Implementing a Magnetic Charge Topology Model for Solar Active Regions

G. Barnes

Colorado Research Associates Division
NorthWest Research Associates, Inc.
3380 Mitchell Lane, Boulder, CO 80301

`graham@cora.nwra.com`

and

D.W. Longcope

Department of Physics,
Montana State University
Bozeman MT 59717

`dana@physics.montana.edu`

and

K.D. Leka

Colorado Research Associates Division
NorthWest Research Associates, Inc.
3380 Mitchell Lane, Boulder, CO 80301

`leka@cora.nwra.com`

ABSTRACT

Information about the magnetic topology of the solar corona is crucial to understanding solar energetic events. One approach to characterizing the topology which has had some success is Magnetic Charge Topology, in which the topology is defined by partitioning the observed photospheric field into a set of discrete sources and determining which pairs are interlinked by coronal field lines. The level of topological activity is then quantified through the transfer of flux between regions of differing field line connectivity. We discuss in detail how to implement such a model for a time series of vector magnetograms, paying particular attention to distinguishing real evolution of the photospheric magnetic flux from changes due to variations in atmospheric seeing, as well as uncorrelated noise. We determine the reliability of our method and estimate the uncertainties in its results. We then demonstrate it through an application to NOAA active region 8210, which has been the subject of extensive previous study.

Subject headings: methods: numerical — Sun: corona — Sun: magnetic fields — Sun: photosphere

1. Introduction

It is now generally believed that solar flares occur through the release of energy stored in the coronal magnetic field. There remains, however, considerable debate over what initiates this release and which factors make some magnetic configurations more likely to flare than others. Theoretical considerations indicate that magnetic reconnection may play a role in either the initiation of a flare or in the energy release. From this hypothesis it would seem that magnetic fields most likely to reconnect would be those most likely to flare. A recent class of reconnection models, called *Magnetic Charge Topology* models, are versatile enough to apply to realistic coronal geometries and thereby test this prediction.

Magnetic charge topology (MCT) models assume that the photospheric magnetic field consists of distinct sources or "charges" (Baum & Bratenahl 1980; Gorbachev & Somov 1988; Priest & Forbes 1989; Lau 1993; Démoulin et al. 1994; Parnell et al. 1994), presumably the manifestations of sub-photospheric flux tubes. Classifying field lines by the sources at their footpoints divides the corona into distinct flux domains separated by *separatrix surfaces*. In this simplified context, magnetic reconnection exchanges the footpoints of a pair of field lines thus transporting flux from one domain to another. Simple geometrical considerations show that this transfer must occur across a *separator* field line, a line of intersection between two separatrices (Baum & Bratenahl 1980; Greene 1988; Gorbachev & Somov 1988; Lau & Finn 1990). In light of this model it would seem that magnetic configurations with more separators would have more opportunity to reconnect and would thus be more likely to produce flares.

In order to apply MCT models to observations, photospheric fields must be somehow represented as a set of distinct sources. This *partitioning* is motivated by the well-known tendency of photospheric flux to concentrate, albeit not always into completely distinct regions. Several investigators have proposed different methods of partitioning and applied these to specific flares (Mandrini et al. 1991; Démoulin et al. 1993; Longcope & Silva 1998). Each of the methods can be classified into one of two basic categories which aim to accomplish different things. Models of the first category place point sources beneath the photospheric surface

($z < 0$) and extrapolate upward to produce the photospheric field. The photospheric image of a single submerged source is a two-dimensional Lorentzian whose width is proportion to the source's depth; the composition of many sources yields a relatively realistic approximation of an observed photospheric field. The magnitude and location of the submerged poles can be varied in order to match a magnetogram pixel-for-pixel, often in a least-squares sense. This partitioning scheme has been applied to numerous flare observations (Gorbachev & Somov 1989; Mandrini et al. 1991; Démoulin et al. 1993; van Driel-Gesztelyi et al. 1994; Bagalá et al. 1995) and good geometrical agreement was found between flare-related signatures, such as Hα or hard X-ray emission, and the locations of separatrices and separators in the model field.

Submerged poles are good at representing the photospheric field and thereby faithfully reproducing the geometry of flare signatures. They are, however, poorly suited to quantitative tasks such as determining the flux connecting a pair of sources, or how that flux changes through reconnection. Of all the flux originating at a given submerged source only a fraction is "real" in the sense that it crosses the photospheric plane into the corona. Moreover, the real fraction depends on the distribution of sources and will change as sources move. Most changes in domain flux are, therefore, artificial and not the result of coronal reconnection. A related philosophical drawback is that separatrices and separators in such a model are defined by connections below the photosphere and not by properties in the "real" portions of field lines.

