Directed dynamical networks applied to SuperMAG ground-based magnetometers

by

Lauren Orr

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Physics

March 2020

Contents

List of Tables

List of Figures

Acknowledgments

I would like to express my gratitude to Prof. Sandra Chapman for her support, patience, encouragement and endless enthusiasm. The opportunities to travel and present my work to so many types of people was much appreciated. I would like to thank my colleagues, particularly those in PS1.16/17 for coffee breaks and always creating a happy atmosphere. I am grateful to our collaborators J.W.Gjerloev and W.Guo for their time, knowledge and encouragement. Thank you to my parents for always supporting me. Lastly thank you to Stacey Cummings and especially David Steele for their constant love, encouragement and spell checking. Phillipians 4 vs 6-7.

Declarations

I declare that the work presented in this thesis is original work carried out at the University of Warwick under the supervision of Prof. S. C. Chapman, unless otherwise stated. All work was carried out between October 2016 to March 2020 and has not been submitted, partially or entirely, for the attainment of a degree in this or any other academic institution. Some parts of this thesis have been published or await publication, as indicated in text.

Abstract

Geomagnetic substorms are a fundamental, global reconfiguration of the magnetosphere during which energy is abruptly transported to the ionosphere where it is dissipated. This fully non-linear and multi-scale process has been observed in-situ for decades, but how it evolves and is spatially structured is a fundamental problem since the various proposed models imply different and conflicting magnetospheric reconfiguration scenarios. We show that the substorm current wedge (SCW) displays large-scale coherent behaviour which puts significant doubt on the recent hypothesis that this current system consists of a series of mesoscale wedgelets. We use techniques from network science to analyse data from > 100 ground-based magnetometers collated by the SuperMAG collaboration. We translate this data into a time-varying directed network. If the canonical cross-correlation between vector magnetic field perturbations, observed at two magnetometer stations, exceeds an event and station specific threshold, they form a network connection. The time lag at which cross-correlation is maximal determines the direction of propagation or expansion of the structure captured by the network connection. If spatial correlation reflects ionospheric current patterns, network properties can test different models for the evolving substorm current system. We obtained the timings for, a consistent picture in which the classic SCW forms. A current system is seen pre-midnight following the SCW westward expansion. Later, there is a weaker signal of eastward expansion. Further, we perform community detection on the network which identifies locally dense but globally sparse groups of connections. We consistently find robust structural change from many small, uncorrelated current systems before substorm onset, to one large spatially-extended correlated system during the expansion phase. All substorms analysed ultimately form a large-scale structure, approximately 10 minutes after onset. This establishes that a single large-scale SCW is central to substorm physics and that substorms do not proceed solely by small-scale wedgelets.

Abbreviations

MHD	Magnetohydrodynamics
SPM	Single particle motion
RHS	Right hand side
IMF	Interplanetry magnetic field
LLBL	Low latitude boundary layer
AE	Auroral electrojet
FAC	Field aligned current
R1/R2	Region 1/2 currents
WTS	Westward travelling surge
SCW	Substorm current wedge
DFBs	Dipolarising flux bundles
DF	Dipolarisation front
BBF	Bursty bulk flow
SEP	Solar energetic particles
CME	Coronal mass ejections
GIC	Geomagnetically induced currents
GPS	Global positioning systems

HF	High frequency
CCA	Canonical correlation analysis
CCC	Canonical cross-correlation
AAFT	Amplitude adjusted Fourier transform
IAAFT	Iterative amplitude adjusted Fourier transform
MLT	Magnetic local time
MLAT	Magnetic latitude

Chapter 1

Introduction

This chapter provides a background to the work in later chapters and outlines why a better understanding of space weather is needed. In section 1.1 an overview of plasma physics is provided, section 1.2 summarises the sun-Earth coupled system and the physical processes that occur at Earth as a result of plasma streaming out of the Sun. Section 1.3 will encapsulate the space weather resulting from these processes, as well as their effect on Earth.

1.1 Properties of Plasmas

In this section a brief overview of basic plasma physics is provided in order to understand how plasma from the Sun can affect the Earth's environment. We will begin with a complete statistical description of a plasma before describing how the relevant properties of a plasma can be usefully described (i) as one or more fluids or (ii) by the motion of a collection of non-interacting charged particles, and the physical properties we can derive from each description.

1.1.1 Definition of a Plasma

A collection of charged particles behaves as a plasma under conditions that we will outline below. A plasma is often described as the "fourth state of matter" [Goldston and Rutherford, 1995]. The complete statistical description of the system is given by the distribution function

$$F_N(\boldsymbol{x_1}, \boldsymbol{x_2}, ..., \boldsymbol{x_N}, \boldsymbol{v_1}, \boldsymbol{v_2}, ..., \boldsymbol{v_N}, t), \tag{1.1}$$

where N is the number of particles in the system and $\int F_N d\boldsymbol{x_1}...d\boldsymbol{x_N}d\boldsymbol{v_1}...d\boldsymbol{v_N} = 1$. The exact position and velocities of each particle are required to calculate F_N, hence plasma is generally described by macro variables which are calculated from a reduced form of F_N [Krall and Trivelpiece, 1973]. A one-particle distribution, $F_{1\alpha}$, is found by integrating F_N over the coordinates and velocities of all but one particle of type α.

The distribution function obey the Liouville equation,

$$\frac{\partial F_N}{\partial t} + \sum_i \left(\frac{\partial F_N}{\partial \boldsymbol{x_i}} \cdot \boldsymbol{v_i} + \frac{\partial F_N}{\partial \boldsymbol{v_i}} \cdot \boldsymbol{a_i} \right) = 0, \tag{1.2}$$

where $\boldsymbol{a_i}$ is the acceleration of particle i due to external and inter-particle forces [Baumjohann and Treumann, 1997]. By integrating the Liouville equation by the coordinates and velocities of all but one particle of type α equation 1.2 becomes,

$$\frac{\partial F_{1\alpha}}{\partial t} + \boldsymbol{v_1} \cdot \frac{\partial F_{1\alpha}}{\partial \boldsymbol{x_1}} + \boldsymbol{a_1^E} \cdot \frac{\partial F_{1\alpha}}{\partial \boldsymbol{v_1}} = \left. \frac{\partial F_{1\alpha}}{\partial t} \right|_c, \tag{1.3}$$

where the RHS is due to collisions and $\boldsymbol{a_1^E}$ is the acceleration of a particle due to external and the average internal forces. The long range electrostatic interactions associated with charged particles means that a 'close encounter' will result in a momentum exchange and particle deflection, essentially the same as a 'collision' between neutral particles [Goldston and Rutherford, 1995]. Equation 1.3 is the kinetic equation of a plasma or the Vlasov equation if the RHS$=0$ [Krall and Trivelpiece, 1973]. Velocity moments of both $F_{1\alpha}$ and equation 1.3 result in several macroscopic variables and equations to describe a conducting fluid containing a single species of charged particles (e.g. ions or electrons).

The continuity equation is given by,

$$\frac{\partial}{\partial t} A + \boldsymbol{\nabla} \cdot A\boldsymbol{V_\alpha} = 0, \tag{1.4}$$

where the macroscopic variable A may represent the number of particles, n_α, the mass density, $\rho_{m\alpha}$, or the charge density, $\rho_{q\alpha}$, and $\boldsymbol{V_\alpha}$ is the average velocity of α particles.

The momentum transfer equation is given by the velocity moments of equation 1.3 multiplied by the momentum. The momentum transfer equation is given by,

$$\rho_{m\alpha} \left(\frac{\partial}{\partial t} + \boldsymbol{V_\alpha} \cdot \boldsymbol{\nabla} \right) \boldsymbol{V_\alpha} - \rho_{q\alpha} \boldsymbol{E} - \frac{\boldsymbol{J_\alpha} \times \boldsymbol{B}}{c} + \boldsymbol{\nabla} \cdot P_\alpha = \int \bar{n} m_\alpha \boldsymbol{v} \left. \frac{\partial F_{1\alpha}}{\partial t} \right|_c d\boldsymbol{v}, \tag{1.5}$$

where $\boldsymbol{J_\alpha} = \rho_{q\alpha}\boldsymbol{V_\alpha}$ is the electric current density, $\bar{n}$ is the number of α particles per unit volume and P_α is the pressure tensor, both for charged α particles. $\boldsymbol{E}$ is the electric

field and $\boldsymbol{B}$ is the magnetic field and they satisfy the Maxwell equations:

$$\boldsymbol{\nabla} \cdot \boldsymbol{E} = \frac{\rho}{\epsilon_0}, \tag{1.6}$$

$$\boldsymbol{\nabla} \times \boldsymbol{B} = \mu_0 \boldsymbol{J} + \epsilon_0 \mu_0 \frac{\partial \boldsymbol{E}}{\partial t}, \tag{1.7}$$

where ρ is the electric charge density, ϵ_0 is the free space permittivity and μ_0 is the magnetic permeability in a vacuum [Krall and Trivelpiece, 1973].

A two-fluid plasma theory will treat the ions and electrons as two coupled, conducting fluids using equations 1.4 with the density, velocity and pressure variables for each species. Each species requires an equation of state to close the system.

A one-fluid plasma uses the combination of the two species' variables. Several approximations can be used to define a simplified set of equations that are commonly used in one-fluid theory. One such simplification is the assumption of quasineutrality. The result is that the total charge density of the system is assumed to be approximately neutral, $\rho_q \approx 0$, on length scales above a certain value known as the *Debye length*, λ_D. The Debye length is the characteristic length scale over which a charged particle will be influenced by the electric field of another charged particle [Baumjohann and Treumann, 1997]. Specifically,

$$\lambda_D \equiv \left(\frac{\epsilon_0 k_B T}{n_e e^2} \right)^{\frac{1}{2}}, \tag{1.8}$$

where k_B is the Boltzmann Constant, T is the temperature, n_e is the electron particle density and e is the electron charge. Within the distances of the order of the Debye length, charged particles arrange themselves to effectively shield any electrostatic fields [Bittencourt, 2013]. Within a sphere of radius λ_D the plasma may be regarded as having fluctuating electric potentials. The *Debye potential* describes the electrostatic potential felt near a test charge particle at radius r and is given by,

$$\phi(r) = \frac{q}{4\pi\epsilon_0 r} e^{-\frac{r}{\lambda_D}}. \tag{1.9}$$

The equation shows that charged particles within a plasma interact when they are at radi smaller than the debye length but beyond λ_D the potential experienced becomes negligible [Bittencourt, 2013]. Therefore quasineutrality is a valid assumption when studying slow motions of fluid elements on scales greater than λ_D and thus we can assume, $\rho_{qe} \approx \rho_{qi}$.

Typical values for the Earth's ionosphere are $n_e = 10^{12}\ m^{-3}$ and $T = 10^3\ K$, such that $\lambda_D = 10^{-3}\ m$ [Bittencourt, 2013]. Further assumptions include neglecting $\frac{m_e}{m_i}$ compared to unity, assuming the plasma has sufficiently low frequency yet a sufficiently high rate of collisions to maintain an isotropic system and $\boldsymbol{\nabla} \cdot \boldsymbol{P} = \boldsymbol{\nabla} p$ where p is the scalar

pressure.

1.1.2 Magnetohydrodynamics (MHD) equations

The resulting equations are known as the magnetohydrodynamic (MHD) equations:

$$\text{Continuity:} \qquad \frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot \rho \boldsymbol{v} = 0, \tag{1.10}$$

$$\text{Momentum:} \qquad \rho \frac{D\boldsymbol{v}}{Dt} = \frac{\boldsymbol{J} \times \boldsymbol{B}}{c} - \boldsymbol{\nabla} p, \tag{1.11}$$

$$\text{Ohm's law:} \qquad \boldsymbol{J} = \sigma_0 \left(\boldsymbol{E} + \frac{\boldsymbol{v} \times \boldsymbol{B}}{c} \right), \tag{1.12}$$

$$\text{Faraday's law:} \qquad \frac{1}{c} \frac{\partial \boldsymbol{B}}{\partial t} = -\boldsymbol{\nabla} \times \boldsymbol{E}, \tag{1.13}$$

$$\text{Ampère's law:} \qquad \boldsymbol{\nabla} \times \boldsymbol{B} = \mu_0 \boldsymbol{J}, \tag{1.14}$$

$$\text{Solenoidal constraint:} \qquad \boldsymbol{\nabla} \cdot \boldsymbol{B} = 0. \tag{1.15}$$

where σ_0 is the electrical conductivity. The subscript for density has been dropped as we have assumed $\rho_q \approx 0$ and hence all future density refers to mass, $\rho_m = \rho$. We note that $\frac{D}{Dt} = \frac{\partial}{\partial t} + \boldsymbol{v} \cdot \boldsymbol{\nabla}$ is the substantial derivative, which gives the rate of change of a quantity at a point moving with the fluid particle. A further equation of state is required to close the system, commonly the assumptions of an incompressible, adiabatic or isothermal fluid, but there are other closure equations [Krall and Trivelpiece, 1973].

Using equations 1.12 to 1.15 the *induction equation* can be derived (e.g. see Priest [2012] for the derivation),

$$\frac{\partial \boldsymbol{B}}{\partial t} = \underbrace{\boldsymbol{\nabla} \times (\boldsymbol{v} \times \boldsymbol{B})}_{(i)} + \underbrace{\eta_0 \nabla^2 \boldsymbol{B}}_{(ii)}, \tag{1.16}$$

where (i) is the *advection* term, (ii) is the *diffusion* term and where $\eta_0 = \frac{1}{\sigma_0 \mu_0}$ is the magnetic diffusivity.

Frozen in Theorem

An important physical interpretation of plasma dynamics can be found from the advection term of the the induction equation (equation 1.16). The following derivation is based on Landau et al. [2013]. We begin by neglecting the diffusion term and expanding

the RHS of equation 1.16.

$$\begin{aligned}\frac{\partial \boldsymbol{B}}{\partial t} =& \boldsymbol{\nabla} \times (\boldsymbol{v} \times \boldsymbol{B}), \\ =& (\boldsymbol{B} \cdot \boldsymbol{\nabla})\boldsymbol{v} + \boldsymbol{v}(\boldsymbol{\nabla} \cdot \boldsymbol{B}) - (\boldsymbol{v} \cdot \boldsymbol{\nabla})\boldsymbol{B} - \boldsymbol{B}(\boldsymbol{\nabla} \cdot \boldsymbol{v}).\end{aligned} \tag{1.17}$$

The *equation of continuity* (equation 1.10) can be expanded to,

$$\frac{\partial \rho}{\partial t} = -(\boldsymbol{v} \cdot \boldsymbol{\nabla})\rho - \rho(\boldsymbol{\nabla} \cdot \boldsymbol{v}). \tag{1.18}$$

We rearrange equation 1.18 and use it to exclude $\boldsymbol{\nabla} \cdot \boldsymbol{v}$ from equation 1.17. Further, we use the solenoidal constraint (equation 1.15) to exclude $\boldsymbol{\nabla} \cdot \boldsymbol{B}$ from equation 1.17.

$$\frac{\partial \boldsymbol{B}}{\partial t} = (\boldsymbol{B} \cdot \boldsymbol{\nabla})\boldsymbol{v} - (\boldsymbol{v} \cdot \boldsymbol{\nabla})\boldsymbol{B} + \frac{\boldsymbol{B}}{\rho}\left(\frac{\partial \rho}{\partial t} + (\boldsymbol{v} \cdot \boldsymbol{\nabla})\rho\right). \tag{1.19}$$

Using the substantial derivative ($\frac{D}{Dt} = \frac{\partial}{\partial t} + \boldsymbol{v} \cdot \boldsymbol{\nabla}$) we can simplify 1.19 and divide through by ρ to ,

$$\frac{1}{\rho}\frac{D\boldsymbol{B}}{Dt} - \frac{\boldsymbol{B}}{\rho^2}\frac{D\rho}{Dt} = \frac{(\boldsymbol{B} \cdot \boldsymbol{\nabla})\boldsymbol{v}}{\rho}. \tag{1.20}$$

By the chain rule

$$\frac{d\rho^{-1}}{dt} = \frac{d\rho^{-1}}{d\rho}\frac{d\rho}{dt} = -\frac{1}{\rho^2}\frac{d\rho}{dt},$$

and by the product rule

$$\frac{d}{dt}\left(\frac{\boldsymbol{B}}{\rho}\right) = \boldsymbol{B}\frac{d\rho^{-1}}{dt} + \frac{1}{\rho}\frac{d\boldsymbol{B}}{dt},$$

we can further simplify equation 1.20. As both the product rule and chain rule hold true for both partial derivative and $\boldsymbol{\nabla}$, they hold true for the substantial derivative and, hence, 1.20 becomes,

$$\frac{D}{Dt}\left(\frac{\boldsymbol{B}}{\rho}\right) = \left(\frac{\boldsymbol{B}}{\rho} \cdot \boldsymbol{\nabla}\right)\boldsymbol{v}. \tag{1.21}$$

The physical interpretation of equation 1.21 is as follows; consider a line which moves with the fluid particle. We let $\delta\boldsymbol{L}$ be a line element and $\boldsymbol{v}$ be the fluid velocity at one end of the element. At the other end of the element, the fluid velocity will be $\boldsymbol{v} + \delta\boldsymbol{L} \cdot \boldsymbol{\nabla}\boldsymbol{v}$. The rate of change at a point moving with the fluid particle is given by,

$$\frac{D}{Dt}(\delta\boldsymbol{L}) = (\delta\boldsymbol{L} \cdot \boldsymbol{\nabla})\boldsymbol{v}. \tag{1.22}$$

Equation 1.22 is of the same form as equation 1.21; therefore, the rate of change of $\frac{\boldsymbol{B}}{\rho}$ is given by an equation of the same form as $\delta\boldsymbol{L}$. Hence, if the line element, $\delta\boldsymbol{L}$, and the magnetic field vector, $\frac{\boldsymbol{B}}{\rho}$, are initially in the same direction, they remain parallel

with one another and their values retain the same ratio. If two infinitely close fluid particles are on the same line of force, at any time, they will always be on the same line of force. The value of $\frac{\boldsymbol{B}}{\rho}$ will remain proportional to the distance between the particles. Therefore, every line of force moves with the fluid particles; the lines of magnetic force are *frozen* in the fluid and move with it [Landau et al., 2013].

1.1.3 Single Particle Motion (SPM)

In section 1.1.1 we outlined how a plasma can be described as a fluid, but the behaviour of the individual particles can also affect large-scale physical processes. SPM neglects the collective behaviour of a plasma but is useful when studying low density plasmas such as those found in the Earth's ring current and in the radiation belts [Baumjohann and Treumann, 1997].

Gyration

The *equation of motion* for a charged particle under the Coulomb (term 1 RHS) and Lorentz forces (term 2 RHS) is given by

$$m\frac{d\boldsymbol{v}}{dt} = q(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}), \tag{1.23}$$

where m is the particle mass and $\boldsymbol{v}$ the particle velocity [Krall and Trivelpiece, 1973]. Charged particles at rest are the sources of the electrostatic field, $\boldsymbol{E}$, which is the origin of the Coulomb force. Charged particles moving with velocity, $\boldsymbol{v}$, are the current elements generating a magnetic field, $\boldsymbol{B}$, which is the origin of the Lorentz force [Baumjohann and Treumann, 1997].

In the absence of an electric field ($\boldsymbol{E} = 0$), the equation of motion reduces to,

$$m\frac{d\boldsymbol{v}}{dt} = q(\boldsymbol{v} \times \boldsymbol{B}). \tag{1.24}$$

From equation 1.24 it can be shown that a particle's kinetic energy, as well as the magnitude of its velocity, are constant in time [Krall and Trivelpiece, 1973].

Again from equation 1.24 the orbit of a charged particle, with a uniform magnetostatic field along the z-axis, $\boldsymbol{B} = B\hat{\boldsymbol{e}}_z$, can be described by,

$$\begin{aligned} x - x_0 &= r_g \sin \omega_g t, \\ y - y_0 &= r_g \cos \omega_g t, \end{aligned} \tag{1.25}$$

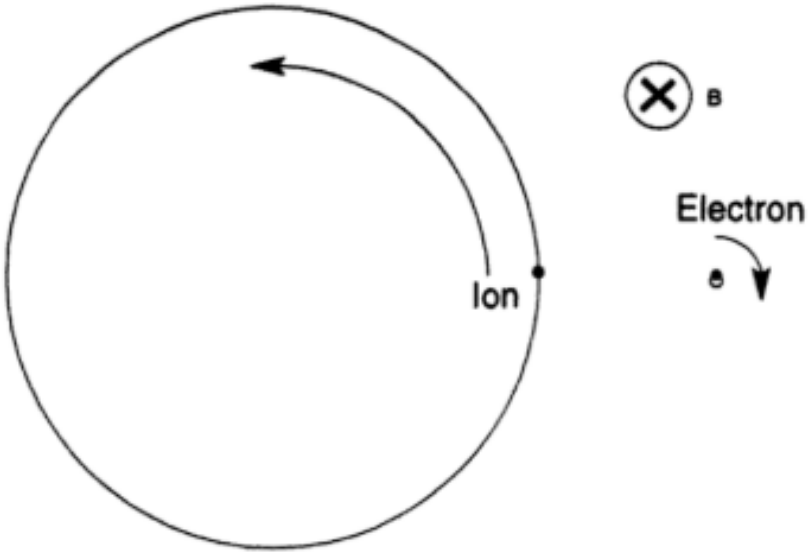

Figure 1.1: The gyration of charged particles around their guiding centres. The ion has a much greater gyroradius than the electron for fixed energy. Figure taken from Goldston and Rutherford [1995].

where r_g is the *gyroradius* or *Larmor radius* as defined by

$$r_g = \frac{v_\perp}{|\omega_g|} = \frac{mv_\perp}{|q|B}, \tag{1.26}$$

$$\tag{1.27}$$

where $v_\perp = (v_x^2 + v_y^2)^{\frac{1}{2}}$ is the constant speed in the plane perpendicular to $\boldsymbol{B}$ [Baumjohann and Treumann, 1997], and

$$\omega_g = \frac{qB}{m}, \tag{1.28}$$

is the *gyro-frequency* or *cyclotron frequency* or *Larmor frequency* [Goldston and Rutherford, 1995].

Equations 1.25 describe a circular orbit which a charged particle follows around the magnetic field. The direction of rotation is dependent on the charge of the particle, q, contained in the gyro-frequency term. The centre, (x_0, y_0), of the orbit is the *guiding centre* and the radius of the orbit is r_g. Figure 1.1 is a schematic of gyration around the guiding centres of an ion and an electron. If plasma phenomena vary on time scales much longer than the gyro-period and spatial scales much larger than the gyroradius, the magnetic field is changing the dynamics of the particle and the gyro-motion affects the plasma behaviour [Goldston and Rutherford, 1995].

The constant velocity, along (parallel to) the magnetic field, results in the trajectory

of the particle being a three-dimensional helix, assuming $v_{\parallel} \neq 0$. An example of this helicoidal trajectory can be seen in figure 1.2. Goldston and Rutherford [1995] describes the "particle gyro-centres as sliding along magnetic field lines, like beads on a wire." The angle between the magnetic field direction and the direction of particle motion is known as the *pitch angle*, α_p, of the helix [Bittencourt, 2013] and is defined as

$$\alpha_p = \tan^{-1}\left(\frac{v_{\perp}}{v_{\parallel}}\right). \tag{1.29}$$

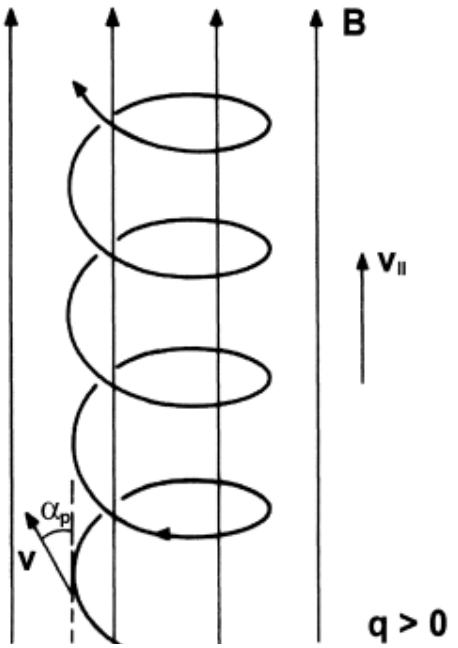

Figure 1.2: The helicoidal trajectory of a positively charged particle's gyration orbit in a uniform magnetic field. Figure adapted from Bittencourt [2013].

Electric and Magnetic Drifts

Electric and magnetic fields induce drifts to the particles gyratory motion.

$$E \times B \text{ Drift:} \qquad \boldsymbol{v_E} = \frac{\boldsymbol{E} \times \boldsymbol{B}}{B^2}, \tag{1.30}$$

$$\text{Polarisation Drift:} \qquad \boldsymbol{v_P} = \frac{1}{\omega_g B}\frac{d\boldsymbol{E}_{\perp}}{dt}, \tag{1.31}$$

$$\text{Gradient Drift:} \qquad \boldsymbol{v_\nabla} = \frac{v_{\perp}^2}{2\omega_g B^2}(\boldsymbol{B} \times \nabla B), \tag{1.32}$$

$$\text{Curvature Drift:} \qquad \boldsymbol{v_R} = \frac{v_{\parallel}^2}{\omega_g R_c^2 B}(\boldsymbol{R}_c \times \boldsymbol{B}), \tag{1.33}$$

$$\text{Diamagnetic Drift:} \qquad \boldsymbol{v_D} = -\frac{\nabla p \times \boldsymbol{B}}{qnB^2}, \tag{1.34}$$

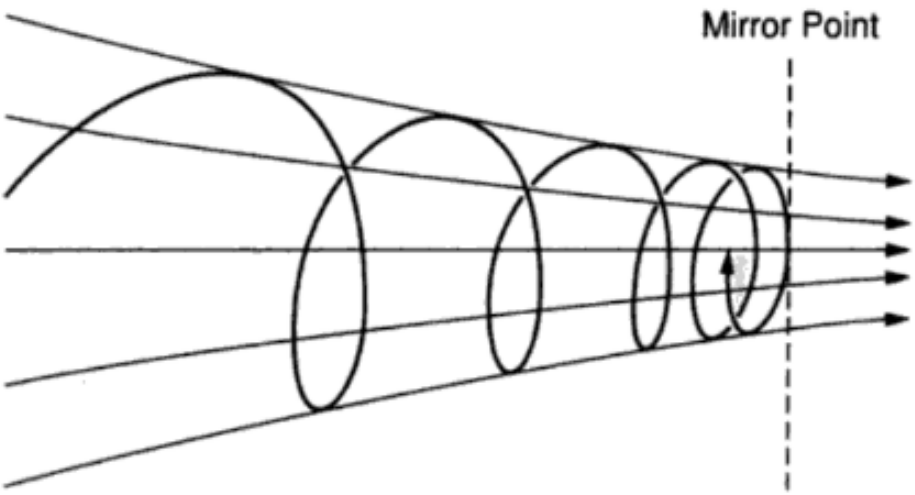

Figure 1.3: The orbit of an ion along a converging magnetic field reflected at its mirror point. Figure from Baumjohann and Treumann [1997].

where $\boldsymbol{R}_c$ is the local radius of curvature and n is the particle density [Baumjohann and Treumann, 1997].

The $E \times B$ drift does not depend of the charge of the particles, thus all particles move perpendicular to $\boldsymbol{B}$ and $\boldsymbol{E}$. When collisions occur between charged particles and neutral particles, the $E \times B$ drift results in an electric current that is perpendicular to both $\boldsymbol{E}$ and $\boldsymbol{B}$. The current is known as the Hall current and is the result of the ion-neutral collision frequency being higher than the electron-neutral collision frequency, causing electrons to move faster than ions [Bittencourt, 2013]. The polarisation, gradient and curvature drifts depend on ω_g (equation 1.28) and therefore depend on the charge of the particles, hence there is a transverse current since ions and electrons are oppositely charged [Baumjohann and Treumann, 1997].

Mirror Force

A *mirror point* is defined as a point along the field line, B_m, where the pitch angle (equation 1.29), $\alpha_m = 90°$. At a mirror point the particle will be reflected as in figure 1.3. The orbit of the particle is a helix around the guiding centre, with decreasing radii with increasing magnetic field. The parallel kinetic energy of the particle decreases with converging magnetic field lines, while the perpendicular kinetic energy increases, until all of the particle's kinetic energy is perpendicular to the drift direction [Baumjohann and Treumann, 1997; Bittencourt, 2013]. This is known as the mirror point, where the parallel motion stops and reverses.

In magnetic fields with converging field lines at either end and a weaker field in the

centre, particles can be mirrored at both ends and become *trapped*. The Earth's magnetic dipole is an example of such a geometry. The trapped energetic particles of the terrestrial magnetic field, within $\sim 6\ R_E$, form the radiation belts. The trapped particles spiral along the field lines, bouncing back and forth between the magnetic poles [Bittencourt, 2013].

The gradient and curvature drifts result in a slow azimuthal drift of the particles, where the ions drift westward and the electrons drift eastwards around the Earth. The associated current is known as the ring current. Further from the Earth the $E \times B$ drift will dominate over the gradient and curvature drifts resulting in a sunward drift [Baumjohann and Treumann, 1997].

1.2 The Sun Earth coupled system

This section summarises the solar terrestrial physics needed to understand the physical processes responsible for space weather. Space weather is the result of coupling between the Sun and the Earth, through the medium of the solar wind. First the solar wind will be summarised, followed by the terrestrial environment and its response to enhanced solar activity.

1.2.1 Solar Wind

The idea that a medium existed between the Earth and the Sun was first suggested when Biermann [1957] pointed out that the observed motions of comet tails would require matter streaming outward from the Sun. Parker [1958] built on this idea and suggested that there was no hydrostatic equilibrium solution for the corona but there must be continuous expansion.

The solar wind is a highly conducting plasma which streams radially out of open-field regions of the Sun, into interplanetary space. Near the Earth it travels at supersonic speeds of $\sim 400 km/s$ but can travel up to $900 km/s$ [Priest, 2012]. It consists of mainly electrons and protons but is $\sim 5\%$ Helium ions and the density of the solar wind decreases with the inverse square of the distance from the Sun [Baumjohann and Treumann, 1997]. Close to the Earth, the solar wind is already quite rarified, with a typical electron density of $n_e \sim 5 \times 10^6 m^{-3}$ and a temperature of $T_e \sim 5 \times 10^4 K$, hence the solar wind is highly conductive [Bittencourt, 2013]. The solar coronal magnetic field is 'frozen' into the streaming solar wind plasma and will be drawn outward by the expanding solar wind flow to eventually become the interplanetary magnetic field (IMF). The IMF direction varies on timescales of minutes [Eastwood et al., 2015].

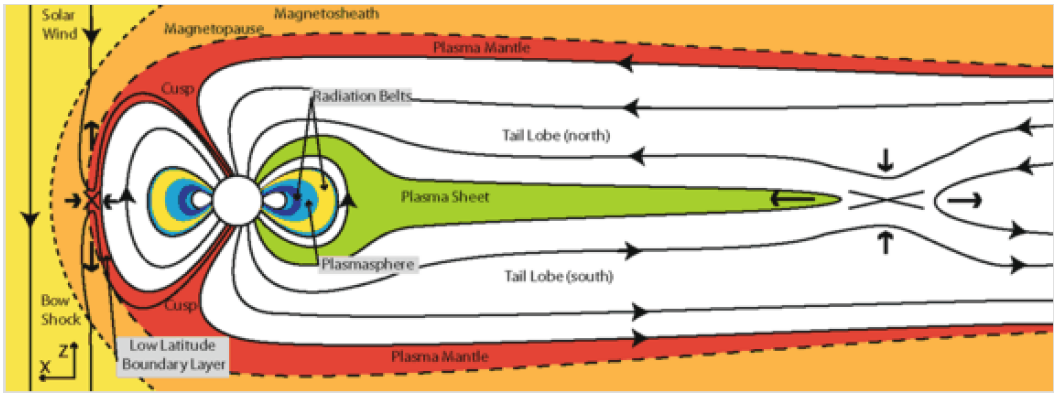

Figure 1.4: A schematic of the Earth's magnetosphere with the sun to the left of the image. From Eastwood et al. [2015].

Figure 1.4 is a large-scale schematic of the terrestrial environment and solar wind coupling. The Earth's dipolar magnetic field acts as a boundary to the solar wind, with most of the plasma and IMF being deflected around the terrestrial field [Bittencourt, 2013]. The bow shock is generated when the solar wind hits the terrestrial field at supersonic speeds. The bow shock acts as a particle accelerator and an energy converter [Eastwood et al., 2015]. This thermalised subsonic plasma forms a region between the bow shock and the outer boundary of the geomagnetic field, called the *magnetosheath* (orange layer in figure 1.4). In the magnetosheath the plasma is compressed and, therefore, is denser and hotter than it was within the solar wind [Baumjohann and Treumann, 1997]. Behind the magnetosheath is the magnetopause (dashed line in figure 1.4), the outer boundary of the geomagnetic field [Bittencourt, 2013].

1.2.2 Magnetosphere

The magnetosphere is the name given to the near Earth field and plasma environment. The magnetosphere protects the planet from much of the solar wind by deflecting the solar wind [Baumjohann and Treumann, 1997]. The dense plasma is mostly deflected around the magnetic field, as the plasma and IMF move together due to the frozen-in property. Only when the terrestrial and interplanetary field have an anti-parallel component can reconnection occur and the plasma infiltrates the magnetosphere, i.e. when the IMF is southwardly directed [Eastwood et al., 2015]. The kinetic pressure of the solar wind plasma distorts the magnetosphere shape, so on average it extends $\sim 10\ R_E$ (Earth radii) on the dayside [Bittencourt, 2013] and to at least 210 R_E on the [Christon et al., 1998], forming a stretched magnetotail. The coupling of terrestrial magnetic field to solar wind interplanetary field creates a highly dynamic, non-equilibrium environment [Baumjohann and Treumann, 1997].

The plasma within the magnetosphere is not evenly distributed but contained in re-

gions of varying temperature and density. The low latitude boundary layer (LLBL) or magnetopause current sheet (red layer in figure 1.4) forms on the dayside beyond the magnetopause. The LLBL forms of the solar wind plasma that enters when the IMF is southwardly directed. It contains mainly solar wind plasma and is a near-permanent feature. The plasma mantle forms at the edge of the magnetotail and can, to some extent, be considered the tailward extension of the LLBL [Eastwood et al., 2015]. The tail lobes contain highly rarified plasma, consisting of magnetic flux with a footpoint on Earth and extending into interplanetary space. A thick plasma sheet is formed mid-plane via reconnection [Dungey, 1961] in the tail. Closed field lines return this plasma towards the Earth to the high-latitude auroral ionosphere.

The radiation belts (yellow and dark blue in figure 1.4) lie on field lines at $\sim 2-6R_E$. Energetic particles (mainly electrons and protons), with high kinetic energy (> 30 keV) but lower density plasma (~ 1 cm^{-3}), becomes trapped within the radiation belts [Baumjohann and Treumann, 1997]. The radiation belts trap the plasma particles by the mirror force discussed in section 1.1.3 and can accelerate ions and electrons to very high energies. Geomagnetic storms often energise the radiation belts by injecting solar wind protons and electrons [Bittencourt, 2013]. The plasmasphere (light blue, figure 1.4) contains a dense, cold plasma of ionospheric origin ($n_e \sim 5\times10^2$ cm^{-3},$T_e \sim 5\times10^3$ K) [Baumjohann and Treumann, 1997].

Figure 1.4 suggests a state of dynamic equilibrium, but this is a driven and dissipating, non-equilibrium system which responds non-linearly to the solar wind driver and is difficult to predict. The solar wind varies continuously on time scales of minutes to hours, even in steady solar wind [Eastwood et al., 2015].

Dungey cycle

The key mechanism for transport of solar wind plasma into the terrestrial system is the Dungey cycle [Dungey, 1961]. This is the open magnetosphere model of the dynamic system that evolves from the interaction of the Earth's magnetic field and the IMF. If there is a southwardly directed IMF line (one of the blue lines in figure 1.5) there will be reconnection (dayside shaded box) between the IMF and the closed terrestrial field, (green) which have footpoints on the earth. This field will then reconfigure into two open field lines (red) and be carried downtail by the solar wind, forming the magnetotail lobes. Reconnection takes place again, in the central plane at $\sim 100-200$ R_E downtail [Baumjohann and Treumann, 1997], (nightside shaded box) forming a stretched, closed terrestrial field line in the magnetotail (green), and an open solar wind field line (blue). As the stretched terrestrial field relaxes, the frozen-in plasma is transported back towards earth. The magnetic field footpoints moving tailward results in a twin-

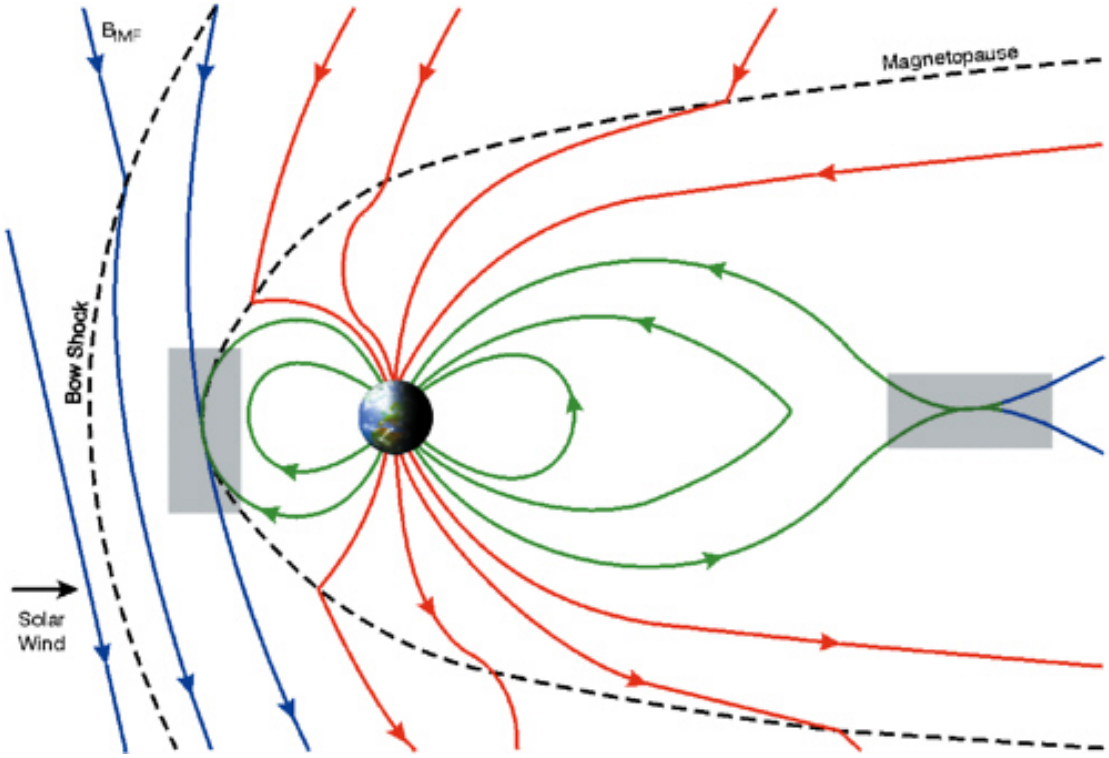

Figure 1.5: The Dungey cycle, with southward IMF lines (blue) reconnecting with Earth's closed field lines (green) at the dayside magnetopause. The resulting open field lines (red) move downtail and reconnect in the distant tail. Figure from [Council, 2004].

cell convective flow of plasma across the ionosphere, resupplying the dayside with flux and plasma [Milan et al., 2007].

Dungey originally proposed the reconnection at the dayside and tail occurred at an equal rate [Dungey, 1961], but they do not have to happen simultaneously. The dayside reconnection rate depends on the magnitude and orientation of the IMF and the solar wind velocity. The nightside reconnection rate is affected by the magnetotail conditions. The tail conditions are affected by the dayside reconnection rate, but it is not an instantaneous response as the open flux take several tens of minutes to propagate to the tail [Milan et al., 2007].

1.2.3 Earth's Atmosphere

The Earth's atmosphere is comprised of several layers (see figure 1.6) but the boundaries depend on pressure and temperature. The troposphere is the lowest layer, from the surface to $\sim 11\ km$; it contains $\sim 90\%$ of the atmosphere's total mass and all living things. Above this is the stratosphere, where ozone molecules form and heat the atmosphere by trapping ultraviolet radiation. It is generally at higher altitudes over the equator and lower altitudes towards the poles. Third is the mesosphere which does not absorb

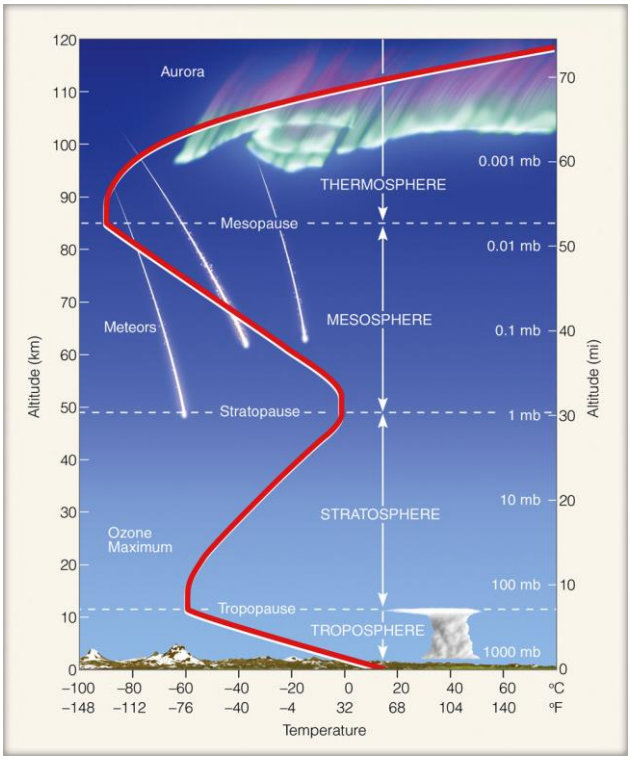

Figure 1.6: Layers of the Earth's Atmosphere with average profile of air temperature (red line) overplotted. From Ahrens and Samson [2010].

a significant amount of radiation but decreases in temperature with altitude. Finally, the thermosphere is the highest and thickest layer, where satellites orbit the Earth and the temperature increases with altitude. The thermosphere contains oxygen molecules which are heated by solar radiation [Ahrens and Samson, 2010]. Overlapping in space with the mesosphere and thermosphere is the ionosphere. At an altitude of $\sim$ 60km from the Earth's surface, it is the ionised component of the Earth's upper atmosphere and the lower boundary of the magnetosphere [Bittencourt, 2013].

The two main sources of ionisation of neutral particles are i) absorption of solar extreme ultraviolet and x-ray radiation, which have maximum ionisation on the dayside near the local noon and equator, and ii) energetic particles from outside the magnetosphere, which penetrate the atmosphere at auroral latitudes [Bittencourt, 2013]. The ionisation is highly altitude dependent due to the chemical composition and density of the atmosphere. The collisions between neutral and ionised ions above $\sim$ 80 km altitude are too infrequent to form a neutral environment, resulting in the permanently ionised atmosphere. At low and mid latitudes the upper ionosphere merges with the lower plasmasphere, the cool and dense plasma layer extending to $\sim 4\ R_E$ [Baumjohann and Treumann, 1997].

At the auroral latitudes there can be precipitation of plasma sheet electrons along the

open magnetic field lines down to the ionosphere, leading to collisions with the particles in the neutral atmosphere. The neutral atmosphere is ionoised and emits photons which are responsible for the auroral light. The aurora are typically observed at the footpoints of the closed magnetic field lines, in the region known as the auroral oval [Baumjohann and Treumann, 1997].

1.2.4 Magnetospheric Currents

Introduction

As described in section 1.1, an electric current describes a flow of charge from one place to another, they are associated with a magnetic field and they combine with the Earth's dipolar magnetic field to form the topology of the magnetosphere [Ganushkina et al., 2018]. A generic expression for the magnetic field produced by moving charges is given by Biot Savart's Law,

$$\boldsymbol{B}(\boldsymbol{r},t) = \frac{\mu_0}{4\pi} q \frac{\boldsymbol{v} \times (\boldsymbol{r} - \boldsymbol{r}'(t))}{|\boldsymbol{r} - \boldsymbol{r}'(t)|^3} \tag{1.35}$$

It calculates the magnetic field produced at position $\boldsymbol{r}$ at time t, by a charge at position $\boldsymbol{r}'(t)$. Magnetospheric physics has debated for years two paradigms for treating magnetospheric dynamics [Alfvén, 1977; Song and Lysak, 1994; Parker, 1996; Lui, 2000]. The debate centres around whether the primary quantities for magnetospheric analysis should be magnetic field and plasma bulk flow, known as the Bu paradigm, or electric field and current density, the Ej paradigm. The Bu and Ej paradigms are respectively related to the concepts of magnetic reconnection and current disruption. The Bu paradigm uses a simplified Ohm's law and Ampere's law to derive the electric field and current density from the primary quantities and is mainly based on the single-fluid theory (MHD). The Ej paradigm derives the magnetic field and the plasma flow through the Biot-Savart law and calculations using the ensemble of single particles, based on a single particle calculation or kinetic analysis. Whilst Parker [1996] asserts that the Ej paradigum could lead to incorrect physics Lui [2000] insists that each paradigm has its limits and merits and the choice of which to use should depend on the magnetospheric phenomenon being studied.

The fluid approach is good as it provides the overall macrodynamics of the magnetospheric and a simplified picture of the influence of the solar wind on the magnetosphere [Parker, 1996], however there is general negligence on the inherent limitations on MHD as they are limited to slow time variations and long spatial scales [Lui, 2000], as described in section 1.1. Phenomena, such as substorm expansion onset, fall beyond the scope of these limitations. According to observations magnetospheric plasmas in general do not obey a simple equation of state such as the ones generally used to close the set of

equations for Bu paradigm [Chew et al., 1956]. Following our assumptions of quasineutrality, Gauss's law (equation 1.6), simplifies to $\boldsymbol{\nabla} \cdot \boldsymbol{E} = 0$, meaning there are no sources or sinks in the electric field, i.e. the system must be closed. Particles are commonly lost through precipitation in the ionosphere and/or as they escape along the magnetic field lines and so magnetospheric plasma does not always form a simple closed system. [Lui, 2000] argues that Ej paradigm is better suited for studying of certain physical processes such as particle acceleration, plasma waves and the breaking of the magnetic field frozen-in condition [Ganushkina et al., 2018] but Parker [1996] states the paradigm is severely limited when applied to dynamical problems. An associated weakness of Ej is that the current closure is in simple, well-defined loops but the currents can have high levels of complexity to their closure [Ganushkina et al., 2018]. Both methods have their benefits and weaknesses [Lui, 2000] but the Bu paradigm is often found useful for addressing magnetospheric physics Ganushkina et al. [2018]; Vasyliūnas [2001].

Magnetospheric currents overview

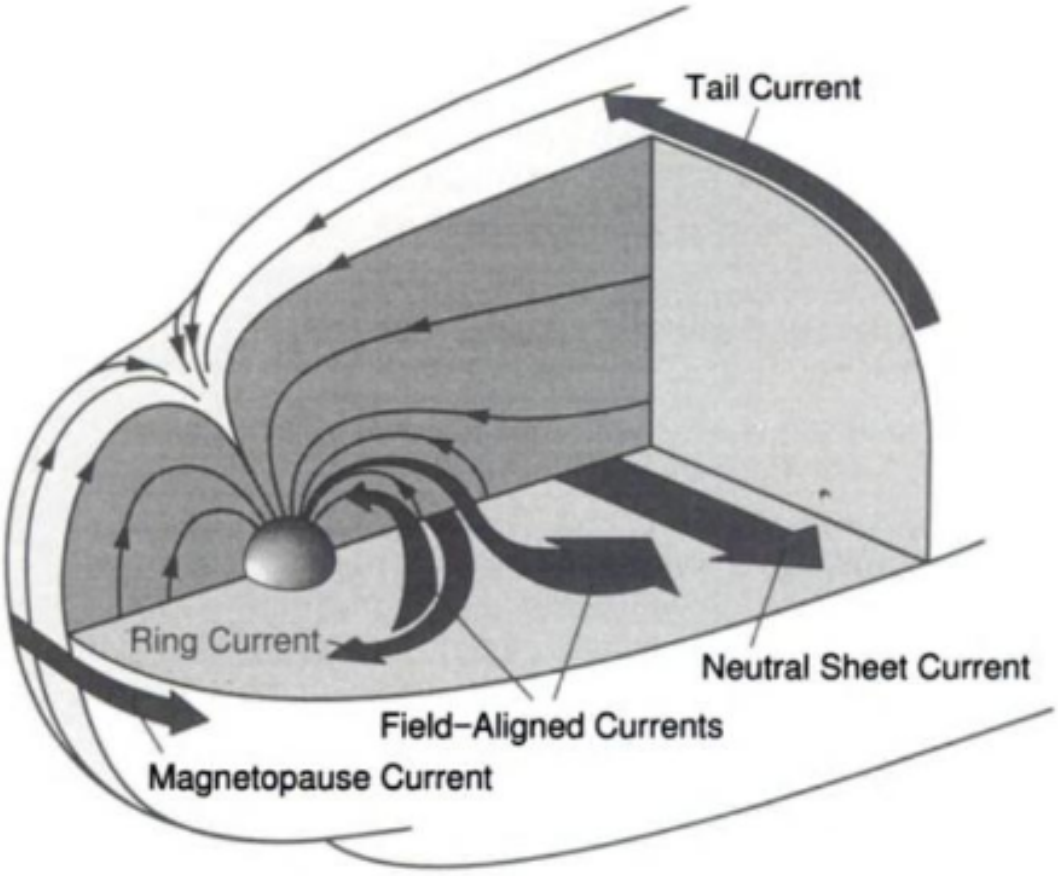

Figure 1.7: Schematic of the large-scale currents flowing in the magnetosphere. Taken from Baumjohann and Treumann [1997].

Figure 1.7 shows the large-scale currents that flow within the magnetosphere. The

magnetopause current is associated with a pressure gradient due to the compression of the dayside terrestrial field and the tail current accompanies the magnetotail magnetic fields [Eastwood et al., 2015]. The neutral sheet current flows westward from dawn to dusk within the plasma sheet due to a gradient drift associated with the weakening magnetic field. It is connected to, and closed by, the tail current [Axford et al., 1965]. Whilst on a global magnetospheric scale this tail current is a stable system, on a macro-scale there are dynamic currents such as dipolorization front currents [Liu et al., 2013; Ganushkina et al., 2018], which we will discuss further in section 1.2.5. In the inner magnetosphere, the ring current flows westward around the Earth at distances of several R_E. The radiation belt particles carry this current. The ring current particles bounce and gyrate around field lines. The protons drift westwards whilst the electrons move eastwards; this drifting in combination with their gyration motion in a region with a pressure gradient results in a net charge transport around the Earth [Baumjohann and Treumann, 1997]. The ring current is a closed system but the current density is generally distributed asymmetric [Ganushkina et al., 2018]. Field-aligned currents (FACs) flow along the magnetic field lines and connect the magnetospheric and ionospheric currents. They are described as region 1 (R1) and region 2 (R2) currents [Iijima and Potemra, 1976] and the current flow is due to polarisation drifts [Stern, 1977].

Within the ionosphere there are several currents, including the auroral electrojets at auroral latitudes, the Sq currents in the dayside at mid-latitude and the equatorial electrojet above the magnetic equator. The Sq current is a result of sunlight and is a permanent feature of the dayside. The Sq currents form two vortices, one in each hemisphere, with their centres at mid latitudes. Where they touch the equator, they flow in the east to west direction, from morning towards evening, and form an equatorial electrojet [Baumjohann and Treumann, 1997].

Auroral Currents

Figure 1.8 from Milan et al. [2017] displays a schematic of a possible current system scenario within the auroral oval. The ionisation from the absorption of solar radiation causes ionospheric conductivity on the dayside [Bittencourt, 2013], but at high latitudes the polar cap is in complete darkness for many months of the year. Energetic particles from the magnetosphere, which are depositing charge in the ionosphere, cause the high ionospheric conductivity in the auroral region [Bittencourt, 2013]. The solar wind plasma is convected along magnetic field lines, producing a horizontal drift of the ionospheric plasma [Cowley, 2000]. The motion of the flux tubes across the polar cap, due to the Dungey cycle, moves the ionospheric footpoint of the flux tube and the plasma tied to it across the polar cap to the nightside. Sunward convection of the magnetic flux tubes footpoints occurs in the high latitude ionosphere leading to the two-cell convection

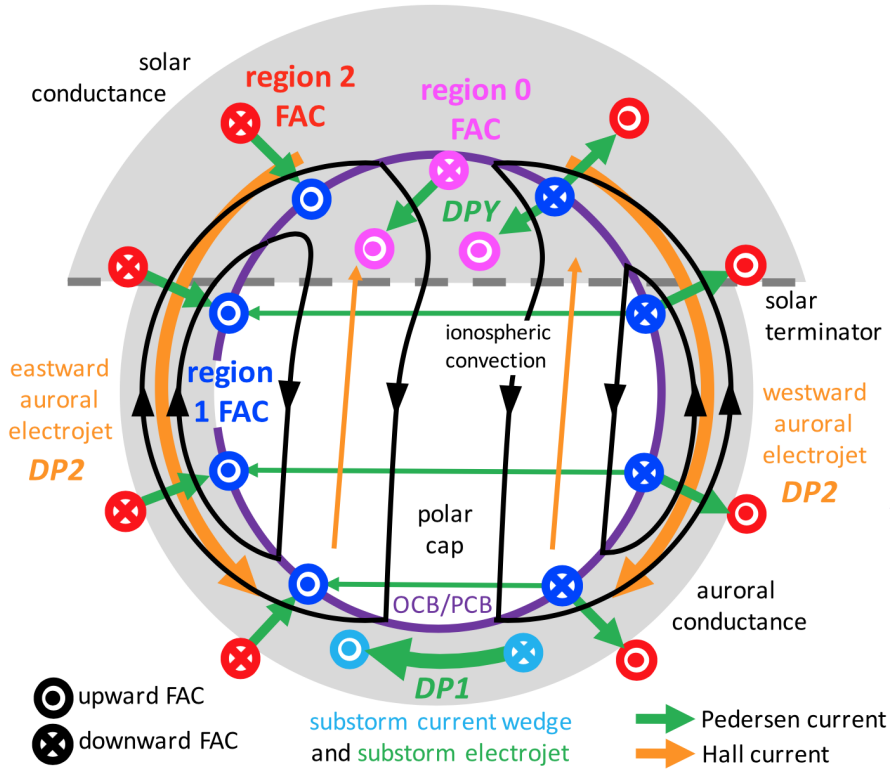

Figure 1.8: A schematic of the current systems in the polar ionosphere of the northern hemisphere. The purple circle is the open/closed field line boundary (OCB). Black arrows are the typical twin-cell ionospheric convection pattern. The region 1 (R1) FACs (blue) and the region 2 (R2) FACs (red) at the poleward/equatorward edge of the auroral zone. Region 0 (R0) FACs (magenta) flow in the cusp throat of the convection pattern. The grey shading represents enhanced ionospheric conductance in the auroral zone and illuminated sections. Pedersen currents (green) flow horizontally between upwards and downwards FACs. Hall currents in the high conductance auroral zones form the eastward and westward electrojets. Substorm current wedge FACs (cyan) and the interconnecting substorm electrojet are present during substorms. From Milan et al. [2017].

pattern around dawn and dusk [Milan et al., 2007].

