, 187, 460

Aly, J. J. 1984, ApJ, 283, 349

—. 1991, ApJL, 375, L61

Arnowitt, R. L., Deser, S., & Misner, C. W. Canonical Analysis of General Relativity (Pergamon Press, Oxford), 127

Arons, J. 2007, eprint arXiv:0708.1050

Arons, J. 2011, in High-Energy Emission from Pulsars and their Systems, ed. D. F. Torres & N. Rea, 165

Arons, J. & Scharlemann, E. T. 1979, ApJ, 231, 854

Asano, E., Uchida, T., & Matsumoto, R. 2005, Publications of the Astronomical Society of Japan, 57, 409

Bai, X.-N. & Spitkovsky, A. 2010, ApJ, 715, 1282

Barnes, C. W. & Sturrock, P. A. 1972, ApJ, 174, 659

Baumgarte, T. W. & Shapiro, S. L. 2003a, ApJ, 585, 930

—. 2003b, ApJ, 585, 921

Beloborodov, A. M. 2008, ApJL, 683, L41

—. 2009a, ApJ, 703, 1044

—. 2009b, ApJ, 703, 1044

Beloborodov, A. M. 2011, in High-Energy Emission from Pulsars and their Systems, ed. D. F. Torres & N. Rea, 299

—. 2012, ArXiv: 1201.0664

Beloborodov, A. M. & Thompson, C. 2007, ApJ, 657, 967

Bertsch, D. L., Brazier, K. T. S., Fichtel, C. E., Hartman, R. C., Hunter, S. D., Kanbach, G., Kniffen, D. A., Kwok, P. W., Lin, Y. C., & Mattox, J. R. 1992, Nature, 357, 306

Biskamp, D. & Welter, H. 1989, Solar Phys., 120, 49

Blandford, R. D. 1976, MNRAS, 176, 465

—. 2002, Lighthouses of the Universe: The Most Luminous Celestial Objects and Their Use for Cosmology: Proceedings of the MPA/ESO/MPE/USM Joint Astronomy Conference Held in Garching, 381

Blandford, R. D. & Znajek, R. L. 1977, MNRAS, 179, 433

Bonazzola, S., Gourgoulhon, E., & Marck, J.-A. 1999, Journal of Computational and Applied Mathematics, 109, 433

Boyd, J. 1992, Journal of Computational Physics, 98, 181

—. 1996, Proceedings of the Third International Conference on Spectral and High-Order Methods, 267

—. 1998, Journal of Computational Physics, 143, 283

Boyd, J. P. 2001, Chebyshev and Fourier Spectral Methods (Mineola, N.Y.; Dover.)

Brandenburg, A. 2003, Advances in Nonlinear Dynamics (Taylor and Francis Group, London), 269

Bürgel, A. & Sonar, T. 2002, International Journal for Numerical Methods in Fluids, 40, 263

Burrows, A. 1987, ApJL, 318, L57

Burrows, A. & Lattimer, J. M. 1986, ApJ, 307, 178

Cai, W., Gottlieb, D., & Shu, C. 1992, SIAM Journal on Numerical Analysis, 29, 905

Camilo, F., Ransom, S. M., Halpern, J. P., Reynolds, J., Helfand, D. J., Zimmerman, N., & Sarkissian, J. 2006, Nature, 442, 892

Camilo, F. et al. 2007, ApJ, 663, 497

Carpenter, M. H., Gottlieb, D., & Shu, C. W. 2003, Journal of Scientific Computing, 18, 111

Chan, C. K., Psaltis, D., & Özel, F. 2005, ApJ, 628, 353

—. 2009, ApJ, 700, 741

Cheng, B., Epstein, R. I., Guyer, R. A., & Young, A. C. 1996, Nature, 382, 518

Cheng, K. S., Ho, C., & Ruderman, M. 1986, ApJ, 300, 500

Cho, J. 2005, ApJ, 621, 324

Contopoulos, I. 2005, A&A, 442, 579

Contopoulos, I., Kazanas, D., & Fendt, C. 1999, ApJ, 511, 351

Cowley, S. C. & Artun, M. 1997, Phys. Rep., 283, 185

Davies, S. R., Coe, M. J., Payne, B. J., & Hanson, C. G. 1989, MNRAS, 237, 973

Deutsch, A. J. 1955, Annales d'Astrophysique, 18, 1

Dimmelmeier, H., Novak, J., Font, J. A., Ibáñez, J. M., & Müller, E. 2005, Phys. Rev. D, 71, 64023