Models from the second category of partitioning use point sources located *on* the photospheric surface, $z = 0$. Since exactly half of the flux from a source enters the corona (the other half entering a mirror corona below the photosphere), such a model is well suited to quantifying the interconnections between sources and possible changes in these interconnections: reconnection. Each photospheric point source matches the leading terms in a multi-pole expansion of the field from the actual (smooth) flux concentration. As such it is reasonably accurate far from the concentration but not at the photosphere where the model field is singular (such singularities are inevitable in multi-pole expansions). Rather than approximating the magnetogram, as submerged poles do, partitioning of this type seeks to identify distinct regions corresponding to the hypothetical flux concentrations.

The examples of magnetogram partitioning found in the literature rely heavily on subjective analysis and are therefore difficult to reproduce, and unsuitable for studying large numbers of active regions. This is particularly true of those from the second, quantitative, category, since they must isolate distinct concentrations (Longcope & Silva 1998). We present here an automatic algorithm for partitioning magnetograms into source regions, with particular emphasis on distinguishing real changes in the photosphere over the course of a time-series of vector magnetogram data from changes due to the effects of noise and atmospheric seeing variations.

Once the partitioning is done, the topological properties of the field can be explored. An extrapolation readily defines the connectivity of the coronal magnetic field. Locating and classifying magnetic null points provides information on the location of separatrix surfaces, and the number of separators. In addition, estimates of the horizontal velocities can be obtained from the motions of each partition.

In this paper, we apply the MCT model to active region NOAA AR 8210 for which on 1998 May 01, a good-quality time series was obtained by the Imaging Vector Magnetograph (Mickey et al. 1996; LaBonte et al. 1999). With a cadence of approximately four minutes, the vector magnetic field was obtained over a $3'$ field of view covering the entire active region, and sampled with $1.1''$ pixels. The 180° ambiguity in the observed transverse component was resolved such that heliographic B_x, B_y, B_z components are used; an additional step was performed to ensure that the resulting ambiguity resolution was least variable over the time series (see Leka & Barnes 2003a, for details). Of note for this analysis, the temporal sequence of vector magnetic field maps are co-aligned on a sub-pixel grid in the image-plane. This is a well-studied region, as it was both the source of a Solar, Heliospheric and Interplanetary Environment (SHINE) Campaign Event[1], and a focus of the Solar MURI project[2], so we can compare some of our findings (e.g., horizontal velocities) to the results of other methods.

2. Partitioning the Photospheric Field

The goal of partitioning is to identify concentrations of photospheric flux which serve as "distinct" anchoring points for the coronal field lines. The vertical component of the photospheric field, $B_z(x, y)$ (distinguished here from the component along the line of sight), defines the distribution of flux to which the coronal field is anchored. It is therefore this function which must be partitioned into *source regions*. Were it actually the case that $B_z(x, y)$ consisted of isolated, unipolar patches surrounded by field-free ($B_z = 0$) photosphere, then partitioning would be quite straightforward (Schrijver et al. 1997b; Kankelborg & Longcope 1999): each isolated patch is a source region. Active region fields, however, do not conform to this simple picture. For instance, typical plage can be considered to be either a single extended source region or countless tiny source regions representing individual field-line foot-

---

[1] `http://www.shinegroup.org/campaign.html`

[2] `http://sprg.ssl.berkeley.edu/home/yanli/public_html/muri/events98.html`

points. Neither extreme gives any practical information about the prospects for magnetic reconnection in the corona. We have therefore developed an algorithm to identify a tractable number of *flux concentrations* which will be taken as source regions in our MCT analysis.

Following the methodology used successfully in the clear-cut case of quiet Sun fields, we identify concentrations as local maxima in the unsigned vertical flux density, $|B_z(x,y)|$. Unfortunately, the vertical field from a magnetogram contains a great deal of fine structure giving rise to many local maxima. Some of this structure is genuine, but some also results from uncorrelated noise. Even more challenging are the variations in the structure caused by changes in the atmospheric seeing. These effects are correlated, and can result in systematic changes in the vertical field which are not genuine. Since one of our primary interests is to consider the time evolution of active regions, it is particularly critical to minimize the impact of the changes due to both random noise and atmospheric seeing. To accomplish this, we construct a "reference" magnetogram for a time-series, apply a partitioning algorithm to it, and partition all the magnetograms in the time series based on this reference partition. This approach greatly reduces the number of random fluctuations in the partitioning between time-steps, and so allows real evolution in the magnetic topology to be identified.