In the ionosphere the electric current density $\boldsymbol{J}$ is assumed to follow Ohm's law (equation 1.12) in the form,

$$\boldsymbol{J} = \boldsymbol{\sigma} \cdot \boldsymbol{E}' = \sigma_P \boldsymbol{E}'_\perp + \sigma_H \hat{\boldsymbol{B}} \times \boldsymbol{E}' + \sigma_\parallel \boldsymbol{E}_\parallel, \tag{1.36}$$

where $\boldsymbol{E}' \equiv \left(\boldsymbol{E} + \frac{\boldsymbol{v} \times \boldsymbol{B}}{c}\right)$, $\hat{\boldsymbol{B}}$ is the unit vector along the ambient magnetic field $\boldsymbol{B}$ and $\boldsymbol{\sigma}$ is the conductivity tensor where σ_P, σ_H, and $\sigma_\parallel$ are the Pedersen, Hall and parallel

conductivities, respectively [Vasyliūnas, 2012]. These conductivities are given by,

$$\sigma_P = \frac{\nu_c^2}{(\nu_c^2 + \omega_g^2)}\sigma_0, \tag{1.37}$$

$$\sigma_H = \frac{\nu_c \omega_g}{(\nu_c^2 + \omega_g^2)}\sigma_0, \tag{1.38}$$

$$\sigma_{\parallel} \equiv \sigma_0 = \frac{n_e e^2}{m_e \nu_c}, \tag{1.39}$$

where ν_c is the collision frequency and ω_g the gyro-frequency (equation 1.28) [Bittencourt, 2013].

The auroral electrojets are a result of the ionisation caused by energetic particle precipitation into the auroral oval, which cause a much higher conductivity than that of the polar cap and carry a current of $\sim 10^6$ A [Baumjohann and Treumann, 1997]. They result in the largest disturbances of the ground magnetic field of all the current systems [Milan et al., 2017]. The auroral electrojets are primarily convective currents, called Hall currents (orange in figure 1.8, discussed in section 1.1.3) travelling from noon, eastwards and westwards towards midnight. The E x B drift causes particles to drift perpendicular to both the electric and magnetic fields. Ion-neutral collisions are more frequent than electron-neutron collisions, hence ions drift more slowly than electrons. The Hall current electrojets create a pattern known as DP2 [Nishida, 1968] which dominates in the dawn and dusk regions.

The electrojets are supplied by downward FACs. The downwards FACs are equatorwards of the auroral oval in the pre-midnight sector and poleward of the auroral oval post-midnight (red and blue in figure 1.8). Upward FACs close the field across the auroral oval and are closed via a Pedersen current (green in figure 1.8) [Milan et al., 2017]. Pedersen currents are the result of ions and electrons being accelerated along the electric field but anti-parallel to each other. The conductivity depends on the rate of collisions compared to the electrostatic force. The poleward FACs are the R1 currents and the equatorward FACs are the R2 currents [Iijima and Potemra, 1976]. Figure 1.9 shows a schematic of how the FACs may be coupled to the magnetosphere. The FACs flow along the magnetic field lines to connect the regions of convective flow within the magnetosphere but, in the collisionless ionosphere, the motions of positive and negatively charged particles result in current flow both parallel (Pedersen) and perpendicular (Hall) to the electric field [Carter et al., 2016; Milan et al., 2017]. The parallel conductivity describes the electric current flow along the magnetic field in the direction of the electric field and do not depend on the strength of the magnetic field [Bittencourt, 2013].

The eastward and westward electrojets meet and overlap, typically in the pre-midnight region, leading to a three-sheet FAC; this region is known as the Harang discontinuity [Baumjohann and Treumann, 1997].

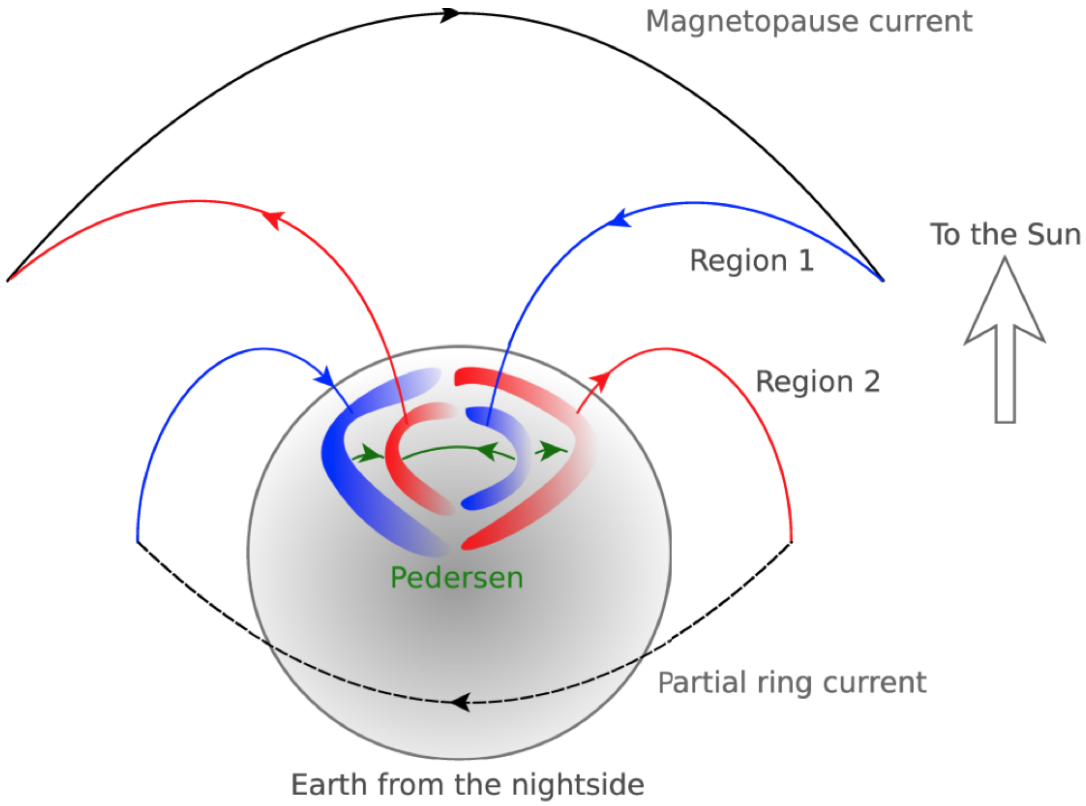

Figure 1.9: The global layout of the magnetospheric current systems of the Earth (northern hemisphere only). The R1 and R2 currents are shown, along with the magnetopause (black), partial ring current (black dashed), and the Pedersen currents (green) that flow across the polar cap. Upward/downward current regions in the polar region and the R1 and R2 currents are coloured red/blue. From Carter et al. [2016].

The region 0 (R0) current (purple in figure 1.8) is poleward of the auroral zone near midday. The polarity of the R0 current is dependent on the B_y component of the IMF. It is upwardly directed in the northern hemisphere if $B_y > 0$ and downwards if $B_y < 0$. R0 is associated with the east to west flows in the dayside polar cap that are caused by magnetic tension forces on newly reconnected field lines [Milan et al., 2017].

Magnetospheric substorm

All of the currents mentioned so far have the implicit assumptions that an equilibrium (balance) exists between the dayside reconnection, tailward motion, nightside reconnection and sunward motion [Eastwood et al., 2015]. Further, it is assumed the rate at which southward directed IMF reconnects with the terrestrial field is constant [Baumjohann and Treumann, 1997]. As discussed in section 1.2.2, the dayside reconnection rate can be different to the nightside reconnection rate. Only a fraction of the flux that is transported tailwards will instantly be reconnected and convected back towards the

dayside [Baumjohann and Treumann, 1997]. The rest of the flux is stored in the tail lobes. When the magnetic flux and associated magnetic energy accumulated becomes too great, ($\sim$ 60 minutes) the tail becomes unstable and the field lines are suddenly reconnected in the tail, causing an explosive release of magnetic energy [Milan et al., 2017]. The time before the release is known as the "growth phase" [McPherron, 1970] and it results in an enhancement of the two-cell pattern and the auroral electrojets [Nishida, 1968].

The substorm onset occurs when the magnetic energy is explosively released into the nightside. This is the start of the expansion phase. The aurora will appear as a localised brightening, before expanding westward and then polewards. The westwards expansion is known as the westward travelling surge (WTS) [Kepko et al., 2015]. Additionally, the ionospheric current is greatly enhanced and the stretched magnetic field in the plasma sheet becomes more dipolar again. After $\sim$ 45 minutes the magnetosphere enters the substorm recovery phase where the aurora will start to fade and retreat poleward, the ionospheric current flow decreases and the field becomes less dipolar [Baumjohann and Treumann, 1997].

The unloading of the stored flux leads to an additional current, the substorm electrojet, with a strongly enhanced westward current flow in the midnight sector [Baumjohann and Treumann, 1997]. The substorm electrojet is the ionospheric part of the substorm current wedge (SCW) and the patterns associated with nightside auroral zone substorm unloading are known as DP1 [Nishida, 1968].

At low latitudes, magnetometers are sensitive to magnetic perturbations associated with the ring currents and the dayside magnetopause [Akasofu and Chapman, 1961], but at auroral latitudes, Hall currents produce perturbations that are observable with ground-based magnetometers [Milan et al., 2017]. It is generally assumed that Pedersen currents produce little ground magnetic perturbations due to the cancellation by the magnetic perturbations produced by the R1 and R2 FACs when they close [Fukushima, 1976], however this is not always a good approximation [Laundal et al., 2016].

Differing models of the large-scale substorm current

The above information on magnetospheric current configuration is the accepted theory of substorms, but there is still much debate around how the current system(s) form and how they are structured. Figure 1.10 shows six different schematics of how the large-scale magnetospheric currents appear during a substorm. These are just an example of some of the many different current configurations that authors have suggested. Since the first schematic picture of the large-scale SCW system [McPherron et al., 1973], there has been a steady stream of competing models [Kamide and Kokubun, 1996; Rostoker,

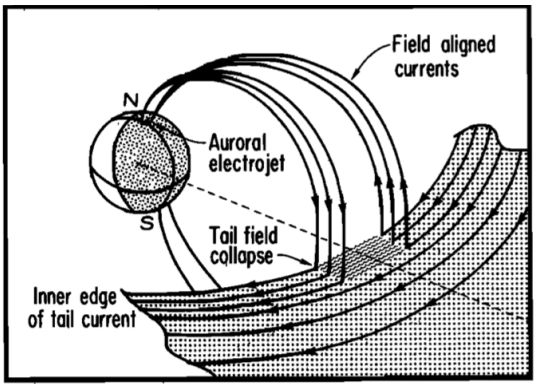

(a) From McPherron et al. [1973].

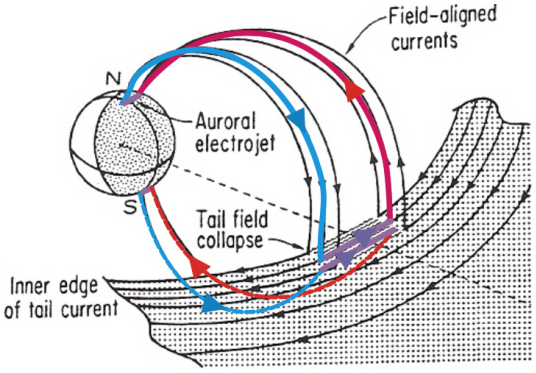

(b) From Ritter and Lühr [2008].

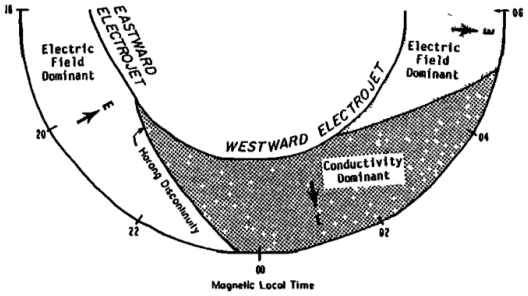

(c) From Kamide and Kokubun [1996].

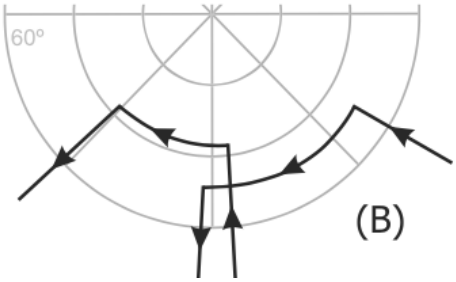

(d) From Gjerloev and Hoffman [2014].

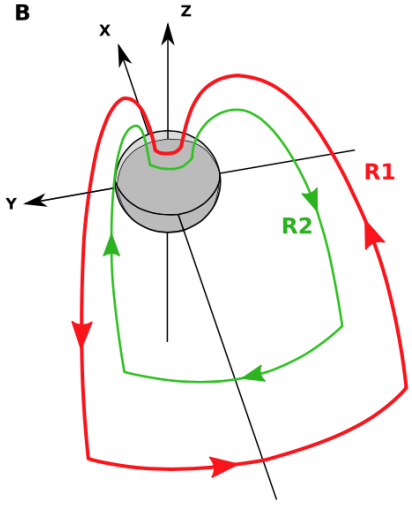

(e) From Sergeev et al. [2014].

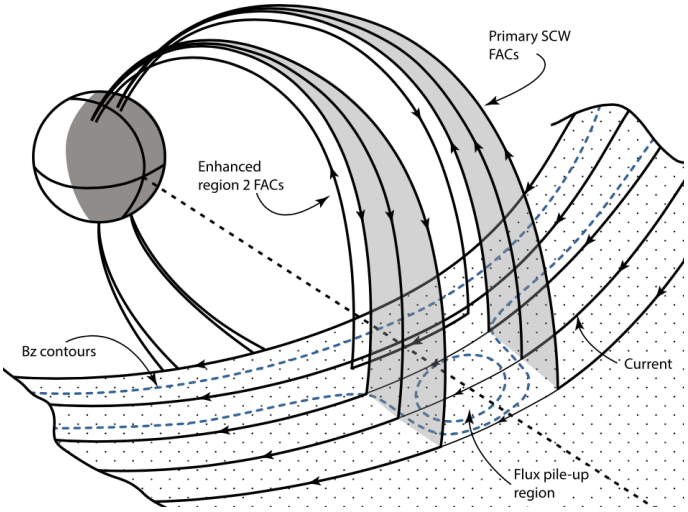

(f) From Kepko et al. [2015].

Figure 1.10: An example of six schematics of the proposed models for the large-scale SCW.

1996; Ritter and Lühr, 2008; Sergeev et al., 2011, 2014; Gjerloev and Hoffman, 2014; Birn and Hesse, 2014].

In the classic scenario [McPherron et al., 1973] (figure 1.10a) the ionospheric segment of the SCW is illustrated as a rapidly appearing, intense westward electrojet, located in the midnight region with postmidnight feeding and premidnight drainage. According to Ritter and Lühr [2008] (figure 1.10b) the observed substorm-related signatures could not be explained by a simple SCW but rather a current reconfiguration scenario combining four elements: i) a decrease of the tail lobe field, ii) a cross-tail current rerouting through the ionosphere, iii) low to mid latitude eastward ionospheric current, driven by R2 FACs and iv) a partial ring current connected to these R2 FACs.

Kamide and Kokubun [1996] (figure 1.10c) proposed a two component westward auroral electrojet, one directly driven (i.e. two-cell convection, DP2 pattern [Nishida, 1968]) and the other the unloading component in the midnight sector wedge (unloading of energy previously stored in the tail). Gjerloev and Hoffman [2014] (figure 1.10d) argue that a two-wedge current system better reflects the ionospheric current systems during an auroral substorm. This current system includes both a bulge current wedge located in the premidnight region, just equatorward of the open-closed field line boundary, and an oval current system postmidnight contained within the auroral oval. Sergeev et al. [2011] (figure 1.10e) suggested an additional R2 polarity FAC, at lower latitudes, be added to the classic SCW cartoon. Kepko et al. [2015] (figure 1.10f) updated the picture of the SCW to include a bounded dipolarisation region between the two original R1 type currents, as well as a typically small R2 type system, earthward of the dipoloarised region.

1.2.5 Mesoscale currents

Although there are many variations on the classic scenario of large-scale current systems [McPherron et al., 1973], most models do not focus on complex, mesoscale, azimuthal substructure. Whilst there has been observational evidence of this substructure [Forsyth et al., 2014], their contribution to the large-scale current have been largely debated with some even suggesting there is no large-scale system but simply many small "wedgelets" [Liu et al., 2018; Birn et al., 2019]. The idea of the SCW comprising of individual wedgelets has recently received significant attention [Birn and Hesse, 2014; Birn et al., 2019; Forsyth et al., 2014; Liu et al., 2015, 2018; Malykhin et al., 2018; Palin et al., 2016; Panov et al., 2016; Merkin et al., 2019]. A summary of the ideas and literature contributing towards the wedgelet scenario will be discussed in this section.

Dipolarising Flux Bundle (DFB)

DFBs are small (typically $< 3R_E$, [Liu et al., 2018]) magnetotail flux tube with a magnetic field which is significantly more dipolar than its background field [Liu et al., 2013]. The leading edge of a DFB shows a sharp enhancement of the north-south component of the magnetic field, B_z, which is known as the dipolarisation front (DF) [Nakamura et al., 2002]. The DF is $\sim 800 - 2000$km thick [Runov et al., 2009; Kepko et al., 2015] and contains a concentrated current sheet. This DFB propagates towards earth, at high speeds (> 100 km/s [Angelopoulos et al., 1992]), from the near-Earth reconnection site ~ 20 R_E downtail [Nagai et al., 1998; Angelopoulos et al., 2013] and becomes associated with a bursty bulk flow (BBF). Theoretically, the association between fast earthward flows and a contracting flux tube was first proposed by Pontius Jr and Wolf [1990]. The earthward motion of a flow burst causes deformation of the flux tube, but they retain their characteristic strong magnetic field and become increasingly more dipolar [Runov et al., 2009; Liu et al., 2014]. The DFBs are drivers of BBFs due to the strong dipolar magnetic field accelerating plasma earthwards [Fu et al., 2011; Liu et al., 2014; Gabrielse et al., 2012; Birn and Hesse, 2013]; they are responsible for $> 65\%$ of BBFs' flux transport [Angelopoulos et al., 1994]. It has been suggested by Liu et al. [2018] that each DFB serves as a "wedgelet", a building block for a large-scale SCW. As DFBs propagate earthward at high speeds, a thermal pressure gradient is established with the maximum pressure ahead of the DF. FACs are associated with this thermal pressure gradient [Liu et al., 2015].

The DFB with associated FACs corresponds to a current configuration that is similar to that of a SCW (see figure 1.11) and it is this configuration that is known as a wedgelet. The width of the FAC in the ionosphere is < 1 hour MLT [Liu et al., 2018]. The R1-sense sheet carries more current than the R2-sense sheet [Liu et al., 2018], corresponding to $\sim 7-36\times10^{-6}$ A/m^2 in the ionosphere [Nakamura et al., 2005; Snekvik et al., 2007; Sun et al., 2013; Forsyth et al., 2008]. For reference this is comparable to the 25×10^{-6} A/m^2 obtained from the ground-based magnetic field perturbations associated with an auroral streamer [Amm et al., 1999]. During substorm expansion, north-south aligned auroral forms (streamers) have been related to the fast bursts in the magnetotail plasma sheet that comprise BBFs [Sergeev et al., 2004; Nakamura et al., 2005; Forsyth et al., 2008]. Streamers are often longitudinally distributed indicating that wedgelets are distributed over several hours of local time [Zesta et al., 2006; Liu et al., 2013].

Another typical mesoscale auroral formation that has been associated with local current closure (wedgelets) are omega bands (see figure 1.12). They are associated with the dawn sector of the SCW and grow over several MLT hours during the recovery phase of the substorm. They are periodic, wave-like undulations of the poleward boundary [Kepko et al., 2015]. Amm [1995] found that $\sim 80\%$ of the FAC within a omega band

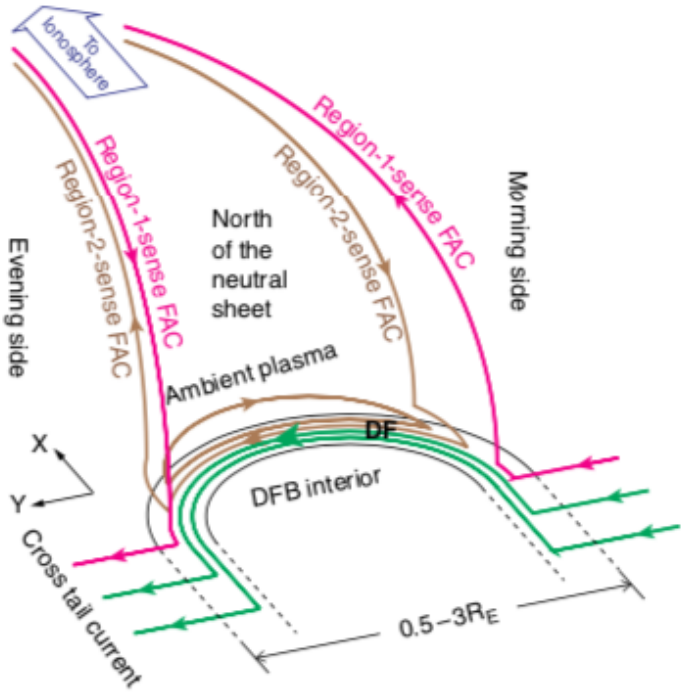

Figure 1.11: A possible closure scheme for the DFB current system. From Liu et al. [2018].

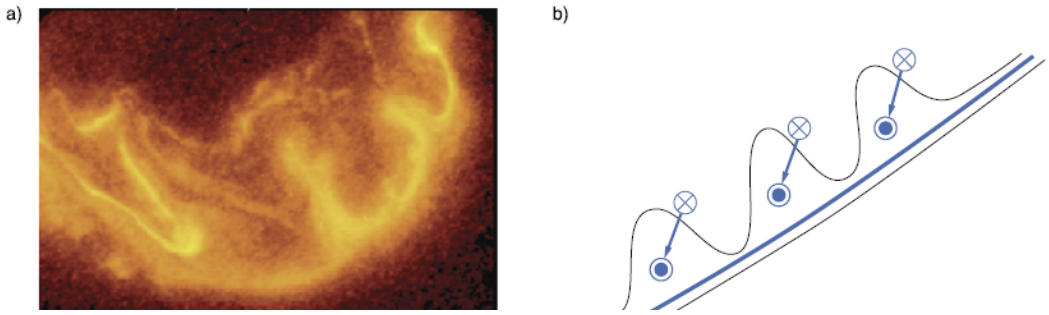

Figure 1.12: Omega bands (a) observed by Polar VIS (b) The structure of the FACs associated, and the geometry of the westward electrojet. From Kepko et al. [2015].

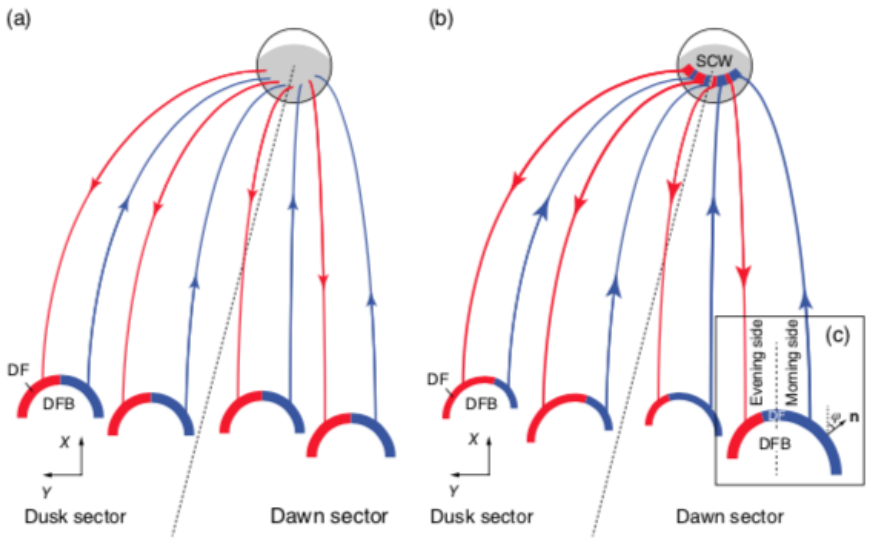

Figure 1.13: A global view from the Northern Hemisphere showing many DFBs (wedgelets) coexisting in the tail. (a) is the symmetric case (b) is the asymmetric case. From Liu et al. [2018].

was closed locally (i.e. did not contribute to the upward, dawn sector current at the end of the westward line current). The implication was that mesoscale sheets of FACs were diverting current from and towards the ionosphere at the scale of a few hundred km [Kepko et al., 2015].

Validity of the SCW consisting of DFBs

The typical length scale of wedgelets is approximately equal to that of the typical SCW width. On the other hand, wedgelets, which have oppositely directed FACs at dusk and dawn, are not approximately equal to the typical SCW magnitude. Side-by-side the upwardly directed FAC of one wedgelet will cancel out the downwardly directed FAC of it's neighbouring wedgelet, making the net current magnitude equal to that of a single wedgelet [Liu et al., 2018]. If the wedgelets are asymmetric in the magnitude of the upwards and downwards FACs, the net current would aggregate. Liu et al. [2015] further examined the asymmetric nature of R1-sense FACs in the DF layer and found the FACs in the dawn (dusk) sector of the tail to be more intense towards (away from) the earth and to cover a larger azimuthal range of the DF than the oppositely directed currents. The conclusion was that "an SCW is effectively a pair of large-scale, R1-sense FACs made up of small, asymmetric pairs of FACs"-[Liu et al., 2018]. See figure 1.13 for a schematic of the proposed substorm current configuration. The schematic shows four wedgelets but Liu et al. [2015] suggests there would be $\sim$ 10 wedgelets. This is based on a wedgelet having a net FAC of $\sim 0.1 \times 10^6$ towards (away from) Earth in the dawn (dusk) sector of the tail and the assumption that the R1-sense pair of a typical

SCW is $\sim 0.5 \times 10^6$ in each sector (dawn or dusk). According to Sergeev et al. [2014] the net SCW is typically $\sim 0.3 \times 10^6$ but Liu et al. [2018] has assumed that $20\% - 60\%$ of the current is contained within the R2-sense FAC pair.

Cluster observations of an event in 2010 by Forsyth et al. [2014] have been used as support for this ~ 10 wedgelets proposal. Liu et al. [2015] takes the ~ 30 individual current sheets seen by Cluster 1 and 4, consisting of a number of upward and downward currents as proof for wedgelets. Further, the observations of wider, upward currents in the dusk sector equated to asymmetric pairs of FACs [Liu et al., 2015]. Whilst accepting that most SCW models neglect the more complicated cross-tail azimuthal substructure, Forsyth et al. [2014] found, by comparing their observations to previous observations of BBF current systems, that their observations did not support the scenario in which BBFs create the small-scale substructure comprising the SCW. Forsyth et al. [2014] suggests that onset instabilities in the inner magnetosphere were the result of the azimuthal substructure.

Current closure

Whilst it is clear that a SCW contains some small/mesoscale structure [Forsyth et al., 2014; Amm, 1995; Zesta et al., 2006; Gjerloev and Hoffman, 2002], the scenario in which wedgelets form the SCW still raises questions. Firstly, how much of a DFB's FACs can eventually reach the ionosphere? The size of the current observed in the tail may be much greater than that which reaches the ionosphere. The depth of penetration depends on entropy reduction [Birn and Hesse, 2013; Dubyagin et al., 2011], so only the most depleted plasma tubes can reach the inner magnetosphere. Secondly, how do asymmetric wedgelets come to be? As of Liu et al. [2018] this was conjecture and no work had been done into determining whether such asymmetric flow shears actually existed.

A further unknown is how do asymmetric wedgelets currents close? The current (either downwards or upwards) is stronger in one direction and can not achieve current closure. Liu et al. [2018] suggests a possible global closure option (see figure 1.14). The current from the redundant FACs of a wedgelet, in one tail sector, will be closed by wedgelets in the other sectors, resulting in a net westward current in the ionosphere. In reality the ionospheric FACs are more complex than simple current sheet pairs [Forsyth et al., 2014] and DFBs have a short life span (a few ~ 1 minutes); they would need to emerge continually during the substorm expansion to match observations.

Alternatively, wedgelets may only contribute to a small section of the substorm; previous observations gave contradictory results on DFB occurrence during different substorm phases [Juusola et al., 2011; Chu, 2015]. According to Kepko et al. [2015] the cumulative

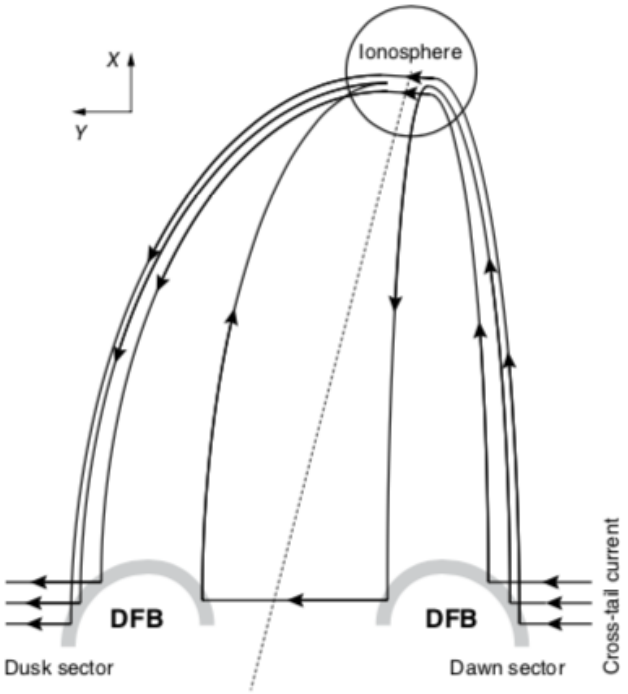

Figure 1.14: A proposed model for current closure for the wedgelet scenario. From Liu et al. [2018].

effect of the flow bursts modify the pressure distribution in the inner magnetosphere and this sustains the SCW. Birn and Hesse [2013] simulated multiple flow bursts at different local times which resulted in the high pressure, dipolarised region expanding azimuthally to resemble the SCW. To address the current closure problem Birn and Hesse [2014] introduced another current loop related to the intense DF current, so that part of the "disrupted" current may not be diverted to the ionosphere.

MHD simulations

Magnetohydrodynamic simulations of magnetotail dynamics have been used in an attempt to build a comprehensive view of the SCW, as it is difficult to obtain experimentally [Kepko et al., 2015]. MHD simulations numerically solve the MHD equations. Insight into the build-up of the currents is increased via simulations, which have produced details of the physical processes that could generate the current wedge and its three-dimensional configuration [Birn and Hesse, 1991; Scholer and Otto, 1991; Raeder et al., 2001; Birn et al., 2019]. Three dimensional quasi-local simulations of tail reconnection were first performed by Birn and Hones Jr [1981].

Scholer and Otto [1991] set the initial conditions as a two-dimensional equilibrium of the tail configuration, with solid wall boundary conditions. In the simulations, reconnection

was initiated from a fixed position in the centre of the plasma sheet by a localised resistivity. The results were that the shear and dipolarisation over a limited dawn-dusk extent was responsible for R1 FACs earthward and tailward of the neutral line. Raeder et al. [2001] noted the key result that substorm dynamics were dependent on the parameterisation of inherently kinetic processes and concluded that the ionosphere is not a "passive recipient of energy from the magnetosphere", but largely controls magnetospheric convection.

Investigations of dipolarisations and their relationship with BBFs have been used by many recent modelling efforts to understand the build up substorm currents and the radiation belts [Cramer et al., 2017; Sorathia et al., 2018; Ukhorskiy et al., 2018; Yang et al., 2015]. Birn and Hesse [2014] investigated the properties of current diversion associated with the SCW, with an emphasis on the current systems contributing to the SCW during various stages. Their conclusion was that flow channels distorted the magnetic field and drove a current system akin to a small SCW (wedgelet), but when these dipolarised fields accumulated near the Earth, the combined system eventually became dominated by a single wide wedge (SCW).

Nakamura et al. [2017] used their simulation as a comparison for their multi spacecraft observations and found that they accurately reproduced the overall characteristics observed. Recently, MHD simulations of magnetotail reconnection, flow bursts and dipolarisation have been used to quantitatively investigate the contributions of individual flow bursts, DFBs and rapid flux transport events in the build up of the SCW, as well as the transport of magnetic flux and energy into the inner magnetotail [Birn et al., 2019]. Birn et al. [2019] concluded there was no single, large current wedge (SCW).

1.3 Space Weather

The magnetospheric and ionospheric response to space weather have already been discussed in detail, but the solar coronal activity which result in enhanced terrestrial current flow have yet to be discussed. Space weather is a collective term to describe the variable conditions of the Sun and near-Earth environment that can influence the performance and reliability of space-based and ground-based technological systems. Major space weather events are regarded as a "low-probability, high-severity" hazard with few major events' time series data available for assessment [Oughton, 2018].

Recently, there has been growing recognition of this natural hazard. The UK released a "Space Weather Preparedness Strategy" in 2015 [Office et al., 2015] and added space weather to the National Risk Register of Civil Emergencies in 2017 [Office, 2017]. In 2016, the United States signed an executive order, titled "Coordinating Efforts to Pre-

pare the Nation for Space Weather Events" [House, 2016]. Although the study of space weather has progressed, there is still need for substantial efforts to further our understanding [Oughton, 2018].

1.3.1 What is Space Weather?

Space weather originates on the solar corona, where eruptive phenomena act as drivers to the solar wind. The three primary forms of solar activity with the potential to impact Earth are:

- Solar Flares - reconnection on the corona results in a rapid release of radiation emissions spanning the entire electromagnetic spectrum. They can release $> 10^{25}$ J of energy in several tens of minutes [Fletcher et al., 2011]. Initial radiation can reach the Earth within 8 minutes of the flare [Donnelly].
- Solar Energetic Particle Events (SEP) - a flux increase above the background noise with energy ranging from $\sim 10^4 - 10^{10}$ eV per nuclei. SEPs consist of electrons and protons but also heavy ions, and they can arrive at Earth within several hours to days [Papaioannou et al., 2018].
- Coronal Mass Ejections (CME)- The solar surface releases an average mass of $\sim 4 \times 10^{11}$ kg of plasma with an average energy of $\sim 5.4 \times 10^{22}$ J. They travel through interplanetary space at speeds of $10^3 - 10^4$ km/s and reach Earth in $1 - 4$ days [Shea and Smart, 2012; Gopalswamy, 2009].

High speed streams in the solar wind can also drive space weather events. Extreme space weather usually involves all three types of solar phenomena but the largest geomagnetic disturbances occur when a CME hits the Earth with a southward magnetic field direction [Webb and Howard, 2012]. The magnetic field in the CME is likely to lead to a geomagnetic storm that can last multiple days [Oughton, 2018].

1.3.2 Impacts of Space Weather

The aurora is a result of ionisation of neutral particles in the ionosphere during even modest IMF interactions with the terrestrial field, but geomagnetic storms result in the aurora extending to much lower latitudes. The Sun's magnetic activity follows an eleven year solar cycle [Baumjohann and Treumann, 1997]. There may be more solar activity during some parts of the solar activity, such as at solar maximum and in declining phase [Juusola et al., 2015], but solar eruptive phenomena can happen at any time. Figure 1.15 outlines how space weather impacts can extend from the Sun to national

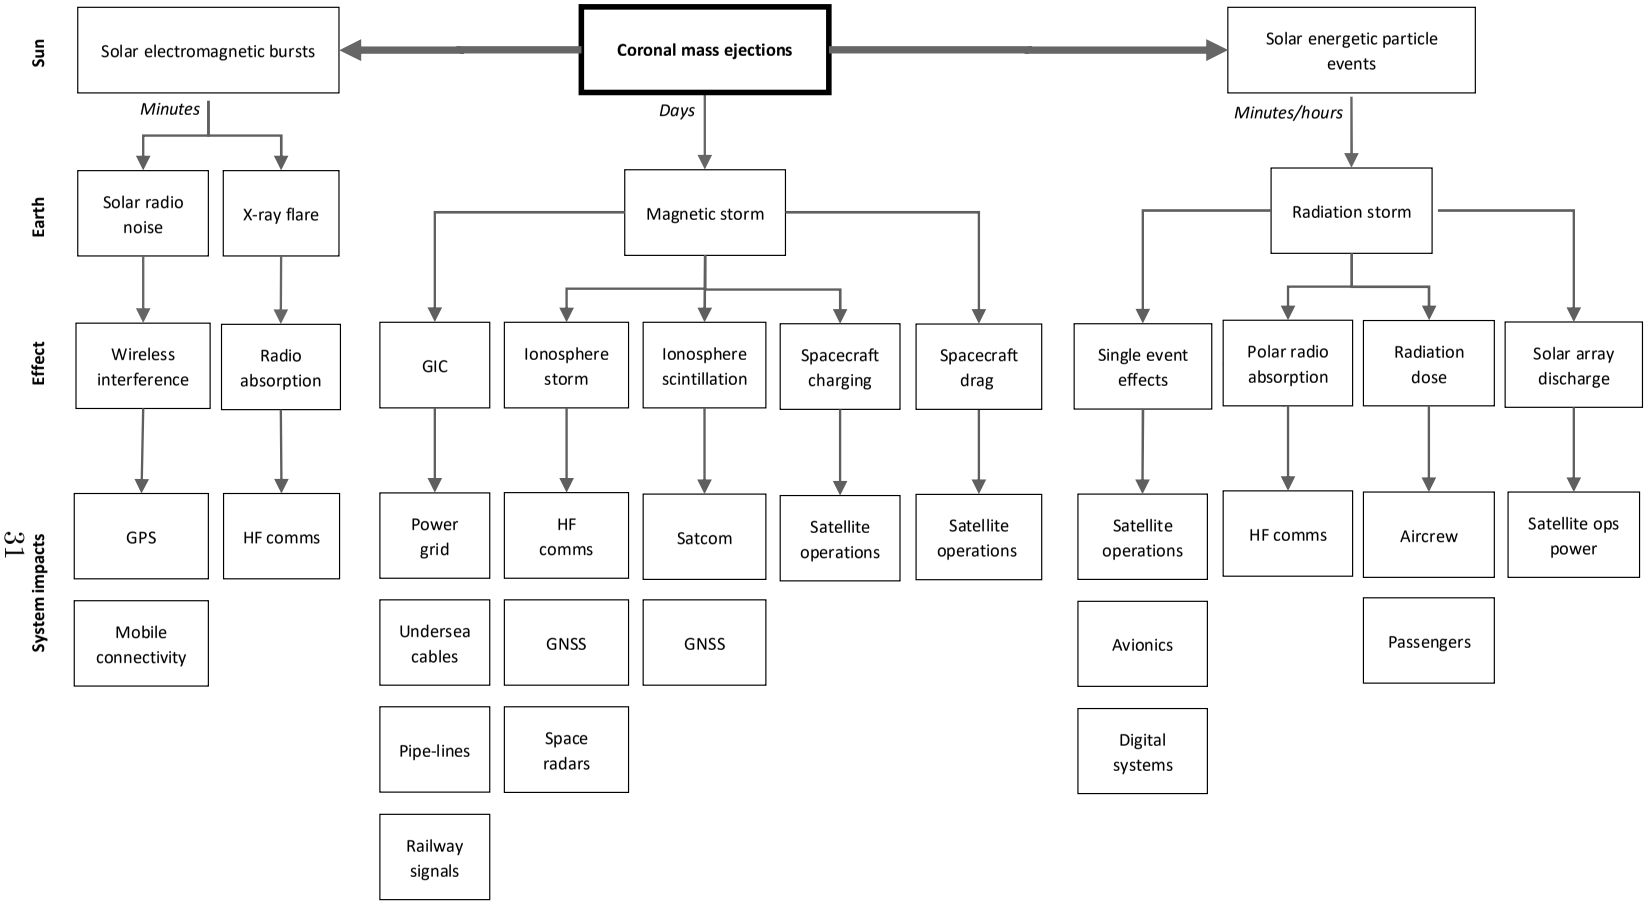

(Oughton 2018, adapted from Hapgood et al. 2012)

Figure 1.15: Space weather impact tree from Oughton [2018].

infrastructure. During a geomagnetic storm, geomagnetically induced currents (GIC) are generated that affect power grids and subject them to excessive thermal heating and voltage instability issues which can result in power loss. GIC can cause immediate or cumulative damage to transformer components [Hutchins and Overbye, 2011]. Rail transportation may be affected by GIC, impacting signalling equipment and tracks. Rail infrastructure can see an increase in the rate of corrosion and pipelines have an increased likelihood of corrosion due to cumulative long-term damage from GIC, however, evidence for long-term corrosion is sparse [Oughton, 2018].

Spacecraft and satellites, including those providing global positioning systems (GPS), are susceptible to space weather, particularly from radiation bursts [Astafyeva et al., 2014]. Short term effects are problems with signal propagation and transmission, causing disruption to other infrastructure systems and economic sectors that rely on communications, navigation and timing services [Oughton, 2018].

Long term affects include spacecraft drag and spacecraft charging. High-frequency (HF) radio communications can be disrupted by electromagnetic radiation bursts. During extreme events, aviation routes may have to be rerouted to avoid high latitudes where there may be HF communications disruption, radiation risk to passengers and crew, as well as problems with GPS navigation systems [Neal et al., 2013; Jones et al., 2005].

According to Oughton et al. [2019] a Carrington-sized 1-in-100-year event without forecasting capability could result in a GDP loss in the UK as high as £15.9 billion. With current forecasting capability this figure drops to £2.9 billion, but additional investment and enhanced forecasting could reduce the loss to £0.9 billion.

Chapter 2

Methods

Provided in the previous chapter was a summary of the physics needed to understand why advancement of space weather research is necessary. In this chapter we will focus on the analysis techniques relevant to this thesis. In section 2.1 we will describe the datasets used in later chapters. Section 2.2 will provide an overview of network analysis and section 2.3 a summary of canonical correlation, surrogates and thresholds.

2.1 Data

Sections 2.1.1, 2.1.2 and 2.1.3 describe the magnetometer observations, magnetic indices and polar VIS data, respectively. All three datasets are used to characterise geomagnetic activity in the auroral zones and will be used in chapters 3-5.

2.1.1 Magnetometer Data

The vast majority of the data used throughout this thesis is magnetometer data, which is freely available at `http://supermag.jhuapl.edu/`. SuperMAG is a worldwide collaboration of organisations that operate more than 300 ground-based magnetometers. The SuperMAG data is provided via an online platform and is made easy to use by rotating the ground magnetic field perturbations into the same coordinate system, with identical time resolution and a common baseline removal approach [Gjerloev, 2012]. The geomagnetic coordinate system is time dependent as the Earth's main magnetic field is constantly changing. SuperMAG rotates the horizontal components of the magnetic field into a local magnetic coordinate system [Gjerloev, 2012]. The magnetic field is given by $\boldsymbol{B} = (B_N, B_E, B_Z)$, where the N-direction is the local magnetic north, E-direction the local magnetic east and the Z-direction is vertically down.

As we explain below, the signal of interest represents a small enhancement on a time-varying baseline which must be removed [Gjerloev, 2012]. The measured magnetic field contains contributions from the Earth's main field, the Sq current system, the FACs, the ring currents, the auroral electrojets and the magnetopause currents. For the purposes of this thesis we are interested in the ionospheric response to space weather events, hence, the perturbations produced by currents flowing in and between the ionosphere and the magnetosphere are of interest, but all other magnetospheric sources should be removed. Whilst it is impossible to uniquely determine the sources of the measured field (Ampere's law, equation 1.7) it is possible to separate sources based on reasonable assumptions e.g. the Earth's main field is slowly varying compared to all other sources. Daily trends can be contributed to the Sq current system and yearly periodic changes are attributed to changes in the Earth's main field and seasonal temperature variations [Gjerloev, 2012].

The type and quality of magnetometers taking the measurements varies across the Earth's surface. Variometers are built and operated by small research budgets while observatories are highly sophisticated instruments built and operated under strict regulations. The cheaper variometers are sufficient for the purpose of studying the magnetospheric-ionospheric currents but we must account for the difference in sensitivity of instruments [Gjerloev, 2012]; this will be addressed in section 2.3.4.

2.1.2 Magnetic Indices

Magnetic indices are used throughout this thesis to compare the network results to a measure of the auroral electrojet. A magnetic index is an estimate of ground magnetic activity, used to quantify physical processes.

The auroral electrojet indices (AE, AU and AL) are used to describe the global magnetic signature of auroral electrojets in the northern hemisphere. They are based on 1 minute resolution readings of the magnetic north component, B_N, of twelve auroral zone (between 65° and 70° magnetic latitude) observatories. The measurements are normalized with a base value, B_{N0}, which represents the average value of B_N during the five days with the least magnetic activity during the month (quiet days) [Davis and Sugiura, 1966; Baumjohann and Treumann, 1997]. The maximum measured and baselined value of B_N, at each time. is then know as AU (upper envelope)

$$AU(t) = \max_i [B_N(t) - B_{N0}]_i, \qquad \text{for } 1 \leq i \leq 12. \tag{2.1}$$

The minimum is AL (lower envelope).

$$AL(t) = \min_i [B_N(t) - B_{N0}]_i, \qquad \text{for } 1 \leq i \leq 12. \tag{2.2}$$

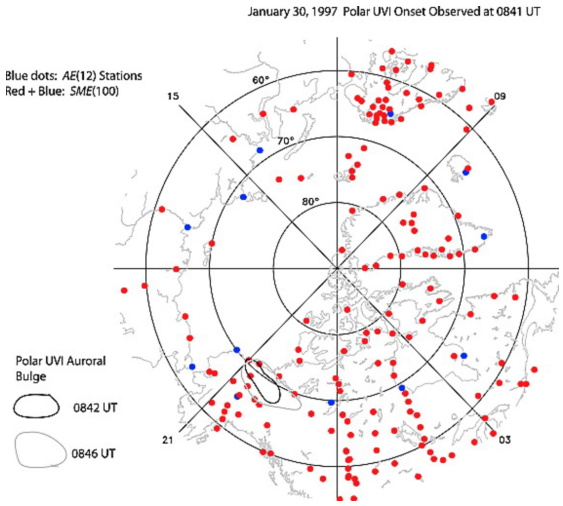

(a) The magnetometers included in the derivation of the AE (blue) and SME (red) indices. From Newell and Gjerloev [2011a].

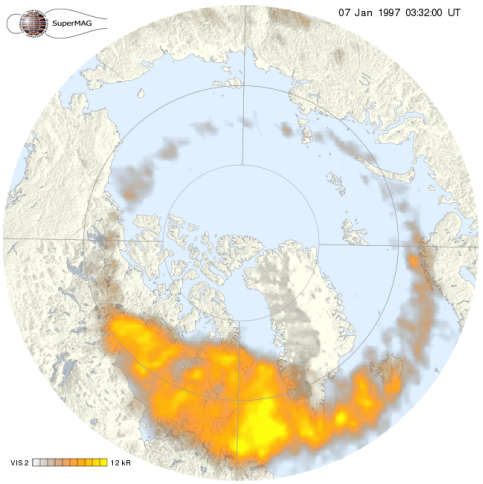

(b) Polar VIS of a substorm on 07/01/1997 at the peak of the substorm. From http://supermag.jhuapl.edu/.

AE is then the difference

$$AE(t) = AU(t) - AL(t). \tag{2.3}$$

SuperMAG have extended these indices to include $\sim$ 110 magnetometers, all between 40° and 80° latitude. Figure 2.1a indicates the magnetometer measurements included in the derivation of AE (blue) and SME (red). The number of magnetometers included in the index varies with the number available. The indices are based on SuperMAG data and, therefore, have undergone error correction, temporal resampling, coordinate system rotation and the baseline has been removed [Newell and Gjerloev, 2011a,b]. The derivation is based on traditional AE index [Davis and Sugiura, 1966] but is not equivalent because of these changes. SMU is the maximum value at each time and SML is the minimum value. SME = SMU − SML.

2.1.3 Polar VIS

The Visible Imaging System (VIS) was flown on the Polar Spacecraft and the catalogue of images has been made available. The camera was equipped with a broad-band filter at far ultraviolet wavelengths and provided wide field-of-view images of observations of Earth [Frank et al., 1995]. Figure 2.1b is an example of the polar VIS image taken during a substorm on 07/01/1997. We have used these images throughout this thesis to provide timings for substorm phases [Gjerloev et al., 2007; Gjerloev and Hoffman, 2014] and comparison to network topology [Orr et al., 2019].

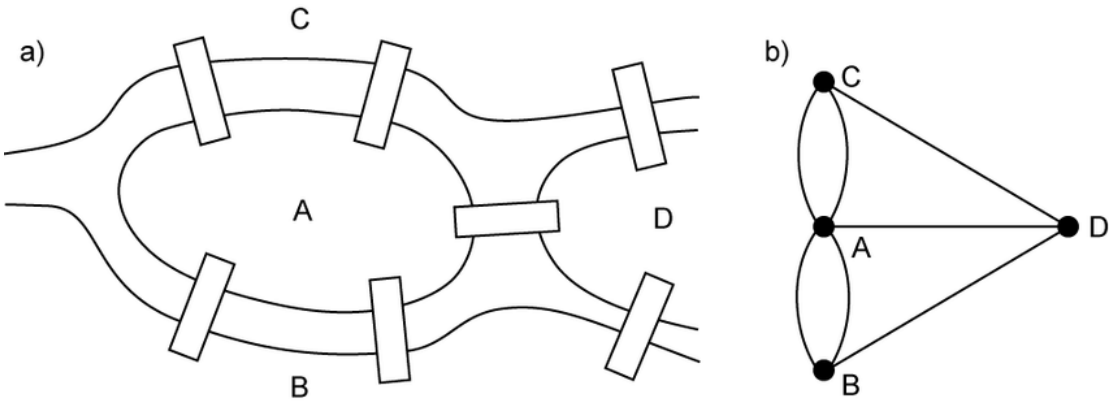

Figure 2.2: The left is a schematic map of Königsberg in 1736. On the right is the resulting graph. The parts of the city are represented as vertices and the bridges are edges.

2.2 Network Analysis

Analysis in chapters 3-5 involves representing magnetometer vector magnetic field perturbation time series as a network and using network parameters to characterise and compare substorm events. In this section a summary of network analysis is provided, along with some useful definition and techniques that will be used in later chapters.

2.2.1 Introduction to graphs

Modern graph theory traces back to 1736 and the problem of the Königsberg bridges (see figure 2.2a for a schematic of the town centre). The problem was to find a route which crossed each bridge once and only once. Euler [1736] simplified the problem by drawing a graph (figure 2.2b), where the different parts of the city (A, B, C and D) became *vertices*, and the bridges between them became connecting lines called *edges*. The key to the problem is in the number of edges per vertex, the *degree*. If the degree of a vertex is even (there are an even number of edges attached to the vertex), one can enter through an edge and leave through another. If the degree of a vertex is odd (there are an odd number of edges attached to the vertex, like all the vertices in figure 2.2b), the only way to cross each edge once, and only once, is if that vertex is the start or end point, and all other vertices have an even degree. If the start and end vertices coincide, no vertex can have an odd degree. Therefore, the problem of crossing each bridge in Königsberg, once and only once has no solution [Caldarelli, 2007; Biggs et al., 1986].

The field of graph theory advanced and evolved into the field of complex networks. Networks have become an important mathematical tool in the analysis of complex systems and have been developed throughout the twentieth century, particularly in the

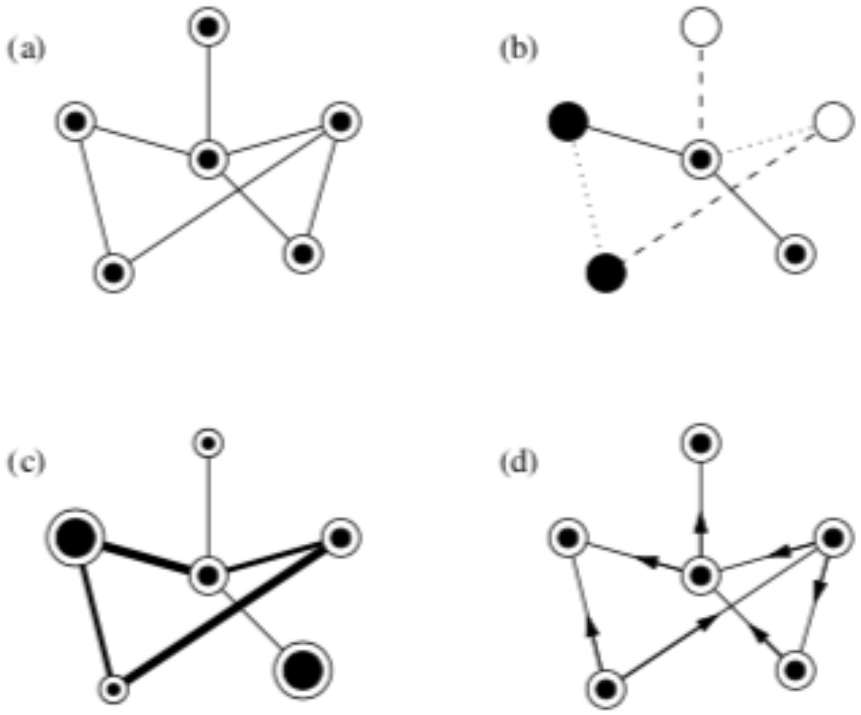

Figure 2.3: Example of the different types of networks: (a) Undirected network (b) a network with a number of discrete vertex and edge types (c) a weighted network (d) directed network. From Newman [2003].

social sciences [Watts and Strogatz, 1998; Newman, 2003; Albert and Barabási, 2002]. Other examples of networks include the Internet, the World Wide Web, neural networks, metabolic networks, food webs, networks of citations between papers e.t.c. [Newman, 2003].

With the recent advance of computer power, networks have developed from small-scale systems, that focused on the properties of individual nodes, to the study of large-scale statistical properties. Previously it was possible to draw a network of tens of nodes and gain understanding of their structure, now networks can contain over a million nodes and analysis is difficult. Statistical methods for quantifying large networks aid describing the network structure [Newman, 2003].

2.2.2 Network basics definitions

Network types

Figure 2.3 shows examples of the different types of networks [Newman, 2003]. In their simplest form, networks are a set of vertices/nodes joined by edges/connections (figure 2.3a) but networks can be more complex. Network nodes or edges may be of different

types, or have different properties (figure 2.3b). For example, in a social network of people, the vertices may represent different genders, nationalities, ages, incomes etc. The edges could represent friendship, professional acquaintance, animosity or geographical proximity. Network edges can be weighted (figure 2.3c), representing how often two people meet, or how much prey a predator eats. Figure 2.3d is a directed network, where an edge has an associated direction. Examples of this type of network include telephone calls or emails, as messages travel in one direction.

Adjacency matrix

An adjacency matrix, $\boldsymbol{A}$, represents the structure of the graph in an $n \times n$ matrix, where n is the number of nodes in the network. Each entry in the matrix, a_{ij}, will be 1 if vertices i and j are connected (have an edge between them), and 0 otherwise. For an undirected graph the adjacency matrix will be a symmetric matrix (a_{ij}=a_{ji}), which generally has the diagonal elements equal to zero (a_{ii}=0), as we do not consider a node to be connected to itself (a 'loop'). The maximum number of possible edges in an undirected graph is $n(n-1)/2$.

For a directed graph, the matrix is no longer symmetric as a_{ij} represents the directed edge from $i \to j$ and a_{ji} represents the directed edge from $j \to i$; i.e. a_{ij} no longer needs to equal a_{ji}. The maximum number of possible edges in an directed graph is $n(n-1)$, if we allow nodes to be connected in both directions ($i \to j$ and $j \to i$), or $n(n-1)/2$, if the vertices can only be connected in one direction.

For a weighted graph, a_{ij} will be the real number representing the weight to the entry, or $a_{ij} = 0$, where there is no edge [Caldarelli, 2007].

Degree

The degree of a vertex is its number of associated edges. The sum of all the degrees in an undirected graph is twice the number of the edges in the graph. The degree, k_i, of a vertex, i, is given by $k_i = \sum_{j=1}^{n} a_{ij}$.