Don, W. S. 1994, Journal of Computational Physics, 110, 103

Don, W. S. & Gottlieb, D. 1998, SIAM Journal on Numerical Analysis, 35, 2370

Don, W. S. & Solomonoff, A. 1995, SIAM Journal on Scientific Computing, 16, 1253

Duncan, R. C. & Thompson, C. 1992, ApJL, 392, L9

Enoto, T., Nakazawa, K., Makishima, K., Rea, N., Hurley, K., & Shibata, S. 2010, ApJL, 722, L162

Fahlman, G. G. & Gregory, P. C. 1981, Nature, 293, 202

Fernández, R. & Thompson, C. 2007, ApJ, 660, 615

Feroci, M., Frontera, F., Costa, E., Amati, L., Tavani, M., Rapisarda, M., & Orlandini, M. 1999, ApJL, 515, L9

Ferrari, V. & Mashhoon, B. 1984, Phys. Rev. D, 30, 295

Fichtel, C. E., Hartman, R. C., Kniffen, D. A., Thompson, D. J., Ogelman, H., Ozel, M. E., Tumer, T., & Bignami, G. F. 1975, ApJ, 198, 163

Finn, J. M., Guzdar, P. N., & Chen, J. 1992, ApJ, 393, 800

Fornberg, B. 1996, A Practical Guide to Pseuospectral Methods (Cambridge, England; Cambridge University Press.)

Frigo, M. & Johnson, S. G. 2005, Proceedings of the IEEE, 93, 216

Gaensler, B. M., Kouveliotou, C., Gelfand, J. D., Taylor, G. B., Eichler, D., Wijers, R. A. M. J., Granot, J., Ramirez-Ruiz, E., Lyubarsky, Y. E., Hunstead, R. W., Campbell-Wilson, D., van der Horst, A. J., McLaughlin, M. A., Fender, R. P., Garrett, M. A., Newton-McGee, K. J., Palmer, D. M., Gehrels, N., & Woods, P. M. 2005, Nature, 434, 1104

Gammie, C. F., McKinney, J. C., & Tóth, G. 2003, ApJ, 589, 444

Gavriil, F. P. & Kaspi, V. M. 2004, ApJL, 609, L67

Gavriil, F. P., Kaspi, V. M., & Woods, P. M. 2002, Nature, 419, 142

Gelb, A. & Tanner, J. 2006, Applied and Computational Harmonic Analysis, 20, 3

Godon, P. 1996, MNRAS, 282, 1107

Godon, P. & Shaviv, G. 1993, Computer Methods in Applied Mechanics and Engineering, 110, 171

Gold, T. 1968, Nature, 218, 731

—. 1969, Nature, 221, 25

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

Goldreich, P. & Reisenegger, A. 1992, ApJ, 395, 250

Goodwin, S. P., Mestel, J., Mestel, L., & Wright, G. A. E. 2004, MNRAS, 349, 213

Gotthelf, E. V. & Halpern, J. P. 2007, Ap&SS, 308, 79

Gottlieb, D. & Gottlieb, S. 2005, Comptes Rendus Mécanique, 333, 3

Gottlieb, D. & Hesthaven, J. S. 2001, Journal of Computational and Applied Mathematics, 128, 83

Gottlieb, D., Shu, C., Solomonoff, A., & Vandeven, H. 1992, Journal of Computational and Applied Mathematics, 43, 81

Gottlieb, D. & Tadmor, E. 1985, Progress and Supercomputing in Computational Fluid Dynamics; Proceedings of U.S.-Israel Workshop, 357