2.1. Constructing a Reference Partition

The time series in which we are interested have a typical duration of at most a few hours, thus we expect that real changes in the magnetic flux concentrations will be of comparatively small amplitude on this timescale. Thus we construct the reference magnetogram by taking the time-averaged value of the vertical field at each (co-aligned) pixel. The next step is to discard those areas of the magnetograms in which the smoothed field is consistent with zero. The uncertainty in the vertical field strength measurement is

$$\sigma^2_{\text{detection}} = (\delta B_\perp)^2 \sin^2\theta + (\delta B_\parallel)^2 \cos^2\theta \qquad (1)$$

where θ is the observing angle, and the intrinsic uncertainties $\delta B_\perp$ and $\delta B_\parallel$ are determined in separate measurements by fitting the distribution functions of quiet Sun areas (Leka & Skumanich 1999). Experience has shown this uncertainty to vary little over a time series for a given active region, although there can be considerable differences among different active regions. We therefore set a threshold field strength of three times the mean detection threshold ($3\bar{\sigma}_{\text{detection}}$) which is typically in the range 50 G to 100 G. Pixels in the reference magnetogram where the vertical field falls below this threshold are excluded from consideration by setting their vertical field strength to zero.

Source regions are now defined by propagating *region labels* downhill from local maxima in $|B_z|$. Working in descending order of $|B_z|$ each pixel inherits the label of the greatest among its eight neighbors; if none are yet labeled then the central pixel is a local maximum and gets a “fresh” label. This type of *gradient based tessellation* scheme has been shown to label each and every pixel uniquely (Schrijver et al. 1997a); all pixels sharing a common label represent a single source region.

The downhill gradient algorithm can, however, produce an unwieldy number of very small regions in finely structured fields, such as plage. In any event, it is unlikely that structuring on scales of $\sim$ 1000 km (a few pixels) can significantly affect the large-scale coronal field. To limit the number of small regions, those with flux less than 1×10^4 G Mm^2 are assigned the same label as a neighbor of the same polarity, provided the neighbor has a common boundary with the small region. Our motivation in merging partitions in this way is to simplify areas where small concentrations of flux are unlikely to contribute to the global topology in an important way. To further simplify regions of plage, we also perform a “saddle point” merging. That is, we determine the value of the vertical field at the saddle point between each pair of neighboring regions. If the field there is within 100 G of the field strength at either of the local maxima defining the neighboring regions, the regions are merged. Once again, the goal here is to identify concentrations of flux, and this technique groups together similar local maxima into a single concentration. These methods are designed to retain regions with polarity opposite to their surroundings, such as δ-spots. The thresholds used were tested on a sample of six active regions in addition to the results presented here.

The saddle point merging concludes the construction of the reference partition. Figure 1 shows the result for NOAA AR 8210 for the time interval from 17:07 UT to 18:07 UT, which has been partitioned into 52 sources. Where the flux is large, such as in the large negative polarity sunspot, all the divisions between flux concentrations have been kept. Likewise, the small positive polarity intrusion located at $(x, y) \approx (-35, -245)$ has been retained. However, the large areas of plage have seen significant merging.

2.2. Matching the Reference Magnetogram

For each magnetogram in a time-series, the partitioning algorithm initially proceeds as in the reference case, by removing from consideration those portions of the magnetogram consistent with zero field, then employing the gradient based tessellation scheme to assign a label to each pixel. This typically results in an extremely large number of distinct labels, due to the amount of random noise in the field. Attempts to reduce the number of distinct labels by smoothing the field, or merging neighboring partitions based on some criteria (flux,