For a directed network there is an 'in' degree and an 'out' degree because generally $a_{ij} \neq a_{ji}$. Likewise, where edges contain directional information, such as east or west, there is both an east degree and a west degree [Dorogovtsev and Mendes, 2013].

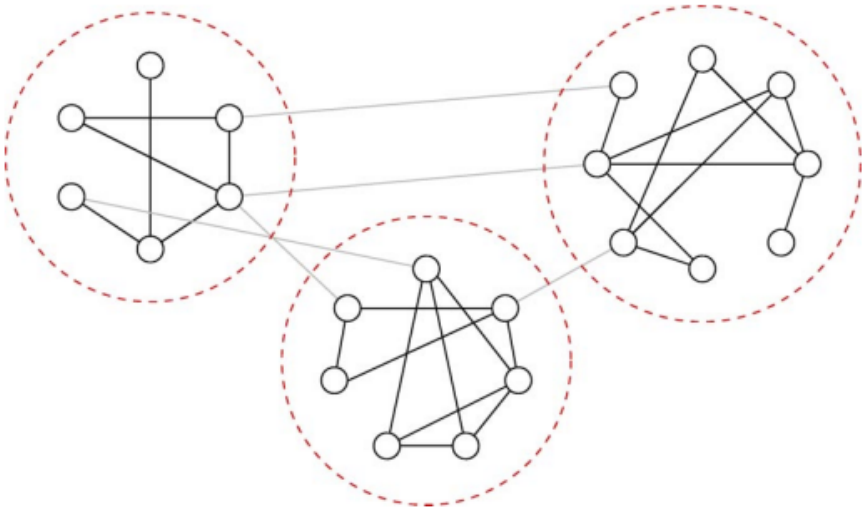

Figure 2.4: An example of a small network with community structure. There are three communities (in dashed circles), with dense internal links, but between which there is a much lower density of external links [Newman and Girvan, 2004].

2.2.3 Communities

In chapter 5 community detection is used to characterise isolated substorms. Real world complex networks are often globally sparse but locally dense, containing groups of vertices called *communities* [Caldarelli, 2007]. Community structure is a property common to many networks, where the network nodes are divided into groups within which the network connections are dense (vertices are highly interconnected), but between which they are sparser (see figure 2.4 for an example).

The ability to find and analyse such groups may provide insight into understanding and visualising the structure of networks [Newman and Girvan, 2004]. Communities not only characterise the graph topologically, but provide information both on the formation of the network and its functionality [Caldarelli, 2007].

In social networks, communities reflect physical communities, e.g. children are likely to have many social interactions within their schools and few between neighbouring schools. In a metabolic network, communities reflect the biological functions of the cell [Ravasz et al., 2002]. In a road network, there will be a high density of roads within cities and a low density of roads between the cities.

There are many algorithms to find if, how and why a network divides into communities. Community structure detection algorithms locate dense subgraphs in directed or undirected graphs, by optimising some criteria, and usually using heuristics [Caldarelli, 2007]. In chapter 5 we use the following six community detection algorithms.

- Edge Betweenness- This algorithm calculates the "edge betweenness", a measure of the shortest paths of the graph, the edge with the highest score is removed and the process repeated [Newman and Girvan, 2004]. The method is a divisive method. It looks for the edges in the network that are responsible for connecting many vertex pairs, the edges that are the most 'between' other vertices, removes them, and repeats this process to divide the network into smaller and smaller components.
- Walk Trap- Based on the idea that short random walks become 'trapped' within the densely connected communities [Pons and Latapy, 2005]. The algorithm uses a measure of similarity between vertices, based on random walks. It can be computed efficiently and it captures the community structure in a network.
- Optimal- This algorithm calculates the optimal community structure of a graph by maximising the modularity measure over all possible partitions [Brandes et al., 2007]. Further information on this method is provided in the following subsection.
- Information Map- Finds the community structure that minimises the expected description length of a random walker trajectory [Rosvall and Bergstrom, 2007].
- Propagating Labels- Labels the vertices with unique labels and updates them by majority voting in the neighbourhood of the vertex [Raghavan et al., 2007].
- Leading Eigenvector- The function locates densely connected subgraphs in a graph, by calculating the leading non-negative eigenvector of the modularity matrix of the graph [Pons and Latapy, 2005].

The edge betweenness community detection is the main algorithm that will be used in chapter 5 and it may be summarised as follows: consider a network comprised of several communities of nodes with a few connections between them. If we consider all possible shortest paths needed to travel between one node and another, the edges that connect between communities will carry the majority of these paths. These edges will have the highest edge-betweenness. The edge-betweenness algorithm [Newman and Girvan, 2004] identifies and successively removes these edges which have the highest edge betweenness, leaving behind sub-networks that are the individual communities. It can be argued that this is an unbiased algorithm which does not predetermine how many communities the network should divided into nor prescribe any other parameter, enabling multi-scale community structures to be recovered by varying the number of edges removed.

Modularity

Modularity is a measure that has been widely used to evaluate how well the community structure has been captured [Clauset et al., 2004; Pons and Latapy, 2005]. To mea-

sure the robustness of the communities formed, the network modularity measures how separate the nodes within the different communities are [Newman, 2004; Newman and Girvan, 2004].

The modularity is expressed as follows. We first divide the network into w communities which defines a $w \times w$ symmetric matrix, $\mathbf{e}$, whose elements e_{xy} are the fraction of all edges in the network that link nodes in community x to nodes in community y. Then e_{xx} is the fraction of the network contained within community x and e_{xy} is the fraction of the network that connects between communities x and y. The fraction of network edges that connect nodes in community x to the rest of the network is $f_x = \sum_{y=1}^{w} e_{xy}$. The modularity, Q, is then given by:

$$Q = \sum_{x=1}^{w} (e_{xx} - f_x^2). \tag{2.4}$$

Generally speaking, modularity is a measure of how many connections are within a region compared to between the regions. It indicates if one division into subgraphs is better than another [Newman, 2003]. The modularity can take any value between negative and positive one, where positive indicates the presence of communities and $Q > 0.3$ is a 'good indicator of significant community structure' [Clauset et al., 2004]. Not all values ($-1 \leq Q \leq 1$) will be available for every combination of nodes and edges but modularity can be compared to that of a randomly generated network as a test of significance. Figure 2.5a shows some possible community partitions with the analogous dendrogram in 2.5b and the corresponding modularity, Q, in 2.5c showing that the best partition consists of two communities [Pons and Latapy, 2005].

The 'optimal' algorithm, which chooses communities based on the maximum modularity, can now be discussed in greater detail. The modularity equation (equation 2.4) reveals an inherent trade-off. The first term is maximised by many edges within a small number of communities, the second is minimised by splitting the graph into many communities, each containing few edges. A clustering with maximum modularity has no community that consists of a single node with a degree of 1. Additionally, a clustering with maximum modularity does not include disconnected communities. These rules mean that if an additional node is connected, the optimal clustering is completely different (see figure 2.6a-b). This property is called non-locality, i.e. a local change affects the global community structure. Figure 2.6c-d shows how adding extra nodes may affect the global community structure. By doubling the graph from figure 2.6c the optimal clustering is completely altered [Brandes et al., 2007].

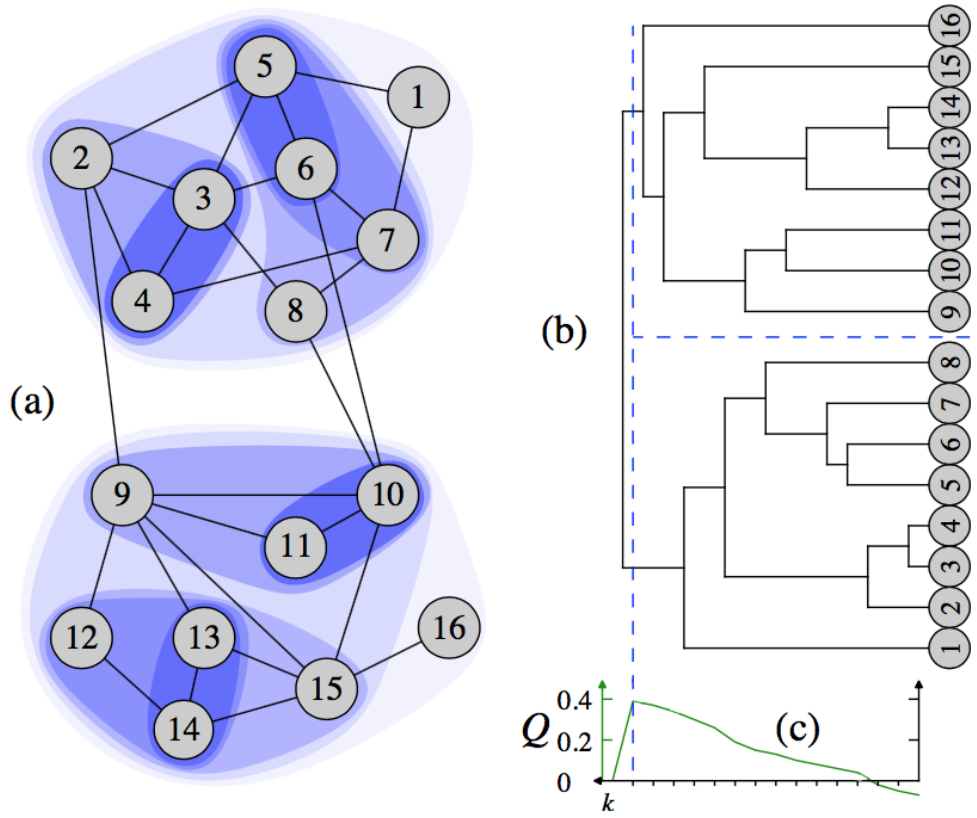

Figure 2.5: (a) An example of community structure found by the walk trap algorithm using random walks of length $t = 3$. (b) The stages of the algorithm encoded as a hierarchical tree (dendrogram), with the blue line marking the community divide with the highest modularity, Q (c) shows the modularity, Q, with the x-axis showing the number of divisions. Adapted from Pons and Latapy [2005].

2.3 Methods and Analytic techniques

We have described the data and network analysis but several analytical techniques are required to calculate a network from the magnetometer data. A measure of association is required to establish an 'edge' between the magnetometer nodes. The canonical cross correlation (CCC) of the magnetic perturbations measured at a magnetometer pair is the method used to decide if the pair are connected. Only the magnetometer pairs with CCC above their station and event specific threshold will be included in the network. Possible methods of establishing an association between magnetometer observations, are described in section 2.3.1 before canonical correlation is described in section 2.3.2 and the method for calculating the threshold is outlined in section 2.3.4. Surrogates are a method of checking statistical significance and will be described in section 2.3.3.

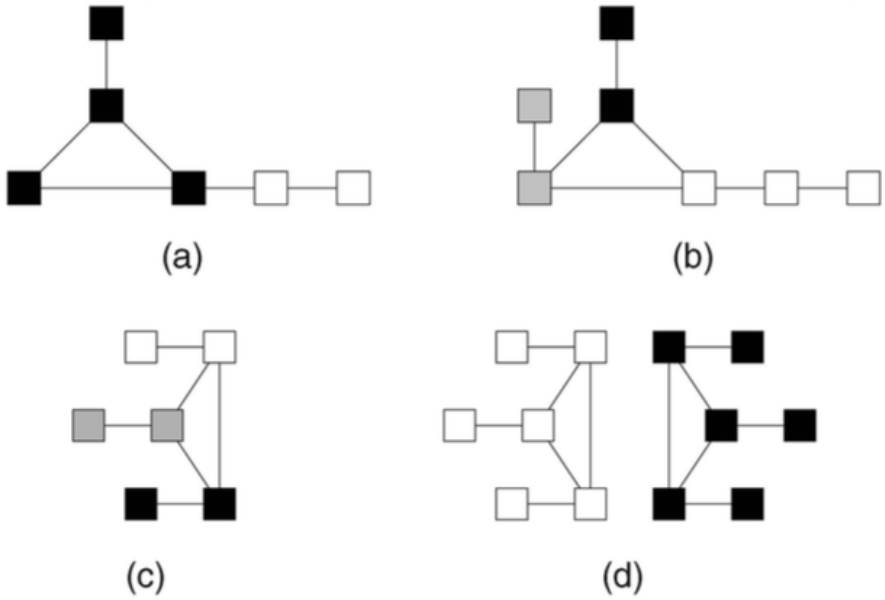

Figure 2.6: (a) and (b) show non-local behaviour, i.e. the addition of a single node changes the entire community structure. (c) and (d) show scaling behaviour. Communities are represented by colour. From Brandes et al. [2007].

2.3.1 Common techniques for quantifying association between datasets

The studies in this thesis required a method of analysis that allowed quantitative characterisation of the vector magnetic field perturbation time series. The Pearson correlation coefficient, r_p, is a measure of the strength of the linear association between two variables [Sedgwick, 2012]. If we have two sets of variables, x and y, the Pearson correlation coefficient, r_p, is given by,

$$\mathrm{Corr}(x, y) \equiv r_p(x, y) = \frac{\mathrm{Cov}(x, y)}{\sqrt{\mathrm{Var}(x)}\sqrt{\mathrm{Var}(y)}}, \tag{2.5}$$

where $\mathrm{Cov}(x, y)$ is the covariance between x and y, and $\mathrm{Var}(x)$ is the variance of x [Krzanowski, 2007]. If we were to plot the two variables against one another on a scatter graph, r_p is a measure of how close the points are to a straight line (linear relationship). The Pearson correlation coefficient can take values between $-1 \leq r_{\leq} 1$ where negative values correspond a negative association between the variables and positive values to a positive association. $r_p = \pm 1$ implies a perfect linear relationship whilst $r_p = 0$ implies there is no discernible linear association [Krzanowski, 2007]. Pearson correlation coefficient could be used, but the data would have to be reduced to one variable, such as one component of the field or simply the magnitude.

Pearson's correlation coefficient is sensitive to outliers in the datasets. Spearman's rank-order correlation is another method that is considered a more appropriate method when the data contains outliers [Chok, 2010]. It is a nonparametric (distribution-free) rank-based version of Pearson's correlation that compares the relationship between two variables to that of a monotonic function [Hauke and Kossowski, 2011]. In mathematics a monotonic function is one which has either a fully positive or negative first derivative [Stover]. Unlike Pearson's correlation coefficient it does not require the relationship between the variables to be linear. Spearman's rank-order correlations first ranks the two sets of variables from 1 to n where n is the number of data points. d is the difference between the ranks of each pair of variables. Spearman's correlation coefficient is then given by,

$$r_s = 1 - \frac{6 \sum_{i=1}^{n} d_i^2}{n(n^2 - 1)} \tag{2.6}$$

Mutual information (MI) is a quantity that measures the relationship between two random variables that are sampled simultaneously. It measures how much information one random variable contains about the other variable [Cover and Thomas, 2012]. If the mutual information between two variables is zero, the two variables are statistically independent [Learned-Miller, 2013]. MI can be used for both linear and non-linear statistical dependence but the method for assigning correct joint probability distributions can be highly data intensive. Mutual information is therefore non-trivial to estimate for quantitative variables whereas correlation coefficients are ideally suited [Song et al., 2012]. Correlation coefficients typically require smaller sample sizes than MI and MI estimation requires certain parameter choices such as the number of bins. Tests of statistical significance, such as the p value, are computationally challenging for MI methods whereas they are quick for correlation. Song et al. [2012] tested various correlation methods against MI as methods of association between network nodes when dealing with gene expression data. They found than correlation outperformed MI in measuring association between pairwise genes.

Fourier decomposition methods are unsuitable as they identify harmonic components but the global response to space weather events is more impulsive than harmonic [Jackel et al., 2001].

Wavelets are mathematical tools for analysing time series with different time and frequency localised processes [Torrence and Compo, 1998]. They are a relatively new concept compared to Fourier analysis, but wavelets can capture temporal (or spatial) variability in spectral character. Where Fourier assumes constant frequency content of a data set in time, wavelets allow for changes and as such, have much better temporal (or spatial) resolution [Cooper and Cowan, 2008]. One method in which time series can be compared using wavelets is cross-wavelet transforms [Torrence and Compo, 1998]. They use a measure of phase correlation between two datasets to show similarities and

differences, but do not consider amplitude. This means there is little sensitivity to noise but large deviations in amplitude do not affect the results. More recently Cooper and Cowan [2008] expanded on this method, describing their work as "semblance filtering"- comparing two datasets based on correlations between their phase angles, as a function of frequency. Decowski and Li [2015] compared two time series of unequal length, using a new test statistic, based on the wavelet transform.

2.3.2 Canonical Correlation Analysis

All methods listed in section 2.3.1 are viable methods of association, but the multivariate nature of magnetometer data (three component vector time series) requires more sophisticated methods [Jackel et al., 2001]. According to the Biot-Savart Law (equation 1.35, the ground magnetic field response to an overhead current, e.g. the auroral electrojets, will be experienced in different components of the field depending on where the magnetometer is in comparison to the current. For example, a magnetometer directly below the east-west directed current system will experience a change in the north-south component, whereas a magnetometer equatorwards and one polewards will have an oppositely directed change in the vertical component [Jackel et al., 2001].

For this reason it is important to use a method of association which considers the three components of the vector as a group, and not as three individual sets of variables. Canonical Correlation Analysis (CCA) is a method of finding the associations between two vector data sets. [Krzanowski, 2007] describes it as 'determining the associations that exist between the *groups* rather than between individual pairs of variables'. It quantitatively explains the variability and similarity between sets. It was first used in Jackel et al. [2001] with respect to magnetometer station pairs. It was used by Dods et al. [2015, 2017] to quantify the spatio-temporal similarities between magnetometer measurements and as the first step for network construction. CCA allows all three vector magnetic field perturbations components to be considered as a group rather than individual pairs.

CCA is a method of measuring the linear relationship between two multidimensional variables. It was first proposed by Hotelling [1933] and has been widely used since. CCA addresses the problem of finding basis vectors for two sets of variables, such that the Pearson's correlation coefficient between the projections of the variable on to these basis vectors is maximised. CCA is applied to two sets of random variables $\boldsymbol{X} = (x_1, x_2, ..., x_k)'$ and $\boldsymbol{Y} = (y_1, y_2, ..., y_l)$ to determine a set of linear combinations of the two, $U_i = \boldsymbol{a_i'X}$ and $V_i = \boldsymbol{b_i'Y}$, where $i = 1, ..k$ provided $k \leq l$. The first canonical correlation component, r_1 by calculating Pearson's correlation coefficient (equation 2.5) between U_1 and V_1, r_1=Corr(U_1,V_1).

U_i and V_i, the set of linear combinations known as the canonical variable, must obey the following conditions:

1. Mutually uncorrelated and have unit variances i.e. $\text{Var}(U_i) = 1$ and $\text{Var}(V_i) = 1$, and $\text{Cov}(U_i, U_j) = 0,\ \forall\ i \neq j$.

2. U_1 and V_1 are the linear combination of $\boldsymbol{X}$ and $\boldsymbol{Y}$ that have the maximum Pearson correlation correlation, while U_2 and V_2 have the second highest correlation, yet are uncorrelated with U_1 and V_1. More generally speaking:

$$r_i = \text{Corr}(U_i, V_i) \geq 0 \qquad \& \quad r_1 \geq r_2 \geq ... \geq rk \geq 0$$

where $\text{Corr}(U_i,U_j)= 0$, $\text{Corr}(V_i,V_j)= 0$ and $\text{Corr}(U_i,V_j)= 0\ \forall i \neq j$.

U_i and V_i are the canonical variables and r_i the ith canonical correlation components between $\boldsymbol{X}$ and $\boldsymbol{Y}$ [Reinsel, 2003].

Calculating U_i and V_i first requires the calculation of $\boldsymbol{a_i}$ and $\boldsymbol{b_i}$ which can be done as follows: Let $\sum_{\boldsymbol{XY}} = \text{Cov}(\boldsymbol{X}, \boldsymbol{Y})$ be the covariance matrix of $\boldsymbol{X}$ and $\boldsymbol{Y}$, and $\boldsymbol{C_X} = \sum_{\boldsymbol{XX}}^{\boldsymbol{-1}} \sum_{\boldsymbol{XY}} \sum_{\boldsymbol{YY}}^{\boldsymbol{-1}} \sum_{\boldsymbol{YX}}$. The eigenvalues and vectors can then be found to satisfy:

$$(\rho_i^2 I \ - \ \boldsymbol{C_X})\ \boldsymbol{a_i} = 0,\ \forall i,$$

where ρ_i^2 and $\boldsymbol{a_i}$ are the ordered eigenvalues and eigenvectors of $\boldsymbol{C_X}$. The vectors $\boldsymbol{a_i}$ are normalized such that $\boldsymbol{a_i}' \sum_{XX} \boldsymbol{a_i} = 1,\ \forall i$. The canonical variables are then calculated as $U_i = \boldsymbol{a_i'X}$.

Similarly,

$$(\rho_i^2 I \ - \ \boldsymbol{C_Y})\ \boldsymbol{b_i} = 0,\ \forall i,$$

The eigenvectors $\boldsymbol{b_i}$ are normalized such that $\boldsymbol{b_i}' \sum_{YY} \boldsymbol{b_i} = 1,\ \forall i$ and the canonical variables are then calculated as $V_i = \boldsymbol{b_i'X}$. It is also given by $\boldsymbol{b_i} = \rho_i^{-1} \sum_{YY}^{-1} \sum_{YX} \boldsymbol{a_i}$ [Reinsel, 2003].

The need for the canonical variables, U_i and V_i to be mutually uncorrelated and uncorrelated with every other canonical variable, V_j and U_j, unless $i = j$, means the correlation coefficient r_i is the 'pure association' between the two groups. This means the quantification of the correlation is only 'between' the groups because the correlations 'within' the groups have been removed [Krzanowski, 2007].

Canonical correlation is used as a measure of similarity between the vector magnetic field perturbation time series measured at a magnetometer pair throughout this thesis.

The Matlab code 'canoncorr' was used to calculated the associations [MATLAB, 2018]. A window of length w minutes will be linearly detrended to remove the best-fit line before calculation of the canonical correlation component, r_1. r_1 takes values between 0 and 1 so does not contain directional information. The relative direction between magnetometers could be obtained from the eigenvectors $\boldsymbol{a_i}$ and $\boldsymbol{b_i}$ but this is not the focus of this thesis.

2.3.3 Surrogates

Geophysical data tend to exhibit a red noise background, where the power spectral density has the form $1/f^2$ [Lancaster et al., 2018]. Jackel et al. [2001] found red noise to be the best approximation for magnetometer observations. Therefore we first need to test whether the dynamics of the system are a consequence of purely random red noise, or if the system evolves in a non-random (deterministic) manner [Lancaster et al., 2018].

Surrogate data testing is a proof by contradiction technique, where a particular property of the data, known as discriminating statistic, is used to compare properties of the original time series to those of a surrogate set. Ideally the surrogate preserves all properties of the original data except one- here, the temporal coherence i.e. the surrogate constructs the network that could occur 'by chance'. Surrogate testing ensures that the results are a true characteristic of the underlying physical system [Lancaster et al., 2018]. Any nonlinear statistic that characterises a time series with a quantifiable parameter can be used to compare the original data to the surrogate data [Schreiber and Schmitz, 1996]. A null hypothesis can then be defined to state that the original data is fully described by the surrogate data, for example the original data is correlated red noise. Rejecting this null hypothesis implies that it is unlikely the original data is just red noise [Lancaster et al., 2018].

There are many ways to construct surrogate data. A random shuffle will simply make a surrogate from a random permutation of the original series [Engbert, 2002]. This random shuffle will destroy the temporal structure of the original signal, but maintain the same mean, variance and histogram distribution. It is a simple method of checking whether the data is just uncorrelated noise, but is an inadequate test for systems with correlated variables e.g. red noise [Lancaster et al., 2018].

Surrogates are needed to maintain the same auto correlation function of the original data but with removed nonlinearity. This would allow surrogates to test the null hypothesis that data was generated by a noisy linear process and can be fully characterized by linear characteristics (such as mean, standard deviation, and autocorrelation). If the null hypothesis cannot be rejected, the data are assumed to be linear and are therefore generated by a linear set of equations driven by noise [Lancaster et al., 2018]. Fourier

transform based surrogates are generally used as autocorrelation can be preserved by preserving the power spectrum.

Amplitude adjusted Fourier transform (AAFT) attempts to preserve the distribution of the original data, by preserving the power spectrum and the amplitude distribution. First, the method scales the data to a Gaussian distribution, then performs a random phase transformation (makes surrogate time series which have the same Fourier spectrum as the re-scaled data) before inversely transforming to de-gaussianise the data and returning it to its original amplitude distribution. Iterative amplitude adjusted Fourier transform (IAAFT) is an improvement of the AAFT to achieve a closer match between the power spectrum and amplitude distribution of the surrogate and original data [Schreiber and Schmitz, 1996; Theiler et al., 1992]. It iteratively fits the amplitudes at each step. At each step the derived time series are reordered until convergence of both the spectral density and amplitude distributions are reached [Kugiumtzis and Tsimpiris, 2010]. The accurate preservation of the distribution and power spectrum means there is a small spread in surrogate values, hence fewer surrogates are required to determine significance [Lancaster et al., 2018].

Once the surrogate data is calculated it is important that the original data and each surrogate are treated in exactly the same way [Lancaster et al., 2018]. Surrogates are used in chapter 3 and 5 to distinguish statistical differences from a IAAFT surrogate.

2.3.4 Thresholds

As discussed in subsection 2.1.1, the type and quality of the magnetometers taking the measurements varies. Additionally, the magnetic field perturbations are affected by the conductivity in that area. Ground conductivity is affected by the local geography. The first two dimensional model of conductivity was a land-sea model where the conductance of land was set at $20S$ and the conductance of sea water was $60S$ [Beggan, 2015]. The siemens, S, is the unit of electric conductance defined by $|S| = |\Omega^{-1}| = |\frac{A}{V}$, where Ω is the ohm, A the ampere and V the volt. More advanced models, including BGS 2012, use a thin-sheet conductivity model that analyses the resistivity of bedrock materials. It finds large ranges in conductivity, with the deeper parts of the Atlantic ocean having conductance of $1600S$. On land, locations including the Northern Highlands are a resistive terrane, as they have lithology consisting of granitic or metamorphic rocks, where as, the southern uplands have a highly conductive upper lithosphere [Beggan, 2015]. As discussed in section 1.2.4, sunlight results in current flow, hence the time of day and season largely affect conductivity.

To generate an informative network for analysis of space weather events, we must consider these factors affecting the observations. We could use one of the available conduc-

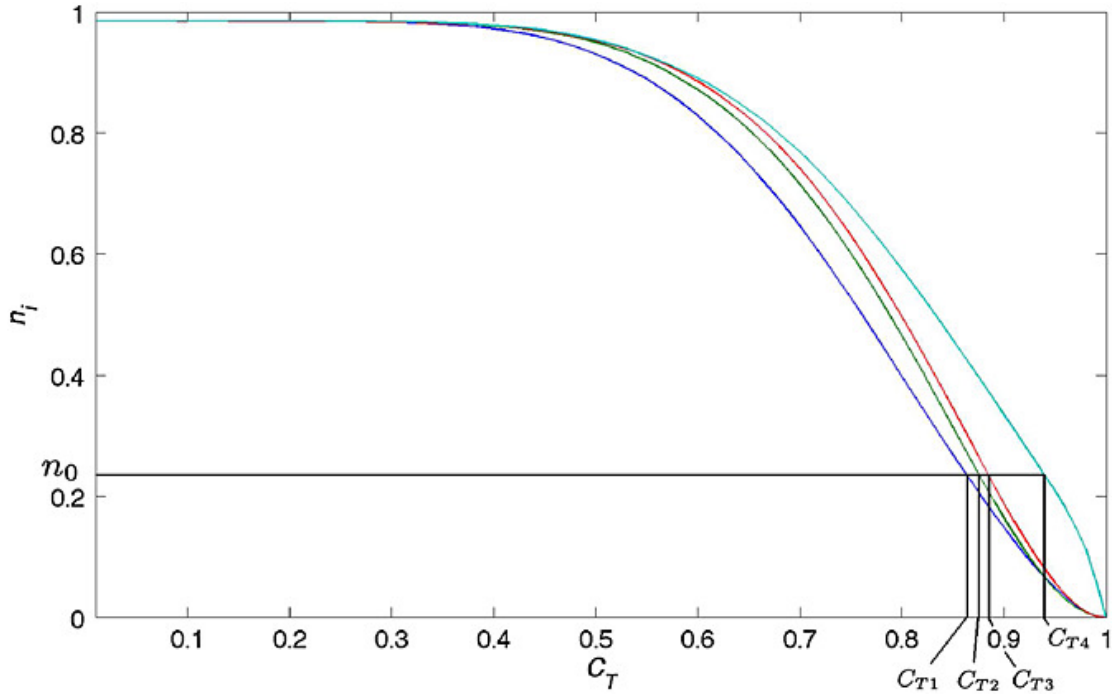

Figure 2.7: The monthly averaged normalized degree for representative stations is plotted as a function of the global threshold, C_T. A given normalized degree for the network, n_0, gives a corresponding threshold for each station, C_{Ti}. Figure from Dods et al. [2015].

tivity models to weight the network and hence, compare observations in resistive areas to conductive areas. However, with many options for 2D surface models and depth models ([Beggan, 2015] considers twenty one different conductivity scenarios) it is unclear which offers the best approximation for surface conductivity. Other considerations, including magnetometer sensitivity, could possibly be dealt with using historical measurement to weight the network observations but this would require constant monitoring of updates and failures of magnetometer hardware.

In Dods et al. [2015, 2017] the observations are normalized using an event and magnetometer specific threshold. A standardised adjacency matrix is constructed such that all magnetometer stations have the same likelihood of being connected to the network over a 28 day period around the event. On average, each magnetometer will then be connected to the same fraction of the network as any other station [Dods et al., 2015]. A normalized degree, $n_i(t)$, is defined as,

$$n_i(t) = \frac{d_i}{N(t) - 1}, \tag{2.7}$$

where d_i is the degree (number of connections (edges) a station i has to the network), and $N(t)$ is the number of active stations at time t. The number of connected magnetometers, d_i, will vary with the threshold, C_T. A magnetometer pair are only connected (to each other and hence the network) if the CCC of their vector magnetic field pertur-

bations exceeds the threshold. The CCC between two magnetometers can vary between 0 and 1. If $C_T = 0$ a magnetometer will always have greater correlation than the threshold; the magnetometer is connected to the network and $n_i = 1$, always. Contrarily, if the threshold is set to $C_T = 1$ a magnetometer will never be connected to the network and $n_i = 0$. Figure 2.7 plots the average normalized degree, n_i, for all thresholds, of four magnetometers. The standardised matrix entries are found by setting a fixed degree, n_0, and finding the corresponding threshold, C_{Ti}, for which each station will be connected to the network for that proportion of the month. The station and event dependent thresholds are then used to determine the adjacency matrix, $\boldsymbol{A}$, where,

$$A_{ij} = \begin{cases} 1, & \text{if} \quad C_{ij} \geq \min(C_{Ti}, C_{Tj}) \quad \text{and} \quad i \neq j \\ 0, & \text{if} \quad i = j \quad \text{or} \quad C_{ij} < \min(C_{Ti}, C_{Tj}). \end{cases}$$

The method used to threshold the network will be used throughout this thesis, as in Dods et al. [2015, 2017]. A problem arises when a magnetometer station has been inactive for a large proportion of the month. For the studies in this thesis we added the additional condition, that if the station had been active for $< 70\%$ of the month it was excluded from the analysis to avoid skewed results.

Chapter 3

Directed Network Methodology

In previous chapters we described the need for improved space weather analyses and outlined the data and techniques we could use for this research. This chapter will outline improvements made to the undirected network methodology outlined in Dods et al. [2015, 2017]. Section 3.1 provides the motivation and method for constructing a directed network. Section 3.2 establishes a common epoch across multiple events and section 3.3 defines a noise level. The work in this chapter underlines the work in chapters 4 and 5.

3.1 Lagged Correlation

Work by Dods et al. [2015] used an undirected network, i.e. it was based on the CCC between two magnetometers with zero time lag, $\tau = 0$, between the two time series measurements.

The first step in any network analysis is to calculate the raw time-varying network. The underpinning methodology for forming the raw network is detailed in Dods et al. [2015, 2017]; Orr et al. [2019] and is summarised here. The magnetometer stations form the nodes of the network and a given pair of nodes are "connected" (have an edge between them) if the CCC of their vector magnetic field perturbation time series exceeds a station and event specific threshold, as described in section 2.3.4 and in Dods et al. [2015]. The time-varying network is calculated at minute resolution, using a sliding leading edge w minute window for the CCC (the time, t, will refer to the last/latest time point of the window).

Section 3.1.1 provides two examples of pseudo time series that are treated as idealised magnetic field perturbation time series in order to demonstrate how the canonical cor-

relation components behave as a function of time lag. Section 3.1.2 takes the time series from a magnetometer pair during an isolated substorm to analyse how the canonical correlation component of real magnetometer data responds to lag. The analysis from sections 3.1.1 and 3.1.2, as well as further inspection of CCC between magnetometer pairs, is then used to identify a robust method of isolating the lag at which the canonical correlation is maximised; this is described in section 3.1.3. The full method is summarised in section 3.1.5.

3.1.1 Pseudo time series constructed with a lag between A and B

Figures 3.1 and 3.3 are pseudo time series that have been constructed to test how the CCC varies with time and lag. The time series have three components (magnetic north, east and down) to represent the vector magnetic field perturbations measured at magnetometers. In each case the two vector time series are randomly generated and independent of each other, except for a window of w minutes. Both time series have the same segment from $t_1 \leq t \leq t_w$ for time series B and $t_1 + \tau_c \leq t \leq t_w + \tau_c$ for time series A, where t is time and τ_c is the lag between the signals.

Random Time Series

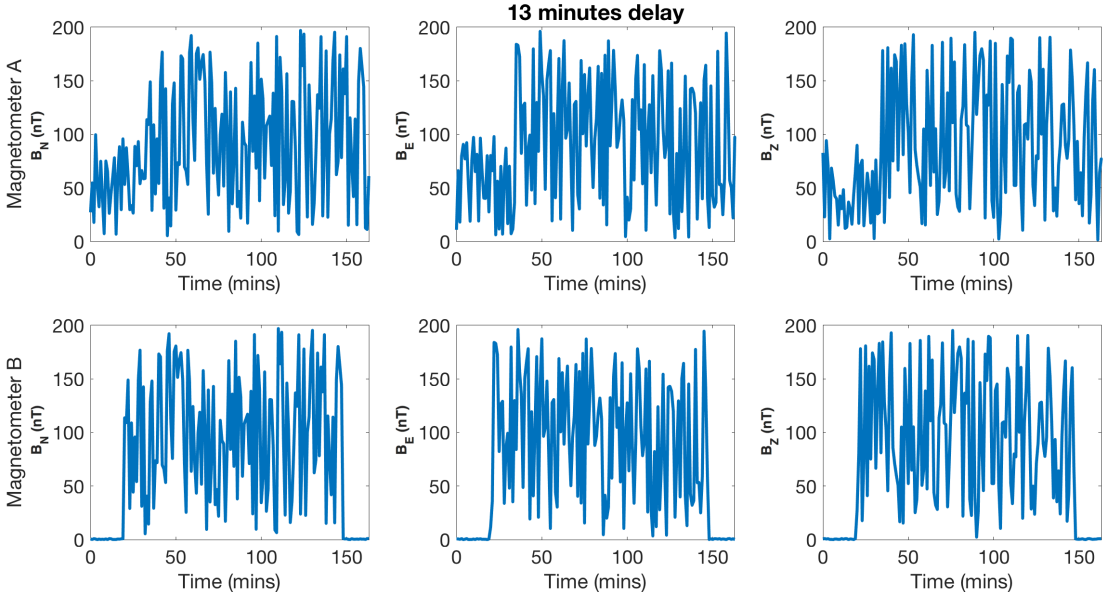

Figure 3.1: Two pseudo vector magnetic field time series. The time series are independent, apart from a $w = 128$ minute window which is the same from $21 \leq t \leq 148$ in B and $34 \leq t \leq 161$ in A, i.e. it has a delay of $\tau = 13$.

The first example is for two time series which have the same signal for a single window, $w = 128$, where A lags B by $\tau_c = 13$ minutes (see figure 3.1).

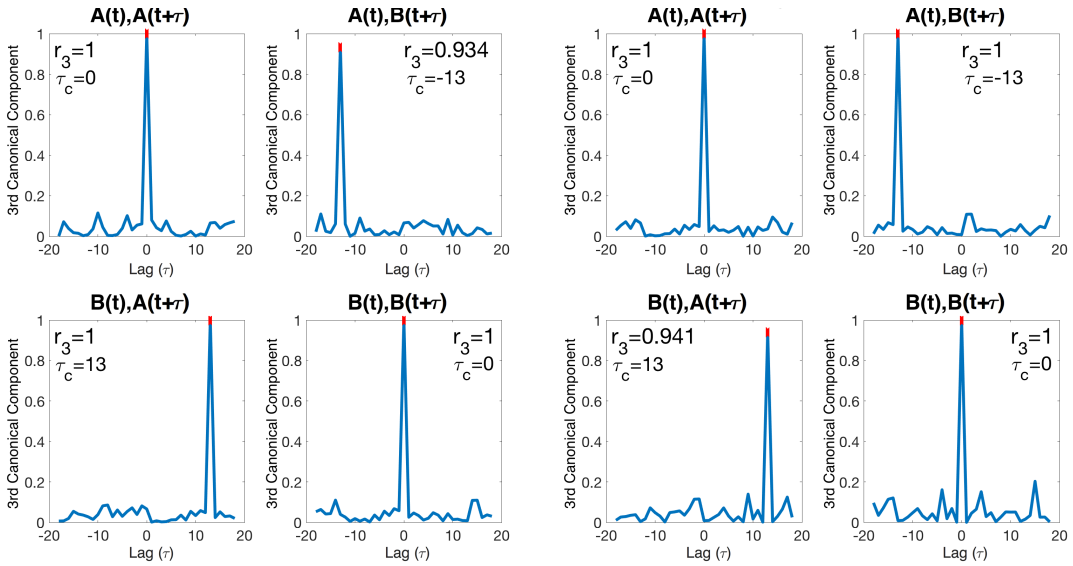

(a) Correlation of A and B using the window $21 \leq t \leq 148$.

(b) Correlation of A and B using the window $34 \leq t \leq 161$ (13 minutes later).

Figure 3.2: Windowed segments of the two pseudo time series (from figure 3.1) are linearly detrended and the canonical correlation components are found, where one time series is lagged with respect to the other. In this example we use the third component, r_3, as the first and second have maximum correlation ($r_1 = r_2 = 1$) even when the two windowed time series are different. 3.2a shows the correlation at the leading edge of the identical section in B, where there is perfect correlation from $B \rightarrow A$ when A is lagged by $\tau = 13$. Likewise 13 minutes later, 3.2b shows perfect correlation from $A \rightarrow B$ when B is lagged by $\tau = -13$. CCC between a time series with itself has $\tau_c = 0$.

In figure 3.2 the two signals are linearly detrended and the CCC is calculated, resulting in a obvious peak in the canonical correlation component value at $\tau = 13$, for B→A $(B(t), A(t+\tau)$, i.e. when the window for A is shifted forward in time by 13 minutes and the window for B remained constant in time). Likewise, there is a peak at $\tau = -13$, for A→B $(A(t), B(t+\tau)$, when the window for B is shifted backwards in time by 13 minutes and the window for A remained constant in time). The two time series are based on random signals, therefore, unless they are perfectly in phase the correlation will be low. When the matching segments of the time series are in phase (figure 3.2a), there is perfect CCC ($r_3 = 1$) from $B \rightarrow A$ when A is lagged by $\tau = 13$ minutes. Likewise, 13 minutes later (figure 3.2b), there is perfect correlation from $A \rightarrow B$ when B is lagged by $\tau = -13$ minutes.

If the network is calculated based on only the zero lag CCC between the time series from figure 3.1, there is not an edge between A and B. However, if the network is calculated based on the lag with the maximum canonical correlation component, A and B will be connected via an edge ($r_3 = 1$ will exceed the station's threshold). In addition, the sign of the lag implies that time series A follows time series B and, based on the magnitude of the lag, it is known that A is 13 minutes behind B.

Although time series A lags behind time series B, the A→B signal (A(t), B(t+lag) in figure 3.2a) results in a high canonical correlation component ($r_3 = 0.934$). When the windowed time series A is compared to the windowed time series B, with the window for B shifted by -13 minutes, both time series segments are made up of 13 minutes of a random signal (independent from each other) plus 115 minutes of the shared signal (the same for each windowed time series); hence the canonical correlation component is high, because 115 (out of 128) minutes are identical and in phase. The sign of the maximal lag, τ_c, is opposite to that in the $B \to A$ direction, hence A lags behind B. The mirror is observed 13 minutes later in figure 3.2b, where a high canonical correlation component ($r_3 = 0.934$) is calculated when A lags 13 minutes behind B.

Combining modulus and sine wave time series

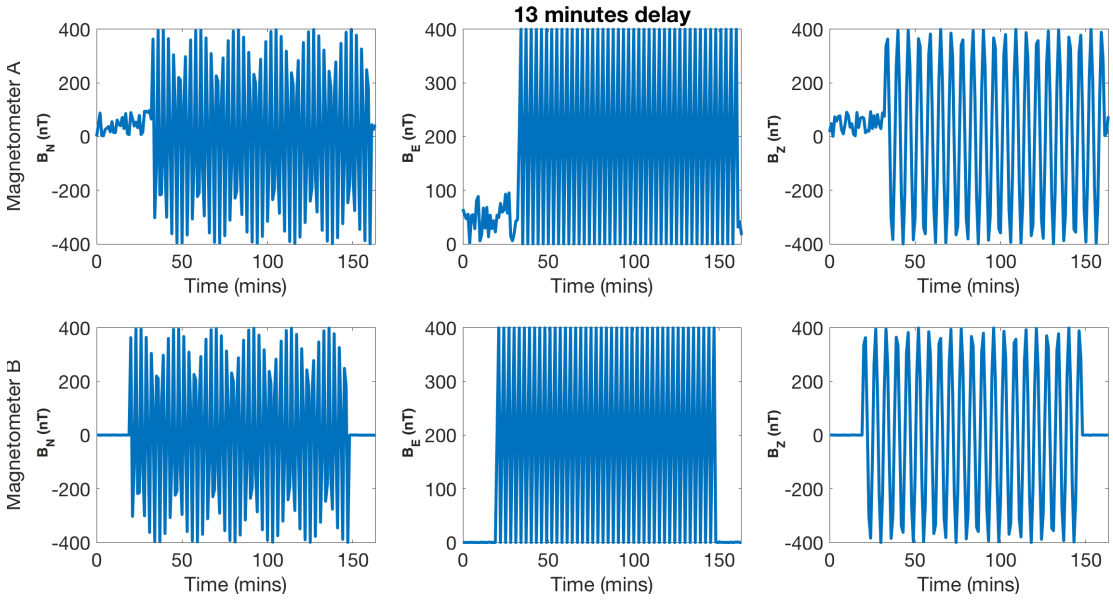

Figure 3.3: Two pseudo magnetic field time series. The time series are independent, apart from a $w = 128$ minute window which is the same from $21 \leq t \leq 148$ in B and $34 \leq t \leq 161$ in A, i.e. it has a delay of $\tau_c = 13$. The identical sections are given by $B_N = 400 \times \sin(2t)$, $B_E = 400 \times (t \times \mod(3))$ and $B_Z = 400 \times \sin(t)$.

The second example is of two random vector time series, with the same $w = 128$ minute window, where A lags B by $\tau_c = 13$ minutes. The matching segments of the time series are given by $B_{North} = 400 \times \sin(2t)$, $B_{East} = 400 \times (t \times \mod(3))$ and $B_{Down} = 400 \times \sin(t)$.

In figure 3.4 there is a peak in the canonical correlation component value at $\tau = \pm 13$ for $B \to A$ and $A \to B$, respectively. Again, when the leading edge of the window is the beginning of the identical section in B, i.e. figure 3.4a, there is perfect correlation from $B \to A$, when A is lagged by $\tau_c = 13$. Likewise, 13 minutes later in the time series, i.e.

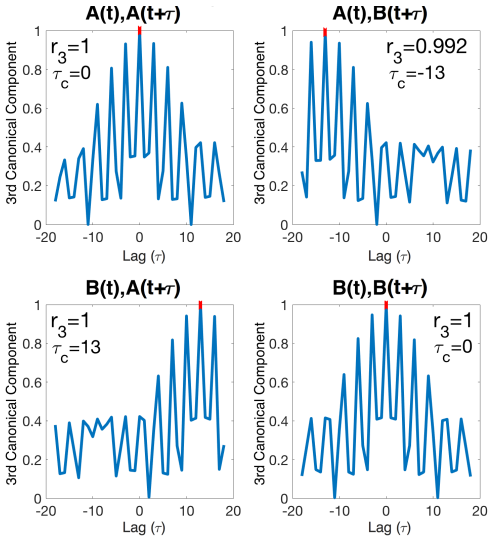

(a) Correlation of A and B using the window $21 \leq t \leq 148$.

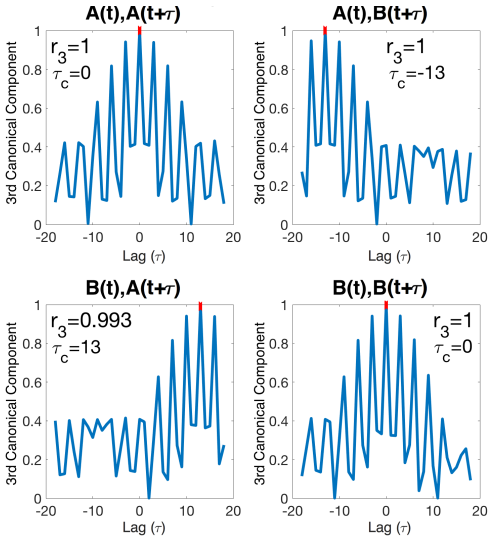

(b) Correlation of A and B using the window $34 \leq t \leq 161$ (13 minutes later).

Figure 3.4: Windowed segments of two pseudo time series (from figure 3.3) are linearly detrended and the canonical correlation components are calculated, where one time series is lagged with respect to the other. In this example we use the third component, r_3, as $r_1 = r_2 = 1$, even when the two windowed time series are different. 3.4a plots the correlation at the leading edge of the identical section in B, where there is perfect correlation from $B \rightarrow A$ when A is lagged by $\tau = 13$. Likewise 13 minutes later, 3.4b plots perfect correlation from $A \rightarrow B$ when B is lagged by $\tau = -13$. CCC between a time series with itself is maximum with $\tau = 0$. The wave like structure is obvious in the correlation, with odd lags having a higher r_3 than even lags. r_3 increases as lag tends toward ± 13.

figure 3.4b, there is perfect correlation from $A \rightarrow B$ when B is lagged by -13 minutes. However, unlike in the random time series in figure 3.1, there is wave like structure in the correlation components, with odd lags having a much larger canonical component when the waves are in phase. The correlation increases as the lags approach τ_c. Depending on the size of the threshold this example may have a canonical component exceeding the threshold at more than one lag.

3.1.2 Magnetometer data examples

Idealised pseudo time series are useful for demonstration of CCC at different lags but the main interest is how magnetometer observations correlate with increasing lag. If magnetometers were always maximally correlated at zero lag it would be unnecessary to incorporate a lagged response into the network. Figure 3.5a plots the vector magnetic field perturbation time series measured at two magnetometer stations during a substorm on the 01/01/1997.

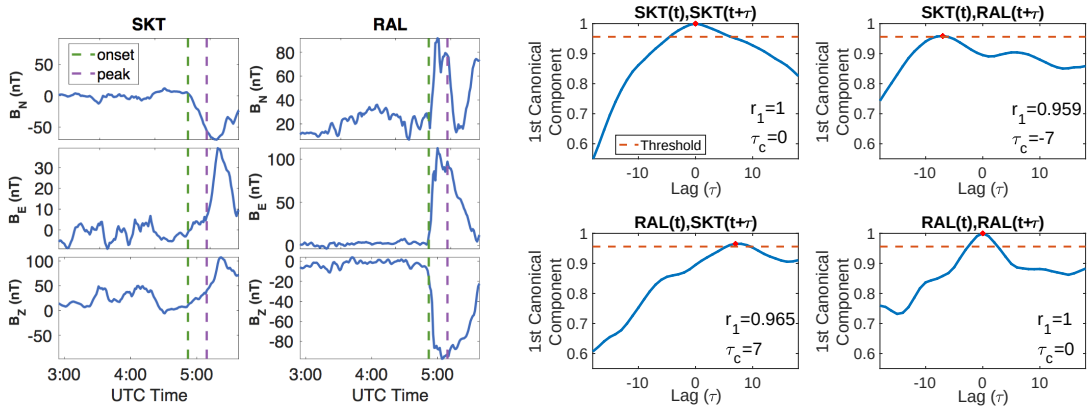

(a) The vector magnetic field perturbation time series measured at two magnetometer stations during a substorm on the 01/01/1997. The dashed green and purple lines mark times of substorm onset and maximum auroral bulge expansion respectively. RAL reaches its maximum field strength before SKT.

(b) A windowed segment of the magnetometer measurements (from figure 3.5a with the window $3:08 \leq t \leq 5:16$) is linearly detrended and the canonical correlation components are found, where one time series is lagged with respect to the other. SKT and RAL have maximum canonical correlation when RAL is windowed 7 minutes before SKT or SKT is windowed 7 minutes ahead of RAL.

Figure 3.5

RAL is the station code of a magnetometer located at Rabbit Lake in central Canada and SKT is the code of a magnetometer in Maniitsoq, Greenland. The magnetometers are separated by 3267 km and yet the magnetic perturbations they pick up are canonically correlated by $\sim 96\%$, with the Greenland magnetometer observing the signal 7 minutes after the Canadian magnetometer. The canonical correlation component exceeds the threshold calculated for the two magnetometers and, hence, there is a network connection with an associated eastward direction and a time lag of 7 minutes. The nodes (magnetometers) would not be connected in a zero lag network as the 1st canonical component, r_1, is below the threshold at $\tau = 0$. Similarly to figures 3.2 and 3.4, the canonical correlation components of RAL(t) with SKT$(t+\tau)$ and SKT(t) with RAL$(t+\tau)$ have little difference.

The normalized number of connections, $\alpha(t)$, is a parameter that is used throughout this thesis, and is defined as:

$$\alpha(t) = \frac{\sum_{i \neq j}^{N(t)} \sum_{j \neq i}^{N(t)} \mathbf{A}_{ij}}{N(t)(N(t)-1)}, \tag{3.1}$$

where $\mathbf{A}$ is the adjacency matrix and $N(t)$ is the number of active magnetometers [Dods et al., 2015].

If using a parameterisation such as $\alpha(t)$, we could choose to include half of the possible connection, such as the top half of the adjacency matrix or only positive lags. The

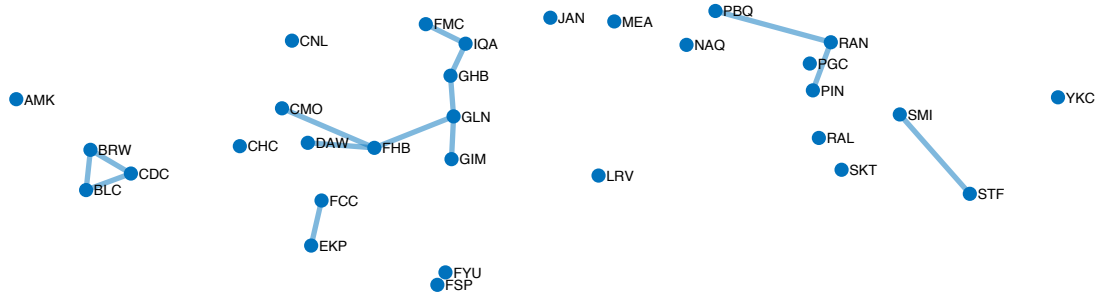

(a) Upper

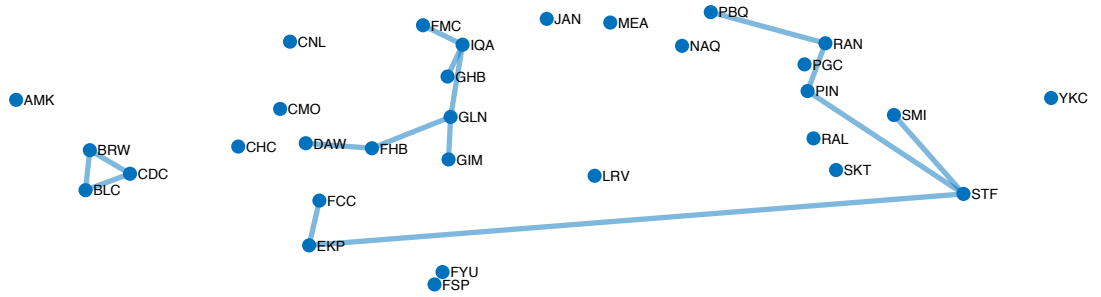

(b) Lower

Figure 3.6: An example network calculated at the lag of maximum correlation, τ_c, for a substorm on 16th March 1997. 3.6a is the upper half of the adjacency matrix and 3.6b is the lower half. The majority of network connections remain the same but the lower half has several extra edges.

choice of top or bottom, or, positive or negative, would result in a difference to the number of connection, and some directional information may be missed. Alternatively, we could choose to include all information, but this may count two edges between a magnetometer pair (e.g. $i \to j$ with positive lag and $j \to i$ with negative lag). However, as $\alpha(t)$ is a normalized parameter, the number of connections (top of equation 3.1) will be divided by twice as many possible connections (bottom of equation 3.1) and hence the maximum will always be $\alpha(t) = 1$ regardless of the choice of including half or all of the adjacency matrix.

When plotting the networks (as in figure 3.6) a double arrow (representing two edges) between two magnetometers does not represent a physical property, so a direction must be selected (i.e top/bottom, positive/negative e.t.c). Figure 3.6 is an example of a network calculated at the lag of maximum correlation, τ_c. If this was an undirected graph the upper and lower networks would be identical but, because of its associated direction, the adjacency matrix will no longer be symmetric; hence the upper and lower graph may be different. The direction and timing information associated with the edges is neglected in figure 3.6.

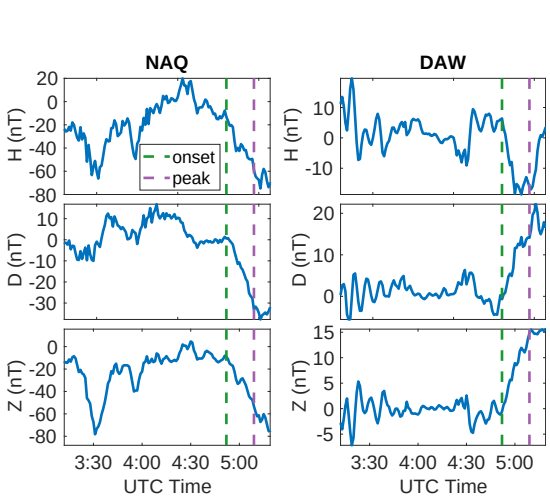

(a) The vector magnetic field perturbation time series measured at two magnetometer stations during a substorm on the 01/01/1997. The dashed green and purple lines mark times of substorm onset and maximum auroral bulge expansion respectively.

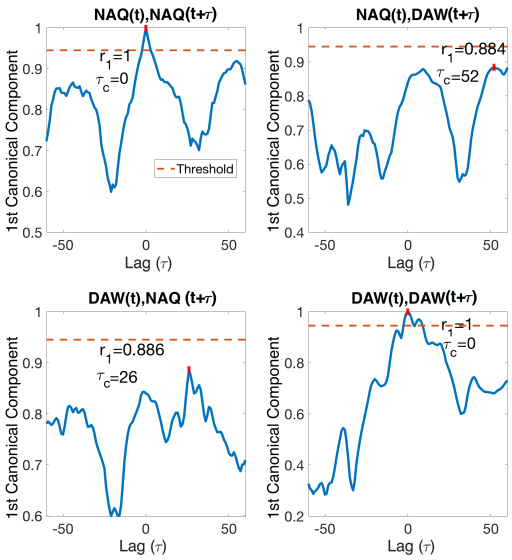

(b) A windowed segment of the magnetometer measurements (from figure 3.7a with the window $3:02 \leq t \leq 5:10$) is linearly detrended and the 1st canonical correlation components are calculated, where one time series is lagged with respect to the other up to ± 60 minutes lag. There are multiple peaks in canonical correlation component at different lags.