Gourgoulhon, E. 1991, A&A, 252, 651

Grandclément, P. & Novak, J. 2009, Living Reviews in Relativity, 12, 1

Gruzinov, A. 1999, eprint arXiv:astro-ph/9902288

—. 2006, eprint arXiv: astro-ph/0604364

Gruzinov, A. 2007, ApJL, 667, L69

Gruzinov, A. 2008, eprint arXiv:0802.1716

—. 2011, eprint arXiv:1101.3100

Halpern, J. P. & Holt, S. S. 1992, Nature, 357, 222

Harding, A. K. 2009, in Astrophysics and Space Science Library, Vol. 357, Astrophysics and Space Science Library, ed. W. Becker, 521

Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Nature, 217, 709

Hou, T. Y. & Li, R. 2007, Journal of Computational Physics, 226, 379

Hurley, K., Boggs, S. E., Smith, D. M., Duncan, R. C., Lin, R., Zoglauer, A., Krucker, S., Hurford, G., Hudson, H., Wigger, C., Hajdas, W., Thompson, C., Mitrofanov, I., Sanin, A., Boynton, W., Fellows, C., von Kienlin, A., Lichti, G., Rau, A., & Cline, T. 2005, Nature, 434, 1098

Hurley, K., Cline, T., Mazets, E., Barthelmy, S., Butterworth, P., Marshall, F., Palmer, D., Aptekar, R., Golenetskii, S., Il'Inskii, V., Frederiks, D., McTiernan, J., Gold, R., & Trombka, J. 1999, Nature, 397, 41

Ibrahim, A. I., Strohmayer, T. E., Woods, P. M., Kouveliotou, C., Thompson, C., Duncan, R. C., Dieters, S., Swank, J. H., van Paradijs, J., & Finger, M. 2001, ApJ, 558, 237

Ibrahim, A. I. et al. 2004, ApJL, 609, L21

If, F., Berg, P., Christiansen, P. L., & Skovgaard, O. 1987, Journal of Computational Physics, 72, 501

Kalapotharakos, C. & Contopoulos, I. 2009, Astronomy and Astrophysics, 496, 495

Kalapotharakos, C., Kazanas, D., Harding, A., & Contopoulos, I. 2011, eprint arXiv:1108.2138

Kaspi, V. M., Gavriil, F. P., Woods, P. M., Jensen, J. B., Roberts, M. S. E., & Chakrabarty, D. 2003, ApJL, 588, L93

Kidder, L. E., Scheel, M. A., Teukolsky, S. A., Carlson, E. D., & Cook, G. B. 2000, Physical Review D (Particles, 62, 84032

Kniffen, D. A., Hartman, R. C., Thompson, D. J., Bignami, G. F., & Fichtel, C. E. 1974, Nature, 251, 397

Komissarov, S. S. 2002, MNRAS, 336, 759

—. 2004a, MNRAS, 350, 427

—. 2004b, MNRAS, 350, 1431

—. 2006, MNRAS, 367, 19

Komissarov, S. S. 2011, MNRAS, 418, L94

Kosloff, D. & Tal-Ezer, H. 1993, Journal of Computational Physics, 104, 457

Koyama, K., Hoshi, R., & Nagase, F. 1987, Pub. Astron. Soc. Japan, 39, 801

Kuiper, L., den Hartog, P. R., & Hermsen, W. 2008, eprint arXiv:0810.4801

Landau, L. D. & Lifshitz, E. M. 1971, The Classical Theory of Fields (Pergamon Press, Oxford)

Li, J., Spitkovsky, A., & Tchekhovskoy, A. 2012, ApJ, 746, 60

Liu, Q. 1997, Microwave and Optical Technology Letters, 15, 158

Low, B. C. 1977, ApJ, 212, 234

—. 1978, ApJ, 224, 668

—. 1986, ApJ, 307, 205

Lyubarskii, Y. E. 1996, A&A, 311, 172

Lyubarsky, Y., Eichler, D., & Thompson, C. 2002, ApJL, 580, L69

Lyutikov, M. 2003, MNRAS, 346, 540

Lyutikov, M. 2006, MNRAS, 367, 1594

Lyutikov, M. & McKinney, J. C. 2011, Phys. Rev. D, 84, 084019

Ma, H. P. 1998, SIAM Journal on Numerical Analysis, 35, 893

Ma, H. P. & Li, H. Y. 2006, Journal of Shanghai University (English Edition), 10, 9