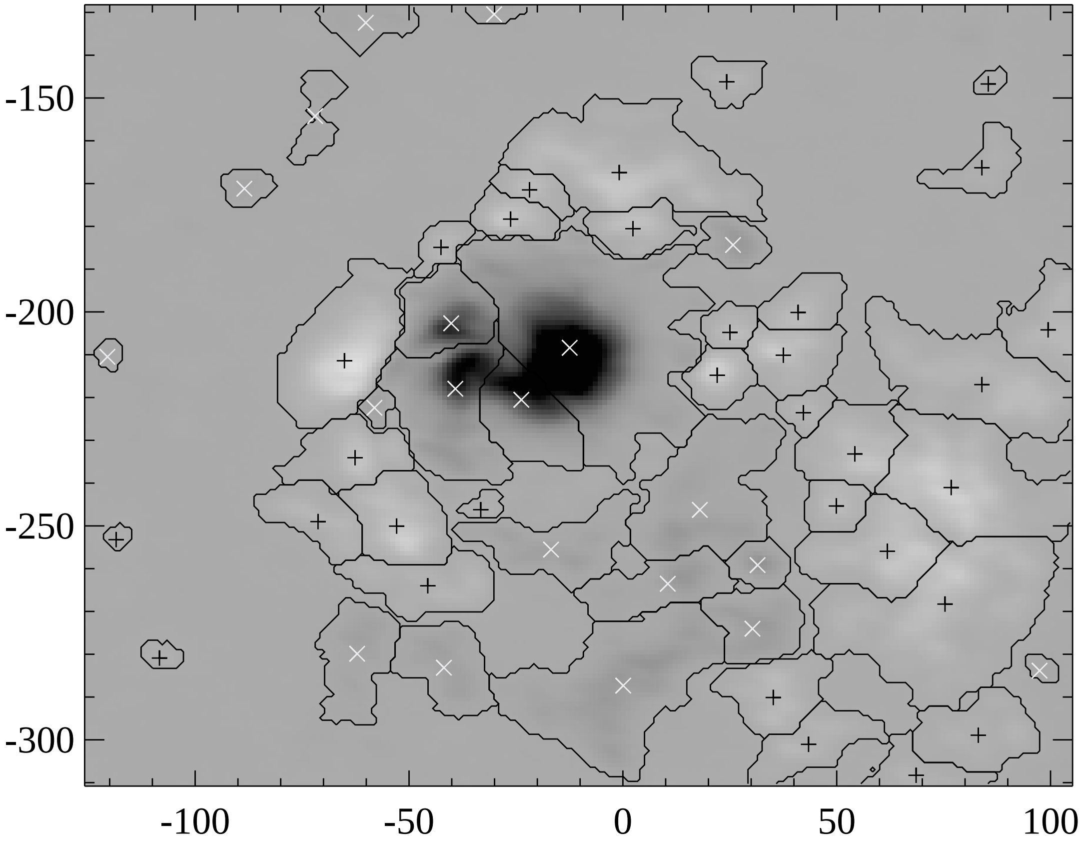

Fig. 1.— Reference magnetogram for NOAA AR 8210, showing the time averaged vertical field, the 52 partitions and the pole locations. Note that our partitioning algorithm retains the divisions between strong flux concentrations in the large negative polarity sunspot, and the small positive polarity intrusion located at $(x, y) \approx (-35, -245)$ while significantly simplifying the large areas of plage. Axes are in Mm in the image plane.

saddle-point, etc.) typically result in a partition which varies greatly from one timestep to the next. Instead, we compare each partition in the magnetogram to the reference magnetogram, and assign to it the label of the reference partition with which it has the greatest overlap in flux. That is, we minimize the difference between the partitioning of each magnetogram in a time series, and the reference magnetogram for that time series.

In Figure 2 we show the partitions of the first (17:07 UT) and last (18:07 UT) magnetograms in our time series for AR 8210. The exact boundary of each partition has changed slightly, but all the partitions are still present, and occupy similar areas of the photosphere. Careful inspection reveals some real changes, such as the motion of the partition at $(x, y) \approx (20, -185)$, labeled `N13`.

3. The Magnetic Topology

According to MCT models, coronal field lines form distinct domains grouped by the source regions of their footpoints. In order to enumerate these domains and describe their topology, we employ a potential-field extrapolation of the photospheric field. The horizontal components of the vector field do typically indicate vertical photospheric currents and non-negligible field twist on both localized and global scales, as demonstrated in Leka & Barnes (2003a). We nevertheless use a potential field for two reasons. First, the topological quantities in question, such as the connectivities, are robust against the existence of modest currents. Secondly, alternatives to potential extrapolations introduce so many practical difficulties that their possible improvements do not seem worthwhile. Constant-α fields, for example, require the introduction of artificial horizontal boundaries since unbounded fields of this type behave unphysically (Nakagawa & Raadu 1972). They can also contain field lines not anchored to the photosphere at all, which cannot therefore be classified in an MCT model. Field lines of this type have never been observed in an active region's corona and so a model which produces them is undesirable. Finally, the actual distribution of vertical currents are rarely, if ever, observed to resemble a constant-α field. Employing a non-linear non-constant-α extrapolation method is beyond the scope of this work. Thus, we use the potential-field extrapolation as a starting point for describing the resultant coronal complexity. That is, the complexity computed using the potential field represents *at least* the complexity in the true corona above these photospheric measurements.