Figure 3.7

3.1.3 Lags of up to 60 minutes

It is clear that the CCC behaves as a function of lag, but the lag at which the correlation between two magnetometers peaks is not always the lag with the highest value of r_1. Figure 3.7 plots the first canonical correlation component, r_1, between two magnetometers, with up to ± 60 minutes lag between the windows. The auto correlation of the time series from NAQ results in $r_1 = 1$ at $\tau = 0$, as expected. However, there are now secondary peaks with significant correlation at $\tau \sim \pm 50$. When looking from NAQ to DAW there is a multiple peak structure with ~ 50 minutes periodicity and the maximum is taken to be at $\tau_c = 50$. The secondary peaks could be an affect of the detrending routine, or of the sinusoidal nature of the wave.

The CCC as a function of lag may have more than one peak and we need a robust method to isolate the peaks. To ensure we are analysing observed signatures of substorm onset, the network will be calculated based on lags between $-15 \leq \tau \leq 15$. The following requirements were identified and verified after extensive inspection of the CCC, $r_1(\tau)$, of a substorm event, with the aim of automatically selecting a real, observed lag of maximum CCC. A CCC component at τ_c must be a true maximum, having a significantly higher magnitude over r_1 given by the lags to either side. The lags

are calculated between $-18 \leq \tau \leq 18$ in order to gauge if τ_c is indeed a peak, $r_1(\tau_c) > r_1(\tau_c \pm 1)$. An additional requirement is that either $r_1(\tau_c \pm 1) \geq r_1(\tau_c \pm 2)$, or $r_1(\tau_c \pm 1) \geq r_1(\tau_c \pm 3)$, or both $r_1(\tau_c) \geq r_1(\tau_c \pm 2)$ and $r_1(\tau_c) \geq r_1(\tau_c \pm 3)$.

The station and event specific threshold continues to be calculated at zero lag. Calculating the threshold (as described in section 2.3.4) is already computationally expensive as r_1 is needed for every station pair for every minute of the 28 days surrounding each event. In order to calculate the CCC at $r_1(\tau_c)$ we would need to calculate $r_1(\tau)$ for all $-18 \leq \tau \leq 18$. As the threshold is by design a quantity to normalize the adjacency matrix, such that each entry has the same likelihood of being $a_{ij} = 1$ for the same fraction of the month, the threshold at zero lag is sufficient.

3.1.4 Establishing Statistical Significance

Statistical significance of network parameters, such as the normalized number of connections, $\alpha(t)$ (equation 3.1), can be established using surrogates, as described in section 2.3.3. The surrogate network is the network that could arise 'by chance', or purely from red noise. The time-series from each magnetometer are Fourier transformed, the phase spectrum is randomised whilst preserving the amplitude spectrum and this is then inverse Fourier transformed to give a surrogate time-series with the same power spectrum as the original signals, but with no time correlation. We used an iterated amplitude adjusted Fourier transform (IAAFT) [Schreiber and Schmitz, 1996] method with the Matlab code supplied by Kugiumtzis and Tsimpiris [2010]. The surrogate time series of each magnetometer is then used to calculate the random phase surrogate of the network parameters using the method as described in this chapter; the surrogate is treated in exactly the same way as the original time series data.

From the network surrogates we can calculate a discriminating statistic, such as the normalised number of connections, $\alpha(t')$. This statistic is a quantification of the likelihood that a connection between a pair of magnetometers could arise by chance- 'a false positive' [Dods et al., 2015]. For clarity we can rewrite equation 3.1 to represent the discriminating statistic as

$$F(t) = \frac{\sum_{i \neq j}^{N(t)} \sum_{j \neq i}^{N(t)} \mathbf{f}_{ij}}{N(t)(N(t)-1)}, \tag{3.2}$$

where $\mathbf{f}$ is the adjacency matrix of the surrogate network, $N(t)$ is the number of active magnetometers and $F(t)$ is the normalised number of connections in the surrogate network (false positives) [Dods et al., 2015]. For each event we repeated this ten times to obtain an average value of the normalised number of false positives, $F(t)$. The discriminating parameter for the random phase surrogate, $F(t)$ is compared to the network parameter, $\alpha(t)$ in figure 3.8. The normalised number of false positive connections is

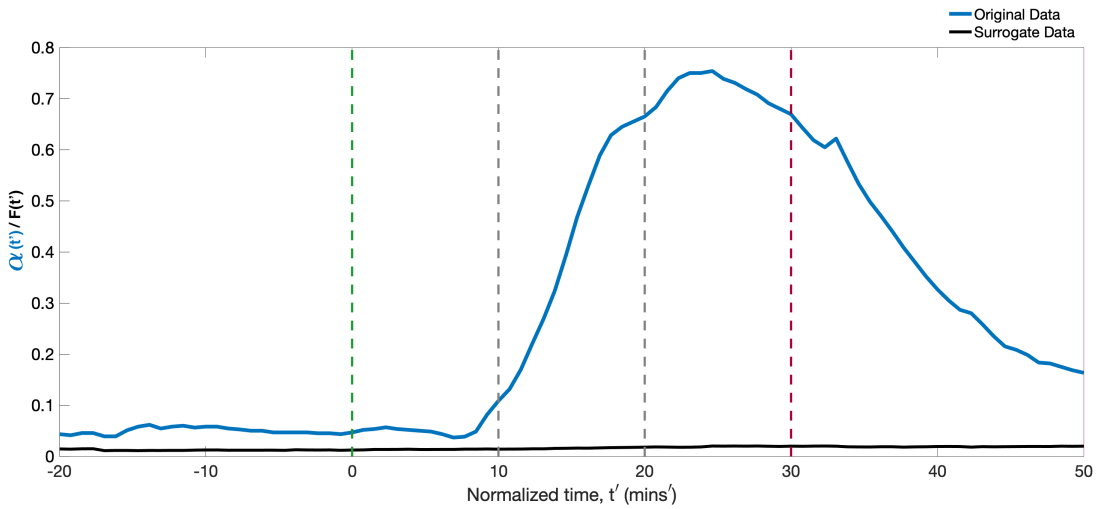

Figure 3.8: The normalized number of connections, $\alpha(t')$ (blue), and the normalized number of false positives, $F(t')$ (black), both within the nightside of an isolated substorm on the 16/03/1997. The surrogate data (false positives) shows no significant response to substorm onset (green dashed line).

~ 0 throughout the substorm whereas the original data clearly response to the substorm onset at $t' = 0$ (N.B. normalized time will be explained in section 3.2). Therefore it is highly unlikely that the response we see in $\alpha(t')$ is simply the result of random noise but it is statistical significant.

3.1.5 Summary

The method for forming a network, at zero lag, is detailed in Dods et al. [2015]. The magnetometer stations form the nodes of the network and a given pair of nodes are *connected* if the CCC of their vector magnetic field perturbation time series exceeds an event and station specific threshold, as specified in Dods et al. [2015] and section 2.3.4. In summary, the CCC is generally calculated over a 128 minute running window of the magnetic field perturbations observed by magnetometer pairs. The data is at minute resolution, giving a 128 point CCC for each station pair, every minute. The 128 minute sliding window is chosen to give sufficient accuracy in the computed cross-correlation function whilst also capturing the large-scale spatial and temporal behaviour of the SCW. Dods et al. [2017] previously demonstrated using model time series that this window length resolves changes on timescales significantly shorter than that of the window, specifically capturing onset where there is a sharp ramp in activity as the SCW forms. A network is calculated for every minute and all times, t, will refer to the leading edge of the window, that is the last time point spanned by the window (i.e. a window spanning time interval $[T, T+127]$ will have network properties plotted at time $t = T + 127$). Each windowed, three-component vector magnetic field time series is (1)

linearly detrended and (2) the CCC is calculated for each station pair then (3) if the correlation between magnetometers i and j exceeds the maximum of the two station thresholds then they are connected and are part of the network. For a network with N active magnetometers, an $N \times N$ adjacency matrix, $\mathbf{A}$, is formed which has $\mathbf{A}_{ij} = 1$, if i and j are connected, and $\mathbf{A}_{ij} = 0$, otherwise. The station specific threshold for each magnetometer station is determined such that the station will be connected to the network for 5% of the month (28 days) surrounding the event. This ensures that all stations have the same likelihood of being connected to the network, independent of their individual sensitivities to an overhead current perturbation, which in turn depends on the individual instrument characteristics and the local time and season dependent ground conductivity.

Dods et al. [2015] constructed the network using just the CCC at zero lag. In this thesis the directed network was formed by considering the lag at which the CCC was maximal, τ_c, up to a lag of ± 15 minutes. The value of the CCC value at lag τ_c is used to determine if the stations are connected (exceeds their threshold) and each connection then also has a direction and timescale of propagation of the observed signal, which is spatially coherent between the two stations. This potentially corresponds to the coherent pattern of time-varying ionospheric currents. The adjacency matrix, $\mathbf{A}$, is not symmetric and the sign of τ_{cij} determines the signal propagation direction for $\mathbf{A}_{ij}$. If the CCC between magnetometer i and j is above the threshold (they are connected), but with $\tau_c < 0$, it can be inferred that the signal originates at j and propagates towards i, i.e. $j \rightarrow i$. If $\tau_c > 0$ the propagation is $i \rightarrow j$. The additional timings and directional information provides a new dimension to the network and the properties are used throughout the following chapters.

3.2 Establishing a common epoch across events for analysis

In this section we will describe an event list of isolated substorms that will be analysed in later chapters. Additionally a method of time normalization will be outlined such that all events can be compared on the same time scale. In this thesis, vector magnetometer time series are analysed with 1 minute time resolution for the full set of magnetometer stations available from the SuperMAG database. This data is processed as in Gjerloev [2012], such that the ground magnetic field perturbations are in the same coordinate system, and have had a common baseline removed. A set of isolated substorm events, occurring between $1997 - 2001$, has been previously identified in Gjerloev et al. [2007]. These events have been selected such that (i) they are isolated single events optically and magnetically; (ii) the onset location is spatially defined; (iii) bulge-type auroral

events; (iv) there is a single expansion and recovery phase (or the end of the event is at the time of a new expansion); (v) the entire bulge region is in darkness to eliminate any terminator effects; and (vi) they are not during magnetic storms ($|Dst| < 30$nT) or prolonged magnetic activity. The requirement for darkness creates biases as the events with the majority of the nightside in darkness are in the months around winter solstice. Excluding daylit stations avoids large differences in ground conductivity between the stations, which would otherwise dominate the CCC. Specific studies require additional selection criteria which will be described where appropriate, but first it is helpful to put all the substorms in a common time frame.

It is well established that substorms vary in duration [Tanskanen et al., 2002; Kullen and Karlsson, 2004]. Out of the 116 events identified by Gjerloev et al. [2007], the time

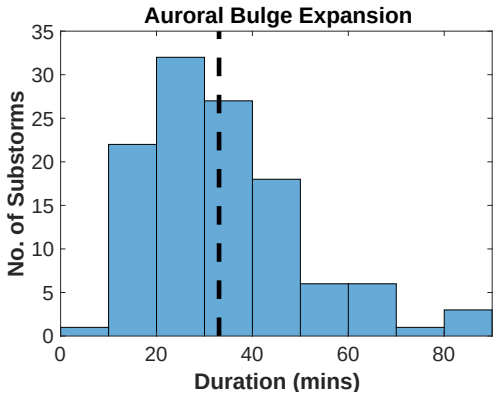

Figure 3.9: Distribution of the duration of auroral bulge expansion for the 116 events identified by Gjerloev et al. [2007] and analysed throughout this thesis.

between onset and peak, both of which are defined from polar VIS images, varies from 8 minutes to 86 minutes (see figure 3.9).

To compare common features across multiple substorms the events need to be mapped onto a single normalized time-base such that, once normalized, all substorms share a common onset time and take the same length of time to evolve from onset to the peak of activity (maximum auroral bulge expansion). Following Gjerloev et al. [2007], the observed event time, t, is related to the normalized time, t', by:

$$t' = \frac{T_E \times (t - t_{onset})}{t_{peak} - t_{onset}}, \tag{3.3}$$

where $T_E = 30$ minutes, approximately the average length of a substorm expansion phase. The onset time is then at $t' = 0$ and the time of peak expansion $t' = 30$. The critical timings for this normalization, t_{onset} and t_{peak}, can be unambiguously identified in these isolated substorm events.

As an example of time normalization, figure 3.10 shows the parameter time series of the

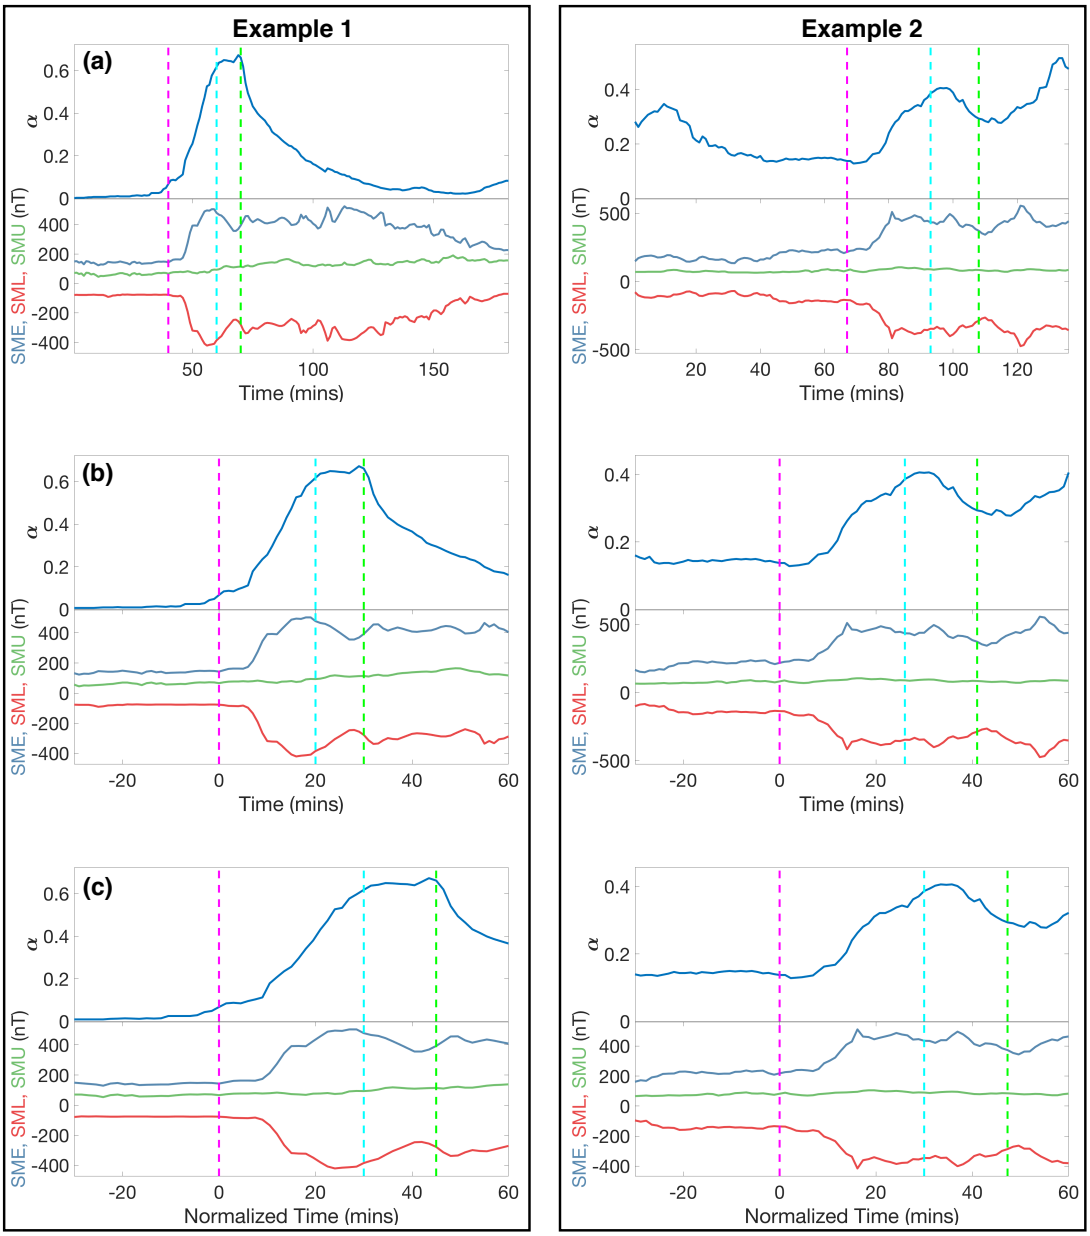

Figure 3.10: A comparison of the parameter time series of two substorms, for the dynamical network calculated at lag of maximum CCC, τ_c, including nightside magnetometers between $60 - 75°$ MLAT. The top half of each y-axis (bright blue line) show $\alpha(t)$ and the bottom is the superMAG auroral electrojet index (blue, green and red lines). Panel (a) shows the parameters for the duration of the two substorms, panel (b) shows 30 minutes before and 60 minutes after onset, and panel (c) shows the normalized time, t'. The two parameter time series increase on similar time scales in panel (c).

normalized number of connections, $\alpha(t)$, for two substorms, one on 25/01/2000 and the second on 18/12/2000.

Figure 3.10 compares the difference between (a) standard time, t, (b) time shifted with respect to onset and (c) normalized time, t' (as described in equation 3.3). The substorm

plots show a 'standard' measure, SME, SML, and SMU, SuperMAG's version of AE, AL and AU. Example 1 has a duration of 3 hours, but the expansion phase is only 20 minutes, therefore the normalized plot is stretched. Example 2 shows a growth phase of 66 minutes before onset; the normalized version does not show much of the growth phase. The expansion phase is 26 minutes, hence normalization does not stretch the parameters as much as example 1. After normalization the two examples have a similar structure.

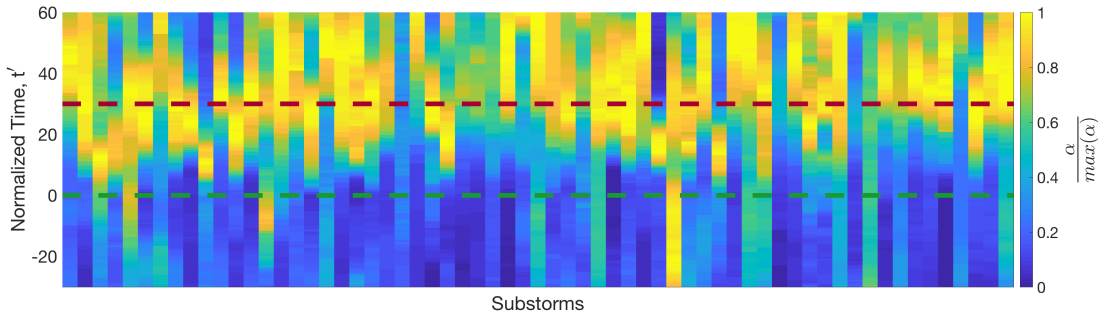

Figure 3.11: Colour plot of $\frac{\alpha}{max(\alpha)}$ for 63 substorms which do not have a preceding substorm, at lag of max cross correlation. The time (y-axis) is normalized as in Gjerloev et al. [2007]. The green and red lines represent onset and peak respectively. We see a brightening shortly after onset in the majority of substorms. There are some substorms which do not brighten significantly after onset and others are high before onset.

To show that time normalization is a useful tool for more than just the examples in figure 3.10, the parameter time series, $\frac{\alpha}{max(\alpha)}$ for 63 substorms which are quiet before onset, is compared. Figure 3.11 is a colour plot showing the normalized time, t' (as described above), for the 63 substorms. $\frac{\alpha}{max(\alpha)}$ generally increases after onset and is low before onset. For this plot, $\frac{\alpha}{max(\alpha)}$ has been used instead of α because substorms tend to show an increase in correlation but this does not necessarily equate to $\alpha = 1$. Time series are displayed in normalized time throughout the following chapters.

3.3 Noise

The aim of this section is to exclude connections from the network which are below the noise level, i.e. we set a minimum threshold where magnetic field perturbations with magnitude below the threshold are not of interest in this study of substorms. Geophysical datasets tend to exhibit a red noise background, where the power spectral density has the form $1/f^2$ [Lancaster et al., 2018]. Jackel et al. [2001] found red noise to be the best approximation for magnetometer observations. Gjerloev [2012] notes the probability distributions of differences between superMAG baselines and official quiet days rarely exceed 20nT. We will define the noise level as where the magnetic field perturbations do not exceed 20nT, in any direction, within the time window, for

all latitudes. Figure 3.12 shows three examples of magnetic field perturbations in a local magnetic coordinate system, the NEZ-system. The top panel has perturbations of $> |100|$ nT in magnitude so it can be assumed to be a real signal. The middle panel shows a time series that does not exceed 20 nT in magnitude; any real signal is below the noise level and any correlation with this signal is excluded from the network. The bottom panel is also has low magnitude, but it does exceed 20 nT in magnitude in the N and Z-direction. It is therefore assumed any correlation with this signal is real. Note this windowed time series corresponds to the onset time ($t' = 0$, pink dashed line in figure 3.13).

Figure 3.13 shows an example of the differences in the network parameter, $\alpha(t)$, when the network is calculated including or excluding the signals below the noise level.

Before the substorm onset, when magnetic field perturbations are generally small, there is a difference between the number of connections from the network calculated with (blue) or without (orange) the perturbations below the noise level. As the substorm expands, inducing larger magnetic perturbations, the difference between the number of connections (i.e. highly correlated vector magnetic field perturbations exceeding a threshold) calculated from each of the network decreases. All network studies in this thesis will exclude magnetic field perturbations that are below the noise level.

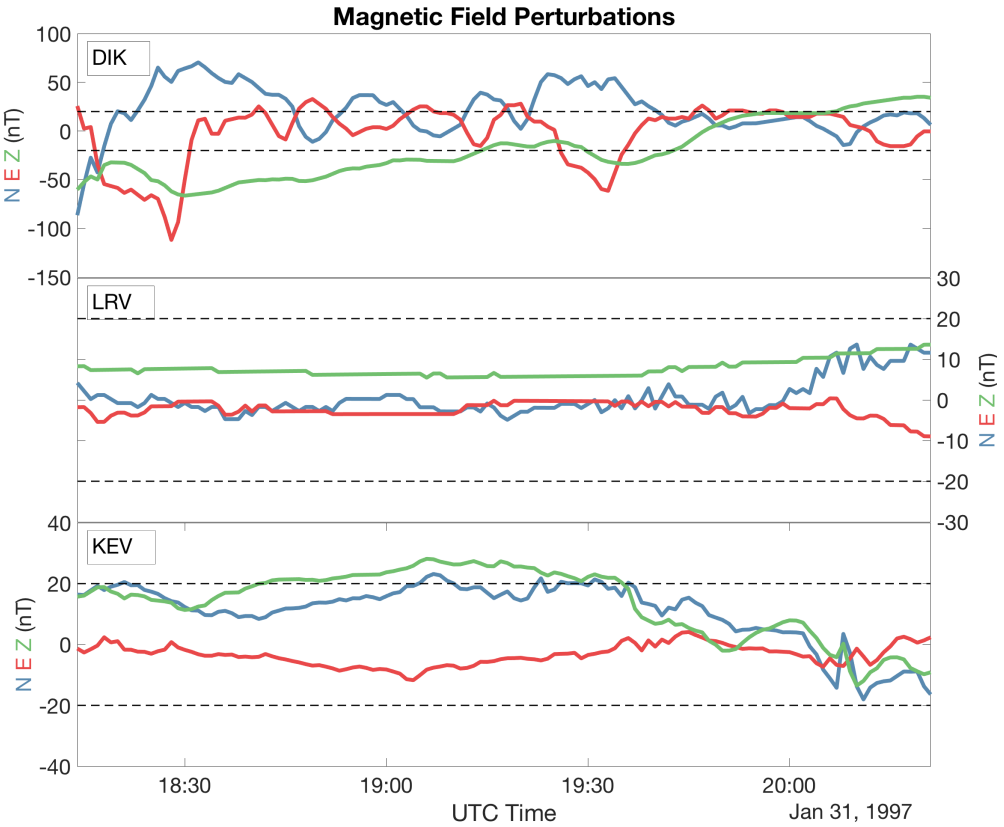

Figure 3.12: The raw magnetic field perturbation time series for three magnetometers during a substorm on the 31/01/1997. The magnetometer codes are in the north-west corners. Each panel plots the three components of the magnetic field in a local magnetic coordinate system. DIK and KEV measure a real signal during this time period as their components exceed 20nT, whereas LRV measures only perturbations below the noise level.

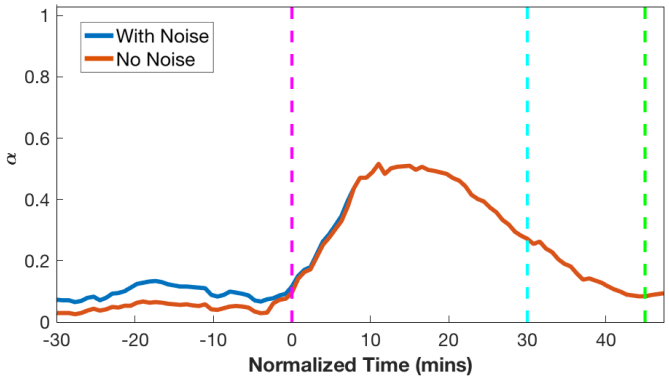

Figure 3.13: $\alpha(t)$ time series for the nightside of the network at lag of maximum CCC, τ_c, for a substorm on the 31/01/1997. The x-axis is normalized time, the y-axis is the magnitude of $\alpha(t)$. The blue line is the network parameter including all magnetometers and the orange is without those magnetometers measuring only perturbations below the noise level. The dashed pink, blue and green lines represent the time of onset/peak/recovery phases, all of which are found using polar VIS images. The exclusion of magnetic field perturbations below the noise level causes a change in the magnitude of α, particularly for values before onset.

Chapter 4

Directed network analysis of isolated substorms using polar VIS boundaries

4.1 Introduction

In the previous chapter we outlined the methodology for constructing a directed network. In this chapter we will use this methodology to quantify the spatio-temporal evolution of the substorm ionospheric current system. If spatial correlation reflects ionospheric current patterns, network properties can test different models (described in chapter 1) for the evolving substorm current system. In this chapter we find, and obtain the timings for, a consistent picture in which the classic substorm current wedge (SCW) forms. A current system is seen pre-midnight following the SCW westward expansion. Later, there is a weaker signal of eastward expansion. Finally, there is evidence of substorm-enhanced convection. The majority of the work in this chapter is from the published paper- "L Orr, SC Chapman, and JW Gjerloev. Directed network of substorms using supermag ground-based magnetometer data. Geophysical Research Letters, 46(12):6268–6278,2019" [Orr et al., 2019].

The organisation of the chapter is as follows. In section 4.2 we describe the methods and the data used to obtain the directed networks. In section 4.3 we highlight a case study of one substorm and present a statistical survey of 86 events which reveals how on average the spatial pattern of correlation evolves as the substorm progresses. We summarise everything up until this point in section 4.4. Section 4.5 contains a additional study which focuses on the directional information.

The proposed current system models from sections 1.2.4 and 1.2.5 point to the outstanding question: what is the average substorm current system morphology that we can quantify and resolve uniquely from the full set of available ground-based magnetometer observations? The goal of this chapter is use the methodology from section 3 to quantify the time-evolving spatial pattern seen across all 100+ magnetometers, in a manner that allows systematic averaging across may substorm events. This will provide a quantitative benchmark to test against model predictions. Key aspects of many of the models, whilst being physically distinctive, are qualitative. Our results place these qualitative predictions in direct contact with the observations, and can thus drive forward the formation of quantitative hypotheses that will allow these models to be distinguished.

The analysis of Dods et al. [2015], only examined the *undirected* network (zero-lag CCC). This was sufficient to reveal the initial formation of the SCW at substorm onset but without directional information could not capture the full spatio-temporal evolution of the current system. In this chapter we construct the networks based on the (often non-zero) time lags at which the CCC between each pair of stations is maximal, to form the substorm *directed* network, which captures the direction of information propagation between network nodes (magnetometers). Looking across a range of CCC lags captures the full pattern of spatial correlation and how it evolves in time. Non-zero CCC lags indicate the time-scale for propagation or expansion of a coherent structure and the sign of the lag gives the direction of propagation or expansion. We construct specific sub-networks to test the hypotheses of different proposed models for how the ionospheric current system evolves. The sub-networks isolate different spatial regions and allow us to test for connections between them. We will focus on spatially well-sampled isolated substorm events and establish network parameters that characterise how the magnetometers collectively respond to the SCW. We have identified 86 events that meet the sampling requirements (this is a subset of the substorm list used in the series of papers by Gjerloev and Hoffman [Gjerloev et al., 2007; Gjerloev and Hoffman, 2014]). We find timings for a pattern in the magnetic field perturbations consistent with the SCW formation at onset which then expands westward to form a coherent current system in the pre-midnight sector. There is additional weaker, eastward expansion of the SCW, followed by coherent correlation patterns spanning the entire nightside.

4.2 Methods

In chapter 3 the methodology for calculating a directed network was outlined. In this section we apply this method to the specifications of the questions posed in section 4.1.

4.2.1 Constructing the dynamical directed network

For each substorm we first construct the full dynamical directed network, which is calculated as described in chapter 3.1, before dividing it into sub-networks that flag spatial correlation within and between specific spatial regions.

These spatial regions are selected to test different proposed models of substorm current patterns. The term 'dynamical' is used here as in the networks literature [Jost, 2007]. Our analysis cannot resolve short-range fine structures that are smaller than, or of the order of the inter-magnetometer spacings, but can test whether long-range spatially correlated patterns exist.

4.2.2 Data and event selection

The substorms used in this study are from the list described in section 3.2 and previously used in Gjerloev et al. [2007]. Here we only use substorms between the years 1997-2000. In addition to the requirements stated in Gjerloev et al. [2007], we require that activity levels are low for a full window of 128 minutes before the substorm onset. Together these selection criteria, along with the requirement for a sufficient number of stations in the spatial regions around onset (described below) give 86 suitably isolated substorms.

Substorms here are translated into normalized time, as described in section 3.2 and specifically by equation 3.3. Onset is always at $t' = 0$ and the time of maximum auroral bulge expansion, known as peak, is always at $t' = 30$. The timings of onset and peak were unambiguously deciphered in Gjerloev et al. [2007]; Gjerloev and Hoffman [2014] from polar VIS images. This time normalization allows the comparison of multiple events.

4.2.3 Sub-networks for specific auroral spatial regions

We construct time-varying directed sub-networks that quantify correlation within and between specific spatial regions in the nightside. These spatial regions are selected for each event as shown in figure 4.1. The network is constructed using stations located between $60 - 75°$ magnetic latitude and within the nightside. Gjerloev and Hoffman individually determined the timings and positions of onset and the east and west ends of the bulge portion of the aurora using polar VIS images [Gjerloev et al., 2007]. The LT of the bulge edges at the time of maximum expansion ($t' = 30$) have been used to define the boundaries of region B.

The study was repeated using the east and west boundaries of the bulge fifteen normal-

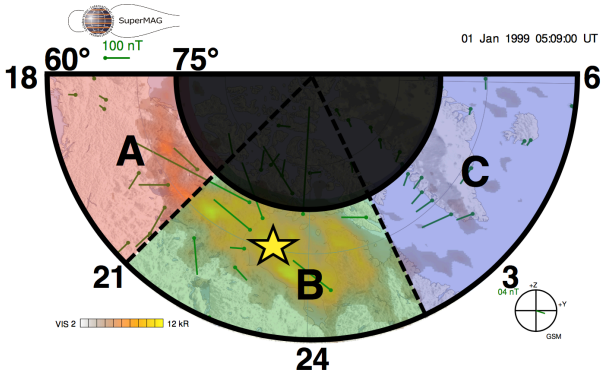

Figure 4.1: SuperMAG polar plot indicating the spatial regions A, B and C for which we obtain sub-networks. All data is from stations between $60 - 75°$ magnetic latitude, within the nightside. The LT boundaries between A, B and C are different for each event and are determined from Polar VIS images; they are separated by the east and west boundaries of the bulge at the time of maximum expansion (dashed lines). The magnetic latitude and local time of onset (again from Polar VIS) for each event, are indicated by the yellow star.

ized minutes before, and after, the maximum expansion phase ($t' = 15$ and $t' = 45$), results are presented in section 4.3.3. This gives slight differences but the overall results and conclusions are unchanged. The SCW is typically six hours of local time in extent [Gjerloev et al., 2007] which corresponds to region B; regions A and C are westwards and eastwards of the SCW respectively.

We will present a detailed study of the sub-networks for a single event and then will compare it to the average sub-network behaviour seen across all 86 isolated substorms. An event was identified which has ≥ 7 magnetometer stations in each of regions A, B and C for the duration of the substorm; this occurs on $01-\text{Jan}-1999$ with onset at $04:52$ UT. It had a relatively short expansion phase, 17 minutes, and a thin SCW, extending over 4.1 hours of LT at the time of peak expansion ($t' = 30$). For the averaged study over 86 events we require at least three magnetometers in a spatial region for it's sub-network to be included in the study. For example, a substorm in which there were ≥ 3 magnetometers in regions A and B, but < 3 in C, will contribute to the average sub-networks behaviour within A and B but not within sub-network C. We repeated the entire analysis with the more restrictive criterion of ≥ 7 magnetometers and found very similar results (see section 4.3.4). One benefit of using network analysis is that we do not require a spatially uniform grid of magnetometers, however, the condition of having ≥ 3 magnetometers per region gives a mean spatial separation distance (within regions) of ~ 1000km.

The spatial regions A, B and C are defined such that the sub-networks are always in the same local time relative to the SCW but, as the earth rotates, the geomagnetic location of the magnetometer stations will vary. This will not affect the properties

of the computed network provided regions A, B, C continue to be well-sampled with stations. However, the number of stations within each region can change. We therefore include a normalization to the number of possible connections as described in section 3.2 by equation 3.1; $\alpha(t', \tau_c)$ then describes the number of normalized connections per normalized time and lag of maximal CCC.

4.3 Results

In this section we present a case study of one substorm and a statistical survey of 86 events which reveals how on average the spatial pattern of correlation evolves as the substorm progresses.

4.3.1 Observed timings of spatial correlation

We now present (in figure 4.2) the directed network for the individual substorm identified above (left column), and the average of all 86 selected substorms (right column). Substorm evolution may vary but the individual substorm is plotted as an example for comparison to the multi-event mean. Having obtained the sub-networks for each region (identified in figure 4.1), we have the normalized number of connections, $\alpha(t', \tau_c)$ (equation 2) within (panels $2-4$), or between the regions A, B and C (panels $5-10$). Analysing connected magnetometers within each region provides timings of the emergence of coherent spatial patterns of correlation in the magnetic field perturbations (at ground level), whilst connections between regions provide information on how these patterns are propagating and/or expanding during the substorm; any inter-region dependencies will also be flagged. If a region was fully connected, the sum over the two plots every magnetometer pair is connected $\alpha = 1$. Since the connections between regions (e.g. A→C and C→A in panel 7 and 8), would be $\alpha = 1$. Hence the range of values for the y-axes for connections between regions (panels 5-10) are half of that within the regions (panels 2-4).

As the networks are constructed using the time delay/lag at which the CCC between each pair of magnetometer stations is maximal, each connection has an associated signed lag, τ_c. We bin the number of connections (α) into ranges of the magnitude of this lag (τ_c). Connections which are at zero lag have no time delay i.e. $\tau_c = 0$ (grey), and connections with an associated direction of propagation/ expansion, from one magnetometer to another, have a range of delays, that is, lags from $1-15$ minutes (blue-red). The sign of the lag indicates a direction of propagation or expansion from one magnetometer location to another; this information is combined with the physical geographical locations of the magnetometers to determine if the propagation/expansion is eastward

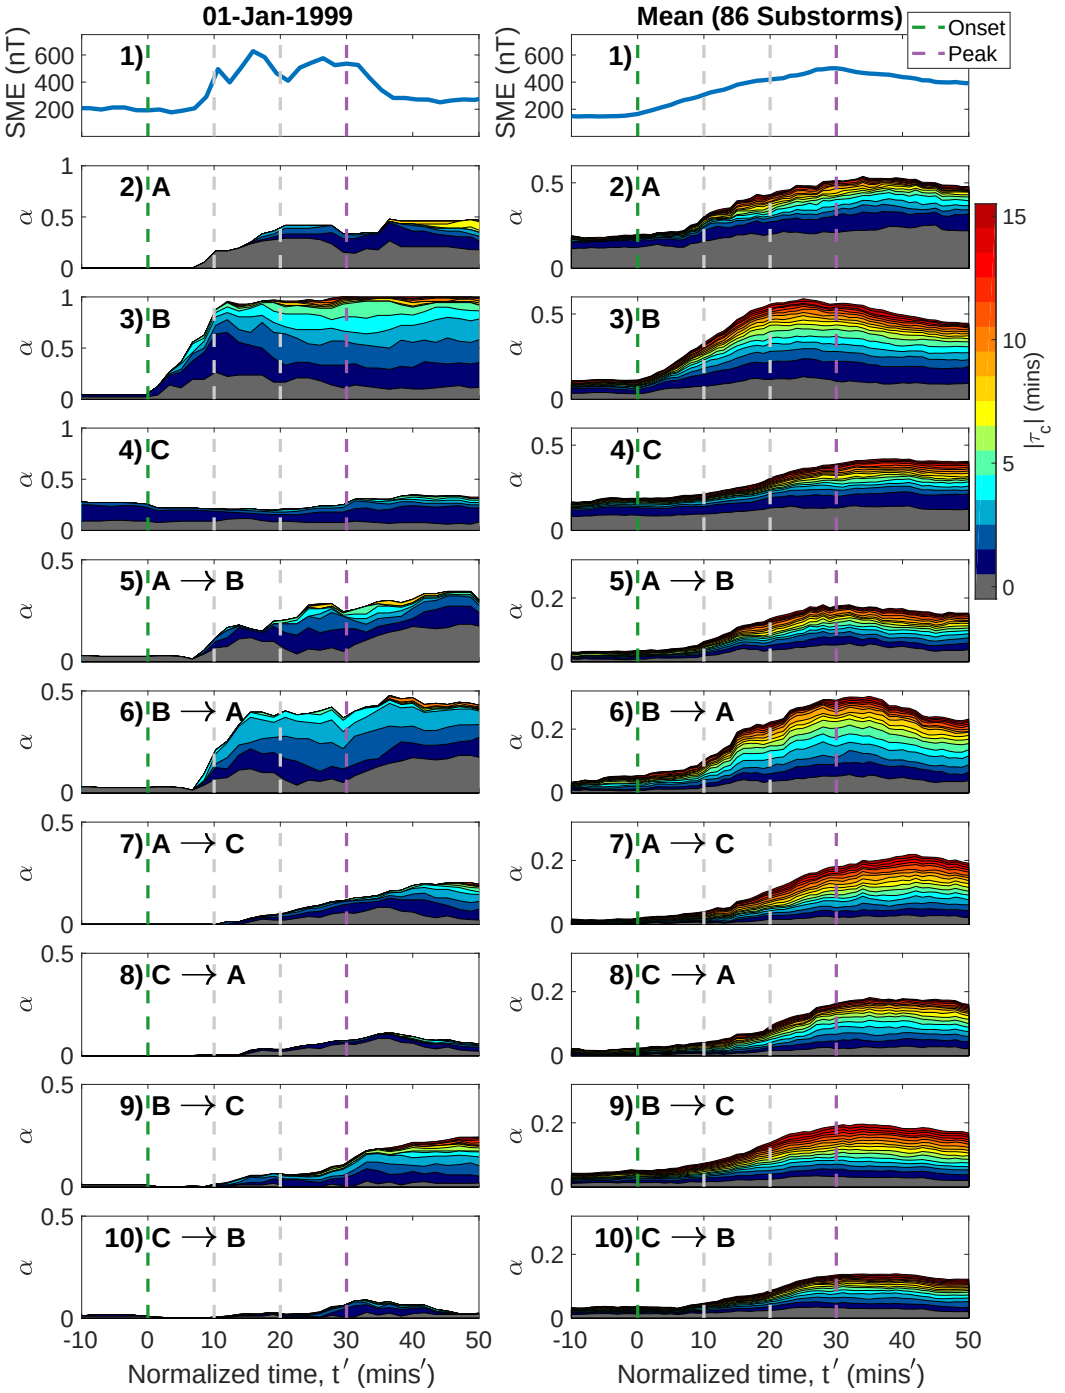

Figure 4.2: The normalized number of connections, $\alpha(t', \tau_c)$, is binned by the lag of maximal CCC, $|\tau_c|$. Each panel stacks vertically (one above the other) $\alpha(t', \tau_c)$ versus normalized time, t', for $|\tau_c| \leq 15$. $|\tau_c|$ is indicated by colour (see colour bar). Panel 1 plots the SuperMAG electrojet index, SME. Panels $2-10$ plot α for connections within and between each of the regions A, B and C (identified in figure 4.1). The left column plots a single event, whereas the right plots the average of 86 events (containing sub-networks with ≥ 3 magnetometers per region). Substorm onset (green dashed line) is at $t' = 0$ and the maximum of the expansion phase (purple dashed line) is at $t' = 30$.

or westward. The connections are separated into different panels for each direction and then binned by the magnitude of the lag. For example, between regions A and C, panel 7 plots the A→C propagation/expansion, eastward, from region A into region C whilst panel 8 plots C→A propagation/expansion, westward, from region C into region A. Connections with a lag $\tau_c = 0$ (indicated in grey) are plotted on both panels 7 and 8 (A→C and C→A) as they simply indicate instantaneous correlation between regions A and C, which have no associated direction and thus, by definition the grey bars are identical on the two plots.

Figure 4.2 stacks the time series of the normalized number of connections, $\alpha(t', \tau_c)$, so that the value of α for each range of $\mid \tau_c \mid$ are plotted one above the other for increasing $\mid \tau_c \mid$. The stacking is such that each independent $\alpha(t', \tau_c)$ is visible. The envelope is then the total (normalized) number of connections over all lags, that is, all $\mid \tau_c \mid \leq 15$. For example, during the individual substorm (left column), we see mostly instantaneous (grey) correlation within A (panel 2) with additional lagged correlation later in the substorm. On the other hand within B (panel 3), at the time of peak expansion (purple dashed line) the network is composed of $\sim 10\%$ instantaneous correlation, $\sim 80\%$ with $1 \leq \mid \tau_c \mid \leq 5$ (fast propagating or expanding) and $< 10\%$ of connections have $\mid \tau_c \mid \geq 6$ (slow expansion or propagation). The plot covers the time interval $-10 \leq t' \leq 50$ normalized minutes where the times of onset, peak expansion and the 10 minute intervals in-between are indicated with vertical dashed lines. The figure presents a summary of time-varying spatial correlation for each sub-network for the duration of the substorm. The full networks for the individual substorm are plotted in section 4.3.5. The SME (SuperMAG electrojet index) for the individual event, and its multi-event average, is plotted in panel 1 of figure 4.2. We can see that although the events are on a normalized time-base, the multi-event average is more smooth and responds less sharply to onset than the individual event.

Prior to onset, the multi-event average shows some spatially coherent connections within each of regions A and C (panels 2 and 4). These connections are mostly instantaneous ($\mid \tau_c \mid = 0$, grey shading) or with 1 minute lag ($\mid \tau_c \mid = 1$, dark blue shading). Importantly these regions are not correlated with each other, so that the number of A→C and C→A connections are small (panels 7 and 8). At onset, in both the individual event and the statistical average, (panel 3) we see that the sub-network within region B has the most prompt and largest response, that is, increase in spatial correlation. For the individual substorm we see a sharp increase in the number of correlated pairs, beginning at onset ($t' = 0$) and increasing to $\sim 100\%$ (all magnetometers within B are highly correlated) at $t' \sim 15$. Likewise, the multi-event average number of correlated magnetometer pairs within the B sub-network begins to increase at onset, but it is smoother and reaches a peak slightly later than in the individual event; in this case, correlation maximises with $\sim 60\%$ connectivity at $t' \sim 25$. The January substorm onset is mainly characterised by

fast propagating ($0-6$ minutes lag) connections whilst in the multi-event mean sub-network $\sim 80\%$ of connections are propagating (non-zero peak lag) throughout. This is consistent with a pattern that is both coherent and propagating and/or expanding. The timings of region B growth are consistent across the majority of substorms observed.

About 10 normalized minutes after onset we can see in panel 6 westward propagation and/or expansion from region B (around onset) into region A (westward of onset), B→A. This coincides with an increase in spatial correlation within region A (panel 2). For both the individual and the multi-event average, the B→A time series (panel 6), that is, the relative increase in the number of connections at different lags, resembles that of the network located wholly within region B (panel 3), except that it occurs ~ 10 normalized minutes later and has about half the magnitude. Within region A (panel 2), $\sim 50\%$ of magnetometers become highly correlated at $20 < t' < 30$. For the individual substorm most of the connections between magnetometers are instantaneous, but for the multi-event average $\sim \frac{2}{3}$ of the increase in the number of connections is at non-zero lag. There is variation between individual substorms as to how spatial correlation between magnetometers within region A develops from $-10 \leq t' \leq 50$, with some substorms having no obvious response to onset. The A→B (panel 5) propagation and/or expansion develops on similar timescales to B →A (panel 6) but there are significantly fewer connections ($20-30\%$ of magnetometers correlated at peak, $t' = 30$) within the multi-event average, with $\sim \frac{1}{3}$ of these connections being instantaneous (zero lag, no direction).

The sub-network for region C (panel 4, east of the SCW) has the smallest response to substorm onset of any region. The January substorm remains moderately connected ($\sim 23\%$) from before onset until long after peak expansion. The multi-event average begins to increase at $10 < t' < 20$ and the region is maximally correlated after peak expansion, $t' > 30$ with $\sim 20\%$ of magnetometer pairs being connected; this pattern of correlation is consistent long into the recovery phase. This is consistent with many of the individual substorms showing little or no response to onset. In panel 9 we see that region B becomes correlated with region C with eastwards propagation and/or expansion (B→C) $\sim 10-20$ normalized minutes after onset, peaking with $\sim 20\%$ magnetometer pairs correlated at $t' > 30$. In panel 10, we can see that for the individual event $< 10\%$ of magnetometers are correlated from C→B, with this small increase only occurring ~ 25 normalized minutes after onset. Correlation increases by $\sim 15\%$ for the multi-event average between $t' \sim 10$ and $t' \sim 30$. Thus the response within region C simply tracks that of the propagation or expansion from B → C, and any propagation from C → B occurs subsequently.

Finally, $\sim 10-20$ normalized minutes after onset, in panels 7 and 8 we see correlation increasing relatively slowly between regions A and C (west and east of the onset location,

respectively), reaching a maximum level of connectedness at $t' \sim 40$, long after the time of peak expansion. There is slightly more eastward propagation (A→C, $> 20\%$ of magnetometers) than westward propagation (C→A, $\sim 10\%$ and $\sim 18\%$ of magnetometer pairs for the individual and multi-event average, respectively). Again, the correlation propagating outwards from C is small (and mostly instantaneous) for the individual event.

4.3.2 Interpretation

If we can interpret coherent patterns of spatial correlation across the distributed SuperMAG magnetometers as the emergence of current systems, the above time-dependent network provides an evolution sequence, with timings, for the substorm current system in the nightside. Our analysis then provides a quantitative measure of spatial coherence as well as the time scales on which evolution occurs. By separating the nightside into three regions we have attempted to isolate the SCW. To use the terminology of Kamide and Kokubun [1996] we have: (A) eastward electrojet; (B) substorm unloading component (SCW) and (C) westward electrojet. Whereas B is associated with the substorm current wedge, or DP1 perturbations, A and C can be related to the general magnetospheric convective system, DP2 [Nishida, 1968], which is enhanced during substorm growth and expansion phases [Milan et al., 2017]. To summarise the above results we identify key time ranges, before and after onset: $0 \leq t'_1 < 10$, $10 \leq t'_2 < 20$, $20 \leq t'_3 < 30$ and $t'_4 \geq 30$ in terms of normalized time, t'. In terms of substorm evolution t'_1 is following onset, t'_2 is expansion phase, t'_3 is near substorm peak and t'_4 is the early recovery phase. The timings are:

- Before onset, the pre- and post-midnight regions A and C each have a relatively weak coherent pattern consistent with convection (DP2); notably A and C are not coherent with each other.
- In t'_1 we first see the formation of a substorm current wedge (SCW/DP1) around onset (correlation within B) which approaches maximum in t'_2.
- In t'_2 there is westward propagation and/or expansion of the SCW west towards the pre-midnight region (A). We see connections (B→A) and at the same time a signature of a coherent current system within region A (correlation within A). This is shortly followed by weaker correlation from A→B, indicating that the entire A-B system is now correlated. These all approach a maximum at t'_3.
- A weaker signal of eastward propagation and/or expansion of the SCW towards the post-midnight region starts in t'_2, and reaches its maximum in t'_4. We see connections (B→C) and on a similar timescale a signature of a coherent current

system within region C (correlation within C), with additional weaker correlation from C→B. The correlation in region C is relatively low.

- The regions eastward and westward of onset (A and C) each have a coherent pattern consistent with enhanced magnetospheric convection (DP2). Later in the substorm there is coherence between regions A and C, beginning well after onset, in t'_2, and reaching maximum correlation in t'_4, that is, only after region B has become correlated with all other regions. This can either reflect direct correlation between A and C, or could simply imply that both A and C are correlated with B.

We can then consider what support these results provide for proposed models for substorm current systems, specifically, models with a single westward electrojet segment [McPherron et al., 1973; Kamide and Kokubun, 1996], a westward and, a lower (but still in the auroral zones) latitude eastward electrojet segment [Ritter and Lühr, 2008; Sergeev et al., 2011, 2014]; two unconnected westward electrojet segments pre- and post-midnight [Rostoker, 1996; Gjerloev and Hoffman, 2014] and finally many small individual segments [Liu et al., 2018]. Importantly, any method for quantifying spatial correlation cannot distinguish between direct correlation (here, A→C) and indirect correlation (here, A→B→C); this indirect correlation may enhance the number of A→C connections relative to B→C. Therefore our results are not inconsistent with multiple separate current systems provided that they are either spatially correlated with each other, or on spatial scales smaller than that of the magnetometer spacing.

The coherent patterns of eastward and westward expansion are in agreement with previous work using synchronous space and ground based magnetometers [Nagai, 1982, 1991]. However, we have not found definitive support for two, or more, distinct and uncorrelated substorm current systems. A lag, $\mid \tau_c \mid > 0$, for B→C connections, implies that C is delayed with respect to B, consistent with a propagation from B to C. Interpreting these results in terms of current components suggest two scenarios for this propagation: i) a single current segment which is expanding from B to C; or, ii) a current segment in B and another in C, where the segment in C is correlated with that in B but is developing with some delay. There is no interpretation of our results which would suggest a scenario where regions B and C are uncorrelated, independent current systems.

If associated entirely with general magnetospheric convection (DP2 system), the pre- and post- midnight (A and C) are directly-driven by the solar wind and must enhance on similar time-scales, although the magnitudes may differ [Kamide and Kokubun, 1996]. We have found that pre-onset, the regions A and C each have coherent, but relatively small, signatures of correlation with little CCC between them. Post-onset, in both the individual event and the average over 86 substorms, the long range east

to west (A→C) correlation patterns only emerge after the growth of the SCW (region B). The growth of spatially coherent patterns appears first in B (the SCW, at onset) followed by A (with correlation between B and A) and later, in C. This suggests that following onset, A and C are not solely attributable to enhanced convection and the presence of contemporaneous B → C and B→ A connections suggests that there may be a combination of contributions from convection enhancement and SCW expansion. Importantly, this does not require that a current segment in A expands or propagates into C.

Finally, if instead of a large scale SCW there were only many small, uncorrelated, individual segments [Liu et al., 2018] we would not expect to find the long-range correlations (A to C) seen here (also in section 4.3.5, figure 4.7). Since we calculate CCC on minute resolution time series, each connection in the network is derived from a 128 minute time window. Thus we cannot resolve short-timescale events such as a large number of small wedgelets each associated with a bursty bulk flow (BBF) in the plasma sheet which have lifetimes of some 10 min. In addition, we cannot resolve structures that are on smaller spatial scales than the inter-magnetometer spacing. If multiple wedgelets are present, their spatial aggregate would give an overall large-scale magnetic amplitude signature mainly at the edges of the region containing the wedgelets, regardless of whether or not the wedglets are spatio-temporally correlated. Here, both spatial and temporal information is used to obtain the cross-correlation so that temporally uncorrelated wedgelets would give no spatially coherent signature of cross correlation at all, whereas if the same wedgelets were temporally correlated, we would find a signature of spatial cross-correlation.

In the analysis and by the organisation of data into three regions, A, B, C, we are quantifying the coherence over these regions. This is over a range in both latitude ($60-75°$) and local time (typically region B is ~ 6 hrs LT). The westward electrojet around onset (B) may not cover all latitudes but our analysis technique is mainly addressing the various SCW models which differ in their local time distribution.