MacDonald, D. & Thorne, K. S. 1982, MNRAS, 198, 345

MacDonald, D. & Thorne, K. S. 1982, MNRAS, 198, 345

Majda, A., McDonough, J., & Osher, S. 1978, Mathematics of Computation, 32, 1041

Mazets, E. P., Cline, T. L., Aptekar', R. L., Butterworth, P. S., Frederiks, D. D., Golenetskii, S. V., Il'Inskii, V. N., & Pal'Shin, V. D. 1999, Astronomy Letters, 25, 635

Mazets, E. P., Golentskii, S. V., Ilinskii, V. N., Aptekar, R. L., & Guryan, I. A. 1979, Nature, 282, 587

McKinney, J. C. 2006a, MNRAS, 367, 1797

—. 2006b, MNRAS, 368, L30

Mead, J. & Renaut, R. 2003, SIAM Journal on Scientific Computing, 24, 143

Mereghetti, S. 2008, The Astronomy and Astrophysics Review, 15, 225

Mereghetti, S. 2008, Astron. Astrophys. Rev., 15, 225

Mereghetti, S., Götz, D., von Kienlin, A., Rau, A., Lichti, G., Weidenspointner, G., & Jean, P. 2005, ApJL, 624, L105

Mereghetti, S. & Stella, L. 1995, ApJL, 442, L17

Merilees, P. 1974, Atmosphere, 12, 77

Mészáros, P. & Rees, M. J. 1997, ApJL, 482, L29

Michel, F. C. 1973, ApJL, 180, L133+

Mikic, Z., Barnes, D. C., & Schnack, D. D. 1988, ApJ, 328, 830

Mikic, Z. & Linker, J. A. 1994, ApJ, 430, 898

Milsom, F. & Wright, G. A. E. 1976, MNRAS, 174, 307

Misner, C. W., Thorne, K. S., & Wheeler, J. A. 1973, Gravitation (W.H. Freeman and Co., San Francisco)

Morini, M. 1983, MNRAS, 202, 495

Muslimov, A. G. & Harding, A. K. 2003, ApJ, 588, 430

Ng, C. S. & Ragunathan, S. 2011, eprint arXiv:1106.0521

Orszag, S. 1970, Journal of Atmospheric Sciences, 27, 890

—. 1974, Monthly Weather Review, 102, 56

Pacini, F. 1968, Nature, 219, 145

Palenzuela, C., Garrett, T., Lehner, L., & Liebling, S. L. 2010, Phys. Rev. D, 82, 044045

Palmer, D. M. 2002, Mem. Societa Astronomica Italiana, 73, 578

Palmer, D. M. et al. 2005, Nature, 434, 1107

Pavan, L., Turolla, R., Zane, S., & Nobili, L. 2009, MNRAS, 395, 753

Pelz, R. 1991, Journal of Computational Physics, 92, 296

Punsly, B. 2003, ApJ, 583, 842

Raadu, M. A. 1971, Astrophys. Space Sci., 14, 464

Roumeliotis, G., Sturrock, P. A., & Antiochos, S. K. 1994, ApJ, 423, 847

Ruderman, M., Zhu, T., & Chen, K. 1998, ApJ, 492, 267

Ruderman, M. A. & Sutherland, P. G. 1975, ApJ, 196, 51

Rudin, L. I., Osher, S., & Fatemi, E. 1992, Physica D: Nonlinear Phenomena, 60, 259

Sarra, S. 2003, Journal of Computational Physics, 186, 630

—. 2006, Numerical Algorithms, 41, 17

Scharlemann, E. T. & Wagoner, R. V. 1973, ApJ, 182, 951

Shu, C. W. & Wong, P. 1995, Journal of Scientific Computing, 10, 357

Spitkovsky, A. 2006, ApJ, 648, L51

Steinolfson, R. S. 1991, ApJ, 382, 677

Sturrock, P. A. 1971, ApJ, 164, 529

—. 1991, ApJ, 380, 655

Sun, Y. S., Zhou, Y. C., Li, S. G., & Wei, G. W. 2006, Journal of Computational Physics, 214, 466