We make the further approximation of using only the two leading order multipoles, the monopole and dipole, to calculate the contribution from each region. This is equivalent to replacing each source region by a magnetic point charge located at its flux-weighted center (see Fig. 1). The flux of the charge is the total flux of all pixels within the partition. The

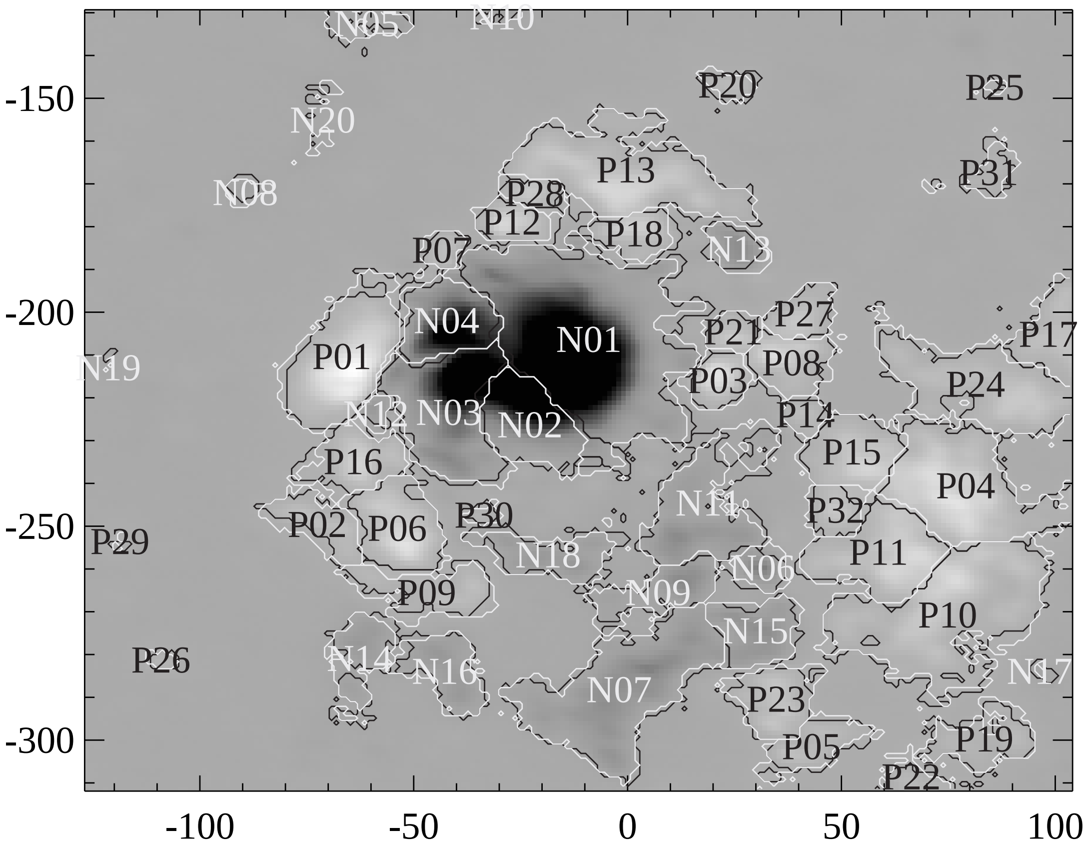

Fig. 2.— Two partitions from the same time series for NOAA AR 8210, at 17:07 UT (black) and 18:07 UT (white) superimposed on the time averaged vertical field. Note that, while the exact boundaries of the partitions have moved somewhat, every partition is still present. Some systematic changes are evident, as in the motion of `N13` in the northwest quadrant, which is likely due to real motion of that flux concentration.

principal advantage of this approximation is that the magnetic field from a collection of coplanar point sources can be rapidly calculated.