4.3.3 Observed timings of spatial correlation with different boundaries

Potential limitations to the technique include sensitivity to the location of the east and west bulge boundaries, which are constant throughout the substorm and therefore may not fully represent fast changes in the time-varying current system. There may also be a spatial coarse-graining effect due to the geographic location of the finite number of magnetometers; there are few near the eastward SWC boundary during the January substorm. To test this we present the same plots, for the January substorm and the

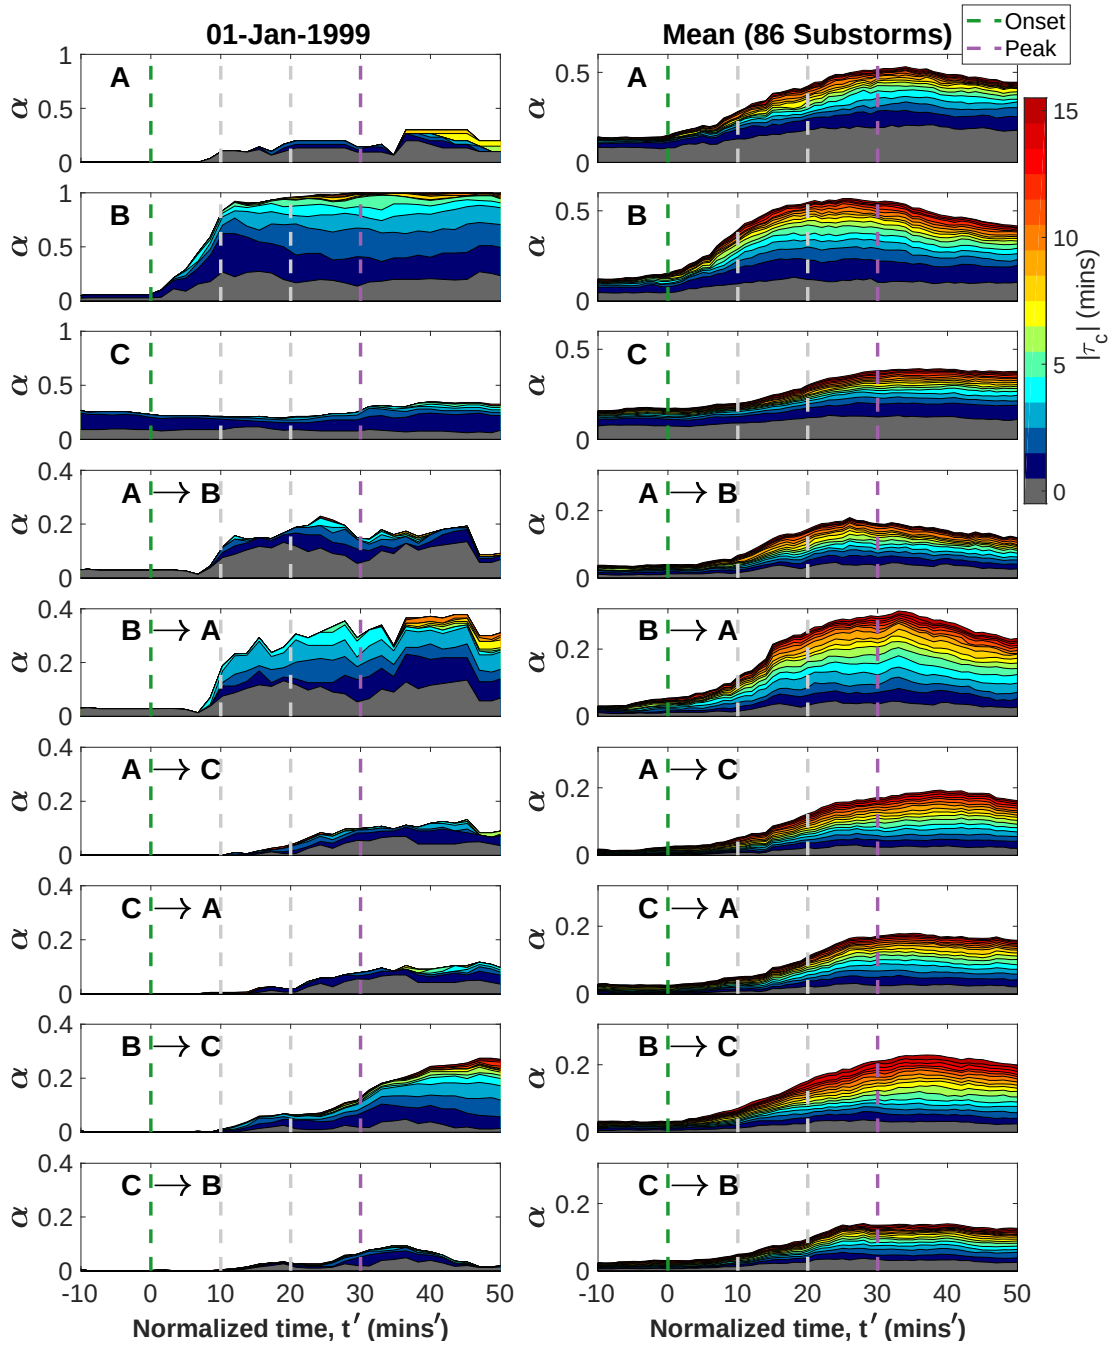

Figure 4.3: The normalized number of connections, $\alpha(t', \tau_c)$, within and between each of the regions A, B and C, where the boundaries of B are identified by polar VIS half way through the expansion phase ($t' = 15$). In the same format as figure 4.2.

event average, but with the east and west boundaries of the bulge at $t' = 15$ and $t' = 45$. These show little change from the results shown in figure 4.2.

Figures 4.3 and 4.4 plot the normalized number of connections in, and between, different sub-regions, as defined in figure 4.1. The figures are in the same format as figure 4.2 but the auroral bulge boundaries have been defined at times $t' = 15$ (for figure 4.3)

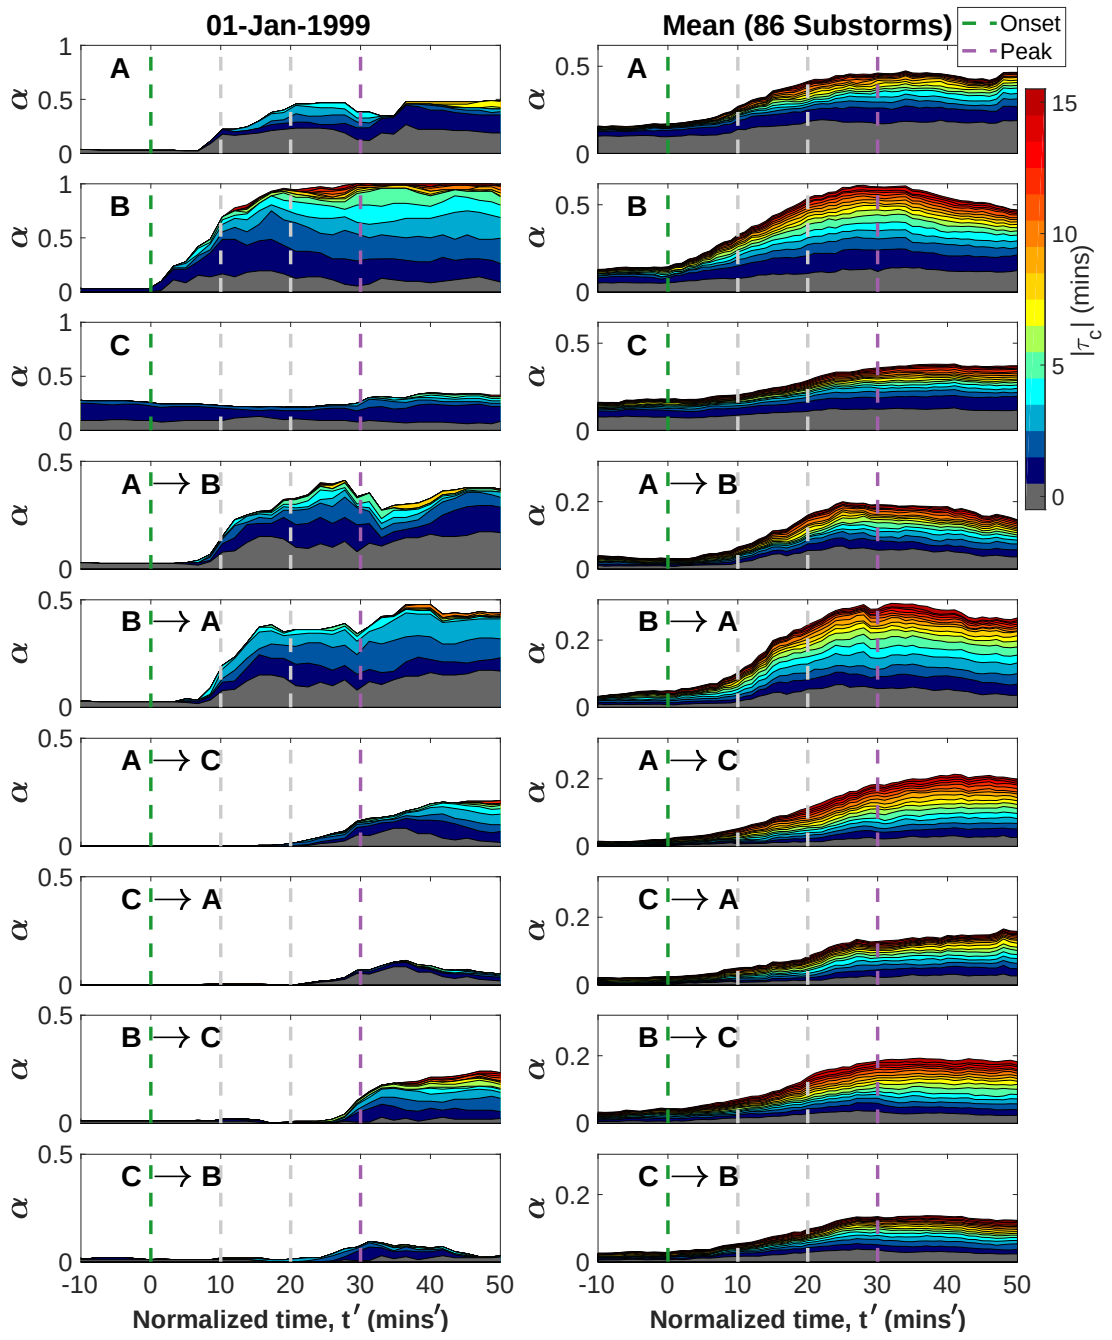

Figure 4.4: The normalized number of connections, $\alpha(t', \tau_c)$, within and between each of the regions A, B and C, where the boundaries of B are identified by polar VIS after peak expansion ($t' = 45$). In the same format as figure 4.2.

and $t' = 45$ (for figure 4.4). The two figures vary slightly, both from each other and figure 4.2, but the main topology remains similar. The numbers of substorms which have ≥ 3 magnetometers in the various regions, and hence are included in the mean $\alpha(t', \tau_c)$ parameter time series for that region, are shown in table 4.1.

Changing the boundaries has an effect on the amount of substorms included in the

Table 4.1: Number of substorms included in the mean $\alpha(t', \tau_c)$ time series (i.e. the number of substorms which have ≥ 3 magnetometers in each region) for auroral bulge boundaries calculated at three different times.

	A	B	C	AB	BC	AC
$t' = 15$	38	46	70	54	57	67
$t' = 30$	33	63	65	52	52	71
$t' = 45$	37	54	64	46	46	66

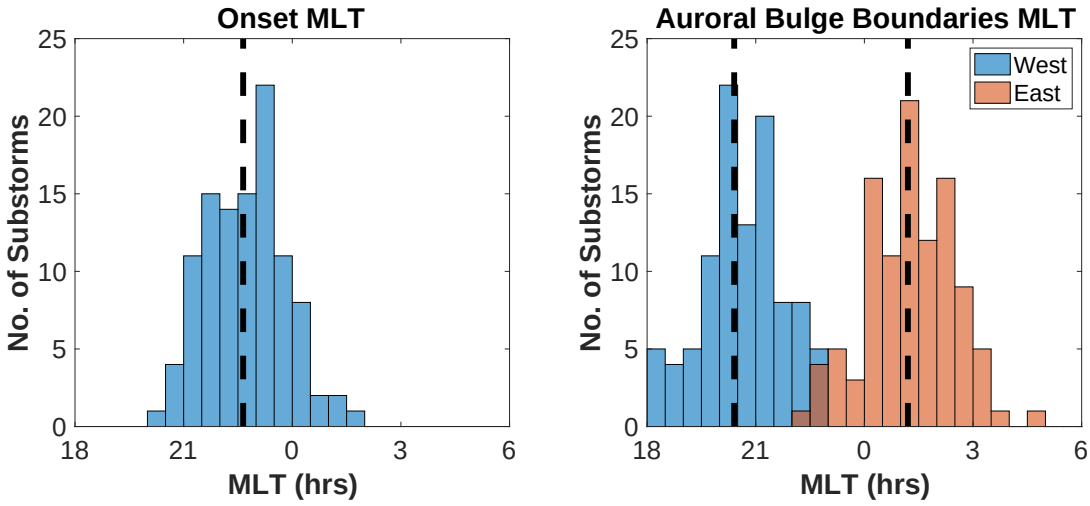

Figure 4.5: Distribution of onset and auroral bulge boundaries (at maximum expansion, $t' = 30$) LT for the substorms used in this chapter.

study, as changing the boundaries excludes or includes substorms based on how many magnetometers they contain in each region (we only include substorms containing ≥ 3 magnetometers in a region). Region A has the fewest number of substorms included. This is unsurprising as the spatial distribution of the westward edge of the auroral bulge (see figure 4.5) is generally at ~ 21 hours LT, meaning region A is generally < 3 hours wide and therefore has a smaller area for magnetometers to be located. On the other hand, region C is bound by the eastward edge of the auroral bulge which is on average at ~ 1.5 hours LT; only seven substorms have their eastward edge at > 3 hours LT. Therefore region C is almost always > 3 hours wide and contains on average more magnetometers than region A. However, the minimum number of substorms included in the average time series parameter for this study is 33 (region A, $t' = 30$), which is still a sufficiently large number for statistical comparison.

For the individual event the eastern boundary only varies by 0.4 hours LT between $t' = 15$ and $t' = 45$ (figure 4.3 and 4.4) so any changes are caused by an individual magnetometer shifting from one side of the boundary to the other throughout the substorm (see figure 4.7). The shift results in $B \rightarrow A$ correlations surpassing the threshold ~ 15 normalized minutes earlier if the eastern boundary was chosen at $t' = 15$ (figure

4.3) as opposed to $t' = 45$ (figure 4.4). The overall coherent patterns of initial westward expansion from the onset location, followed by later eastward expansion, remain consistent with those observed in figure 4.2. The mean number of connections (right column, Figures 4.3 and 4.4) remains robust to the choice in expansion wedge boundaries by showing little changes.

4.3.4 Repeated study with different number of required magnetometers

In section 4.2.3 we set the criteria that there must be ≥ 3 magnetometers in a region for it to be included in the event average. Here we repeated the analysis with the more restrictive criterion of ≥ 7 magnetometers and found very similar results. Figure 4.6 plots the mean normalized number of connections, $\alpha(t', \tau_c)$, for substorms with ≥ 3 magnetometers per region (left column) and that of substorms with ≥ 7 magnetometers in a region. The overall shape of the time-series remains very similar. If we only include the regions of a substorm with ≥ 7 magnetometers there are slightly fewer longer lagged connections (slow propagating), particularly in the regions east and west of onset (A and C) but the overall change is minimal; the same is true between regions. If we include all regions which have ≥ 3 magnetometers the sample size will be larger (see table 4.2). The similarity between of the left and right column proves how robust the mean $\alpha(t', \tau_c)$ parameter is to substorm expansion, even when the left column can contain more than twice as many substorms (see table 4.2, region B).

Table 4.2: Number of substorms included in the mean $\alpha(t', \tau_c)$ time series for each magnetometer number requirement.

	A	B	C	AB	BC	AC
≥ 3 magnetometers	33	63	65	52	52	71
≥ 7 magnetometers	16	30	44	31	28	56

4.3.5 Network maps of spatial correlation

Here we plot the network maps for the isolated substorm event to highlight the importance of the spatial coverage and geographical locations of the highly correlated magnetometer pairs. The detailed network maps of the event which occurred on 01−Jan−1999, for the times represented by the vertical dashed lines in figure 4.2, between the times of onset and peak auroral bulge expansion, are provided in figure 4.7. The geographical location of the sub-networks within regions A, B and C (identified in figure 4.2) can be overplotted on a map centred at magnetic north. The time evolution of the directed networks for these events are binned into propagation speed intervals of $\tau_c = 0$ (left

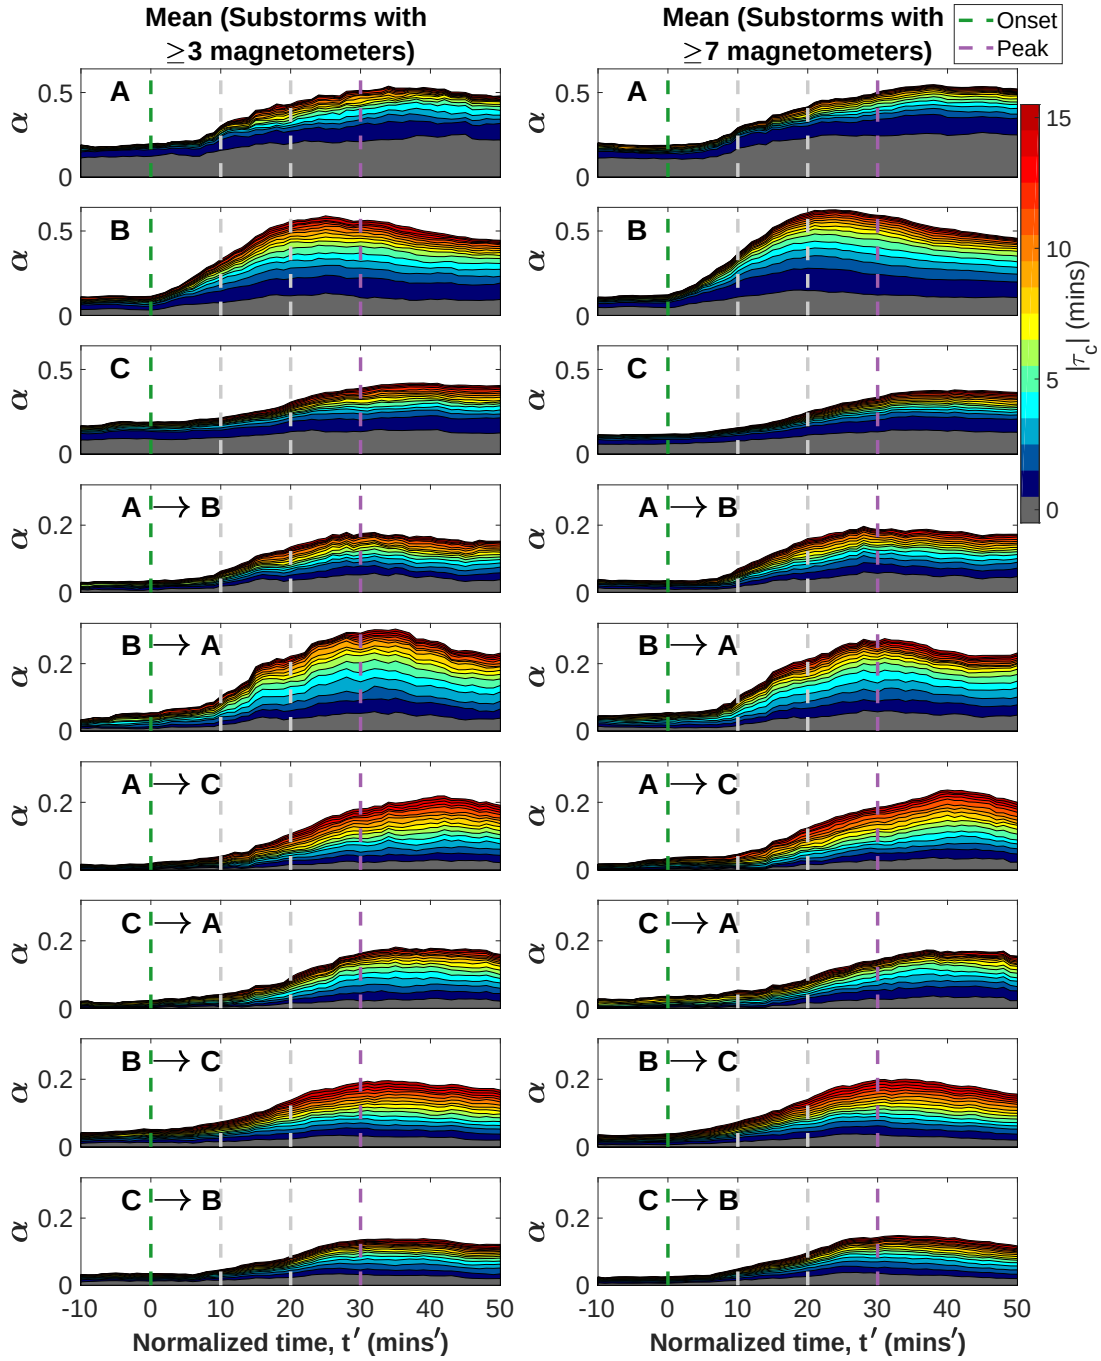

Figure 4.6: The normalized number of connections, $\alpha(t', \tau_c)$, within and between each of the regions A, B and C. Substorms with ≥ 3 (left, the same as the right hand column of figure 4.2) or ≥ 7 (right) magnetometers in a region are included in the mean time series. In the same format as figure 4.2.

column) and $1 \leq |\tau_c| \leq 15$ (right column). The network maps plot the full directed network obtained from all available magnetometer stations in the nightside and located between $60 - 75°$ magnetic latitude with the regions A, B and C overplotted. In the maps we look, in detail, at the network at four different instances in time, from on-

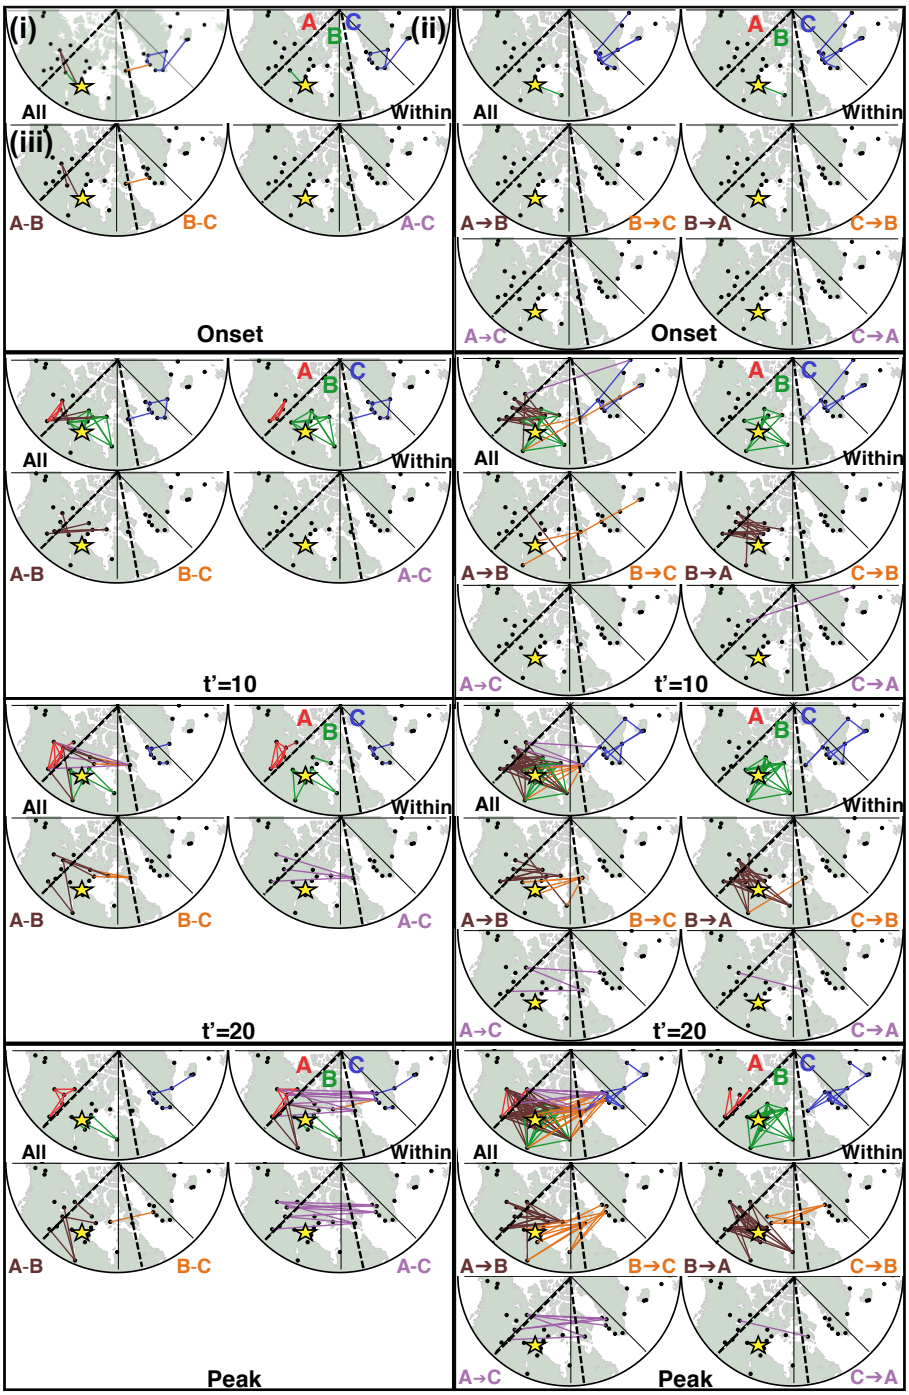

Figure 4.7: The directed network maps of the substorm occurring 01-01-1999, separated into instantaneous ($\tau_c = 0$, left) and propagating ($1 \leq |\ \tau_c\ | \leq 15$, right) connections. Maps are in LT with midnight at the bottom. Dots show the geographical location of active magnetometers. Coloured lines plot network connections. Networks are calculated at minute resolution but here we show ten normalized minute intervals from onset, $t' = 0$, until peak, $t' = 30$. (i) at each time, plots the entire network, (ii) the sub-networks strictly within regions and the remainder are the connections that cross the expansion wedge boundaries (as noted on maps). Stars and dashed lines indicate the onset and auroral bulge boundaries (at the time of peak expansion, $t' = 30$) locations respectively.

set to peak, in intervals of ten normalized minutes. The zero lagged connections (left) correspond to the grey time series plotted in figure 4.2, whilst the propagating connections are represented by the coloured time series. The maps provide more information about the location of the magnetometers in LT with respect to the onset location and expansion wedge boundaries at the time of peak expansion (found using polar VIS). The network in figure 4.7 is equivalent to the adjacency matrix right of the diagonal, i.e. $\mathbf{A}_{ij}$, where $j > i$.

There is a clear change in network structure as the substorm expands. At onset there are connections in the dusk sector, within region C, but few in or between other regions. We then see an increase in correlation around the onset location, within region B, and westward propagation and/or expansion from B→A at $t' = 10$. At $t' = 20$ we begin to see propagation/expansion in the eastward direction from B→C. At $t' = 30$ there are many connections in all directions including both instantaneous and lagged connections between A and C; there is correlation across the entire nightside.

In section 4.3.1 we saw that region C had the smallest increase in connectivity in response to substorm onset, with $\sim 20\%$ of magnetometer pairs within C being highly correlated throughout the substorm. What we did not plot in the parameter time series from figure 4.2 is that the eastward polar VIS boundary is at 1.1 LT, leaving region C with both a larger spatial area and greater number of magnetometers than in regions A or B. Additionally, only a single magnetometer was located between $0 - 2$ LT. It is possible that due to this magnetometer spacing region C does not observe the increased magnetic field perturbations relating to the SCW. Region C, in the case of this substorm, may only be experiencing the enhancements from the DP2 convection centred around dusk and dawn, which are consistent from the growth phase of the substorm, before onset. The maps highlight the importance of the spatial coverage and geographical locations of the highly correlated magnetometer pairs.

4.4 Summary I

In this section we summarise the content of this chapter up until this point. The following section will look in more detail at the directional information of the data set. In sections 4.1-4.3 we used the full set of SuperMAG ground-based magnetometer observations of isolated substorms to quantify the time evolution of patterns of spatial correlation. If the observed pattern of spatial correlation between magnetometer observations captures ionospheric current patterns then we can directly test different models for substorm ionospheric current systems. We have obtained the first directed networks for isolated substorms. Each connection in the network indicates when the maximum CCC between the vector magnetic field perturbations seen at each pair of

magnetometers exceeds an event and station specific threshold. The maximum of the CCC corresponding to each connection in the network can occur at a non-zero time lag. The resulting *directed* network then contains information, not only on the formation of coherent patterns seen by multiple magnetometers, but also on the propagation and/or expansion of these spatially coherent structures.

To gain insight on the ionospheric current system during a substorm, we obtained specific time-varying sub-networks from the data which isolate specific physical regions. These regions are west (A), within (B) and east (C) of the bulge boundaries for each substorm (obtained from polar VIS images at the time of peak expansion). We presented both a study of an individual event, which has at least seven magnetometers in each of these regions for the duration of the substorm, as well as the average of the network properties of 86 substorm events. If the observed pattern of spatial correlation between magnetometer observations captures ionospheric current patterns, we find the following sequence of events in terms of key time ranges after onset: $0 \leq t'_1 < 10$, $10 \leq t'_2 < 20$, $20 \leq t'_3 < 30$ and $t'_4 \geq 30$ (t' is normalized time [Gjerloev et al., 2007]):

- Pre- onset, the pre- and post-midnight regions A and C each have a relatively weak coherent pattern consistent with general magnetospheric convection (DP2) and are not coherent with each other.
- A dominant substorm current wedge (SCW) forming around the onset location (within region B) at the time of onset, t_1, which reaches maximum spatial correlation at t_2, half way through the expansion phase.
- This is followed by a westward expansion of this SCW (starting at t_2, with peak at t_3) contemporaneous to and coherent with a current system in the pre-midnight region (within A).
- An additional weaker eastward expansion of the SCW (starting slowly at t_2 with peak at t_4). The signal of a self-contained current post-midnight (region C) is relatively weaker and occurs late in the substorm. The enhancement of C is delayed with respect to that of A.
- Following the SCW expansion, A and C are coherent with each other, but at the same time are coherent with the SCW. This is consistent with a combination of convection and expansion of the SCW.

These conclusions are drawn from the averaged network over 86 isolated substorms. Although the overall spatio-temporal timings revealed by this network analysis are reasonably consistent between individual events and the 86 event average for the formation of a SCW around onset (B) and it's expansion both east (B→C) and west (B→A), the exact timings of the current system evolution varies. Variability between events could

be intrinsic or could relate to the observing conditions, such as differing magnetometer spatial coverage or the static choice of location for region boundaries.

4.5 Directional information

In sections 4.1-4.3 we used the full set of SuperMAG ground-based magnetometer observations of isolated substorms to quantify both the formation and expansion of substorm current systems in terms of spatial correlation. In this section will look in more detail at the directional information of the data set and expand on the sequence of timings in which the SCW forms by analysing the angle at which signals propagate and/or expand in time.

4.5.1 Average information flow direction

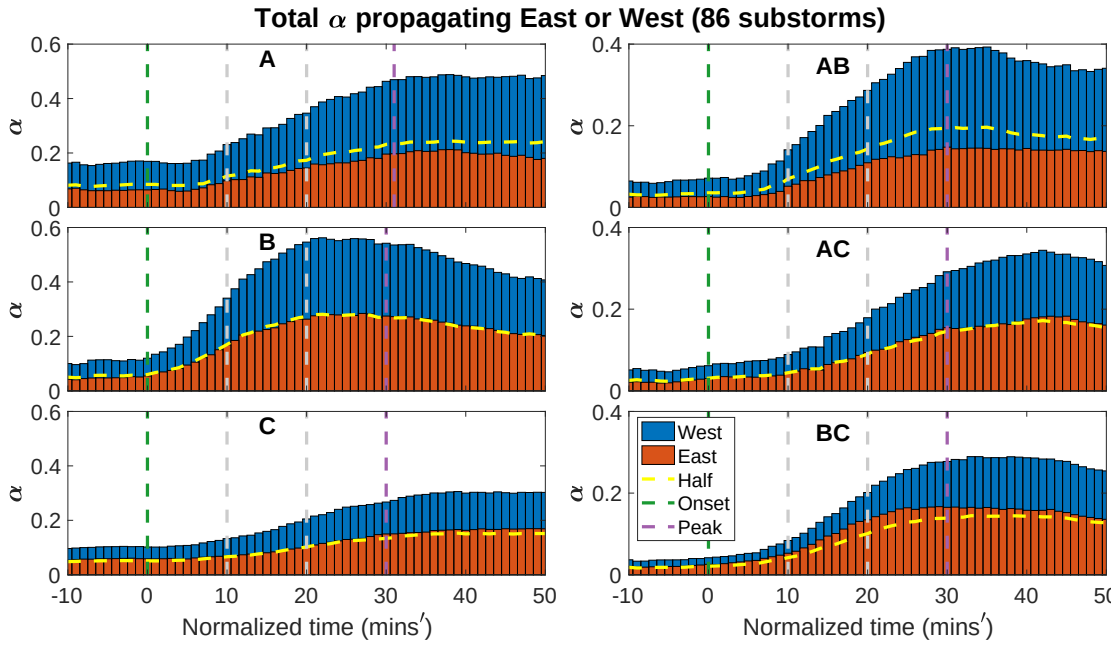

Figure 4.8: The normalized number of connections at each normalized time divided into propagation/expansion East or West. There are clearly more connections going westward from $B \rightarrow A$ and eastward from $B \rightarrow C$ after onset. Between regions A and C, and within regions B and C, there are as many connections propagating Eastwards as Westwards. There is slightly more propagation westward within A.

Directed networks provide information about current system propagation and/or expansion in terms of highly correlated magnetometer measurements, often with a time lag of maximal CCC. Therefore, when two magnetometers are connected within a network it can be useful to know in which direction, in terms of magnetic latitude and magnetic

local time, that connection is propagating in time. Figure 4.8 plots the proportion of the normalized number of connections, $\alpha(t')$, which are propagating/expanding east or west. There are clearly more connections going westwards from $B \rightarrow A$, i.e. region B observes the perturbations first and region A observes a correlated perturbation at a later time; in the same way there are more connections eastwards from $B \rightarrow C$ after onset. Within regions B and C, and between regions A and C, there are as many connections propagating eastwards as westwards. There is slightly more propagation westward within A. Figure 4.8 is a simple way to compare the proportion of eastward to westward propagation. To further analyse the directional and lagged information we can use a polar plot.

4.5.2 Conversion to polar coordinates

We first convert the magnetic coordinates into the polar coordinate system such that the connections for a given region propagate from a reference point, along a line, at an angle towards a point on a plane. The reference point is the magnetometer, i, that measures the (correlated) windowed time series segment prior to the magnetometer, j (i.e. $i \rightarrow j$). The angle between the two magnetometer stations is called the azimuth. It will be denoted as Θ_{ij} and is calculated as follows:

The difference in MLT is denoted as $\Delta\theta_{ij}$ and is calculated such that the distance is always the shortest route possible,

$$\Delta\theta_{ij} = \begin{cases} \theta_j - \theta_i, & \text{if } \|\theta_j - \theta_i\| \leq 12. \\ \theta_j - \theta_i - 24, & \text{if } \theta_j - \theta_i > 12. \\ \theta_j - \theta_i + 24, & \text{if } \theta_j - \theta_i < -12. \end{cases} \tag{4.1}$$

The difference in magnetic latitude is denoted $\Delta\phi_{ij}$ and calculated as,

$$\Delta\phi_{ij} = \phi_j - \phi_i. \tag{4.2}$$

Both $\Delta\theta_{ij}$ and $\Delta\phi_{ij}$ can be positive or negative where we take the positive direction as j being further east/north than i. The angle (azimuth) between the normalized differences is Θ_{Tij}, where $-\frac{\pi}{2} \leq \Theta_{Tij} \leq \frac{\pi}{2}$ and is calculated by,

$$\tan\Theta_{Tij} = \frac{\Delta\phi_{ij}}{|\Delta\theta_{ij}|} \times \frac{12}{15}. \tag{4.3}$$

In order to view the direction of lagged expansion/propagation for all magnetometers

we translate Θ_{Tij} to Θ_{ij}, where $-\pi \leq \Theta_{ij} \leq \pi$,

$$\Theta_{ij} = \begin{cases} \Theta_{Tij}, & \text{if } \Delta\theta_{ij} \geq 0. \\ \pi - \Theta_{Tij}, & \text{if } \Delta\theta_{ij} < 0 \text{ and } \Delta\phi_{ij} > 0. \\ -(\Theta_{Tij} + \pi), & \text{if } \Delta\theta_{ij} < 0 \text{ and } \Delta\phi_{ij} \leq 0. \end{cases} \tag{4.4}$$

If the azimuth, Θ_{ij}, has a negative lag we can conclude that the propagation is from $j \to i$ and hence,

$$\Theta_{ji} = \begin{cases} \Theta_{ij} + \pi, & \text{if } \Theta_{ij} < 0. \\ \Theta_{ij} - \pi, & \text{if } \Theta_{ij} \geq 0. \end{cases} \tag{4.5}$$

Negatively lagged connections can then be included in the polar coordinate system with the positively lagged connections. We have excluded zero lag connections from this study as they contain no directional information. When the directional information has been converted to polar coordinates all network connections can be viewed as originating from one point and propagating/expanding outward from that point. Figure 4.9 is an example of this polar coordinate system for the time of peak expansion of the example substorm (01/01/1999).

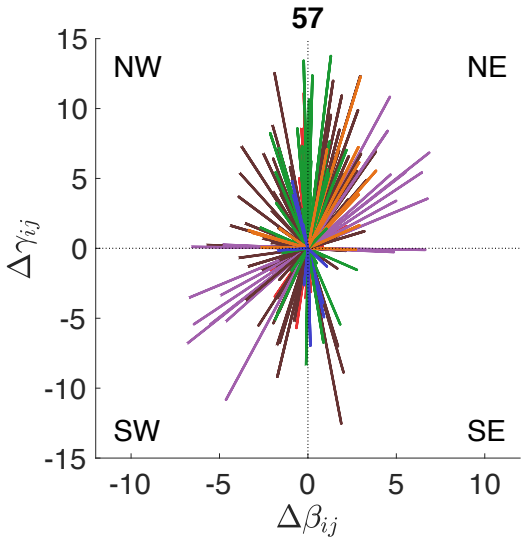

Figure 4.9: The connections (edges) from the network at lag of maximal correlation from the 01/01/1999 substorm at the time of peak expansion, in polar coordinates. The x coordinate system is the difference in LT ($\Delta\phi_{ij}$) and the y coordinate system is the difference in latitude ($\Delta\theta_{ij}$). The connections are the same as those in figure 4.7 at peak but they are in magnetic coordinates.

Figures of the type of figure 4.9 are not informative in terms of analysis. From looking at a single snapshot of an isolated substorm, we see many connections in all directions.

In order to compare multiple substorms we can bin values of the azimuth, Θ_{ij}, and create a simplified plot in the form of a polar histogram where the radius of each bin represents the proportion of the connections with that range of values of Θ_{ij}.

From the 86 substorms with good spatial coverage we can analysis the direction connections are propagating/expanding for multiple events by using the polar histograms. Figures 4.10-4.11 shows the mean direction of propagation/expansion of the 86 isolated substorms. The azimuthal bins are $\Theta = \frac{\pi}{8}$ wide and the radius of each bin is the mean normalized number of connections within that range of Θ_{ij}. We show polar histograms in 10 normalized minutes intervals from onset, $t' = 0$, until peak, $t' = 30$. Additionally, we have lag information so the range of directions at different lags are colour coded. The polar histograms at each lag have been stacked, one on top of each other, with the colour representing the lag, where $1 \leq \tau_c \leq 15$.

4.5.3 Average Azimuth of Multiple Events

In figure 4.10 we see that at onset, on average substorms have few connections (few magnetometers are correlated above a threshold) in all three regions. All of the azimuthal bins contain $< 2.5\%$ of the possible network connections. Ten normalized minutes later the correlation within region B has increased and we now have up to $\sim 10\%$ of possible connections propagating in the bins with azimuth in the north direction. At $t' = 20$ there are further connections within region B, with several bins containing $> 5\%$ of possible network connections, but there is now also correlation with southwards propagation within region A ($\sim 4\%$ of possible connections) and predominantly in the northwards direction within region C ($\sim 3.5\%$ of possible connections). By the time of peak expansion, $t' = 30$, the radius of the polar histogram representing region B has increased from a maximum of $\sim 2\%$ to $\sim 7\%$ of possible network connections. Both numbers are relatively small but recall we have excluded zero lag connections which account for $\sim \frac{1}{5}$ of the average network within region B at peak (see figure 4.2), the maximum normalized number of connections was ~ 0.6 (again see figure 4.2) and here the total normalized number of connections has been divided (un-uniformly) into 16 bins of width $\frac{\pi}{8}$. To put the polar histograms in the context of the global increase in correlation, and hence connections, we have included a smaller circle within the large, of which the radius is the total normalized number of connections (with lags $1 \leq \tau_c \leq 15$) divided by the number of bins (16), i.e. the radius that all bins would have if the propagation was uniform in every direction. From this we see that the north and south propagation/expansion is significantly more than that of the average.

It is important to remember that magnetometers are in fixed locations. By taking an average over many substorms we have measurements from many locations and not

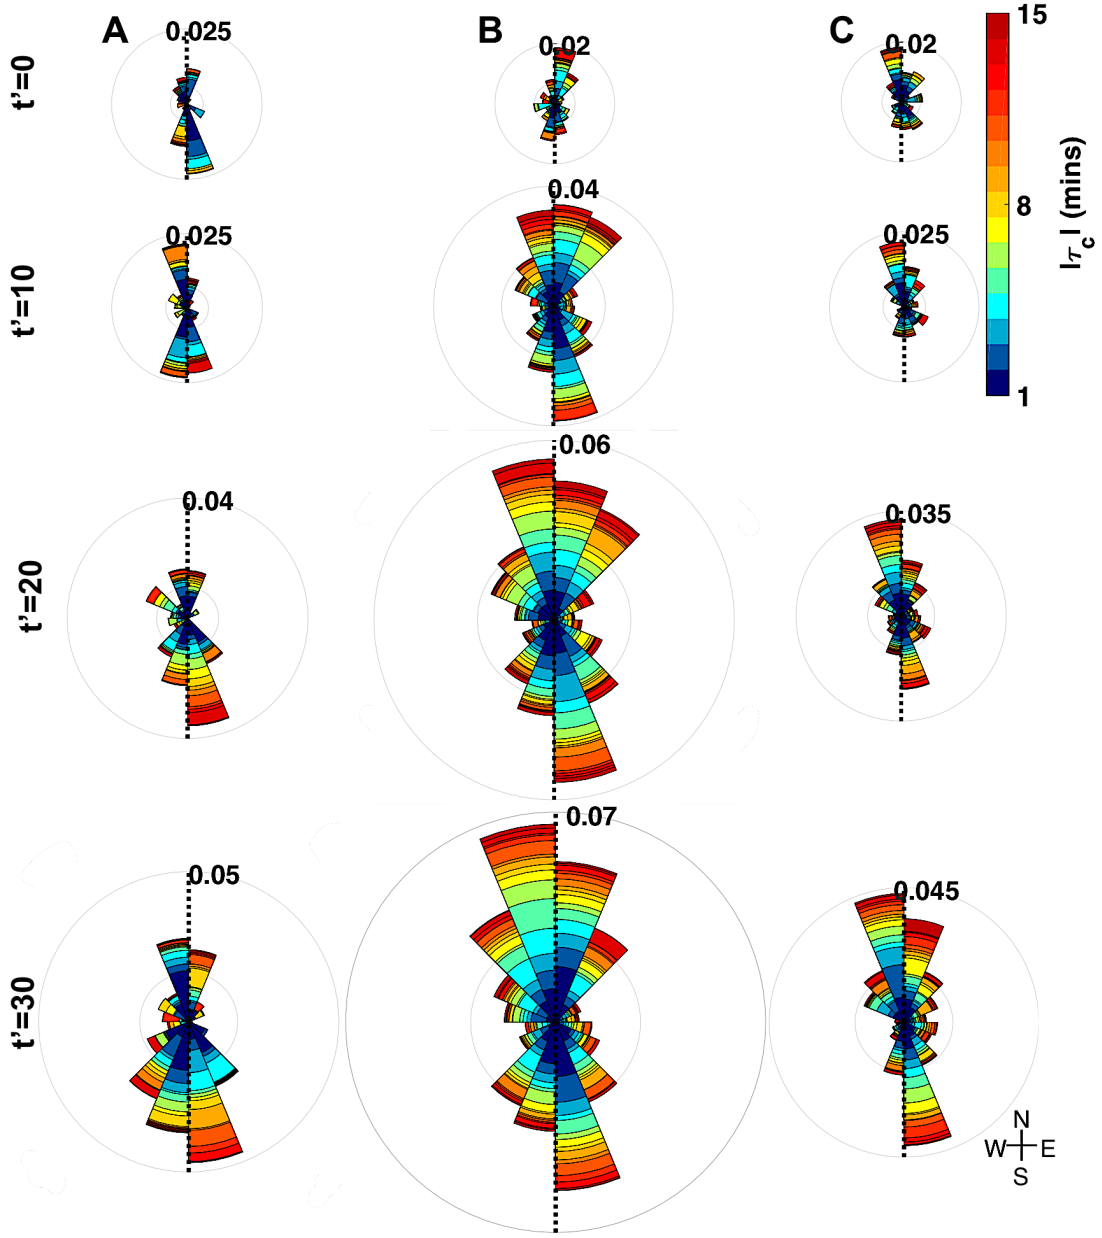

Figure 4.10: Stacked polar histograms of the mean angle between the two connected magnetometer stations, Θ_{ij}, i.e. the direction of propagation/expansion, for 86 substorms. The bins are $\Theta = \frac{\pi}{8}$ wide and the radius of the histograms is the mean normalized number of connections contained within that Θ_{ij} bin. The relative scale of each polar histogram is consistent throughout. The smaller circle within the larger circle, at each time, represents the normalized number of connections if the network was uniformly propagating in all directions. The top panel of polar histograms represents the average network at onset and the bottom the average network at maximum auroral bulge expansion. The columns are the polar histograms of directional information contained in each regions (A, B, C). The colour is the proportion of connections with each lag. Most connections within regions propagate in the north or south direction.

just $\sim 30-50$ fixed points, but there may be biases towards a preferred direction. However, one of the most interesting features that we can see from this type of plot is a clump of bins with larger radius, such as the southward trend in region A. This feature is independent of magnetometer positional bias as, by design, the network can have propagation from $i \rightarrow j$ or $j \rightarrow i$. At $t' = 30$ all four bins with the azimuth in the southward direction exceed the circle of uniform propagation, accounting for $> 60\%$ of the connected network.

Within the regions the MLT extent between two magnetometers is limited by the region boundaries and therefore, the latitudinal distance could be proportionally much greater than the longitudinal. When we look between regions this becomes less significant as there can be greater longitudinal distance between magnetometers. In figure 4.11 there are much clearer preferences in direction.

In figure 4.11, again we see an increase in the normalized number of connections contained within any azimuthal bin between onset ($t' = 0$) and the time of maximum expansion ($t' = 30$). This is a feature discussed in section 4.3 and as such needs no further discussion. The new information is the azimuthal direction of the increased number of connections. The direction of propagation/expansion between regions A and B has more westwardly directed signals (correlated with a time lag associated with magnetometer in the west) then eastwardly directed; the magnetic field signature is first measured in region B and several minutes later in region A. Again this westward expansion was discussed following figure 4.2, but here we have much more detail about the exact direction. There are few connections in the bins closest to the north-south line; this is to be expected as for connections to flow between regions A and B they must have some latitudinal difference. The two bins with the largest radius are those containing connections with SW propagation/expansion ($\sim 5\%$ of possible connections in each bin). The azimuth bins gradually decrease in radius as the propagation direction becomes more westwardly ($\sim 3.5\%$) and then northwardly ($\sim 2.5\%$). This south-westwardly trend is consistent with the southwardly propagating signals within region A (figure 4.10). Propagation in the SW direction has longer time lags than those in the NW direction; $\sim 50\%$ of the connections in the SW direction have $\tau_c \geq 6$ minutes compared to $\sim 30\%$ in the NW direction. This implies that a coherent current system expands first NW and then, to a large extent, SW from the onset location.

Between regions A and C there is little directional difference, temporally or spatially, with an increase in connections propagating/expanding in both westwardy and eastward directions.

Between regions B and C we see mainly eastward propagation from $B \rightarrow C$. The maximum bin having a radius representing $\sim 3.5\%$ of the possible connections compared to $\sim 5\%$ between regions A and B, and regions A and C. The largest bin is in the SE

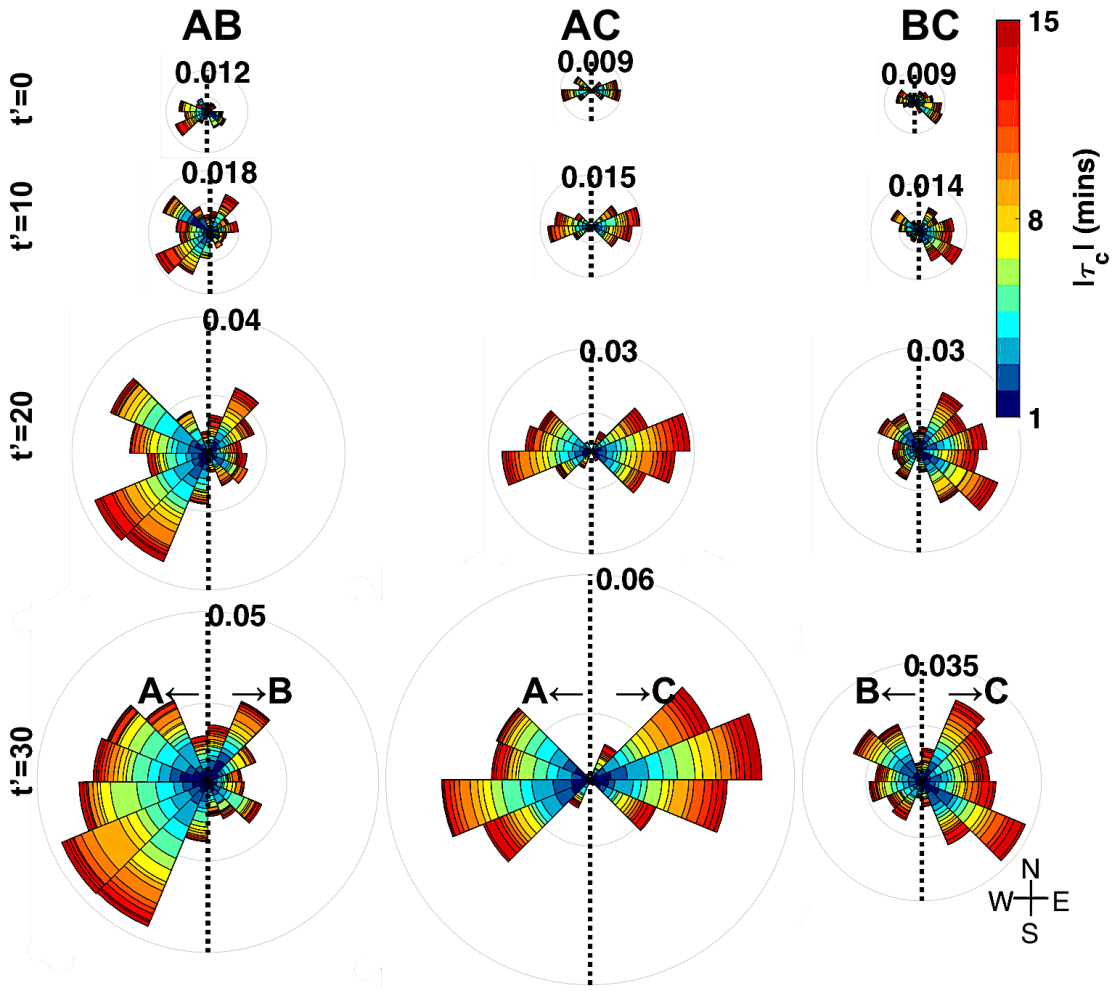

Figure 4.11: Stacked polar histograms of the mean azimuth between the two connected magnetometer stations, Θ_{ij}, i.e. the direction of propagation/expansion, for 86 substorms. The bins are $\Theta = \frac{\pi}{8}$ wide and the radius of the histograms is the mean normalized number of connections contained within that Θ_{ij} bin. The relative scale of each polar histogram is consistent throughout. The smaller circle within the larger, at each time, is the normalized number of connections if the network was uniformly propagating in all directions. The top panel of polar histograms represents the average network at onset and the bottom the average network at maximum auroral bulge expansion. The columns are the polar histograms of directional information contained between each region (A-B, A-C, B-C). The colour is the proportion of connections with that lag. It is clear that more connections propagate from $B \rightarrow A$ and from $B \rightarrow C$ but eastward and westward propagating between regions A and C are approximately equal.

direction but all the other bins in the E direction are approximately equal so we cannot claim a southwards preference in propagation. Likewise the proportion of each bin with each value of $|\tau_c|$ is approximately equal.

4.5.4 Summary II

In section 4.5 we expanded on the observed patterns in spatial correlation between magnetic field perturbations and how they can capture ionospheric current patterns by looking at the angle at which these patterns propagate/expand in time. By converting the magnetic coordinates of the magnetometers into polar coordinates we can analyse what direction coherent patterns propagate in time. Again we divide the nightside into the regions A (west of the auroral bulge), B (around the onset location) and C (east of the auroral bulge) for each substorm; this allows us to study how current system behaviour varies across the nightside, as well as compare multiple events.

When we study the mean spread of the azimuth (angle of propagation) for the 86 events we find the same sequence of events as found in section 4.3, but with the full range of angles at which correlated signals propagate and/or expand:

- A dominant SCW forming around the onset location (within region B) at the time of onset, $t' = 0$, with predominantly northward propagation.
- This is followed by a westward, and predominantly south-westward, expansion of this SCW ($t' = 10$, with peak at $t' = 30$) contemporaneous to and coherent with a current system in the pre-midnight region (within A), where propagation and/or expansion is southward.
- An additional weaker eastward expansion of the SCW (starting slowly at $t' = 10$ with peak at $t' = 30$). The signal of a self-contained current post-midnight (region C) is relatively weaker and occurs late in the substorm ($t' = 20$ and $t' = 30$). There is not a preferred north or south direction of propagation post-midnight.
- Following the SCW expansion, A and C are coherent with each other, but at the same time are coherent with the SCW. This is consistent with a combination of convection and expansion of the SCW. Propagation is approximately equivalent in the east and west directions.

The expansion westward of the onset location (B→A) implies that a coherent current system has faster expansion in the northwest direction but a slower, and greater or more consistent expansion towards the southwest ($\sim 2\times$ as many southward as northward). West of the auroral bulge this extension towards lower latitudes is evident ($> 60\%$ of connection propagating southward). Propagation both across the auroral bulge (between regions A and B) and eastward from the onset location (B→C) are uniformly spread between north and south.

Chapter 5

Community structure of isolated substorms

5.1 Introduction

In chapter 4 [Orr et al., 2019] we made the first application of directed networks to the full set of 100+ ground-based magnetometers collated by SuperMAG which allowed us to test these ideas on a large set of isolated substorm events. To objectively test the hypothesised large-scale SCW models listed in chapter 1, we used the raw network properties to resolve timings of propagation/expansion of the current wedge within and between three predefined spatial regions and found timings of a consistent sequence in which the classic SCW forms. In particular this required the polar VIS data to be analysed by hand to determine the region boundaries for each event, and necessarily makes the assumption that these regions will unambiguously demarcate the various features of the dynamically evolving current system. We now want to move to a fully automated methodology with the goal of finding a parameter for the network that characterises the full spatio-temporal pattern. The majority of work in this chapter is from the paper submitted to Nature Communications, "L Orr, SC Chapman, JW Gjerloev and W. Guo. Network Community Structure of Substorms using SuperMAG Magnetometers".

The chapter is organised as follows. Section 5.2 will give an overview of the method, including the data used and an explanation of how to calculate the networks and their communities. Section 5.3 presents an example of a single substorm and a statistical survey across all 41 isolated substorm events. A summary of the material up until that point is contained within section 5.4 with some additional examples of community structure in section 5.5.

In this chapter, we perform community detection on the dynamical networks constructed from all magnetometers collaborating with the SuperMAG initiative for the first time. In the context of ionospheric current systems, communities are subsets of the 100+ ground-based magnetometer observations that are more strongly correlated with each other than they are with the rest of the network. The method can identify whether one or more community (i.e. current system) exists and how this changes in time, and importantly, once the raw network has been constructed, it does not rely on any assumptions or other inputs.

We perform this analysis across 41 isolated substorm events at one minute temporal resolution and we find a robust and consistent configuration of ionospheric substorm currents evolving as the substorm progresses. Multiple discrete current systems that are present before onset are found to progressively merge into a single coherent SCW. This merging occurs over 10-20 minutes following onset and characterises the peak expansion phase of the substorm. Our analysis reveals consistently, across many events, a final, fully expanded SCW state which is a single spatially-extended correlated system, consistent with a single SCW. Our results underline the central role of time-dynamics in models for the substorm current system, since an observational 'snapshot' of the system after onset, but before the SCW has fully formed, would suggest multiple current systems which may indeed differ from one event to another. This may resolve much of the controversy surrounding models for the substorm ionospheric current system.

5.2 Methods

In chapter 3 the methodology for calculating a directed network was outlined and in section 2.2.3 community detection methods were described. In this section we apply these methods to the directed network.