Tadmor, E. 1989, SIAM Journal on Numerical Analysis, 26, 30

—. 1990, Computer Methods in Applied Mechanics and Engineering, 80, 197

—. 1993, Numerical Methods for Fluid Dynamics IV", Proceedings of the 1992 Conference on Numerical Methods for Fluid Dynamics, 69

Tadmor, E. & Tanner, J. 2005, IMA Journal of Numerical Analysis, 635

Tanner, J. 2006, Mathematics of Computation, 75, 767

Thompson, C. & Blaes, O. 1998, Physical Review D (Particles, 57, 3219

Thompson, C. & Duncan, R. C. 1993, ApJ, 408, 194

—. 1995a, MNRAS, 275, 255

—. 1995b, MNRAS, 275, 255

Thompson, C., Duncan, R. C., Woods, P. M., Kouveliotou, C., Finger, M. H., & van Paradijs, J. 2000, ApJ, 543, 340

Thompson, C., Lyutikov, M., & Kulkarni, S. R. 2002, ApJ, 574, 332

Thompson, D. J., Fichtel, C. E., Kniffen, D. A., & Ogelman, H. B. 1975, ApJL, 200, L79

Thorne, K. S., Price, R. H., & MacDonald, D. A. 1986, Black Holes: The Membrane Paradigm (Yale University Press, New Haven)

Timokhin, A. N. 2006, MNRAS, 368, 1055

Tomczyk, S., McIntosh, S. W., Keil, S. L., Judge, P. G., Schad, T., Seeley, D. H., & Edmondson, J. 2007, Science, 317, 1192

Tóth, G. 2000, Journal of Computational Physics, 161, 605

Uchida, T. 1997, Phys. Rev. E, 56, 2181

Uzdensky, D. A. 2002, ApJ, 574, 1011

Vandeven, H. 1991, Journal of Scientific Computing, 6, 159

Viganò, D., Parkins, N., Zane, S., Turolla, R., Pons, J. A., & Miralles, J. A. 2012, Journal of Physics: Conference Series, 342

Viganò, D., Pons, J. A., & Miralles, J. A. 2011, A&A, 533, A125

Wasserman, I. & Shapiro, S. L. 1983, ApJ, 265, 1036

Wolfson, R. 1989, ApJ, 344, 471

—. 1995, ApJ, 443, 810

Wolfson, R. & Low, B. C. 1992, ApJ, 391, 353

Woods, P. M., Kouveliotou, C., Finger, M. H., Göğüş, E., Wilson, C. A., Patel, S. K., Hurley, K., & Swank, J. H. 2007, ApJ, 654, 470

Woods, P. M., Kouveliotou, C., Gavriil, F. P., Kaspi, V. M., Roberts, M. S. E., Ibrahim, A., Markwardt, C. B., Swank, J. H., & Finger, M. H. 2005, ApJ, 629, 985

Woods, P. M., Kouveliotou, C., Göğüş, E., Finger, M. H., Swank, J., Smith, D. A., Hurley, K., & Thompson, C. 2001, ApJ, 552, 748

Woods, P. M. & Thompson, C. Soft gamma repeaters and anomalous X-ray pulsars: magnetar candidates, ed. Lewin, W. H. G. & van der Klis, M., 547–586

Woods, P. M. et al. 1999, ApJL, 524, L55

—. 2002, ApJ, 576, 381

Yang, B., Gottlieb, D., & Hesthaven, J. 1997, Journal of Computational Physics, 134, 216

Yang, W. H., Sturrock, P. A., & Antiochos, S. K. 1986, ApJ, 309, 383

Yu, C. 2011, MNRAS, 411, 2461

Zenitani, S. & Hesse, M. 2008, ApJ, 684, 1477

Zenitani, S. & Hoshino, M. 2007, ApJ, 670, 702