3.1. Mapping the Connectivity

The field's connectivity is quantified by the domain matrix, ψ_{ij}, giving the net flux connecting source i to source j. From the approximate coronal field described above we estimate the domain matrix using a Monte-Carlo technique. N_i field lines are initiated in random directions from source i and followed to their opposite end. The number of these field lines that end at opposing source j will be denoted m_{ij}. Since the field lines are initialized at random m_{ij} will be a Poisson random variable with mean $N_i\psi_{ij}/\Phi_i$ (Φ_i is the flux of source i). Similarly, of the N_j field lines initialized at source j, m_{ji} will terminate at i. The probability that the independent Poisson variables, m_{ij} and m_{ji}, will assume a particular pair of values, for a specified domain flux, is the product of their individual probabilities

$$p(m_{ij}, m_{ji}|\psi_{ij}) \;=\; \frac{1}{m_{ij}!m_{ji}!}\frac{\psi_{ij}^{m_{ij}+m_{ji}}}{(\Phi_i/N_i)^{m_{ij}}(\Phi_j/N_j)^{m_{ji}}}\,e^{-\psi_{ij}/\psi_0}. \tag{2}$$

where $\psi_0 = 1/(N_i/\Phi_i + N_j/\Phi_j)$.

Using Bayes' theorem, and following the common practice of assuming the prior distribution of ψ_{ij} to be uniform (D'Agostini 2003), the conditional probability (2) can be converted to the probability of ψ_{ij} for a given pair of observations. After normalization this probability is

$$p(\psi_{ij}|m_{ij}, m_{ji}) \;=\; \frac{1}{(m_{ij}+m_{ji})!}\frac{\psi_{ij}^{m_{ij}+m_{ji}}}{\psi_0^{m_{ij}+m_{ji}+1}}e^{-\psi_{ij}/\psi_0}. \tag{3}$$

The peak of this probability curve yields the maximum likelihood estimate for the domain flux

$$\psi_{ij} \;=\; \frac{m_{ij}+m_{ji}}{(N_i/\Phi_i + N_j/\Phi_j)}. \tag{4}$$

This expression also applies to cases where one of the domains, say j, is ∞ from which no field lines are initialized: $N_j = m_{ji} = 0$.

When two sources are found to be unconnected, $m_{ij} = m_{ji} = 0$, we take the domain to be absent altogether: $\psi_{ij} = 0$. An extant domain with flux ψ_{ij} can, however, occasionally be erroneously overlooked by our Monte Carlo method. This will happen with a probability

$p(0,0|\psi_{ij})$ given by equation (2), which we express as the probability of missing an extant domain

$$P_{\rm miss}(\psi_{ij}) \;=\; e^{-\psi_{ij}/\psi_0}. \tag{5}$$

Therefore we can expect to find at least 95% of all domains with flux $\psi_{ij} \geq 3\psi_0$.

The threshold flux, $3\psi_0$, can vary between source pairs since ψ_0 depends on the flux of, and the number of lines initiated at, each source. To obtain a single threshold, $\psi_c \equiv 3\psi_0$, for all source pairs, we initiate a number of field lines N_i proportional to the flux of that source: $N_i = 3\Phi_i/2\psi_c$ (actually the next largest integer). With this choice we can regard ψ_c as the nominal threshold for 95% yield. Domains connecting to infinity, however, have $\psi_0 = \Phi_i/N_i = (2/3)\psi_c$; these will be found with 95% only when $\psi_{ij} > 2\psi_c$.

In practice, we are able to select ψ_c sufficiently small that the statistical uncertainty in determining the domain fluxes is typically small compared with that due to the uncertainty in the locations and magnitudes of the sources (see §4). For the calculations presented here, $\psi_c = 0.15\,{\rm G\,Mm}^2$. This resolution limit dictates that approximately 2.5×10^7 field lines be integrated; $N_i = 1896$ are initiated at the smallest source `N19`.

Figure 3 shows the domain matrix elements larger than ψ_c as lines whose thickness reflects ψ_{ij} in our model of AR 8210. As expected, the large central sunspot has many connections to the surrounding plage. However, it is interesting to note that even tiny, isolated sources, such as source `P26` in the southeast corner of the field of view, sometimes have multiple connections. Because the field of view is not flux balanced, there are also a number of connections to "infinity", that is, to flux concentrations outside of the field of view.

Figure 4 summarizes the sizes of the 99 different connections found by the Monte Carlo integration. At least in this case, the size distribution is notably bimodal. There is a population of large domains, $\psi_{ij} > 70\,{\rm G\,Mm}^2$, comprising 99.98% of all flux. The smaller population consists of 17 domains with fluxes ranging from $55\,{\rm G\,Mm}^2$ down to $\psi_{ij} = (2/3)\psi_c$; this connection is so small that it would be missed 14% of the time. We discuss below the evidence that at least two more domains this small or smaller have avoided detection.