5.2.1 Community detection and network parameters

The underpinning methodology for forming the raw directed, time-varying network is detailed in chapter 3, as well as in Dods et al. [2015, 2017]; Orr et al. [2019]. Once the raw network has been calculated we can utilise the network properties, such as the tendency to form communities. Community detection is a method of locating groups of nodes within the network which are locally dense with connections but which have sparse inter-group connections. Communities characterise the meso-scale topology of the network [Newman and Girvan, 2004] and provide insight on its formation and functionality [Caldarelli, 2007]. We apply community detection algorithms to the raw directed time-varying networks obtained from SuperMAG observations of substorm events to char-

acterise the network structure in terms of communities. There are many community detection algorithms that identify optimal dense subgraphs in directed or undirected graphs (summarised in section 2.2.3). We used a variety of community detection algorithms in the igraph package in R [Csardi and Nepusz, 2006] to determine community structure.

Results using the edge betweenness algorithm [Newman and Girvan, 2004] are highlighted here but we have verified that our results are robust against the choice of algorithm; results using different algorithms, including the 'optimal' [Brandes et al., 2007], walk trap [Pons and Latapy, 2005], information mapping [Rosvall and Bergstrom, 2007], leading eigenvector [Pons and Latapy, 2005] and label propagation [Raghavan et al., 2007] algorithms are reported in section 5.5.2 (Figures 5.19-5.23).

Modularity is a measure that has been widely used to evaluate how well the community structure has been captured [Clauset et al., 2004; Pons and Latapy, 2005] and is summarised in section 2.2.3 (equation 2.4).

The modularity parameterizes to what extent the network is characterised by many, separate, communities or one dominant community. If the magnetometer signals become more strongly correlated to each other, forming a single community, the modularity will tend to zero as $Q \rightarrow e_{xx} - e_{xx}^2 \rightarrow 0$ as $0 < e_{xx} < 1$. For a network comprised of smaller, interconnected communities, Q will be finite. In order to compare many events, we then normalize the modularity such that in the quiet interval before each individual substorm, the maximum value of $Q = 1$.

For the networks formed from the SuperMAG set of ground-based magnetometers, the network communities identify the spatially coherent perturbations from ionospheric current systems. The change in the modularity as the substorm progresses identifies how these current systems are changing. The modularity will be maximum when spatially localized coherent perturbations cause a group/community of magnetometers to be internally correlated, but not cross-correlated with other groups/communities of magnetometers. If all auroral latitude magnetometers are highly correlated due to a large-scale ionospheric current system overhead, the modularity will be near-zero. A schematic of the current system models we would expect from high or low modularity is contained in Figure 5.1.

The overall network response is parameterized by the normalized number of connections, $\alpha(t)$ (equation 3.1).

We have performed two tests of the significance of our results. We checked the statistical significance of the modularity of a given substorm by constructing a random phase surrogate for that event. The random phase surrogate provides an estimate of the network

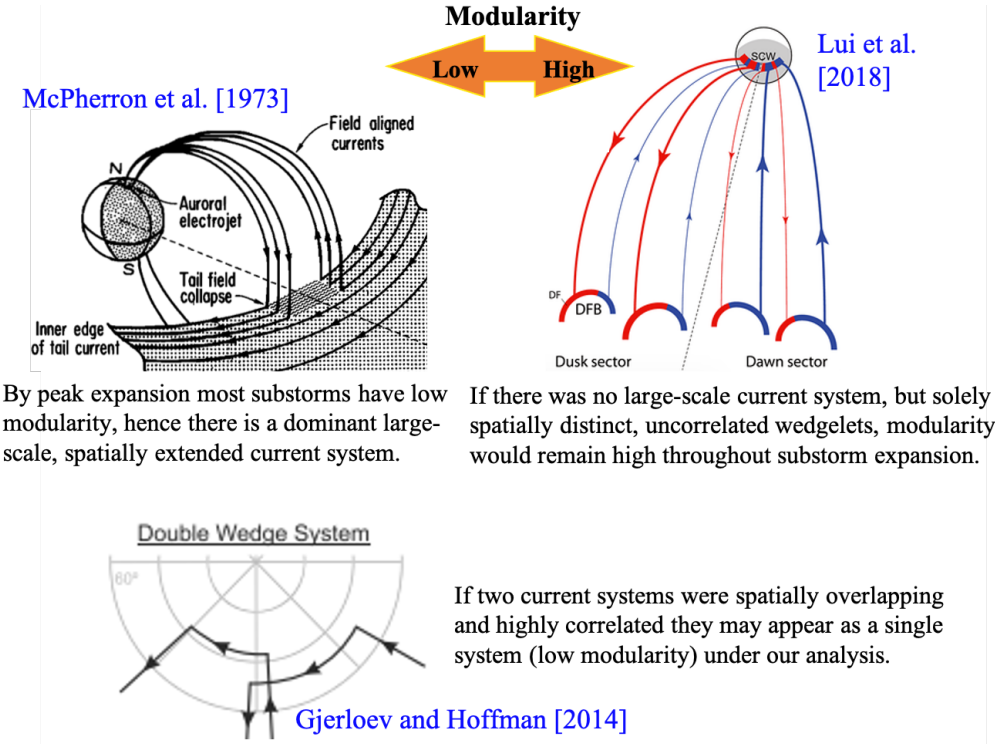

Figure 5.1: A schematic showing our interpretation of high or low modularity. Modularity will be near-zero when all auroral latitude magnetometers are highly correlated due to a large-scale ionospheric current system overhead, e.g. McPherron et al. [1973]. The modularity will be maximum when spatially localized coherent perturbations cause a group/community of magnetometers to be internally correlated, but not cross-correlated with other groups/communities of magnetometers e.g. Liu et al. [2018]. If current systems were highly correlated and not spatially distinct e.g. Gjerloev and Hoffman [2014] they could appear as a single system (low modularity) under our analysis.

properties that could arise 'by chance' in a given set of data and is summarised in section 2.3.3. For each event, the time-series from each magnetometer are Fourier transformed, the phase spectrum is randomized whilst preserving the amplitude spectrum and this is then inverse Fourier transformed to give a surrogate time-series with the same power spectrum as the original signals, but with no time correlation (for an example see Tindale et al. [2018]). We used an iterated amplitude adjusted Fourier transform (IAAFT) Schreiber and Schmitz [1996] method with the Matlab code supplied by Kugiumtzis and Tsimpiris [2010]. The surrogate time series of each magnetometer is then used to calculate the random phase surrogate of the network parameters using the method as described in this chapter; the surrogate is treated in exactly the same way as the original time series data. For each event we repeated this ten times to obtain an average value of the random phase modularity. The random phase surrogate is indicated on plots of our results. We also compared the modularity with that of random networks generated according to the Erdős-Rényi model Erdős and Rényi [1959], constructed with the same (time dependent) number of nodes, this is shown in figure 5.29.

5.2.2 Data and event section

The substorms used in this study are from the list of 116 events described in section 3.2 and previously used in Gjerloev et al. [2007]. We also require that the nightside is quiet for at least one CCC window (127 minutes) before the substorm onset so that the network calculated at the time of substorm onset is not contaminated with previous activity, rejecting events where the SML index [Newell and Gjerloev, 2011a] exceeded 25% of its maximum value (at the peak of the substorm) in the 127 minutes before the start of the substorm. To maintain good magnetometer coverage of the nightside during the substorm we require two or more magnetometers in each three hour window of the nightside (e.g. ≥ 2 magnetometers in each segment of $18-21$, $21-24$, $00-03$ and $03-06$ hours of local time). 41 substorms fulfil these criteria. We will present as an example a substorm with excellent magnetometer spatial coverage alongside an epoch analysis of the normalized modularity of the 41 substorms. Further substorm examples (figures 5.6-5.17) and an epoch analyses of other substorms (figures 5.24-5.26) are included in sections 5.5 and 5.5.2.

5.2.3 A single normalized time-base for the events

To compare multiple substorms we first map each event onto a common normalized time-base such that, once normalized, all substorms share a common onset time and take 30 normalized minutes to develop from onset to the maximum expansion. The method for time normalization, t', developed in Gjerloev et al. [2007] is described in

section 3.2.

5.3 Results

In this section we present a case study of one substorm to show how the community structure of the network changes throughout the substorm expansion and a statistical survey of 41 events which reveals how on average the nightside becomes globally correlated during substorm expansion.

5.3.1 Community structure of a single substorm

Figure 5.2 is an example of how the community structure of the network varies during a substorm for a single event. The network is calculated for a substorm on the 16th March 1997 which has excellent magnetometer spatial coverage of the entire nightside, 18-6 hours of magnetic local time (MLT), between 60-75 degrees magnetic latitude (MLAT). Throughout the substorm there are always ≥ 7 magnetometers in each 4 hours of MLT (i.e. in each of MLT windows $18-22$, $22-02$ and $02-06$), as well as ≥ 4 magnetometers in each 3 hours of MLT (i.e. in each of MLT windows $18-21$, $21-00$ e.t.c.). In all panels of figure 5.2 the abscissa plots normalized time (equation 3.3) so that the onset of the substorm is at $t'=0$ (green dashed line) and the time at which the auroral bulge has reached its maximum expansion is $t'=30$ (red dashed line).

Panels 1 and 2 of figure 5.2 visualise the overall importance in the network, and physical location, of the network communities as a function of time. At each time, each community is indicated by a circle in panel 1. The edge betweenness algorithm does not pre-define how many communities there should be so the number of circles at each time is completely unconstrained and changes during the substorm. The ordinate of panels 1 and 2 plot the MLT, $\bar{\theta_x}(t')$, and MLAT, $\bar{\phi_x}(t')$, of the physical centroid (mean) of the magnetometer stations within each community at each time, t'. The size of the circle reflects the normalized number of connections, $\alpha(t')$ (equation 3.1) contained within the community whilst the colour reflects the proportion of connections with each time lag, $|\tau_c|$. The network has been constructed by identifying connections using the CCC lag at which the CCC between each pair of magnetometers is at its peak so that each connection has an associated lag, τ_c. The magnitude of the lag $|\tau_c|$ indicates whether connections within a given community are formed rapidly (zero CCC lag i.e. $\tau_c = 0$ (grey)) or whether they are associated with propagation and/or expansion (non-zero CCC lags from $1-15$ minutes (blue-red)) [Orr et al., 2019]. The dashed lines overplotted are the east and west edges of the auroral bulge (MLT) in panel 1 and the onset location (MLAT) in panel 2 as seen in Polar VIS auroral images [Gjerloev et al., 2007].

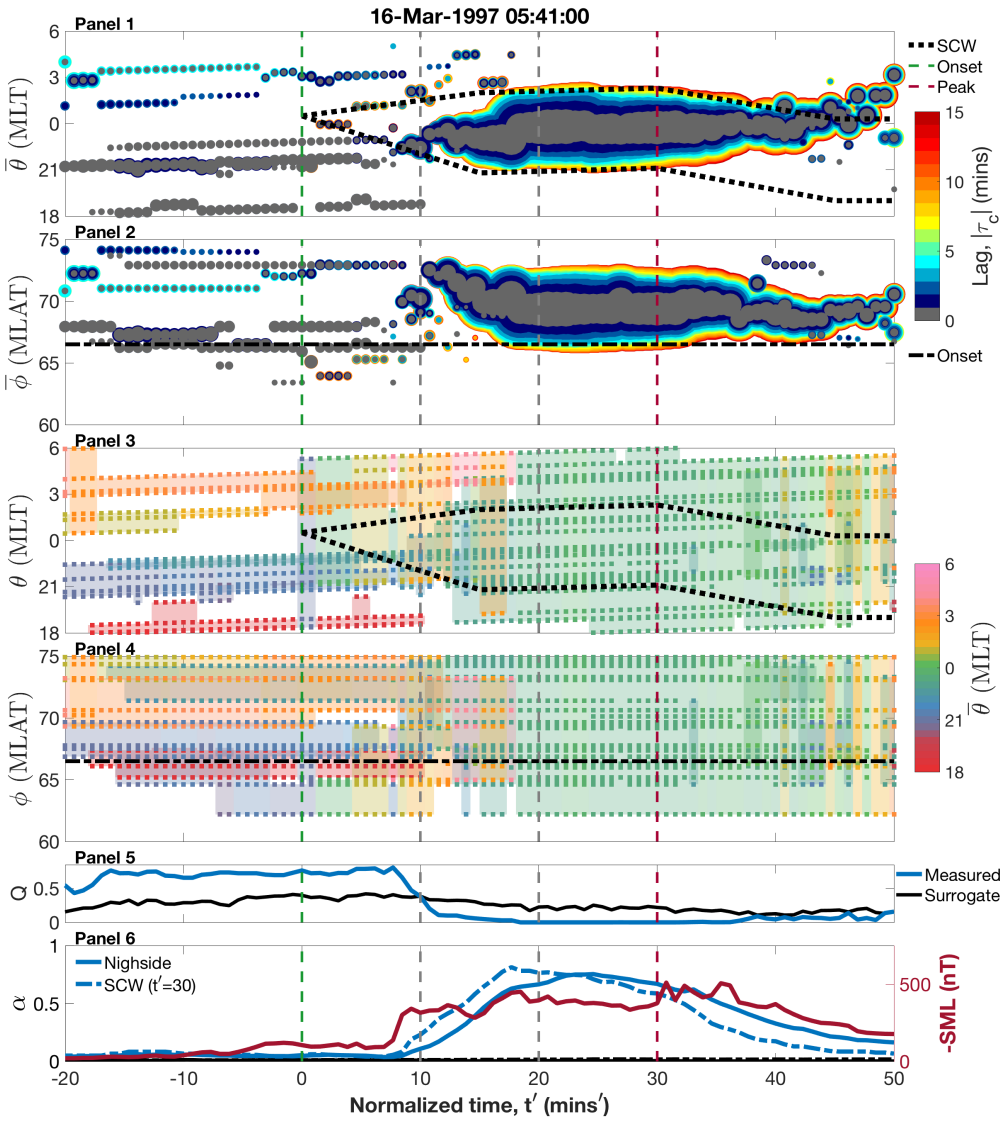

Figure 5.2: The community structure of a substorm on the 16/03/1997. The abscissa of all panels is normalized time ($t' = 0$ is onset (green line) and $t' = 30$ (purple line) is the time of maximum auroral bulge expansion). Panels 1-2 plot the community structure where the size of the circle reflects the normalized number of connections within the community, the ordinate plots the mean MLT/MLAT of the community, $\bar{\theta}_x(t')$ and $\bar{\phi}_x(t')$, and the colour indicates the proportion of connections with each lag, $|\tau_c|$. The dashed lines overplotted are the edges of the auroral bulge (MLT) and the onset location (MLAT), found using Polar VIS. Panels 3-4 show the spatial extent of each community, where the dots are the specific location of the magnetometers and the shading is the extent. colour represents the mean MLT of the stations contained within each community, $\bar{\theta}_x(t')$. Panel 5 plots the modularity, Q, (blue line) and the random phase surrogate result (black line) which is high if the network is densely connected within communities and sparsely connected between them, and low when the system is globally densely connected. Panel 6 plots the normalized number of connections, $\alpha(t')$, both within the nightside and within the SCW, as well as their surrogates. The right ordinate plots (negative) SML.

From before onset until approximately 10 normalized minutes after there are ~ 5 small communities. Each community contains few connections (few magnetometers are highly correlated) and are mostly instantaneous (zero lag, there is no delay of the signal between pairs of magnetometers). These communities spread throughout the nightside at all MLT. These persist beyond onset until $t' \sim 10$ where there is a clear transition in structure. The small communities begin to merge into one large community. By $t' \sim 20$ the small communities have been completely replaced by a single large structure containing many highly connected magnetometers (the circle radius is now large representing a large fraction of the available network). This single community contains about half of connections at < 2 minute lag and half at longer lags, consistent with expansion and/or propagation of the structure. The MLT, $\bar{\theta_x}$, of the centroid of this single large community is located well within the auroral bulge. At $t' \sim 40$ this large structure begins to breakdown as the auroral bulge shrinks.

Panels 3 and 4 of figure 5.2 plot the location of all individual magnetometers that are connected to the network at each time. Each dot represents a single magnetometer and each colour indicates a distinct community, the colour used to label each community corresponds to the MLT of the centroid ($\bar{\theta}(t')$) of that community at each time. In panel 3 the dashed lines are again the east and west edges of the auroral bulge (MLT) and onset MLAT determined from Polar VIS images. There are few, if any, magnetometers in the network below 65° MLAT (ϕ) during the substorm. Before onset we see that the multiple small communities are mostly spatially separated in MLT (θ) and each only involves 3-4 magnetometers. The communities are slightly more overlapping in MLAT (ϕ). By $t' \sim 20$ the entire network is one community, spanning the entire nightside.

Panel 5 plots the Modularity, Q, a measure of the proportion of the network connections between communities compared to connections contained within communities (equation 2.4). If Q is large the network is highly densely connected within communities and sparsely connected between them, whereas if Q is low the system is globally densely connected, that is, dominated by one main community. The modularity for the event (blue line) is compared to that of its random phase surrogate (black line) and we can see that in the event, there is a clear drop in modularity at $t' \sim 10$ from spatially separated localised communities ($Q \sim 0.8$) to one dense global community ($Q \sim 0$) which is clearly distinct from the behaviour of the surrogate which shows little change during the event.

Panel 6 shows the (negative) SML time series, an index of substorm electrojet enhancement, alongside the normalized number of connections, $\alpha(t')$ (equation 3.1) contained both within the nightside (solid blue) and within the auroral bulge boundaries at the time of maximum expansion, $t' = 30$. Both $-$SML and $\alpha(t')$ begin to increase just before $t' = 10$, just as the network begins to merge into one community and drop in modularity, and maximises at $t' \sim 20$, when the entire network is a single highly connected global

structure. The random phase surrogate remains $\sim$ zero during the substorm.

The detailed evolution of the spatio-temporal current system from onset to peak expansion is found to vary between events. We show several examples in section 5.5 (see figures 5.6, 5.9, 5.12 and 5.15) and the modularity before onset varies across these five examples (e.g. $\sim$ 0.4 in 5.6 and $\sim$ 0.8 in figure 5.2). However, we consistently find a transition from many distinct communities to a single coherent community at peak expansion.

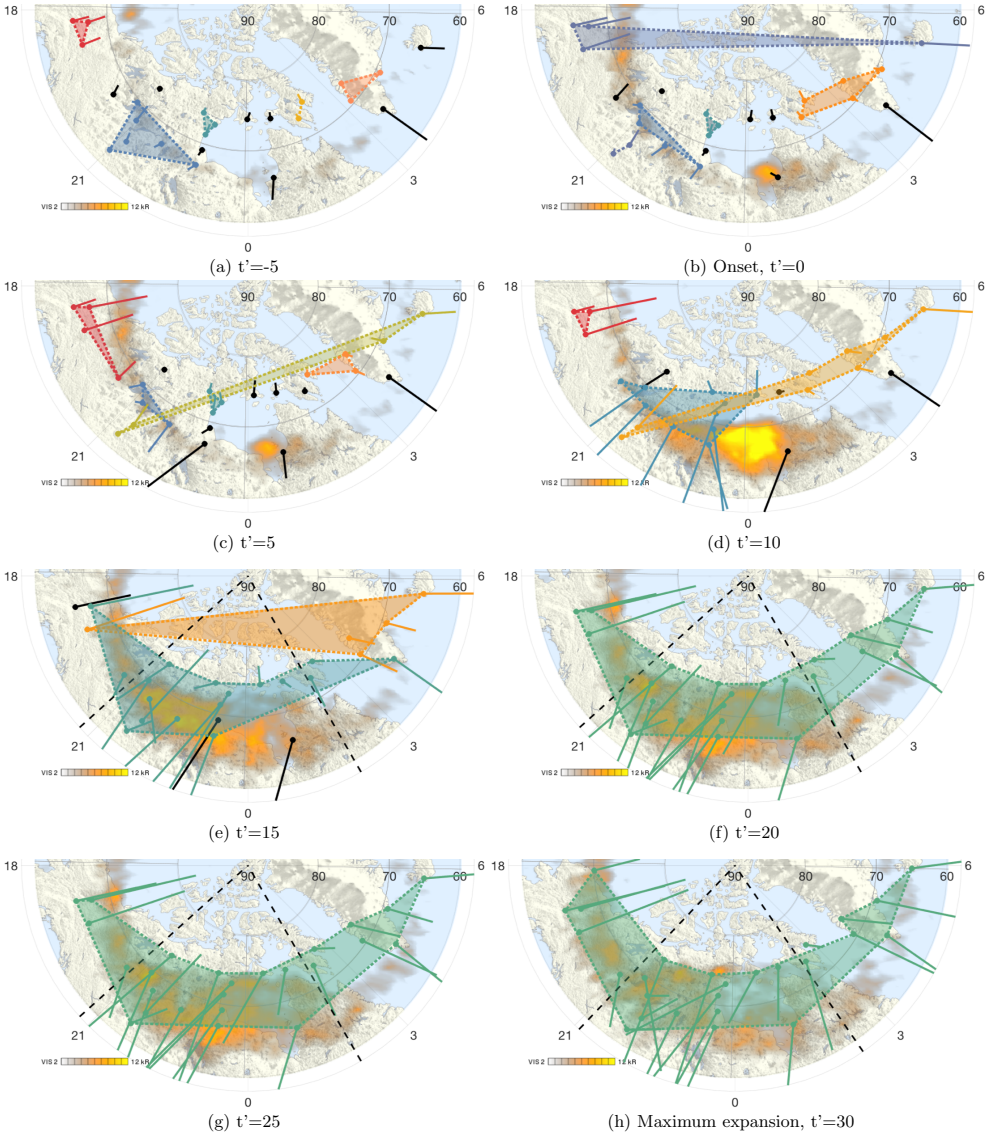

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.3: The magnetic field perturbation vectors (B_N and B_E) are coloured by the the MLT of the centroid ($\bar{\theta}(t')$) of each community for a substorm on 16/03/1997. The colours match those of panels 3-4 in figure 5.2. Each subplot represents a snapshot of the nightside community structure in five normalized minutes from before onset to the time of maximum expansion, corresponding to the times in figure 1. The circles represent ground magnetometers in MLT with the line representing the $B_{N,E}$ vector. Black circles represent nodes that are not connected to the time-varying network. The dashed lines the locations of the auroral bulge as defined in [Gjerloev et al., 2007].

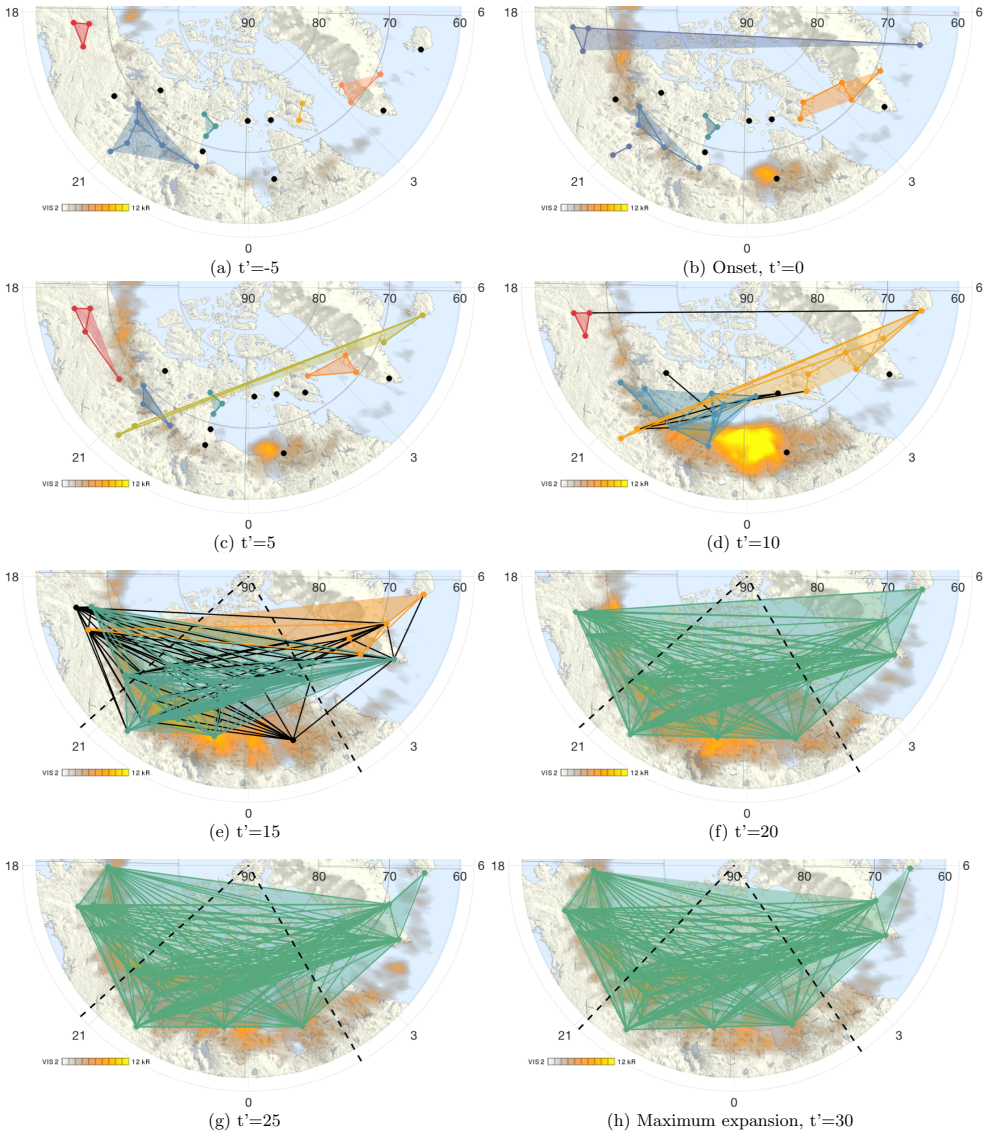

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.4: The community structure of a substorm on the 16/03/1997 which is plotted in the same format as figure 5.3 except that all network connections are plotted (and magnetometer vectors are not plotted). Connections are colour coded for community, black connections are connections between communities. There is a clear change in community structure from $t' \sim 10$, with many more inter-community connections forming at $t' \sim 15$, and merging into a single community at $t' \sim 20$.

We can see this in the physical maps of individual events. Figure 5.3 contains eight snapshots of the same substorm, corresponding to the times indicated by dashed lines in figure 5.2 (see also section 5.5 figures 5.7, 5.10, 5.13 and 5.16)). The snapshots are SuperMAG polar plots of the nightside overplotted with images of Polar VIS. The polar plots are overlayed with the magnetometer locations and their associated magnetic field perturbation vectors (B_N and B_E). The magnetometers which are not part of any community, at each time, are coloured black whilst the magnetometers contained in a community are coloured by the the MLT of the centroid ($\bar{\theta}(t')$) of each community. The colours match those of panels 3-4 in figure 5.2. Each community is also shaded

and surrounded with a dashed line. The full network is shown in figure 5.4. We see from figure 5.3a that just before onset there are five communities, all of which are small and spatially separate, with the largest community only containing six magnetometers. The same picture is seen in figure 5.4a. These communities only begin to incorporate larger numbers of magnetometers at about $t' = 10$ (figure 5.3d and 5.4d) but the communities are still relatively separate (few black connections between groups). There is increased activity around the auroral bulge, 10 magnetometers with dense connectivity form a community centred at $\sim$ 22.5 MLT (blue) and a second community formed of 9 magnetometers spans from $\sim$ 21.5 MLT to $\sim$ 5.5 MLT (orange). Magnetometers that are not included in these communities are still included in the network (see figure 5.4d). At this time the modularity decreases in figure 5.2. At $t' = 15$ figure 5.3e all magnetometers that are within the network are within these two communities. There are still 3 magnetometers which are not contained within the communities but they are highly connected to the network (see figure 5.4e). Although the network forms two communities they are highly inter-connected, as is expected with low modularity. By $t' = 20$ (figure 5.3f) all available magnetometers are connected to the network and are now contained in one single community, suggesting the entire global system is highly connected (see figure 5.4f). This configuration is unchanged across the maps at $t' = 25$ and $t' = 30$ figure 5.3g-h, consistent with figure 5.2.

The overall dynamics of this event is then a coalescing of multiple, small communities into two, and then one single global community. We have analysed other events (see section 5.5 for a further four examples) and find that there is always a transition from initially many small, spatially separated communities at onset to one single global system at the peak of the expansion phase. However there is a great deal of variety in how these small communities coalesce depending on the substorm. For examples see figures 5.6-5.17 where all begin as several communities, sometimes overlapping in MLT, and tend towards a global system.

5.3.2 Community structure of multiple substorms

We can see how robust this transition, from unstructured to structured, is by comparing the time evolution of the modularity of the networks of the set of 41 substorms, which have been selected using the above criteria to be quiet before onset and to have good magnetometer spatial coverage of the nightside, each has at least 2 magnetometers in each three hour MLT sector. Figure 5.5 plots the normalized modularity, Q_N for these 41 events. We use normalized modularity, Q_N, as a parameter for community structure, where $Q_N \to 1$ when the network has multiple, separate sub-networks of magnetometers with many connections within but few between and $Q_N \to 0$ when the network is globally dense (either a single community or several large-scale communities with many

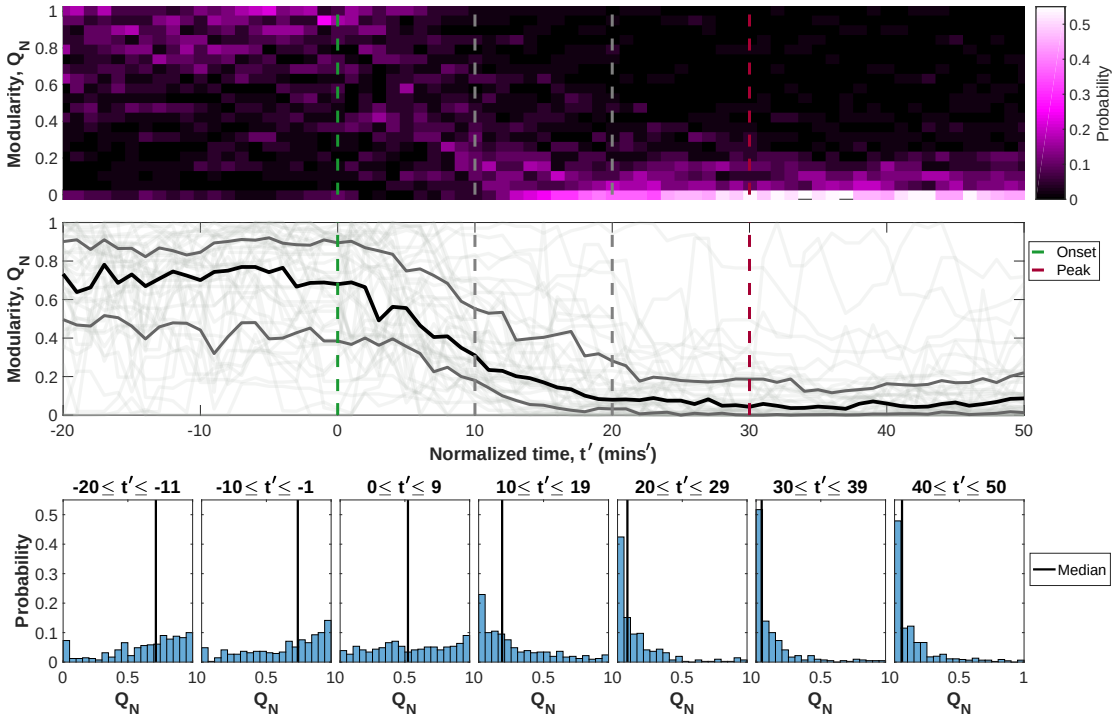

Figure 5.5: The normalized modularity, Q_N, of the set of 41 substorms that have 2 or more magnetometers in four even LT sectors of the nightside and are quiet before onset. The top two panels have normalized time as the abscissa. The top panel ordinate bins Q_N at each normalized time and the colour indicates the probability (count of substorms with Q_N/total number of substorms). The centre panel plots Q_N of all substorms as a function of normalized time with the median overplotted in black and the 25% and 75% quantiles in grey. The bottom panel plots the normalized histograms of Q_N of the events aggregated over 10 minute intervals as time progresses. The median is overplotted.

connections between them). There is a value of the modularity for each substorm at each minute in time and the top panel plots as a function of time the modularity probability (count of substorms with Q_N/total number of substorms), that is, the fraction of the substorms which have normalized modularity within each $Q_N = 0.05$ bin, indicated by colour. We have not included data where there was less than five connections in the network, i.e. we imposed a $m \geq 5$ criteria on the modularity data to avoid $Q = 0$ simply because there are no connections.

There is a clear pattern of high modularity before onset which drops to low modularity from $t' \sim 20$. This is also shown in the centre panel which plots an overlay of the modularity time series for each substorm. The median modularity and the 25th and 75th quantiles are overplotted on this panel. This again shows a clear transition from high to low modularity that takes place in the (normalized) time between onset and expansion peak. Finally, the bottom panel plots histograms of the probabilities of modularity from all 41 events aggregated across $10 - 11$ normalized minute time windows. There is a clear transition from a right shifted distribution, with median ~ 0.7 before onset, to an approximately uniform distribution from $0 \leq t' \leq 9$, to a highly peaked left shifted

distribution from $t' \sim 10$, with $\geq 50\%$ of substorms having $Q_N \leq 0.1$. By the time of maximal auroral bulge expansion less than a quarter of the substorms have modularity over 0.2. This suggests that whilst before the substorm, and up to 10 minutes following onset, there can be a variety of community structures which give a broad spread in modularity values, once the expansion phase is reached, almost all the events are at low modularity, consistent with a single spatially coherent community.

We find the same overall behaviour looking over a wider set of events (see figures 5.24-5.25), however more active conditions before the substorm make the pattern less clear. If we only consider substorms with large values of SML in the 127 minutes preceding onset, we can no longer isolate such a clear transition in the modularity distribution (see figure 5.26). Figures 5.19-5.23 show the same figure but using five alternative algorithms for community detection. The same pattern of decreasing modularity is repeated in all five algorithms and futher, is repeated across a range of cross-correlation thresholds (see figures 5.27-5.28). If the threshold is set too high ($< 1\%$ of magnetometers "connected"averaged over a month) there would be too few network connections to see any pattern (hence 5.27 contains more outliers). Figure 5.29 shows that the modularity, Q, derived from the substorm events has more structure (higher Q) than random networks and for a given number of connections, m, the data explores a broad range of modularity values, Q, and vice versa.

5.4 Summary I

In this section we summarise all material covered in sections 5.1-5.3.2. Further examples of community structure will be discussed in the following section. We have used well-established network science techniques to analyse data from > 100 ground-based magnetometers. We translated this data into a time-varying directed network, based on CCC of the vector magnetic field perturbations measured at each magnetometer pair [Dods et al., 2015, 2017] and performed community detection on the network. Communities are locally dense but globally sparse groups of connections in the network [Newman and Girvan, 2004], identifying emerging coherent patterns in the current system as the substorm evolves. We consistently find robust structural change from many small, uncorrelated groups of magnetometers before substorm onset, to one large spatially-extended correlated system during the expansion phase.

We have shown that the SCW consistently, over 40+ events, displays large-scale correlated behavior and this is inconsistent with the recent hypothesis that this current system consists of multiple distinct mesoscale wedgelets [Birn and Hesse, 2014; Birn et al., 2019; Liu et al., 2015, 2018; Malykhin et al., 2018; Palin et al., 2016; Panov et al., 2016; Merkin et al., 2019].

As our technique is based on correlated magnetic field perturbations from magnetometes we cannot resolve structures that are on spatial scales smaller than that of the inter-magnetometer spacing, nor on short temporal scales of less than 1 minute, thus we cannot resolve short-timescale, small-scale events such as individual bursty bulk flows or wedgelets. However, unless the wedgelets are spatially and temporally correlated with each other, we would not expect to see the spatially coherent signature of global cross-correlation seen across the auroral bulge. We cannot rule out the scenario in which two or more spatio-temporally correlated currents would appear as one community because they would not be able to be resolved by ground-based magnetometers. Nonetheless, the spatially-extended communities we observe cannot be obtained by having many, small, spatially localised wedgelets which are each internally correlated, but not cross-correlated with each other. All 40+ substorms analysed here ultimately form a single large scale current wedge structure. The structural shift from multiple to a single global current system occurs approximately 10 minutes after onset. It excludes models in which the current system is solely comprised of individual wedgelets, a fully correlated large-scale SCW is required.

The structural transition from multiple to a single global current system is not instantaneous, it occurs approximately over the 10 minutes after onset. We have found examples where this transition is a direct coalescence of multiple small communities into a single global community. We also found examples where multiple small communities first coalesce into two or more large communities which then merge into the single global community at expansion peak. This emphasises the need for an understanding of the dynamical evolution of substorms which may resolve the controversy surrounding models for the substorm current system. If we were to only look at a snapshot of the nightside's magnetic activity within the first ten minutes of a substorm we may indeed see multiple structures but as the substorm evolves we clearly see one underlying spatially extended current system.

This work introduces a new parameter for the spatio-temporal pattern captured by the full set of ground-based magnetometers- the network modularity. From $\sim$ 20 minutes into the substorm expansion phase over 75% of the substorms analysed here have extremely low normalized modularity (< 0.2), indicating a highly correlated, large-scale global system. The modularity provides a quantitative spatio-temporal response benchmark for MHD simulations and SCW models. For example, Birn and Hesse [2014], which concluded that flow channels distorted the magnetic field and drove a current system akin to a small SCW (wedgelet), but their accumulation near the Earth led to a system dominated by a single, large-scale SCW would be an interesting simulation to compare to our analysis. Would their simulated large-scale SCW be correlated on spatial scales bigger than that of the wedgelets? If the wedgelet accumulation was simply side-by-side uncorrelated, spatially distinct wedgelets we would expect our analysis to

identify small, distinct communities with a high modularity. If the wedgelet accumulation involved reconnection of magnetic field and the diversion of the current system across multiple wedgelets our analysis would identify a large-scale, correlated system with low modularity.

5.5 Community structure of additional substorms

Here we show a further four substorm examples, as well as the epoch analyse of the modularity for different sets of events and under different network or community conditions. We find that there is always a transition from initially many small, spatially separated communities at onset to one single global system at the peak of the expansion phase. However there is a great deal of variety in how these small communities coalesce depending on the substorm.

5.5.1 Individual Events

Figures 5.6-5.17 show examples of the community structure for a further four substorms. We present the sequence of plot as shown in figures 5.2-5.4, maintaining the same format and using the same network detection method, namely edge betweenness [Newman and Girvan, 2004]. The four additional events are a substorm on the 07/01/1997 (figures 5.6-5.8), 06/09/1997 (figures 5.9-5.11), 06/01/1998 (figures 5.12-5.14) and 20/01/1998 (figures 5.15-5.17). All four of the substorms have good coverage during the event with ≥ 5 magnetometers in each 3 hours MLT for the duration of the substorm.

07/01/1997 event

Figures 5.6-5.8 show an example of the community structure of the network during the substorm on the 07/01/1997. The figures are in the same format as figures 5.2-5.4. In figure 5.6, we see many individual communities from $-20 \leq t' \leq 10$ (panels 1-2). There is a larger community in the dusk sector, at $\hat{\theta} \sim 20$ hours MLT, during this time period which contains ~ 7 magnetometers. Modularity (panel 5) fluctuates from $-20 \leq t' \leq 10$ but exceeds the modularity expected from a random network (surrogate). At $t' \sim 10$ there is an increase in the number of connections, $\alpha(t')$, as well as the auroral electrojet indicies, -SML, (panel 6), which correspond with a sharp decrease in modularity (panel 5) and the merging of communities (panels 1-4). Initially the merging is near and westward of the onset location (blue/purple coloured communities in panel 3 between $18 \leq \hat{\theta} \leq 1$) but include stations that were located further east by $t' \sim 15$. By $t' \sim 15$, modularity has dropped to $Q \sim 0$ as the one community dominates, suggesting the

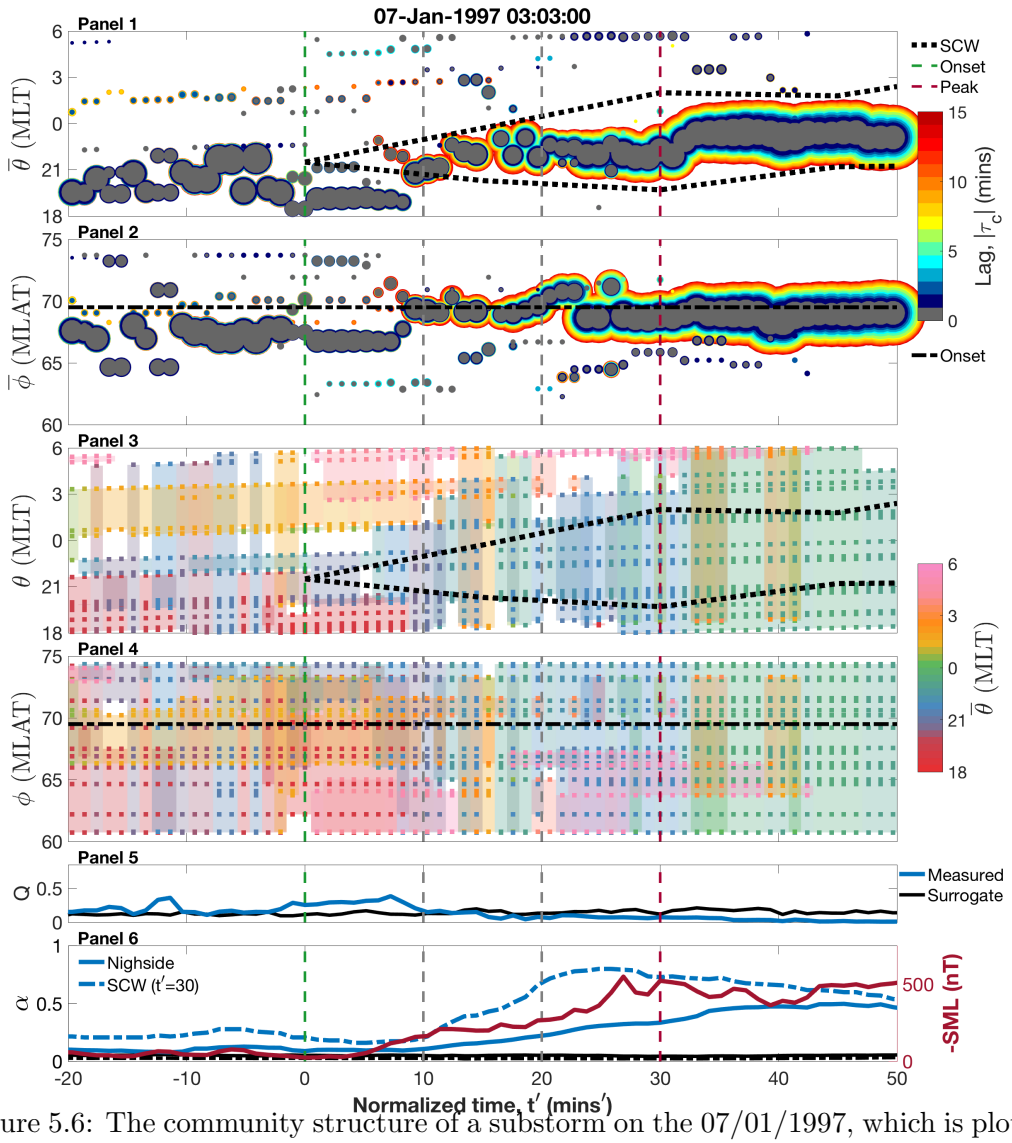

Figure 5.6: The community structure of a substorm on the 07/01/1997, which is plotted in the same format as in figure 5.2 (which shows a different event, 16/03/1997). Following onset, figures 1 and 2 show that there is a dominant community but additionally several smaller communities which persist during the substorm. The modularity plotted in panel 5 drops from ~ 0.5 before the event to almost zero by $t' \sim 20$. This plot used the edge betweenness community detection algorithm.

entire nightside is highly correlated by $t' \sim 15$. In the dawn sector, at $\hat{\theta} \sim 6$, a small community of $2-3$ magnetometers is maintained during the substorm. Looking at the snap shots of the network in figures 5.7-5.8, the communities detected change spatially throughout the substorm but the correlation appears to be centred around the onset location (5.8e-h) with fewer connections to the magnetometers closest to dawn ($\hat{\theta} \sim 6$). This may suggest that the SCW has extended further west than east.

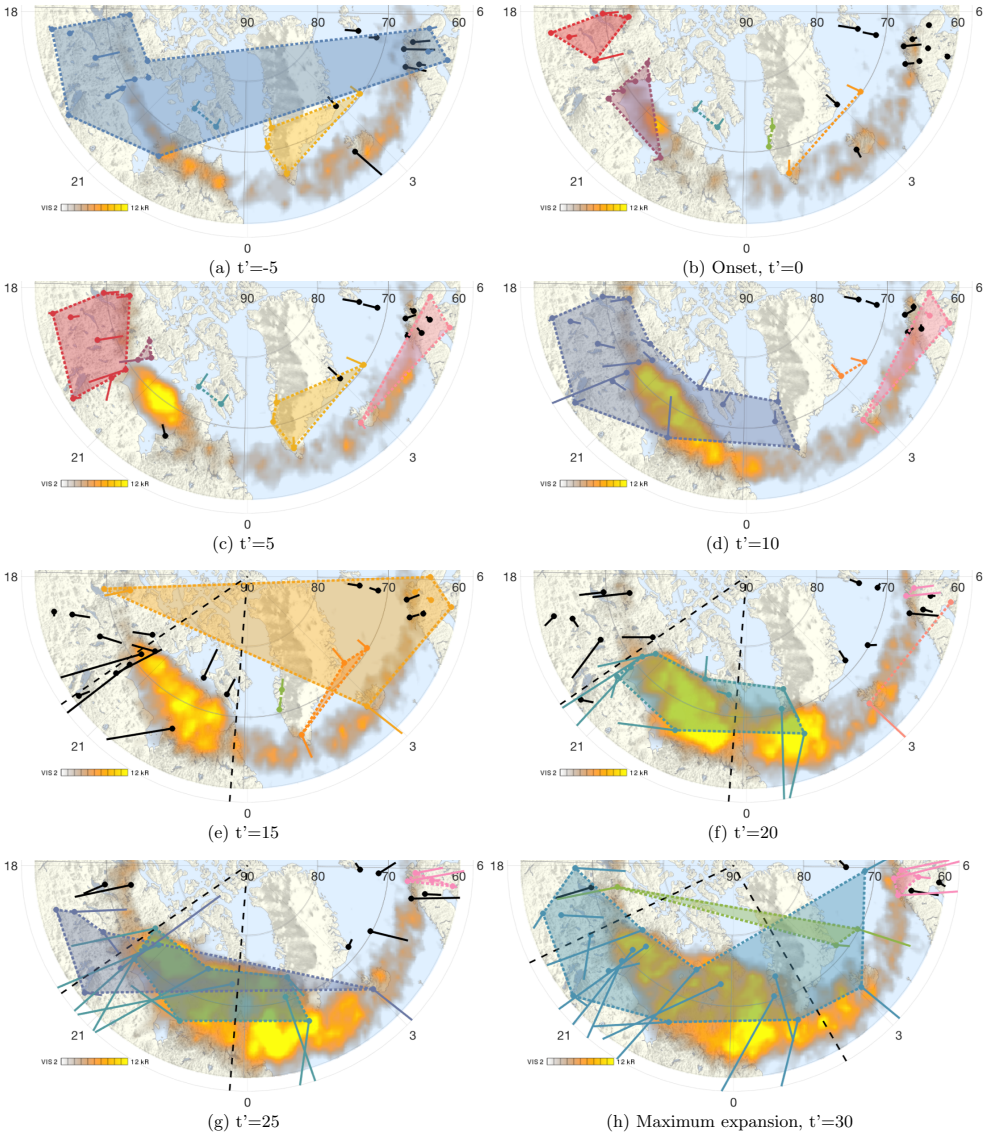

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.7: The community structure of a substorm on the 07/01/1997, which is plotted in the same format as in figure 5.3 (which shows a different event, 16/03/1997). Several smaller communities are maintained for the duration of the substorm but the network tends towards one main community. This plot used the edge betweenness community detection algorithm.

06/03/1997 event

Figures 5.9-5.11 show an example of the community structure of the network during the substorm on the 06/03/1997. The figures are in the same format as figures 5.2-5.4. In figure 5.9, we see several communities before onset, each of which contains only $2-5$ magnetometers. Most are localised with spatial extent of <2 hours local time. At $t' \sim 5$ we begin to see communities merging to contain the magnetometers within the auroral bulge boundaries. Although there are still several communities in the dawn side during the substorm, the modularity drops to $Q \sim 0$ from $t' \sim 10$ implying that the communities are highly connected. If we refer to figure 5.11 we can see that there

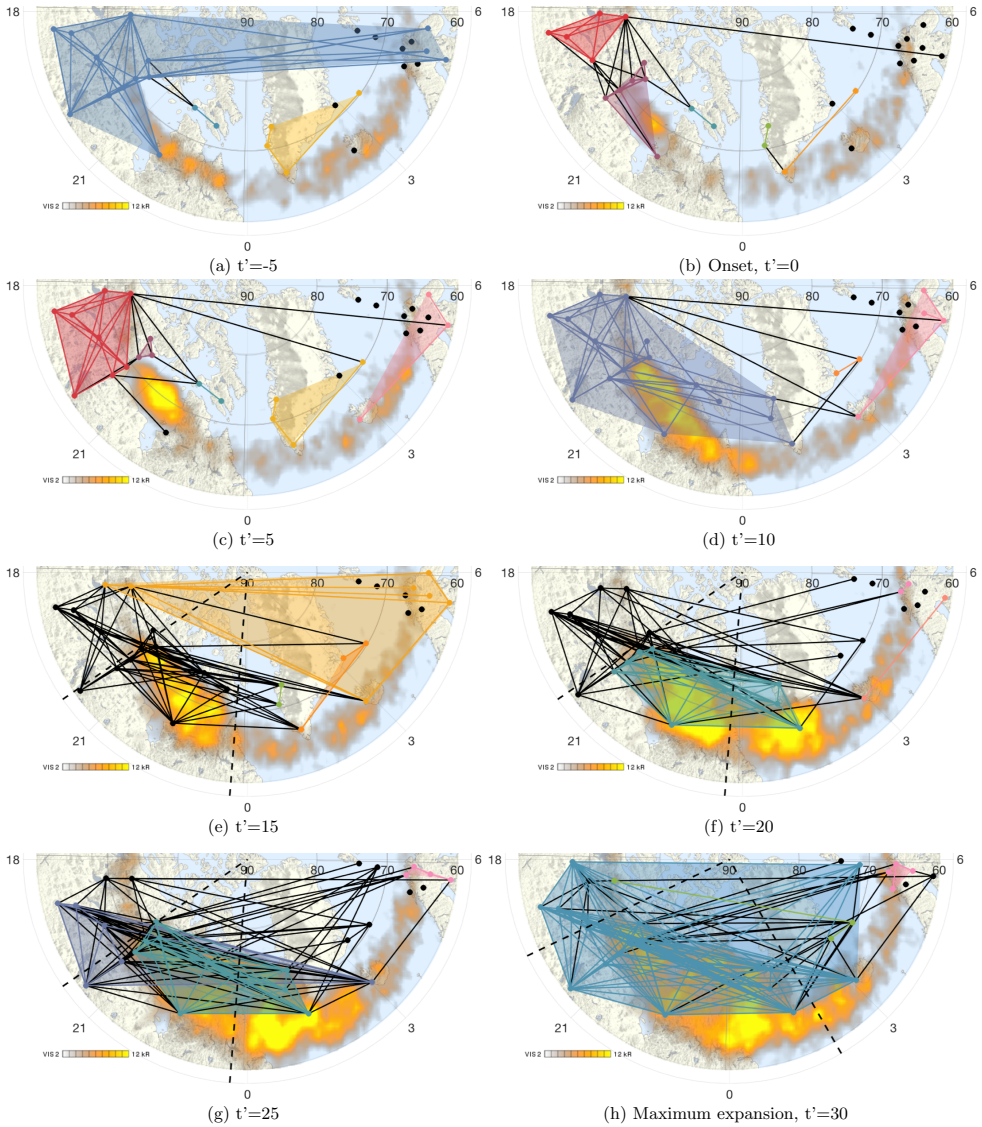

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.8: The community structure of a substorm on the 07/01/1997, which is plotted in the same format as in figure 5.4 (which shows a different event, 16/03/1997). The smaller communities transition into a dominant community, with several smaller communities persisting during but they are connected via many inter-community edges. This plot used the edge betweenness community detection algorithm.

are many connections between the magnetometers in the dusk and dawn, between the different communities. The two magnetometers located in Iceland and West Greenland only ever have one connection with the main network; they do not have much correlation with the rest of the nightside.

06/01/1998 event

Figures 5.12-5.14 show an example of the community structure of the network during the substorm on the 06/01/1998. The figures are in the same format as figures 5.2-5.4.

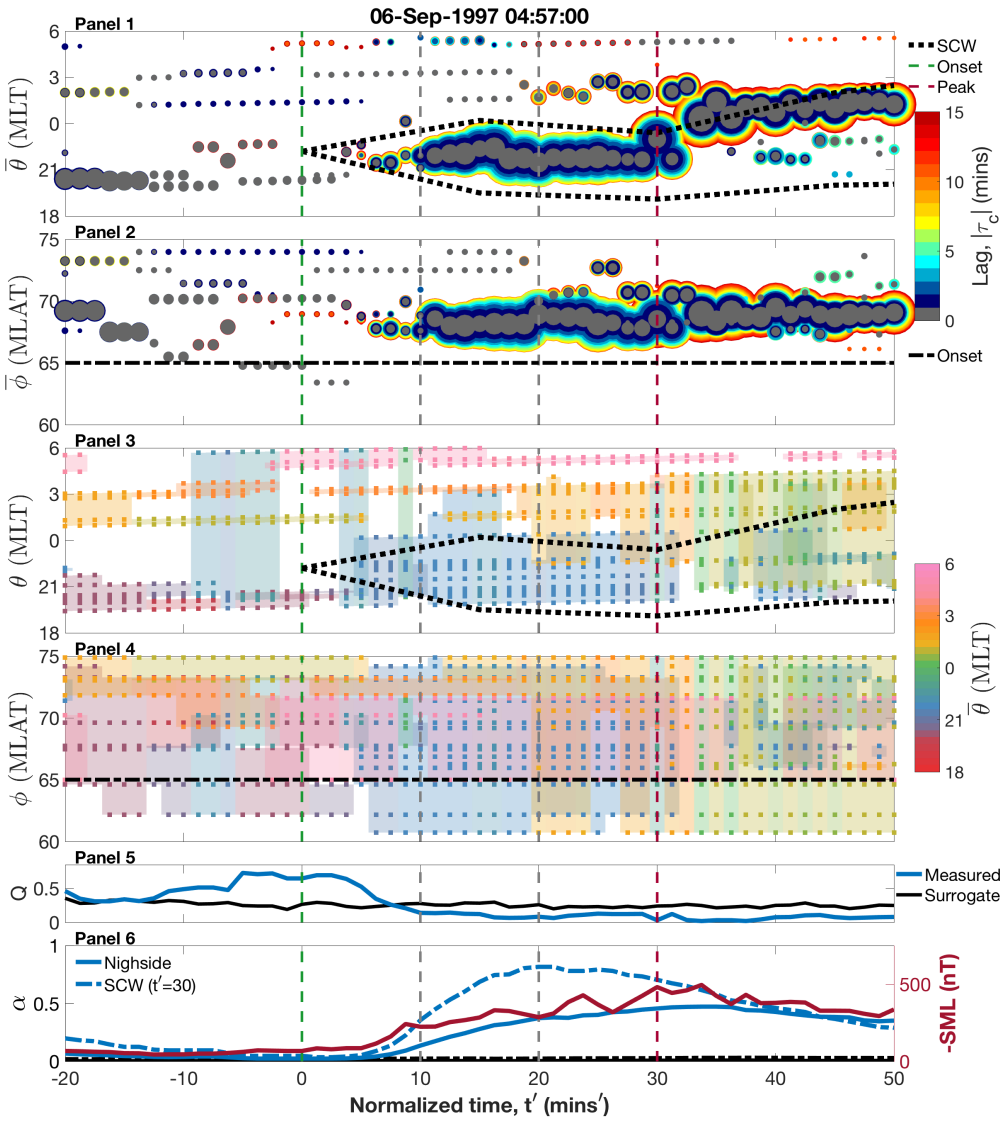

Figure 5.9: The community structure of a substorm on the 06/09/1997, which is plotted in the same format as in figure 5.2 (which shows a different event, 16/03/1997). Following onset, figures 1 and 2 show that there is a dominant community but, additionally, several smaller communities which persist during the substorm. The modularity plotted in panel 5 drops from ~ 0.7 before the event to almost zero by $t' \sim 10$. This plot used the edge betweenness community detection algorithm.

In figure 5.12, we see several communities before onset each containing a small number of magnetometers (< 7). The communities often have a large spatial extent in terms of local time (panel 3) but there are few connections between they given the modularity is $Q \sim 0.5$. After onset the modularity begins to drop as the communities merge into one main structure. Several magnetometers near dawn maintain a separate community during the substorm. From figure 5.14 we can see that the magnetometers near dawn (pink) become connected to the network later in the substorm ($t' \sim 20$) and the three on Svalbard have few connections with the rest of the network, even at the maximum extent of the substorm.

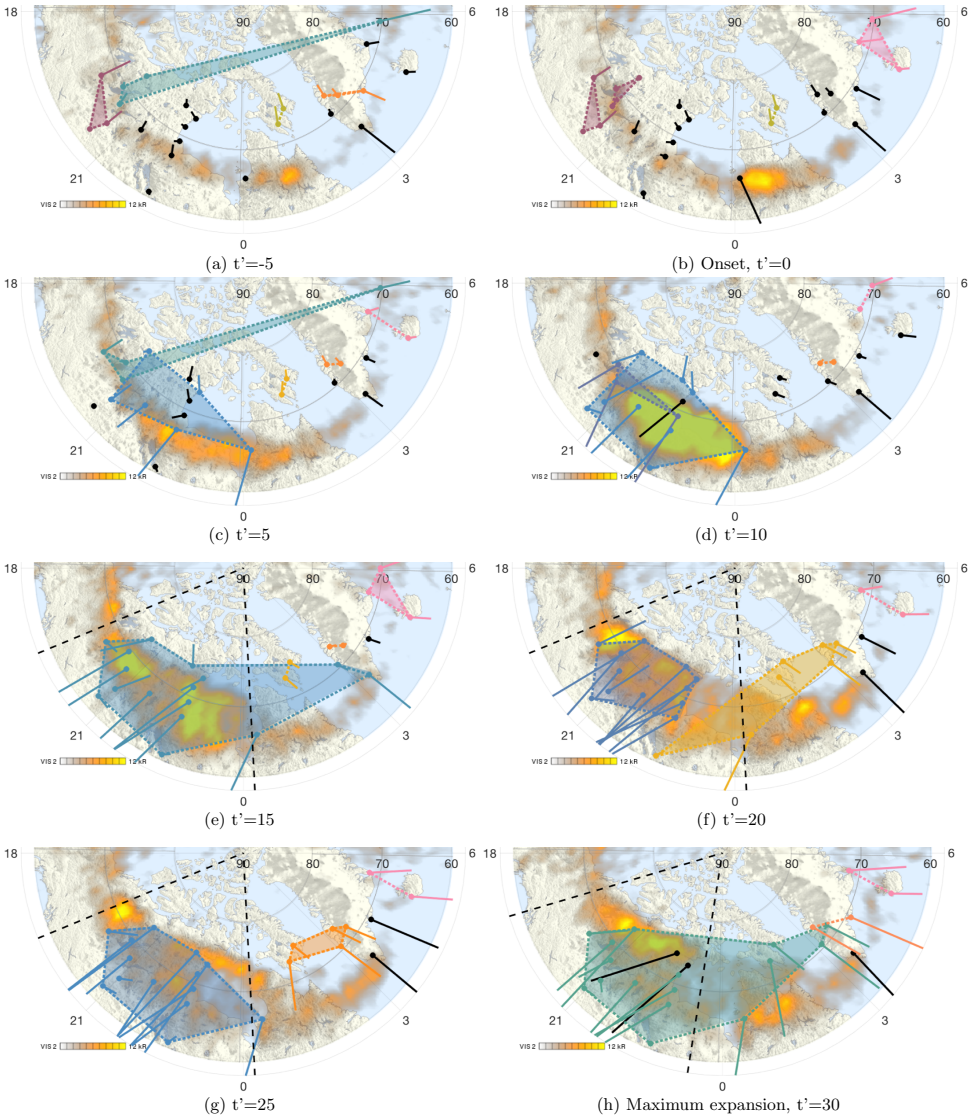

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.10: The community structure of a substorm on the 06/09/1997, which is plotted in the same format as in figure 5.3 (which shows a different event, 16/03/1997). Several smaller communities persist during the substorm but the network tends towards one main community. This plot used the edge betweenness community detection algorithm.

20/01/1998 event

Figures 5.15-5.17 show an example of the community structure of the network during the substorm on the 20/01/1998. The figures are in the same format as figures 5.2-5.4. In figure 5.15, we again see a pattern of small communities before onset which tend towards one main group from $t' \sim 10$. The modularity begins to drop slightly before onset quantifying that the communities are becoming more inter-connected. Again we see a small $2-3$ station community, at dawn, is maintained during the expansion phase but the community centred within the expansion wedge contains the majority of nightside magnetometers. From $t' \sim 15$ modularity is $Q \sim 0$ meaning the entire nightside is highly correlated. From figure 5.17 we see that even the dawn side community,

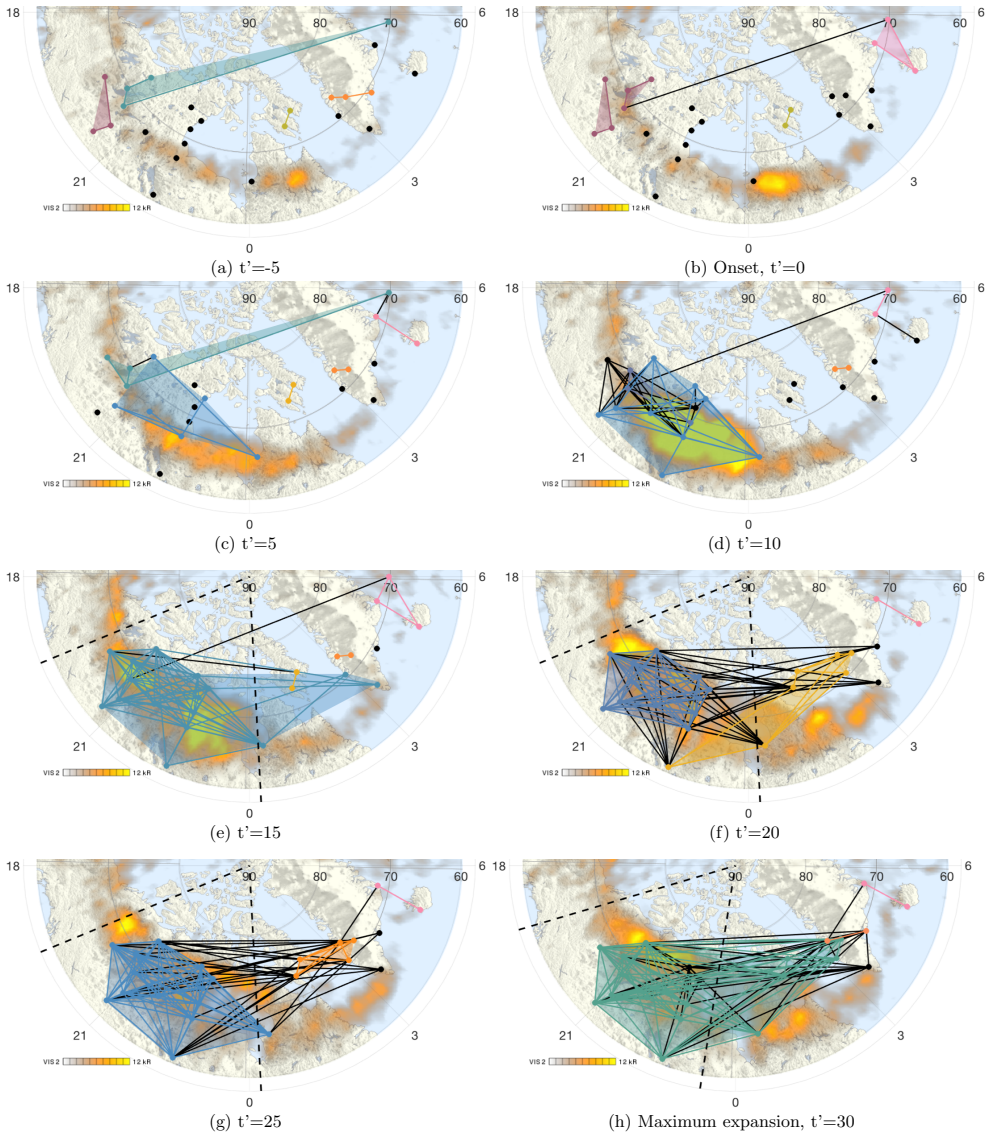

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.11: The community structure of a substorm on the 06/09/1997, which is plotted in the same format as in figure 5.4 (which shows a different event, 16/03/1997). The smaller communities transition into a dominant community, with several smaller communities persisting during the substorm but they are connected via many inter-community edges. This plot used the edge betweenness community detection algorithm.

that is maintained during the auroral expansion, is highly connected to the rest of the magnetometers.

Summary II

In this section we summarise the results from comparing multiple isolated substorms. Figures 5.6-5.17 show that the behaviour displayed in the substorm on the 16/03/1997 is representative of an isolated substorm. The five examples of substorms in this chapter were chosen because they had the best spatial spread of magnetometer during the events.

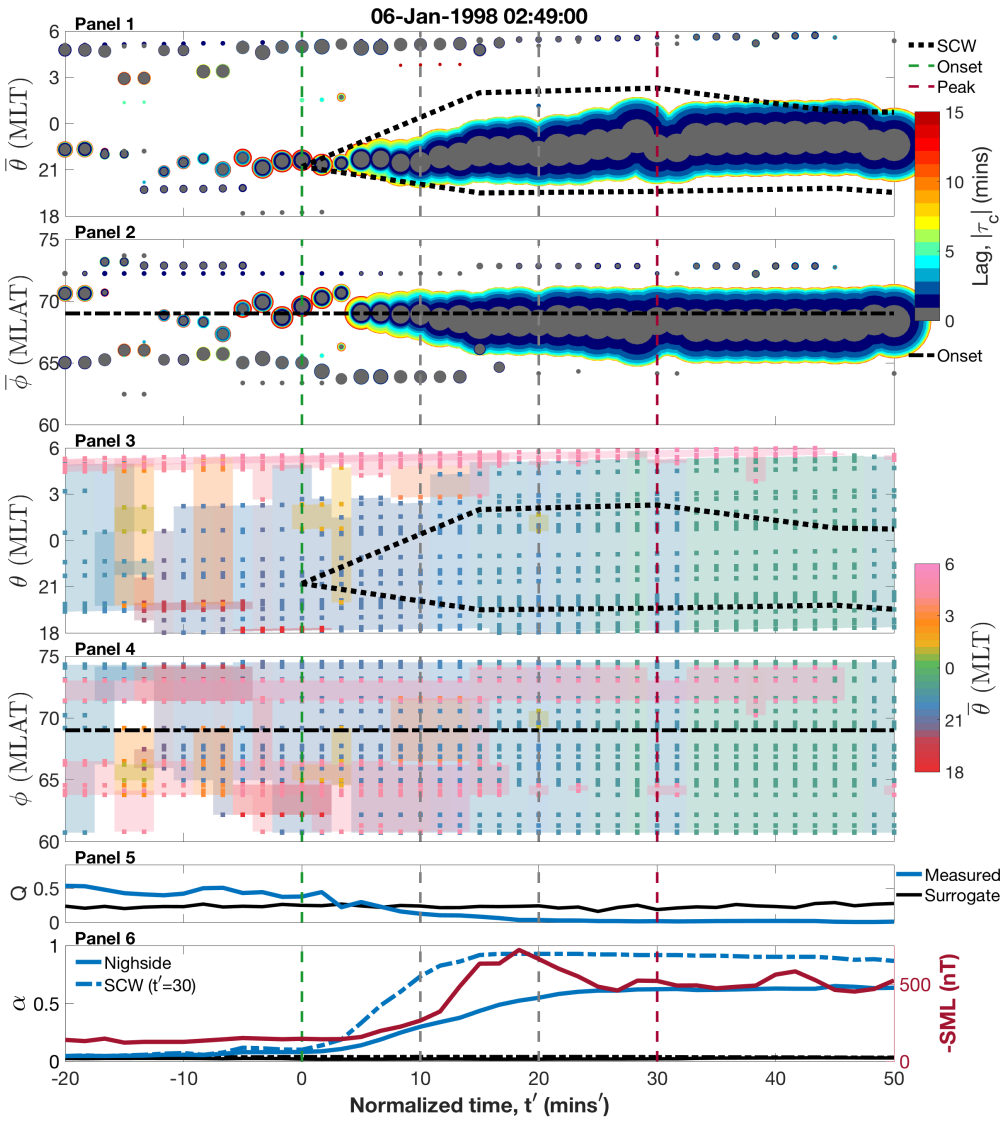

Figure 5.12: The community structure of a substorm on the 06/01/1998, which is plotted in the same format as in figure 5.2 (which shows a different event, 16/03/1997). Following onset, figures 1 and 2 show that there is a dominant community but also a single smaller community (at $\sim$ 6 MLT) which persist during the substorm. The modularity plotted in panel 5 drops from $\sim$ 0.5 before the event to almost zero by $t' \sim 20$. This plot used the edge betweenness community detection algorithm.

They are the five substorms from the list of events which were quiet before onset and had the highest number of magnetometers per three hours of MLT (i.e. $18-21$, $21-24$ etc). Although it appears the magnetometers in the dawn sector, towards 6 MLT, are the least likely to be correlated with the rest of the nightside, it is clear that each of the substorms forms a global community structure.

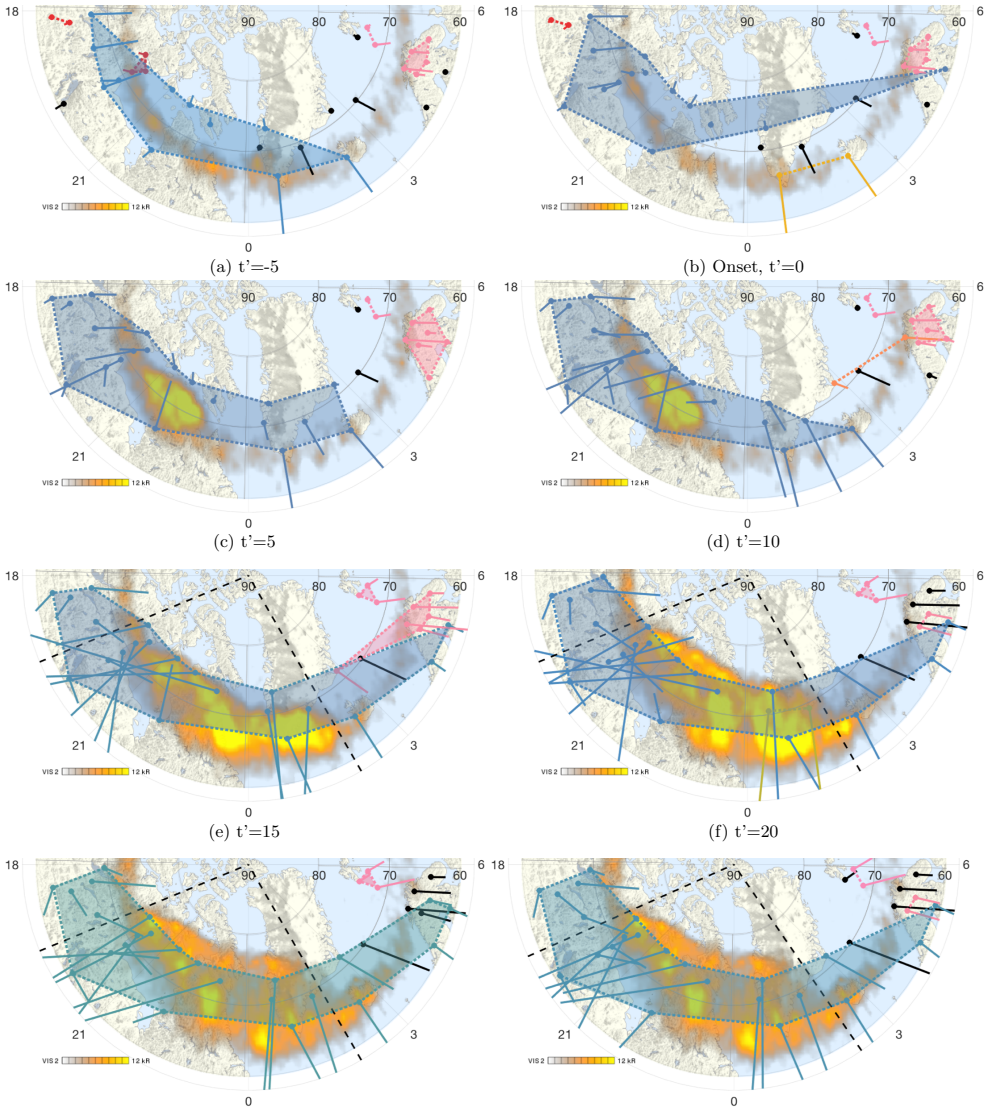

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.13: The community structure of a substorm on the 06/01/1998, which is plotted in the same format as in figure 5.3 (which shows a different event, 16/03/1997). A few smaller communities are maintained during expansion (at $\sim$ 6 MLT) but the network tends towards one main community. This plot used the edge betweenness community detection algorithm.

5.5.2 Modularity of multiple isolated substorms

Unnormalized Modularity

Figure 5.18 shows the unnormalized modularity, Q, (in figure 5.5 the normalized modularity, Q_N, is overlayed for multiple substorms) using the same detection algorithm, threshold and set of substorms as were analysed within section 5.3. The maximum modularity reached is $\sim$ 0.8 and before onset the distribution spread is much greater than that in normalized version (figure 5.5). Ten normalized minutes into the expansion phase there is the same leftwards shift in the distribution. No substorms have modular-

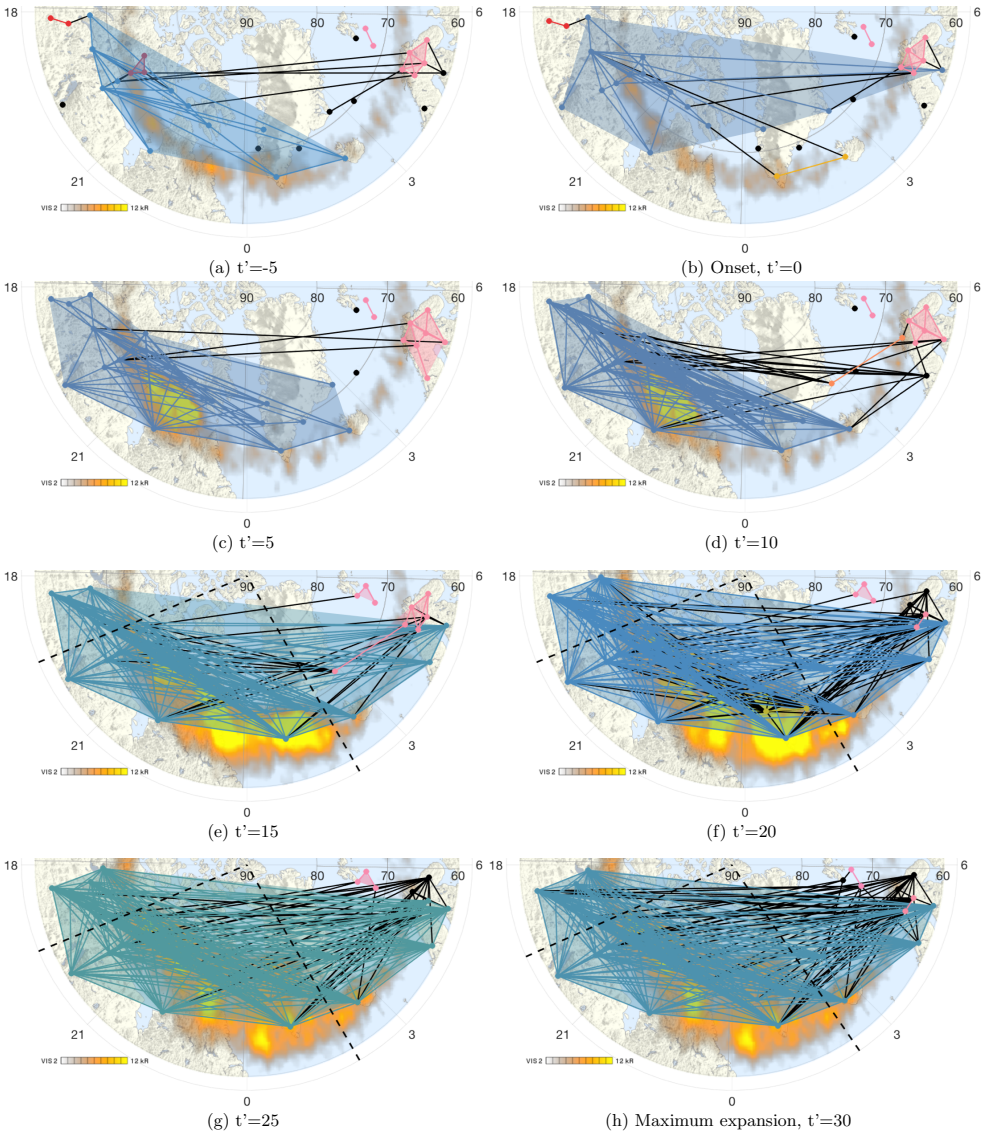

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.14: The community structure of a substorm on the 06/01/1998, which is plotted in the same format as in figure 5.4 (which shows a different event, 16/03/1997). The main community expands eastward to form a large global community, with several smaller communities around dawn but between the communities there are many connections. This plot used the edge betweenness community detection algorithm.

ity of $Q > 0.6$ and more than 75% have modularity under $Q \sim 0.4$. By $t' \sim 20$ the mean modularity is < 0.1. The unnormalized modularity decreases in modularity throughout the substorm expansion, implying increasing network structure and a shift from local to global correlation during the substorm expansion.

Changing the community detection methods

Figures 5.19-5.23 plot the normalized modularity, Q_N, obtained using the optimal Brandes et al. [2007], walk trap Pons and Latapy [2005], information mapping Rosvall and

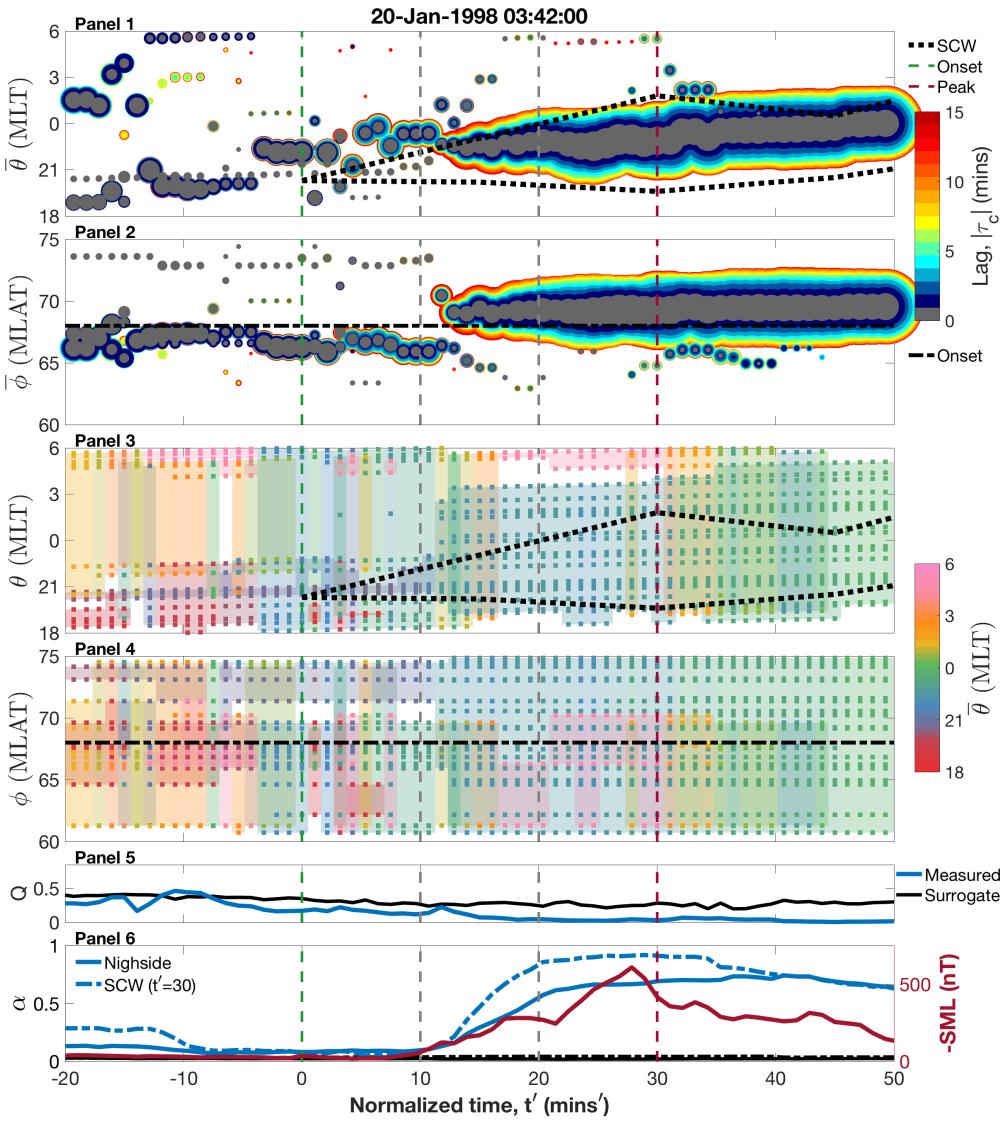

Figure 5.15: The community structure of a substorm on the 20/01/1998, which is plotted in the same format as in figure 5.2 (which shows a different event, 16/03/1997). Following onset, figures 1 and 2 show that there is a dominant community but also a single smaller community which persists during expansion. The modularity plotted in panel 5 drops from ~ 0.5 before the event to almost zero by $t' \sim 20$. This plot used the edge betweenness community detection algorithm.

Bergstrom [2007], leading eigenvector Pons and Latapy [2005] and label propagation Raghavan et al. [2007] methods for community detection, respectively, for the networks calculated from the set of events used in section 5.3.

All five methods show the same pattern of high normalized modularity before onset, with a general decrease between $0 \leq t' \leq 20$ and a low modularity for the remainder of the substorm. The optimal, walk trap and leading eigenvector vector methods (figure 5.19, 5.20 and 5.22) give similar results with a change in the mean of the normalized modularity, from $Q_N \sim 0.8$, before onset to $Q_N \sim 0.2$ at peak expansion. The distributions of Q_N for both the right shifted ($-20 \leq t' \leq 0$) and left shifted ($20 \leq t' \leq 50$)

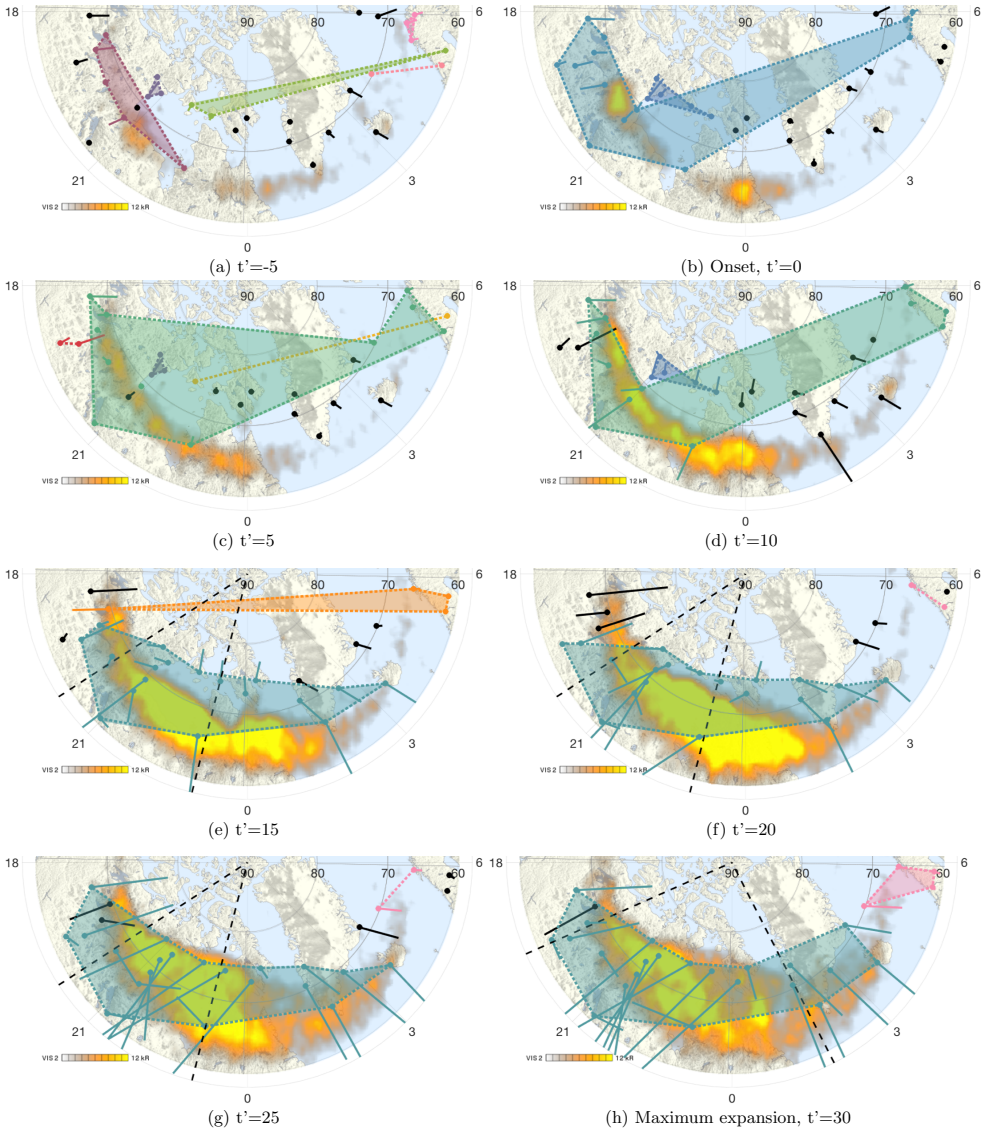

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.16: The community structure of a substorm on the 20/01/1998, which is plotted in the same format as in figure 5.3 (which shows a different event, 16/03/1997). The smaller communities transition into a dominant community but there remains a single smaller community (at ~ 6 MLT) which persists during. This plot used the edge betweenness community detection algorithm.

extremes are broad and never exceeds a probability of > 0.2 for any bin. The label propagating method (figure 5.23) tends towards a lower mean normalized modularity throughout with a shift from $Q_N \sim 0.6$ before onset to $Q_N \sim 0$ from half way through the expansion phase ($t' \sim 15$). The distribution of modularity is wider before onset with $\sim 15\%$ of the substorms having $Q_N \leq 0.1$ before onset. However, by the time of peak expansion $> 85\%$ of substorms sampled have $Q_N < 0.05$. The information mapping method (figure 5.21) likewise has a wider range of values before onset but $> 75\%$ of the substorms have $Q_N > 0.4$. After onset there is a clear decrease in modularity with $\sim 75\%$ of substorms having modularity $Q_N < 0.2$ by peak ($t' = 30$).

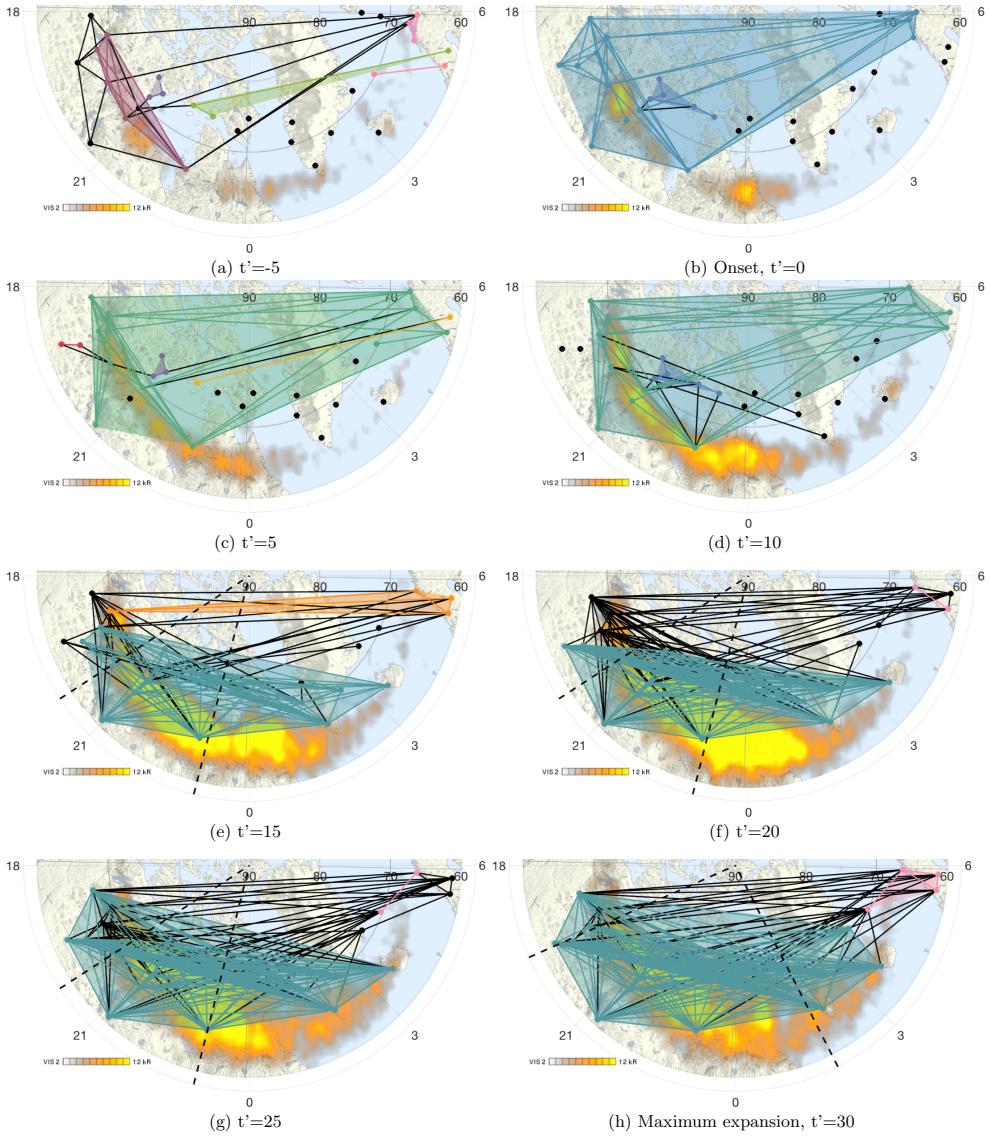

(a) t'=-5 (b) Onset, t'=0

(c) t'=5 (d) t'=10

(e) t'=15 (f) t'=20

(g) t'=25 (h) Maximum expansion, t'=30

Figure 5.17: The community structure of a substorm on the 20/01/1998, which is plotted in the same format as in figure 5.4 (which shows a different event, 16/03/1997). The smaller communities transition into a dominant community, with a single smaller community ($\sim$ 6 MLT) persisting during the substorm. Between the communities there are many connections. This plot used the edge betweenness community detection algorithm.

Despite the slight differences in community structure calculated using each of the six methods (including the Edge Betweenness algorithm used in figures 5.2-5.4), the observed trend from high to low modularity, during the substorm expansion, is independent of the detection algorithm.

Changing the set of substorms

Figures 5.24-5.26 show the normalized modularity, Q_N, of three different sets of substorms. The three figures were plotted in the same format as figure 5.5, using the edge

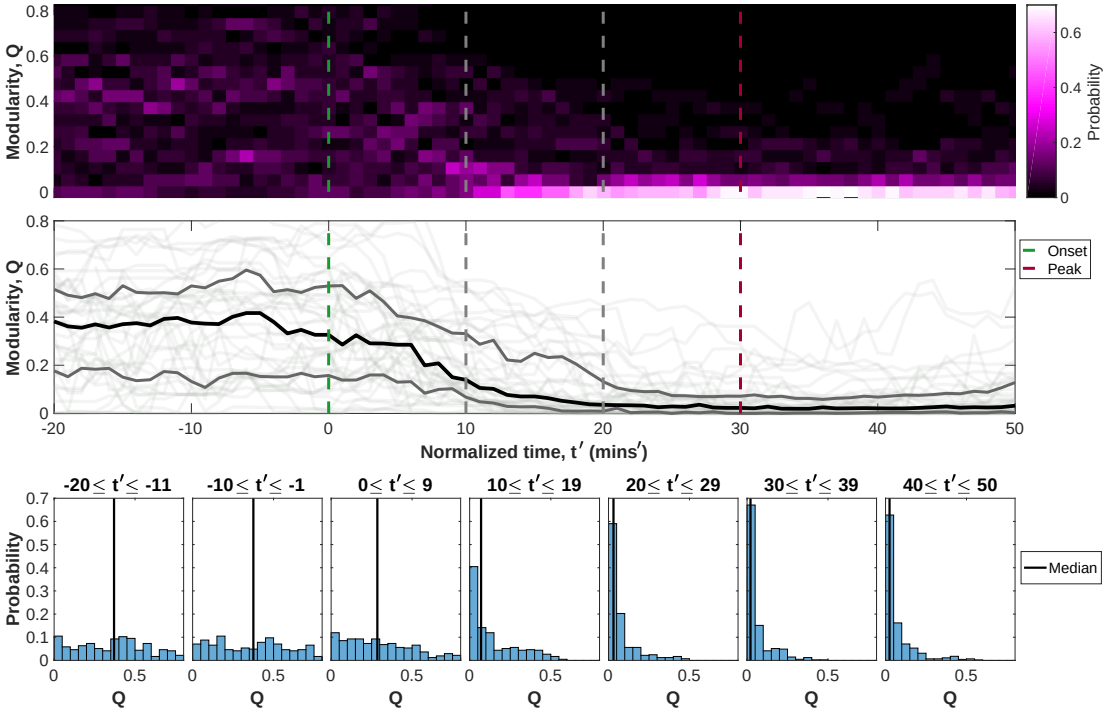

Figure 5.18: The unnormalized modularity, Q, in the same format as figure 5.5 (which showed the normalized modularity). This plot used the edge betweenness community detection algorithm to calculated the modularity for the 41 isolated substorms. All the substorms have a different starting value of modularity but all tend to near zero during substorm expansion.

betweenness detection method. Figure 5.24 contains 75 isolated substorms from the list of 116 substorms, between the years 1997 – 2001 [Gjerloev et al., 2007; Gjerloev and Hoffman, 2014]. The 75 substorms have two or more magnetometers in each three hour wedge of the nightside but there are no additional requirements beyond those specified in [Gjerloev et al., 2007; Gjerloev and Hoffman, 2014]. Therefore, some of the substorms are not as quiet before onset as the 41 used in figures 5.2-5.23. Before onset $\sim 12\%$ of the substorms have normalized modularity $Q_N < 0.05$ compared to $\sim 8\%$ when only substorm which were quiet from 127 minutes before onset were considered. Before onset the mean modularity is $Q_N \sim 0.6$ compared to $Q_N \sim 0.7$ in figure 5.5. By peak 75% of the substorms have $Q_N < 0.4$ in figure 5.24 compared to $Q_N < 0.2$ in figure 5.5. Additionally, the distribution of values in normalized modularity is much wider during the substorm. These differences imply that increased magnetic activity in the 127 minute before onset interferes with the pattern of increasing community structure during the substorm. One explanation could be previous substorms or pseudo breakups have caused high correlation between magnetic field perturbations. However, with this larger sample size we still see the pattern of an increasingly correlated global system during the substorms.

Figure 5.25 uses the set of 11 substorm events which have slightly increased magnetic perturbations before onset, specifically where the SML in the window before onset is

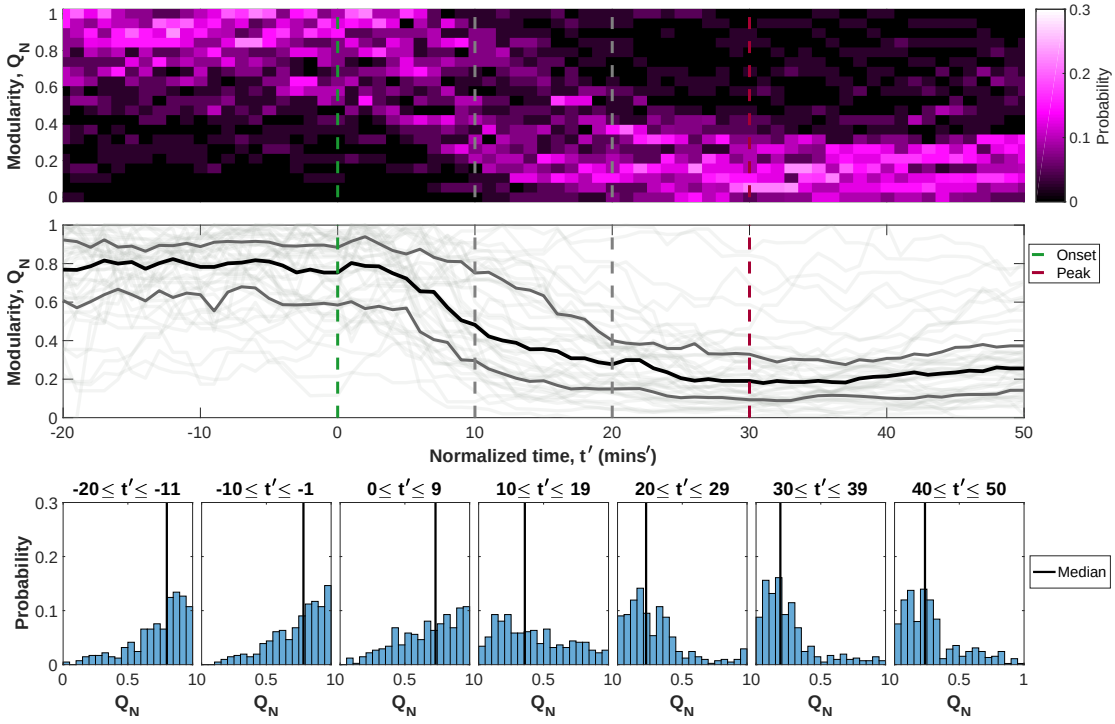

Figure 5.19: The normalized modularity, Q_N, calculated using the optimal community detection algorithm, in the same format as figure 5.5 (calculated using the edge betweenness algorithm). This plot was calculated for the 41 isolated substorms. There is a clear transition from high to low modularity during the expansion phase.

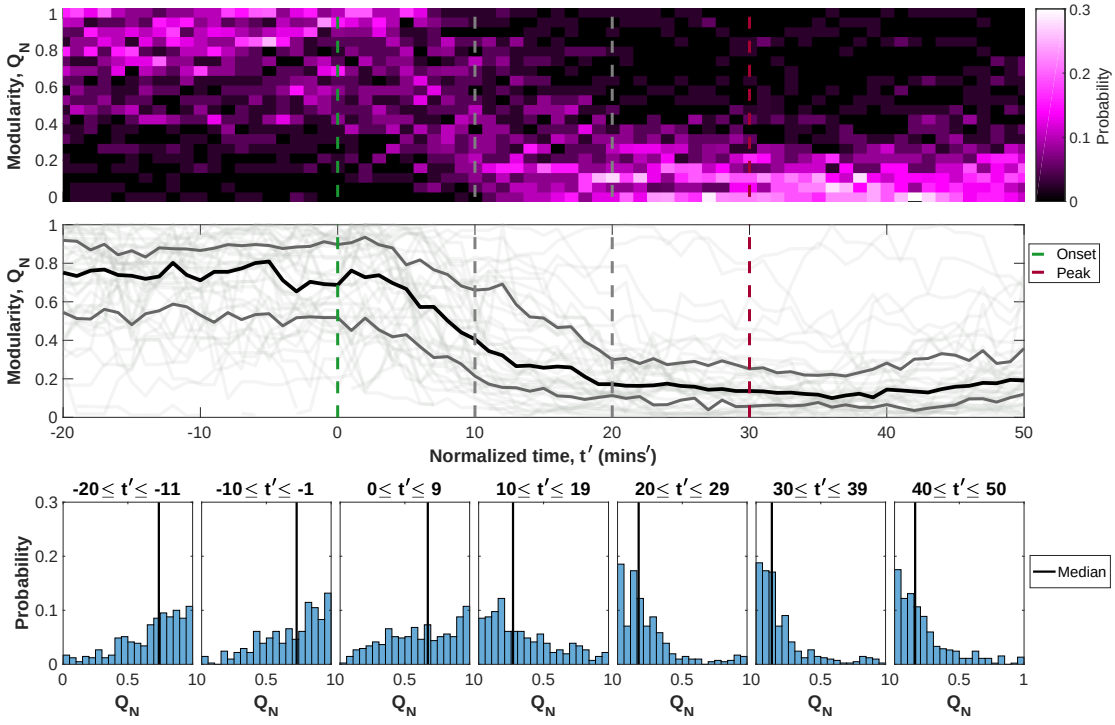

Figure 5.20: The normalized modularity, Q_N, calculated using the walk trap community detection algorithm, in the same format as figure 5.5 (calculated using the edge betweenness algorithm). This plot was calculated for the 41 isolated substorms. There is a clear transition from high to low modularity during the expansion phase.

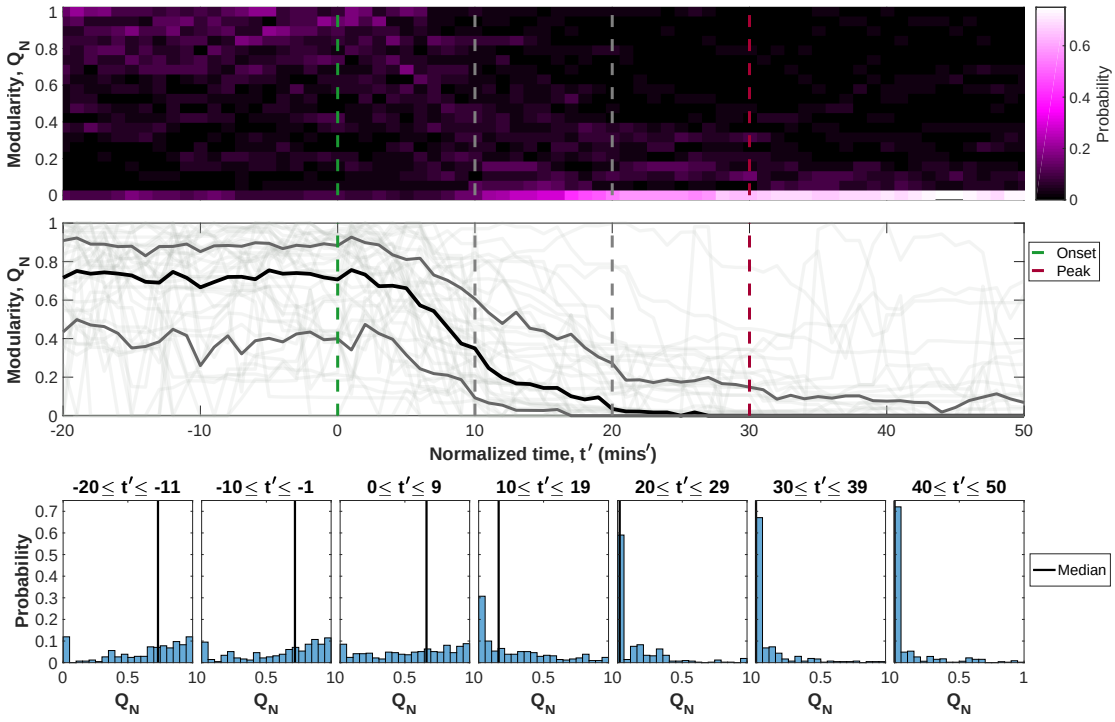

Figure 5.21: The normalized modularity, Q_N, calculated using the information mapping community detection algorithm, in the same format as figure 5.5 (calculated using the edge betweenness algorithm). This plot was calculated for the 41 isolated substorms. There is a clear transition from high to low modularity during the expansion phase.

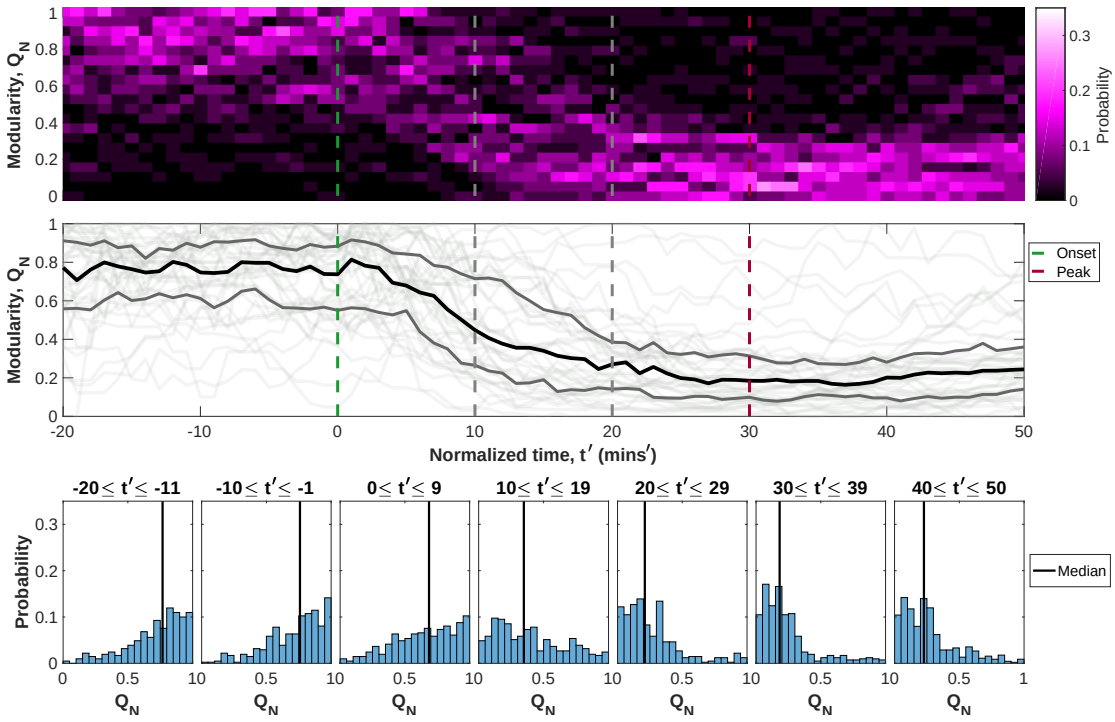

Figure 5.22: The normalized modularity, Q_N, calculated using the leading eigenvector community detection algorithm, in the same format as figure 5.5 (calculated using the edge betweenness algorithm). This plot was calculated for the 41 isolated substorms. There is a clear transition from high to low modularity during the expansion phase.

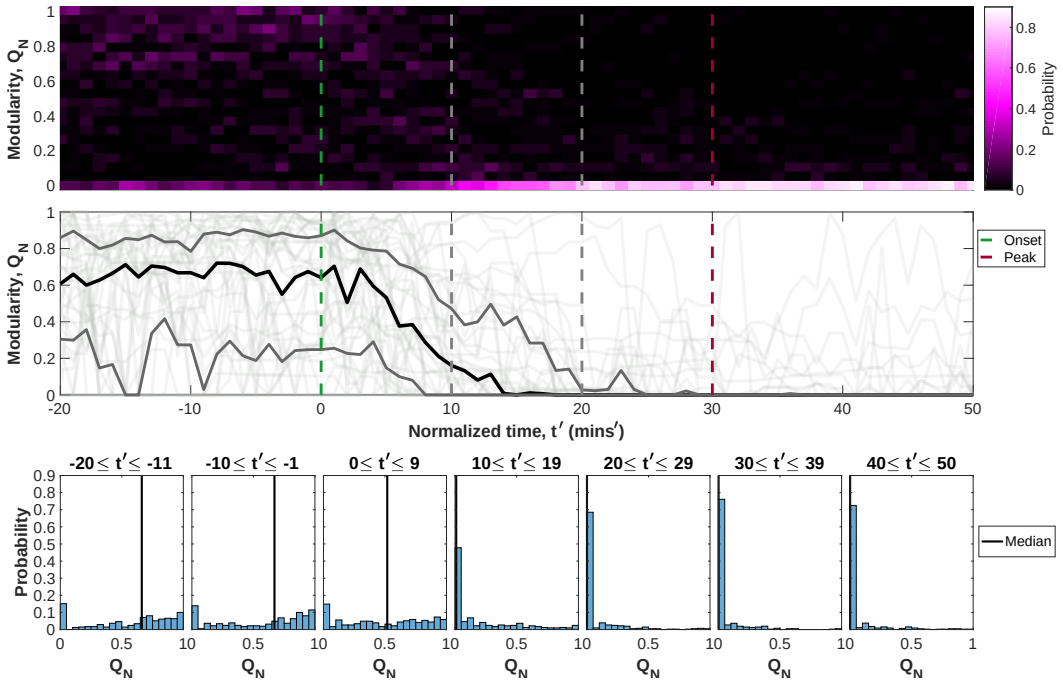

Figure 5.23: The normalized modularity, Q_N, calculated using the label propagating community detection algorithm, in the same format as figure 5.5 (calculated using the edge betweenness algorithm). This plot was calculated for the 41 isolated substorms. There is a clear transition from high to low modularity during the expansion phase.

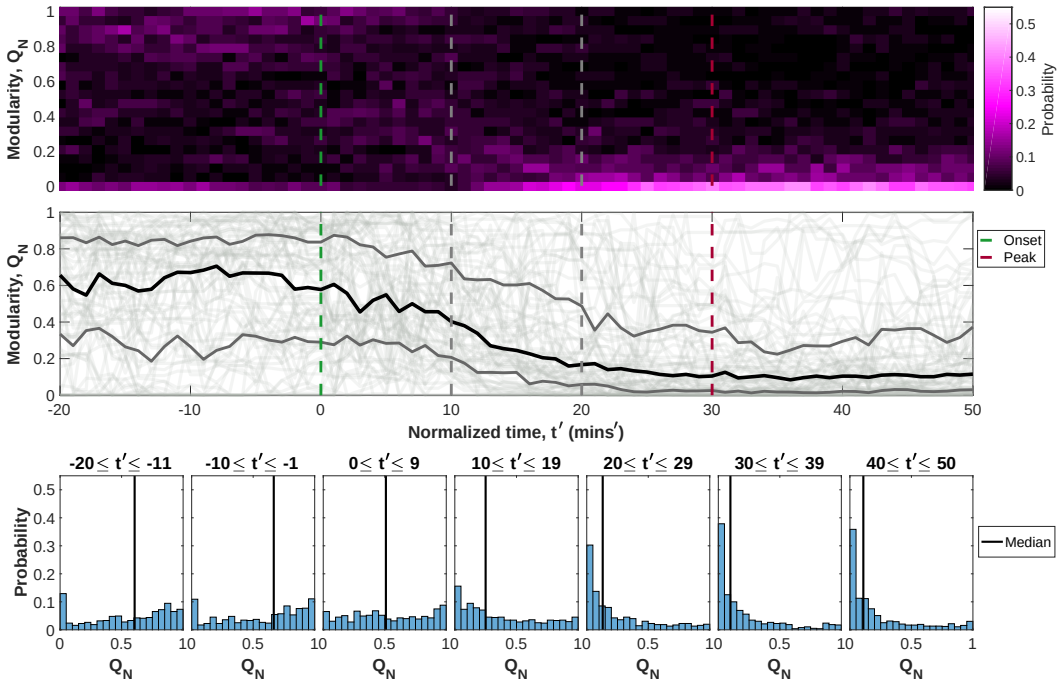

Figure 5.24: The normalized modularity, Q_N, of 75 substorms, in the same format as figure 5.5 (which contained 41 substorms which were quiet before onset). This plot used the edge betweenness community detection algorithm. The transition from high to low modularity during the expansion phase is clear, but there are more outliers, particularly those with low modularity before onset.