Had the Monte Carlo calculation been done with fewer field lines, by setting ψ_c higher, an even larger portion of this small population would have been missed. The population of 82 larger domains would still be found, with 99% certainty, as long as $\psi_c < 45\,{\rm G\,Mm}^2$, which would require the integration of only 1.5×10^5 field lines. In as much as these are the field lines likely to play a significant role in solar activity, it is reasonable to use such a high threshold and thereby disregard the smaller population.

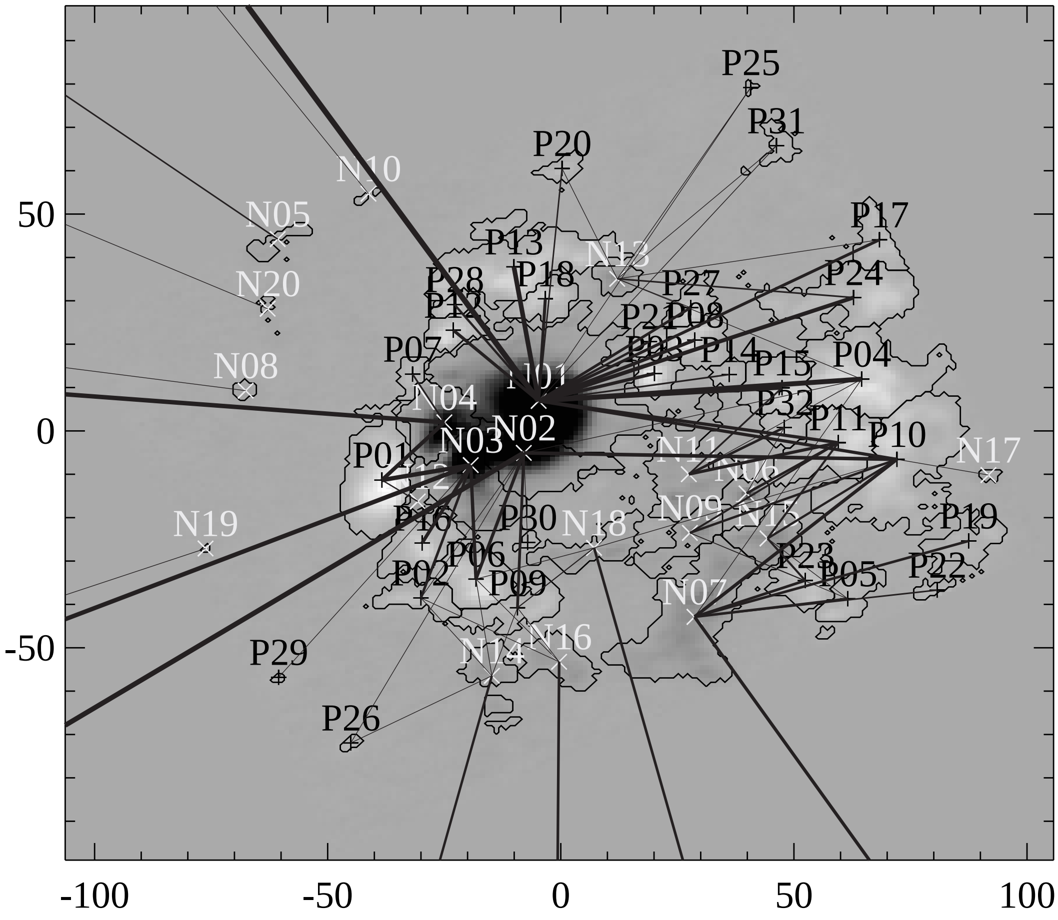

Fig. 3.— Magnetogram of NOAA AR 8210, taken at 17:07 UT, showing in heliographic coordinates the vertical field, the pole locations, the boundaries of the partitions defining the pole locations and the connections between sources (black lines). The thickness of a line indicates the amount of flux in the domain. Lines leading off the edge of the figure are connections to “infinity”, that is, to flux concentrations outside of the field of view.