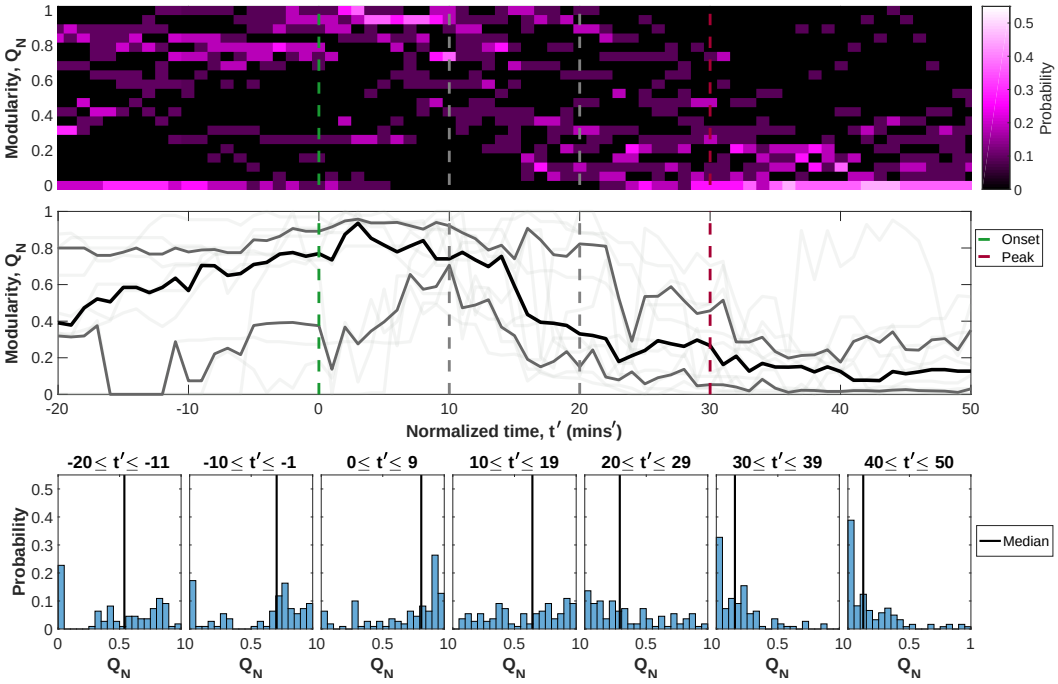

Figure 5.25: The normalized modularity, Q_N, of 11 substorms with SML> 25% of that at the time of peak expansion during the 127 minute window before onset but < 50%, in the same format as figure 5.5 (which contained 41 substorms which were quiet before onset (SML < 25%)). This plot used the edge betweenness community detection algorithm. The transition from high to low modularity is clear but there are more outliers, particularly those with low modularity before onset.

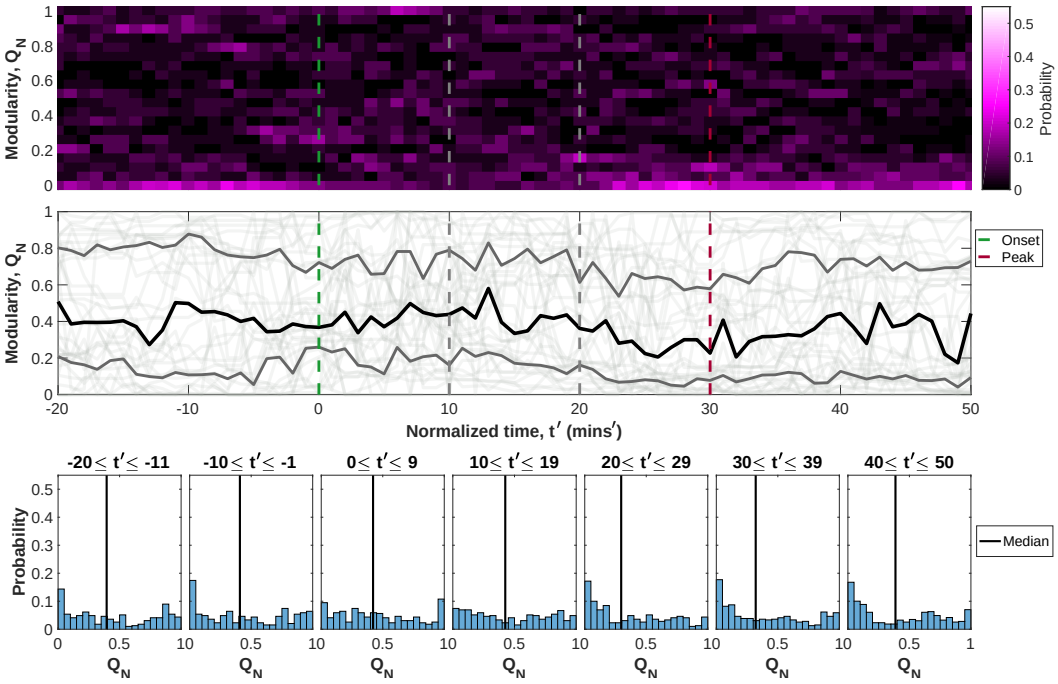

Figure 5.26: The normalized modularity, Q_N, of the substorms (23) that have good nightside coverage and are NOT quiet before onset, in the same format as figure 5.5 (which only contained the 41 substorms which were extremely quiet for 127 minutes before onset). This plot used the edge betweenness community detection algorithm. The transition from high to low modularity is not observed for the non-isolated substorms.

between $25-50\%$ of the SML measured at the peak of the substorm. Recall figures 5.2-5.23 used only the substorms which were "extremely" quiet for 127 minutes before onset (SML $< 25\%$). The heightened magnetic activity results in increased correlation across the nightside, leading to lower modularity (the mean $Q_N \sim 0.4$ at $t' = -20$) before onset. As substorm onset approaches, less of the magnetic perturbations from the previous substorm/ pseudo-breakup are included in the correlation window and the modularity begins to rise again (i.e. the nightside becomes less structured and returns to separate communities). By onset the mean is $Q_N = 0.8$ but $\sim 25\%$ of the substorms still have small modularity (< 0.3). By peak more than 50% of the substorms have $Q_N < 0.2$ and, therefore, resemble a globally correlated system. Increased magnetic perturbations before onset clearly have an affect on the community structure.

Figure 5.26 uses the 23 substorms which had a very disturbed ionosphere before onset i.e. the SML reached $> 50\%$ of its maximum extend before the substorm onset. Many of these events had a substorm in the 128 minutes preceding the substorm of interest and hence the network already had many connections and did not have a clear increased network parameter response at onset. In figure 5.26 the modularity response is completely different to that seen in figure 5.5. Here we see little pattern in the modularity.

Isolated substorms with at least 127 minutes of low SME before onset are a requirement for high to low modularity, otherwise the increased correlation reflected in the community structure.

Changing the choice of threshold for network calculation

All previous figures in this chapter (5.2-5.26) had a threshold set at 5%. The threshold is set such that each magnetometer would be connected to the network for 5% of the month surrounding the event in question. A magnetometer pair are only then connected to the network if their CCC exceeds the threshold (the lower of the two stations' thresholds). In figures 5.27 and 5.28 the threshold has been changed such that fewer and more connections are included in the network. For figure 5.27 the threshold for the network is set at 1% i.e. magnetometers need to be more highly canonically cross correlated to be included in the network and there will be less magnetometer pairs which exceed the threshold. For figure 5.28 the threshold for the networks is set at 10% i.e. magnetometers need a lower component of CCC to be connected and more magnetometers pairs will exceed the threshold. Here we do not see the individual connections but the effect increasing or decreasing the sensitivity has on the overall modularity. In figure 5.27 the pattern of decreasing mean Q_N during the substorm expansion is observed but significant variety in Q_N values compared to in figure 5.5. The distribution still changes from right shifted before onset to left shifted from half way through the expansion phase,

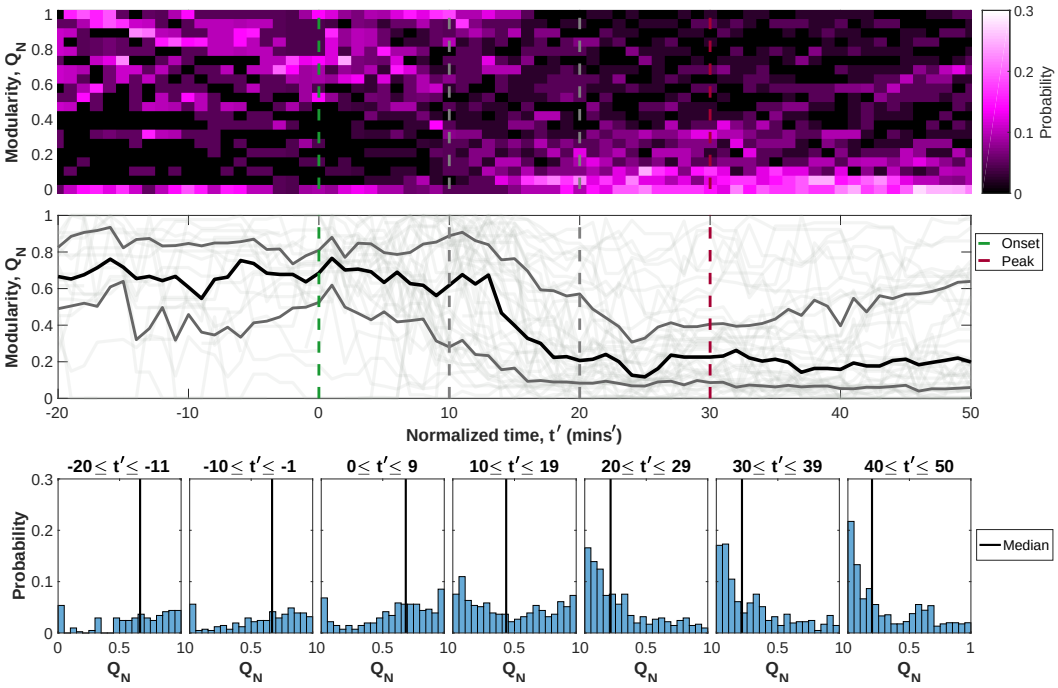

Figure 5.27: The normalized modularity, Q_N, where the network cross-correlation threshold has been increased such that station pairs are correlated for only 1% of the month surrounding the event, hence there are less connected magnetometers compared to figure 5.5 (for which the threshold was set at 5%). The figure is in the same format as figure 5.5. The modularity has been calculated using the edge betweenness community detection algorithm for the 41 isolated substorms. The transition from high to low modularity is still evident, but there are a greater proportion of outliers. Before onset many of the substorms with zero modularity are due to the network containing few connections.

but there are many more outliers. The small number of network connections clearly has an effect on the community structure.

In figure 5.28 there are less outliers. There is a clear decrease from a mean of $Q_N \sim 0.6$ before onset to $Q_N \sim 0$ at peak expansion. If we have too many connections (i.e. threshold is set at $\gg 10\%$ and the CCC required to exceed the threshold is much lower) the network would become saturated and community structure would be challenging to decipher. However, in 5.28 we still see a consistent pattern of increasing global structure during the substorm expansion.

5.5.3 Comparison of modularity to that of a random network.

Figure 5.29 shows how the modularity scales with the number of connections and the number of nodes. The modularity from a random network with n nodes and m connections is plotted in red with the modularity from the magnetometer network overplotted. The colour of the magnetometer network points reflects the normalized time at which the measurement was observed. For each number of nodes, n, and each number of con-

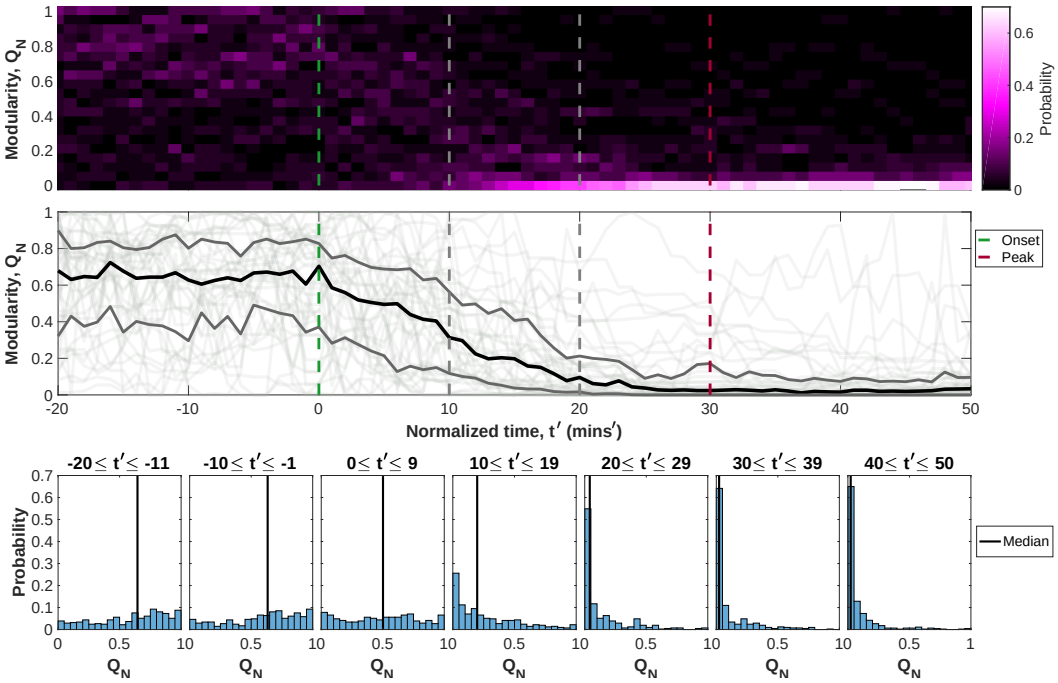

Figure 5.28: The normalized modularity, Q_N, where the network cross-correlation threshold has been lowered such that station pairs are correlated for 10% of the month surrounding the event, hence there are more connected magnetometers compared to figure 5.5 (for which the threshold was set at 5%). The figure is in the same format as figure 5.5 . The modularity has been calculated using the edge betweenness community detection algorithm for the 41 isolated substorms, as in the main text. The transition from high to low modularity is clear.

nections, m, 100 random networks were calculated. The modularity of the community structure was then calculated for each network and plotted against the number of connections. There is a clear exponential decay of the modularity, Q, with the number of connections, m. The coloured modularities are chosen from the 41 events used in figure 5.5 and binned by the number of nodes (active magnetometers). The modularity is then plotted against the number of connections. The modularity from the networks derived from the substorm events shows more structure (higher Q) than the random networks and, for a given m, the data explores a broad range of Q and vice versa. There is still a clear upper bound to the value of Q per m but it is much higher than that of random networks. The patterns of decreasing modularity observed throughout this chapter are not simply a result of the increasing numbers of connections.

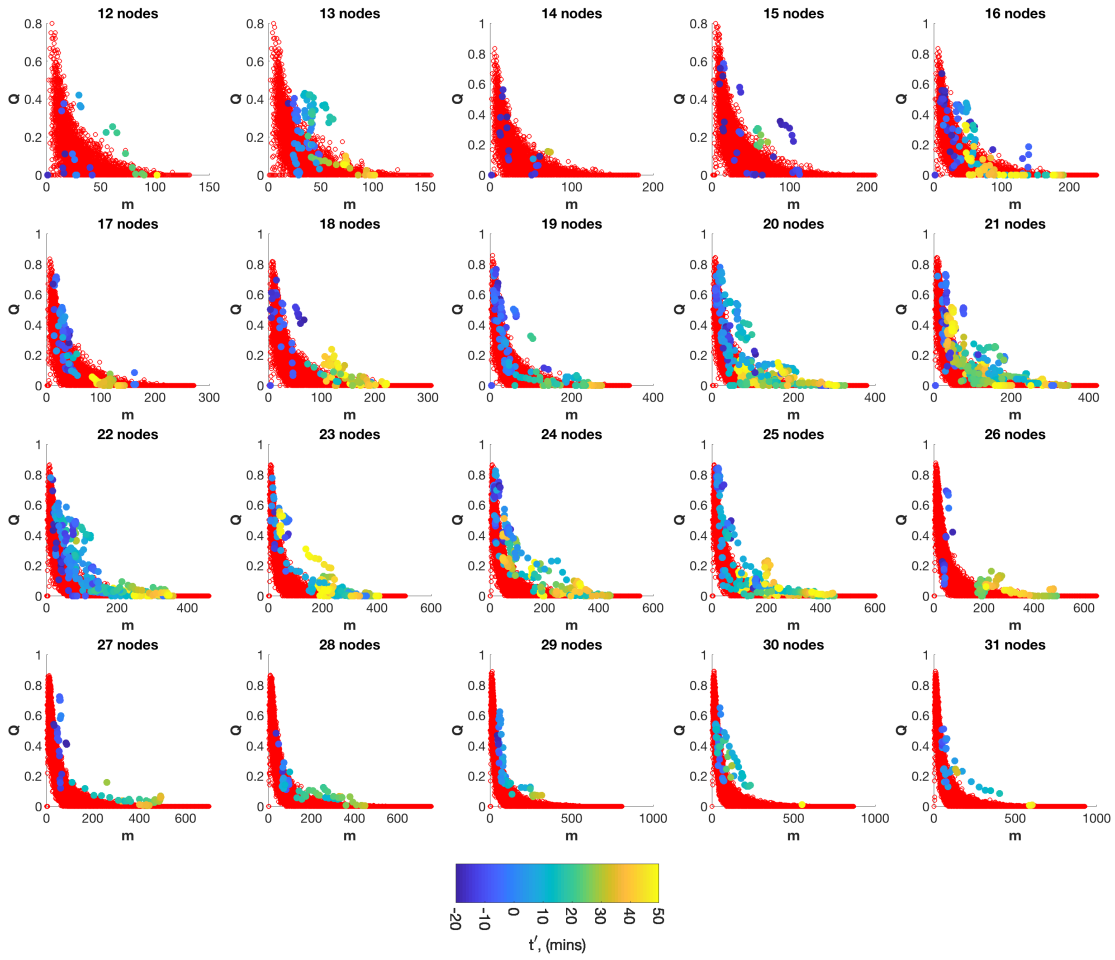

Figure 5.29: Modularity, Q, plotted versus the number of connections, m. Each panel selects times from all 41 events when the network is comprised of a specific number of nodes (normalized time indicated by colour). The edge betweenness algorithm has been used for community detection as in main text. The red circles are the modularity obtained from randomly generated networks ($m \times 100$). The plot shows random networks have a range of modularities (a measure of how separated the communities are) and there is a threshold in Q, m space. The modularity from the networks derived from the substorm events shows more structure (higher Q) than the random networks and for a given m the data explores a broad range of Q and vice versa.

Chapter 6

Conclusions

6.1 Summary

The magnetosphere is a driven, dissipative, out of equilibrium system that couples nonlinear physical processes across multiple spatial and temporal scales. The substorm is a fundamental, global reconfiguration of the magnetosphere during which energy is abruptly transported to the ionosphere where it is dissipated. Central to this phenomenon are the auroral electrojets: large-scale electrical currents flowing in ionosphere, that are part of a global three-dimensional current system. This current system is known as the substorm current wedge (SCW). How the SCW is evolving and how it is spatially structured is a fundamental problem and there are many conflicting models [McPherron et al., 1973; Kamide and Kokubun, 1996; Sergeev et al., 2011; Gjerloev and Hoffman, 2014; Liu et al., 2018].

Heliospheric, and in particular, magnetospheric observations are becoming a data analytical challenge as we move from a data-poor to a data-rich era. This requires novel approaches to the data analysis and this thesis presents an example of this. We have for the first time obtained the time-varying directed network that represents the observations from 100+ ground-based magnetometers and this captures the perturbation from the substorm current system. If the canonical cross-correlation between vector magnetic field perturbations, observed at two magnetometer stations, exceeds an event and station specific threshold, they form a network connection. The time lag at which cross-correlation is maximal determines the direction of propagation or expansion of the structure captured by the network connection. If spatial correlation reflects ionospheric current patterns, network properties can test different models for the evolving substorm current system. This offers an entirely new way to synthesize multi-point magnetospheric observations into a form that readily tests physical model hypothesis.

We calculated specific time-varying sub-networks from the magnetometer data which isolate specific physical regions. These regions are west, within and east of the auroral bulge boundaries for each of 86 substorms, where the boundaries were obtained from polar VIS images at the time of peak expansion. We obtained the timings for, a consistent picture in which the SCW forms. A current system is seen pre-midnight following the SCW westward expansion. Later, there is a weaker signal of eastward expansion. The westward expansion extends towards lower latitudes. These spatio-temporal timings are concluded from the averaged network over 86 isolated substorms.

This analysis required polar VIS data to be analysed by hand to determine the region boundaries for each event, and makes the assumption that these regions will unambiguously demarcate the various features of the dynamically evolving current system. We then moved to a fully automated methodology with the goal of finding a parameter for the network that characterises the full spatio-temporal pattern. We perform community detection on the network which identifies locally dense but globally sparse groups of connections. We consistently find robust structural change from many small, uncorrelated current systems before substorm onset, to one large spatially-extended correlated system during the expansion phase. All substorms analysed ultimately form a large-scale structure, approximately ten minutes after onset. This establishes that a single large-scale SCW is central to substorm physics and that substorms do not proceed solely by small-scale wedgelets.

6.2 Future Work

The parameter modularity provides a quantitative benchmark for global models of the magnetosphere, these models in turn feed into space weather forecasting and future data collection. The wider implications could include modularity as a parameter to characterise the current pattern. It has the potential to be used as an index, like AE or SME, to describe the global magnetic signature of auroral electrojets. Having a quantifiable characterisation of the ground-based response to substorms, such as the modularity, allows us to quantify the correlation with the solar wind and the magnetospheric state. A lagged correlation with solar wind parameters, such as the B field, the plasma density or the velocity, from upstream monitors, such as WIND [Lepping et al., 1995] and ACE [Smith et al., 1998], could provide advanced warning of when space weather is approaching. Data from industries effected by space weather, e.g. aviation, GPS or satellite e.t.c. could be incorporated to discover if there is a quantifiable difference between events which cause disruption to their infrastructure and those that do not.

The methodology in this thesis could be applied to other data sets such as the AMPERE (Active Magnetosphere and Planetary Electrodynamics Response Experiment), a set of

magnetometers attached to spacecrafts in low earth orbit observing magnetic perturbations [Anderson et al., 2000]. SuperMAG is in the process of releasing magnetometer observations at 1 second temporal resolution (analysis in this thesis was at 1 minute temporal resolution). The network methodology could easily be adapted to this higher resolution. The network could be applied to national grid infrastructure to monitor how GIC propagate through the system.

The network can be applied to a wider set of substorms using lists provided by SuperMAG or the SOPHIE technique for determining 'Substorm Onsets and Phases from Indices of the Electrojet' [Forsyth et al., 2015]. Other space weather phenomena such as geomagnetic storms and IMF turnings [Dods et al., 2017] could also be analysed. The modularity parameter could be used to characterise the differences between isolated substorms and substorms within storms. Perhaps there is a precursor to large storms that could be characterised as a prediction of incoming space weather.

Canonical cross correlation has both magnitude and direction. So far we have only looked at magnitude but we can use the direction to infer transient current patterns. In principle this could resolve the direction of the electojet (eastward/westward). Communities is one commonly used network property that we have applied to the substorm network, but there are others, such as hierarchical incoherence and the source nodes of interactions (basal nodes) [Pagani et al., 2019], that may yield more information on current pattern propagation.

Bibliography

C Donald Ahrens and Perry J Samson. *Extreme weather and climate*. Cengage Learning, 2010.

Syun-Ichi Akasofu and Sydney Chapman. The ring current, geomagnetic disturbance, and the van allen radiation belts. *Journal of Geophysical Research*, 66(5):1321–1350, 1961.

Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.

Hannes Alfvén. Electric currents in cosmic plasmas. *Reviews of Geophysics*, 15(3): 271–284, 1977.

O Amm, A Pajunpää, and U Brandström. Spatial distribution of conductances and currents associated with a north-south auroral form during a multiple-substorm period. In *Annales Geophysicae*, volume 17, pages 1385–1396. Springer, 1999.

Olaf Amm. Direct determination of the local ionospheric hall conductance distribution from two-dimensional electric and magnetic field data: Application of the method using models of typical ionospheric electrodynamic situations. *Journal of Geophysical Research: Space Physics*, 100(A11):21473–21488, 1995.

Brian J Anderson, Kazue Takahashi, and Bruce A Toth. Sensing global birkeland currents with iridium® engineering magnetometer data. *Geophysical Research Letters*, 27(24):4045–4048, 2000.

V Angelopoulos, CF Kennel, FV Coroniti, R Pellat, MG Kivelson, RJ Walker, CT Russell, W Baumjohann, WC Feldman, and JT Gosling. Statistical characteristics of bursty bulk flow events. *Journal of Geophysical Research: Space Physics*, 99(A11): 21257–21280, 1994.

V Angelopoulos, A Runov, X-Z Zhou, DL Turner, SA Kiehas, S-S Li, and I Shinohara. Electromagnetic energy conversion at reconnection fronts. *science*, 341(6153):1478–1482, 2013.

Vassilis Angelopoulos, Wo Baumjohann, CF Kennel, F Vo Coroniti, MG Kivelson, R Pellat, RJ Walker, H Lühr, and G Paschmann. Bursty bulk flows in the inner central plasma sheet. *Journal of Geophysical Research: Space Physics*, 97(A4):4027–4039, 1992.

E Astafyeva, Yu Yasyukevich, A Maksikov, and I Zhivetiev. Geomagnetic storms, superstorms, and their impacts on gps-based navigation systems. *Space Weather*, 12(7): 508–525, 2014.

W Ian Axford, Harry E Petschek, and George L Siscoe. Tail of the magnetosphere. *Journal of Geophysical Research*, 70(5):1231–1236, 1965.

Wolfgang Baumjohann and Rudolf A Treumann. *Basic space plasma physics*. World Scientific, 1997.

Ciarán D Beggan. Sensitivity of geomagnetically induced currents to varying auroral electrojet and conductivity models. *Earth, Planets and Space*, 67(1):24, 2015.

L Biermann. Solar corpuscular radiation and the interplanetary gas. *The Observatory*, 77:109–110, 1957.

Norman Biggs, E Keith Lloyd, and Robin J Wilson. *Graph Theory, 1736-1936*. Oxford University Press, 1986.

J Birn and M Hesse. The substorm current wedge and field-aligned currents in mhd simulations of magnetotail reconnection. *Journal of Geophysical Research: Space Physics*, 96(A2):1611–1618, 1991.

J Birn and M Hesse. The substorm current wedge in mhd simulations. *Journal of Geophysical Research: Space Physics*, 118(6):3364–3376, 2013.

J Birn and M Hesse. The substorm current wedge: Further insights from mhd simulations. *Journal of Geophysical Research: Space Physics*, 119(5):3503–3513, 2014.

J Birn, J Liu, A Runov, L Kepko, and V Angelopoulos. On the contribution of dipolarizing flux bundles to the substorm current wedge and to flux and energy transport. *Journal of Geophysical Research: Space Physics*, 2019.

Joachim Birn and Edward W Hones Jr. Three-dimensional computer modeling of dynamic reconnection in the geomagnetic tail. *Journal of Geophysical Research: Space Physics*, 86(A8):6802–6808, 1981.

José A Bittencourt. *Fundamentals of plasma physics*. Springer Science & Business Media, 2013.

Ulrik Brandes, Daniel Delling, Marco Gaertler, Robert Gorke, Martin Hoefer, Zoran Nikoloski, and Dorothea Wagner. On modularity clustering. *IEEE transactions on knowledge and data engineering*, 20(2):172–188, 2007.

Guido Caldarelli. *Scale-free networks: complex webs in nature and technology*. Oxford University Press, 2007.

J. A. Carter, S. E. Milan, J. C. Coxon, M.-T. Walach, and B. J. Anderson. Average field-aligned current configuration parameterized by solar wind conditions. *Journal of Geophysical Research: Space Physics*, 121(2):1294–1307, 2016. ISSN 2169-9402. doi: 10.1002/2015JA021567. URL http://dx.doi.org/10.1002/2015JA021567. 2015JA021567.

GF Chew, ML Goldberger, and FE Low. The boltzmann equation an d the one-fluid hydromagnetic equations in the absence of particle collisions. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 236(1204):112–118, 1956.

Nian Shong Chok. *Pearson's versus Spearman's and Kendall's correlation coefficients for continuous data*. PhD thesis, University of Pittsburgh, 2010.

SP Christon, TE Eastman, T Doke, LA Frank, G Gloeckler, H Kojima, S Kokubun, ATY Lui, H Matsumoto, RW McEntire, et al. Magnetospheric plasma regimes identified using geotail measurements: 2. statistics, spatial distribution, and geomagnetic dependence. *Journal of Geophysical Research: Space Physics*, 103(A10):23521–23542, 1998.

Xiangning Chu. *Configuration and generation of substorm current wedge*. PhD thesis, UCLA, 2015.

Aaron Clauset, Mark EJ Newman, and Cristopher Moore. Finding community structure in very large networks. *Physical review E*, 70(6):066111, 2004.

GRJ Cooper and DR Cowan. Comparing time series using wavelet-based semblance analysis. *Computers & Geosciences*, 34(2):95–102, 2008.

National Research Council. *Plasma physics of the local cosmos*. National Academies Press, 2004.

Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

SWH Cowley. Magnetosphere-ionosphere interactions: A tutorial review. *Magnetospheric Current Systems, Geophys. Monogr. Ser*, 118:91–106, 2000.

William D Cramer, Joachim Raeder, FR Toffoletto, M Gilson, and Bei Hu. Plasma sheet injections into the inner magnetosphere: Two-way coupled openggcm-rcm model results. *Journal of Geophysical Research: Space Physics*, 122(5):5077–5091, 2017.

Gabor Csardi and Tamas Nepusz. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695, 2006. URL http://igraph.org.

T Neil Davis and Masahisa Sugiura. Auroral electrojet activity index ae and its universal time variations. *Journal of Geophysical Research*, 71(3):785–801, 1966.

Jonathan Decowski and Linyuan Li. Wavelet-based tests for comparing two time series with unequal lengths. *Journal of Time Series Analysis*, 36(2):189–208, 2015.

J Dods, SC Chapman, and JW Gjerloev. Characterizing the ionospheric current pattern response to southward and northward imf turnings with dynamical supermag correlation networks. *Journal of Geophysical Research: Space Physics*, 122(2):1883–1902, 2017.

Joe Dods, Sandra C Chapman, and Jesper W Gjerloev. Network analysis of geomagnetic substorms using the supermag database of ground-based magnetometer stations. *Journal of Geophysical Research: Space Physics*, 120(9):7774–7784, 2015.

Richard Frank Donnelly. *Solar-terrestrial predictions proceedings*, volume 2.

Sergei N Dorogovtsev and José FF Mendes. *Evolution of networks: From biological nets to the Internet and WWW*. OUP Oxford, 2013.

S Dubyagin, V Sergeev, S Apatenkov, V Angelopoulos, A Runov, R Nakamura, W Baumjohann, J McFadden, and D Larson. Can flow bursts penetrate into the inner magnetosphere? *Geophysical Research Letters*, 38(8), 2011.

James W Dungey. Interplanetary magnetic field and the auroral zones. *Physical Review Letters*, 6(2):47, 1961.

JP Eastwood, H Hietala, G Toth, TD Phan, and M Fujimoto. What controls the structure and dynamics of earth's magnetosphere? *Space Science Reviews*, 188(1-4): 251–286, 2015.

Ralf Engbert. Testing for nonlinearity: the role of surrogate data. *Chaos, Solitons & Fractals*, 13(1):79–84, 2002.

Paul Erdős and Alfréd Rényi. On random graphs. *Publicationes mathematicae*, 6(26): 290–297, 1959.

Leonhard Euler. Commentarii academiae scientiarum petropolitanae. *Solutio problematis ad geometriam situs pertinentis*, 8:128–140, 1736.

Lyndsay Fletcher, Brian R Dennis, Hugh S Hudson, Sam Krucker, Ken Phillips, Astrid Veronig, Marina Battaglia, Laura Bone, Amir Caspi, Qingrong Chen, et al. An observational overview of solar flares. *Space science reviews*, 159(1-4):19, 2011.

C Forsyth, AN Fazakerley, IJ Rae, CE J. Watt, K Murphy, James A Wild, Tomas Karlsson, Robert Mutel, CJ Owen, Robert Ergun, et al. In situ spatiotemporal measurements of the detailed azimuthal substructure of the substorm current wedge. *Journal of Geophysical Research: Space Physics*, 119(2):927–946, 2014.

C Forsyth, IJ Rae, JC Coxon, MP Freeman, CM Jackman, J Gjerloev, and AN Fazakerley. A new technique for determining substorm onsets and phases from indices of the electrojet (sophie). *Journal of Geophysical Research: Space Physics*, 120(12): 10–592, 2015.

Colin Forsyth, M Lester, SWH Cowley, I Dandouras, AN Fazakerley, RC Fear, HU Frey, Adrian Grocott, A Kadokura, E Lucek, et al. Observed tail current systems associated with bursty bulk flows and auroral streamers during a period of multiple substorms. 26(1):167–184, 2008.

LA Frank, JB Sigwarth, JD Craven, JP Cravens, JS Dolan, MR Dvorsky, PK Hardebeck, JD Harvey, and DW Muller. The visible imaging system (vis) for the polar spacecraft. *Space Science Reviews*, 71(1-4):297–328, 1995.

Huishan S Fu, Yuri V Khotyaintsev, Mats André, and Andris Vaivads. Fermi and betatron acceleration of suprathermal electrons behind dipolarization fronts. *Geophysical Research Letters*, 38(16), 2011.

Naoshi Fukushima. Generalized theorem for no ground magnetic effect of vertical currents connected with pedersen currents in the uniform-conductivity ionosphere. *Rep. Ionos. Space Res. Jpn.*, 30(1):35–40, 1976.

C Gabrielse, V Angelopoulos, A Runov, and DL Turner. The effects of transient, localized electric fields on equatorial electron acceleration and transport toward the inner magnetosphere. *Journal of Geophysical Research: Space Physics*, 117(A10), 2012.

N Yu Ganushkina, MW Liemohn, and S Dubyagin. Current systems in the earth's magnetosphere. *Reviews of Geophysics*, 56(2):309–332, 2018.

JW Gjerloev. The supermag data processing technique. *Journal of Geophysical Research: Space Physics*, 117(A9), 2012.

JW Gjerloev and RA Hoffman. Currents in auroral substorms. *Journal of Geophysical Research: Space Physics*, 107(A8):SMP–5, 2002.

JW Gjerloev and RA Hoffman. The large-scale current system during auroral substorms. *Journal of Geophysical Research: Space Physics*, 119(6):4591–4606, 2014.

JW Gjerloev, RA Hoffman, JB Sigwarth, and LA Frank. Statistical description of the bulge-type auroral substorm in the far ultraviolet. *Journal of Geophysical Research: Space Physics*, 112(A7), 2007.

Robert J Goldston and Paul Harding Rutherford. *Introduction to plasma physics*. CRC Press, 1995.

Nat Gopalswamy. Coronal mass ejections and space weather. In *Climate and Weather of the Sun-Earth System (CAWSES): Selected Papers from the 2007 Kyoto Symposium*, pages 77–120. Terrapub, 2009.

Jan Hauke and Tomasz Kossowski. Comparison of values of pearson's and spearman's correlation coefficients on the same sets of data. *Quaestiones geographicae*, 30(2): 87–93, 2011.

Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.

White House. Executive order—coordinating efforts to prepare the nation for space weather events., 2016.

Trevor R Hutchins and Thomas J Overbye. The effect of geomagnetic disturbances on the electric grid and appropriate mitigation strategies. In *2011 North American Power Symposium*, pages 1–5. IEEE, 2011.

Takesi Iijima and Thomas A Potemra. The amplitude distribution of field-aligned currents at northern high latitudes observed by triad. *Journal of Geophysical Research*, 81(13):2165–2174, 1976.

Brian J Jackel, Paul Eglitis, Eric F Donovan, Ari T Viljanen, Don D Wallis, Leroy L Cogger, and Hermann J Opgenoorth. Observations of highly correlated near-simultaneous magnetic field perturbations at contraposed ground stations. *Journal of Geophysical Research: Space Physics*, 106(A11):25857–25872, 2001.

JBL Jones, RD Bentley, R Hunter, RHA Iles, GC Taylor, and DJ Thomas. Space weather and commercial airlines. *Advances in Space Research*, 36(12):2258–2267, 2005.

Jürgen Jost. Dynamical networks. In *Networks: from biology to theory*, pages 35–62. Springer, 2007.

L Juusola, N Østgaard, E Tanskanen, N Partamies, and K Snekvik. Earthward plasma sheet flows during substorm phases. *Journal of Geophysical Research: Space Physics*, 116(A10), 2011.

L Juusola, A Viljanen, M Van De Kamp, EI Tanskanen, H Vanhamäki, N Partamies, and K Kauristie. High-latitude ionospheric equivalent currents during strong space storms: Regional perspective. *Space Weather*, 13(1):49–60, 2015.

Y Kamide and S Kokubun. Two-component auroral electrojet: Importance for substorm studies. *Journal of Geophysical Research: Space Physics*, 101(A6):13027–13046, 1996.

Larry Kepko, RL McPherron, O Amm, S Apatenkov, Wolfgang Baumjohann, J Birn, M Lester, R Nakamura, Tuija I Pulkkinen, and V Sergeev. Substorm current wedge revisited. *Space Science Reviews*, 190(1-4):1–46, 2015.

Nicholas A Krall and Alvin W Trivelpiece. Principles of plasma physics. *American Journal of Physics*, 41(12):1380–1381, 1973.

Wojtek J Krzanowski. *Statistical principles and techniques in scientific and social research*. Oxford University Press on Demand, 2007.

Dimitris Kugiumtzis and Alkiviadis Tsimpiris. Measures of analysis of time series (mats): A matlab toolkit for computation of multiple measures on time series data bases. *arXiv preprint arXiv:1002.1940*, 2010.

Anita Kullen and Tomas Karlsson. On the relation between solar wind, pseudobreakups, and substorms. *Journal of Geophysical Research: Space Physics*, 109(A12), 2004.

Gemma Lancaster, Dmytro Iatsenko, Aleksandra Pidde, Valentina Ticcinelli, and Aneta Stefanovska. Surrogate data for hypothesis testing of physical systems. *Physics Reports*, 748:1–60, 2018.

Lev Davidovich Landau, JS Bell, MJ Kearsley, LP Pitaevskii, EM Lifshitz, and JB Sykes. *Electrodynamics of continuous media*, volume 8. elsevier, 2013.

Karl Magnus Laundal, JW Gjerloev, Nikolai Østgaard, Jone Peter Reistad, Stein Haaland, Kristian Snekvik, Paul Tenfjord, Shin-ichi Ohtani, and SE Milan. The impact of sunlight on high-latitude equivalent currents. *Journal of Geophysical Research: Space Physics*, 121(3):2715–2726, 2016.

Erik G Learned-Miller. Entropy and mutual information. *Department of Computer Science, University of Massachusetts, Amherst*, 2013.

RP Lepping, MH Acũna, LF Burlaga, WM Farrell, JA Slavin, KH Schatten, F Mariani, NF Ness, FM Neubauer, YC Whang, et al. The wind magnetic field investigation. *Space Science Reviews*, 71(1-4):207–229, 1995.

J Liu, V Angelopoulos, A Runov, and X-Z Zhou. On the current sheets surrounding dipolarizing flux bundles in the magnetotail: The case for wedgelets. *Journal of Geophysical Research: Space Physics*, 118(5):2000–2020, 2013.

Jiang Liu, V Angelopoulos, Xu-Zhi Zhou, and A Runov. Magnetic flux transport by dipolarizing flux bundles. *Journal of Geophysical Research: Space Physics*, 119(2): 909–926, 2014.

Jiang Liu, V Angelopoulos, Xiangning Chu, Xu-Zhi Zhou, and Chao Yue. Substorm current wedge composition by wedgelets. *Geophysical Research Letters*, 42(6):1669–1676, 2015.

Jiang Liu, V Angelopoulos, Zhonghua Yao, Xiangning Chu, Xu-Zhi Zhou, and A Runov. The current system of dipolarizing flux bundles and their role as wedgelets in the substorm current wedge. *Electric Currents in Geospace and Beyond*, pages 323–337, 2018.

ATY Lui. Electric current approach to magnetospheric physics and the distinction between current disruption and magnetic reconnection. *GEOPHYSICAL MONOGRAPH-AMERICAN GEOPHYSICAL UNION*, 118:31–40, 2000.

Andrey Y Malykhin, Elena E Grigorenko, Elena A Kronberg, Rositza Koleva, Natalia Y Ganushkina, Ludmila Kozak, and Patrick W Daly. Contrasting dynamics of electrons and protons in the near-earth plasma sheet during dipolarization. In *Annales Geophysicae*, volume 36, pages 741–760. Copernicus GmbH, 2018.

MATLAB. *9.7.0.1190202 (R2019b)*. The MathWorks Inc., Natick, Massachusetts, 2018.

Robert L McPherron. Growth phase of magnetospheric substorms. *Journal of Geophysical Research*, 75(28):5592–5599, 1970.

Robert L McPherron, Christopher T Russell, and Michel P Aubry. Satellite studies of magnetospheric substorms on august 15, 1968: 9. phenomenological model for substorms. *Journal of Geophysical Research*, 78(16):3131–3149, 1973.

Viacheslav G Merkin, Evgeny V Panov, KA Sorathia, and A Ukhorskiy. Contribution of bursty bulk flows to the global dipolarization of the magnetotail during an isolated substorm. *Journal of Geophysical Research: Space Physics*, 2019.

SE Milan, G Provan, and Benoît Hubert. Magnetic flux transport in the dungey cycle: A survey of dayside and nightside reconnection rates. *Journal of Geophysical Research: Space Physics*, 112(A1), 2007.

Stephen Eric Milan, Lasse Boy Novock Clausen, John Charles Coxon, Jenny A Carter, M-T Walach, K Laundal, Nikolai Østgaard, P Tenfjord, J Reistad, Kristian Snekvik, et al. Overview of solar wind–magnetosphere–ionosphere–atmosphere coupling and the generation of magnetospheric currents. *Space Science Reviews*, 206(1-4):547–573, 2017.

T Nagai, M Fujimoto, Y Saito, S Machida, T Terasawa, R Nakamura, T Yamamoto, T Mukai, A Nishida, and S Kokubun. Structure and dynamics of magnetic reconnection for substorm onsets with geotail observations. *Journal of Geophysical Research: Space Physics*, 103(A3):4419–4440, 1998.

Tsugunobu Nagai. Observed magnetic substorm signatures at synchronous altitude. *Journal of Geophysical Research: Space Physics*, 87(A6):4405–4417, 1982.

Tsugunobu Nagai. An empirical model of substorm-related magnetic field variations at synchronous orbit. *Washington DC American Geophysical Union Geophysical Monograph Series*, 64:91–95, 1991.

R Nakamura, W Baumjohann, B Klecker, Y Bogdanova, A Balogh, H Reme, JM Bosqued, I Dandouras, JA Sauvaud, K-H Glassmeier, et al. Motion of the dipolarization front during a flow burst event observed by cluster. *Geophysical Research Letters*, 29(20):3–1, 2002.

R Nakamura, O Amm, H Laakso, NC Draper, M Lester, Adrian Grocott, B Klecker, IW McCrea, A Balogh, H Reme, et al. Localized fast flow disturbance observed in the plasma sheet and in the ionosphere. In *Annales Geophysicae*, volume 23, pages 553–566, 2005.

Rumi Nakamura, Tsugunobu Nagai, Joachim Birn, Victor A Sergeev, Olivier Le Contel, Ali Varsani, Wolfgang Baumjohann, Takuma Nakamura, Sergey Apatenkov, Anton Artemyev, et al. Near-earth plasma sheet boundary dynamics during substorm dipolarization. *Earth, Planets and Space*, 69(1):129, 2017.

Jason J Neal, Craig J Rodger, and Janet C Green. Empirical determination of solar proton access to the atmosphere: Impact on polar flight paths. *Space Weather*, 11 (7):420–433, 2013.

PT Newell and JW Gjerloev. Evaluation of supermag auroral electrojet indices as indicators of substorms and auroral power. *Journal of Geophysical Research: Space Physics*, 116(A12), 2011a.

PT Newell and JW Gjerloev. Substorm and magnetosphere characteristic scales inferred from the supermag auroral electrojet indices. *Journal of Geophysical Research: Space Physics*, 116(A12), 2011b.

Mark EJ Newman. The structure and function of complex networks. *SIAM review*, 45 (2):167–256, 2003.

Mark EJ Newman. Fast algorithm for detecting community structure in networks. *Physical review E*, 69(6):066133, 2004.

Mark EJ Newman and Michelle Girvan. Finding and evaluating community structure in networks. *Physical review E*, 69(2):026113, 2004.

Atsuhiro Nishida. Coherence of geomagnetic dp 2 fluctuations with interplanetary magnetic variations. *Journal of Geophysical Research*, 73(17):5549–5559, 1968.

Cabinet Office. National risk register of civil emergencies-2017 edition, 2017.

Cabinet Office, Innovation Department for Business, and Skills. Space weather preparedness strategy, 2015.

L Orr, SC Chapman, and JW Gjerloev. Directed network of substorms using supermag ground-based magnetometer data. *Geophysical Research Letters*, 46(12):6268–6278, 2019.

Edward J Oughton. The economic impact of critical national infrastructure failure due to space weather. In *Oxford Research Encyclopedia of Natural Hazard Science*. 2018.

Edward J Oughton, Mike Hapgood, Gemma S Richardson, Ciarán D Beggan, Alan WP Thomson, Mark Gibbs, Catherine Burnett, C Trevor Gaunt, Markos Trichas, Rabia Dada, et al. A risk assessment framework for the socioeconomic impacts of electricity transmission infrastructure failure due to space weather: An application to the united kingdom. *Risk Analysis*, 39(5):1022–1043, 2019.

Alessio Pagani, Guillem Mosquera, Aseel Alturki, Samuel Johnson, Stephen Jarvis, Alan Wilson, Weisi Guo, and Liz Varga. Resilience or robustness: identifying topological vulnerabilities in rail networks. *Royal Society open science*, 6(2):181301, 2019.

Laurianne Palin, Hermann J Opgenoorth, Karin Ågren, Tatjana Zivkovic, VA Sergeev, MV Kubyshkina, A Nikolaev, K Kauristie, M van de Kamp, O Amm, et al. Modulation of the substorm current wedge by bursty bulk flows: 8 september 2002—revisited. *Journal of Geophysical Research: Space Physics*, 121(5):4466–4482, 2016.

Evgeny V Panov, Wolfgang Baumjohann, Richard A Wolf, Rumi Nakamura, Vassilis Angelopoulos, James M Weygand, and Marina V Kubyshkina. Magnetotail energy dissipation during an auroral substorm. *Nature physics*, 12(12):1158, 2016.

Athanasios Papaioannou, Anastasios Anastasiadis, Ingmar Sandberg, and Piers Jiggens. Nowcasting of solar energetic particle events using near real-time coronal mass ejection characteristics in the framework of the forspef tool. *Journal of Space Weather and Space Climate*, 8:A37, 2018.

EN Parker. The alternative paradigm for magnetospheric physics. *Journal of Geophysical Research: Space Physics*, 101(A5):10587–10625, 1996.

Eugene N Parker. Dynamics of the interplanetary gas and magnetic fields. *The Astrophysical Journal*, 128:664, 1958.

Pascal Pons and Matthieu Latapy. Computing communities in large networks using random walks. In *International symposium on computer and information sciences*, pages 284–293. Springer, 2005.

DH Pontius Jr and RA Wolf. Transient flux tubes in the terrestrial magnetosphere. *Geophysical research letters*, 17(1):49–52, 1990.

Eric Ronald Priest. *Solar magnetohydrodynamics*, volume 21. Springer Science & Business Media, 2012.

JRLM Raeder, RL McPherron, LA Frank, S Kokubun, G Lu, T Mukai, WR Paterson, JB Sigwarth, HJ Singer, and JA Slavin. Global simulation of the geospace environment modeling substorm challenge event. *Journal of Geophysical Research: Space Physics*, 106(A1):381–395, 2001.

Usha Nandini Raghavan, Réka Albert, and Soundar Kumara. Near linear time algorithm to detect community structures in large-scale networks. *Physical review E*, 76(3): 036106, 2007.

Erzsébet Ravasz, Anna Lisa Somera, Dale A Mongru, Zoltán N Oltvai, and A-L Barabási. Hierarchical organization of modularity in metabolic networks. *science*, 297(5586):1551–1555, 2002.

Gregory C Reinsel. *Elements of multivariate time series analysis*. Springer Science & Business Media, 2003.

Patricia Ritter and Hermann Lühr. Near-earth magnetic signature of magnetospheric substorms and an improved substorm current model. In *Annales Geophysicae*, volume 26, pages 2781–2793. Copernicus GmbH, 2008.

Gordon Rostoker. Phenomenology and physics of magnetospheric substorms. *Journal of Geophysical Research: Space Physics*, 101(A6):12955–12973, 1996.

Martin Rosvall and Carl T Bergstrom. Maps of information flow reveal community structure in complex networks. *arXiv preprint physics.soc-ph/0707.0609*, 2007.

A Runov, V Angelopoulos, MI Sitnov, VA Sergeev, J Bonnell, JP McFadden, D Larson, K-H Glassmeier, and U Auster. Themis observations of an earthward-propagating dipolarization front. *Geophysical Research Letters*, 36(14), 2009.

Manfred Scholer and Antonius Otto. Magnetotail reconnection: Current diversion and field-aligned currents. *Geophysical Research Letters*, 18(4):733–736, 1991.

Thomas Schreiber and Andreas Schmitz. Improved surrogate data for nonlinearity tests. *Physical Review Letters*, 77(4):635, 1996.

Philip Sedgwick. Pearson's correlation coefficient. *Bmj*, 345:e4483, 2012.

VA Sergeev, K Liou, PT Newell, S-I Ohtani, MR Hairston, and F Rich. Auroral streamers: Characteristics of associated precipitation, convection and field-aligned currents. In *Annales Geophysicae*, volume 22, pages 537–548, 2004.

VA Sergeev, NA Tsyganenko, MV Smirnov, AV Nikolaev, HJ Singer, and W Baumjohann. Magnetic effects of the substorm current wedge in a "spread-out wire" model and their comparison with ground, geosynchronous, and tail lobe data. *Journal of Geophysical Research: Space Physics*, 116(A7), 2011.

VA Sergeev, AV Nikolaev, NA Tsyganenko, V Angelopoulos, AV Runov, HJ Singer, and J Yang. Testing a two-loop pattern of the substorm current wedge (scw2l). *Journal of Geophysical Research: Space Physics*, 119(2):947–963, 2014.

MA Shea and DF Smart. Space weather and the ground-level solar proton events of the 23rd solar cycle. *Space science reviews*, 171(1-4):161–188, 2012.

Charles W Smith, Jacques L'Heureux, Norman F Ness, Mario H Acuna, Leonard F Burlaga, and John Scheifele. The ace magnetic fields experiment. In *The advanced composition explorer mission*, pages 613–632. Springer, 1998.

Kristian Snekvik, Stein Haaland, Nikolai Østgaard, Hiroshi Hasegawa, R Nakamura, Taku Takada, L Juusola, O Amm, F Pitout, H Rème, et al. Cluster observations of a field aligned current at the dawn flank of a bursty bulk flow. 2007.

Lin Song, Peter Langfelder, and Steve Horvath. Comparison of co-expression measures: mutual information, correlation, and model based indices. *BMC bioinformatics*, 13 (1):328, 2012.

Yan Song and Robert L Lysak. Alfvénon, driven reconnection and the direct generation of the field-aligned current. *Geophysical research letters*, 21(17):1755–1758, 1994.

Kareem A Sorathia, Aleksandr Y Ukhorskiy, Viacheslav G Merkin, Joseph F Fennell, and Seth G Claudepierre. Modeling the depletion and recovery of the outer radiation belt during a geomagnetic storm: combined mhd and test particle simulations. *Journal of Geophysical Research: Space Physics*, 123(7):5590–5609, 2018.

David P Stern. Large-scale electric fields in the earth's magnetosphere. *Reviews of Geophysics*, 15(2):156–194, 1977.

Christopher. Stover. Monotonic function. URL http://mathworld.wolfram.com/MonotonicFunction.html. Visited on 29/08/2020.

WJ Sun, SY Fu, GK Parks, J Liu, ZH Yao, QQ Shi, Q-G Zong, SY Huang, ZY Pu, and T Xiao. Field-aligned currents associated with dipolarization fronts. *Geophysical Research Letters*, 40(17):4503–4508, 2013.

E Tanskanen, TI Pulkkinen, HEJ Koskinen, and JA Slavin. Substorm energy budget during low and high solar activity: 1997 and 1999 compared. *Journal of Geophysical Research: Space Physics*, 107(A6), 2002.

James Theiler, Stephen Eubank, André Longtin, Bryan Galdrikian, and J Doyne Farmer. Testing for nonlinearity in time series: the method of surrogate data. *Physica D: Nonlinear Phenomena*, 58(1-4):77–94, 1992.

E Tindale, SC Chapman, NR Moloney, and Nicholas W Watkins. The dependence of solar wind burst size on burst duration and its invariance across solar cycles 23 and 24. *Journal of Geophysical Research: Space Physics*, 123(9):7196–7210, 2018.

Christopher Torrence and Gilbert P Compo. A practical guide to wavelet analysis. *Bulletin of the American Meteorological society*, 79(1):61–78, 1998.

Aleksandr Y Ukhorskiy, Kareem A Sorathia, Viacheslav G Merkin, Mikhail I Sitnov, Donald G Mitchell, and Malamati Gkioulidou. Ion trapping and acceleration at dipolarization fronts: High-resolution mhd and test-particle simulations. *Journal of Geophysical Research: Space Physics*, 123(7):5580–5589, 2018.

Vytenis M Vasyliūnas. Electric field and plasma flow: What drives what? *Geophysical research letters*, 28(11):2177–2180, 2001.

Vytenis M Vasyliūnas. The physical basis of ionospheric electrodynamics. In *Annales Geophysicae*, volume 30, pages 357–369. Copernicus GmbH, 2012.

Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684):440, 1998.

David F Webb and Timothy A Howard. Coronal mass ejections: observations. *Living Reviews in Solar Physics*, 9(1):3, 2012.

Jian Yang, Frank R Toffoletto, Richard A Wolf, and Stanislav Sazykin. On the contribution of plasma sheet bubbles to the storm time ring current. *Journal of Geophysical Research: Space Physics*, 120(9):7416–7432, 2015.

E Zesta, L Lyons, C-P Wang, E Donovan, H Frey, and T Nagai. Auroral poleward boundary intensifications (pbis): Their two-dimensional structure and associated dynamics in the plasma sheet. *Journal of Geophysical Research: Space Physics*, 111 (A5), 2006.