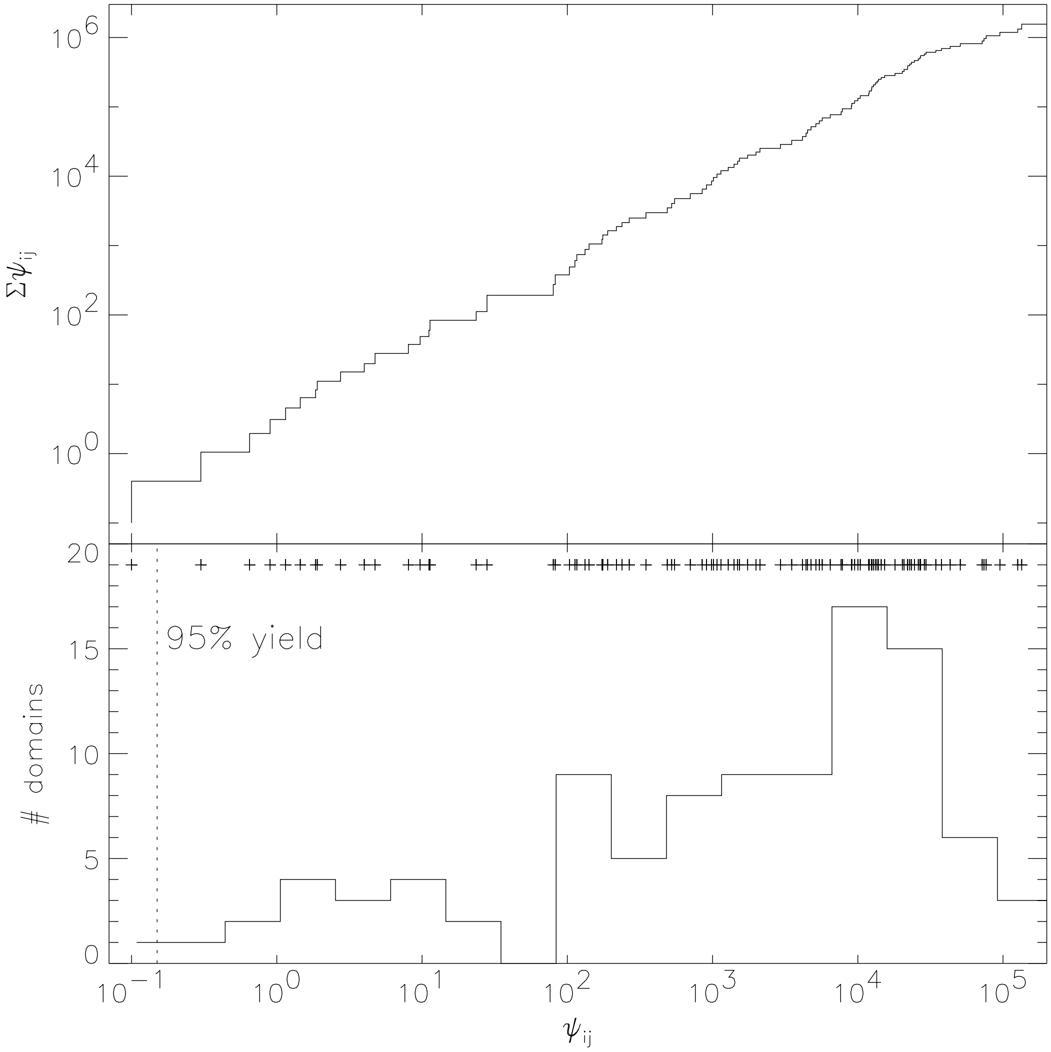

Fig. 4.— A summary of the sizes of the 99 connectivities found by Monte Carlo integration. The bottom panel shows a histogram of the regions with ψ_{ij} in units of G Mm^2 on a logarithmic scale. The dotted vertical line shows $\psi_c = 0.15\,\mathrm{G\,Mm}^2$, which is the limit for a 95% probability of detection. Pluses along the top show the actual values. The top panel shows the cumulative flux in the connectivities. It reaches its maximum, $1.5 \times 10^6\,\mathrm{G\,Mm}^2$ at the right.

Choosing the number of field lines to be proportional to the source flux, as we have done, results in an approximately constant absolute statistical error in the domain fluxes. For questions which make use of the domain fluxes, and in particular for questions which deal with *changes* in the domain fluxes, such as estimates of the rate of reconnection, it is more important to have an upper limit on the absolute error in the domain fluxes. If the goal is to fully determine the magnetic topology by enumerating *all* of the domains, then it will generally be more efficient to assign a fixed number of field lines to every source, irrespective of its flux. This will resolve domains accounting for more than a fixed fraction of their smallest footpoint. Such an approach can lead to an estimate of the number of separator field lines present (see §3.2).

3.2. Magnetic Null Points, Separatrices and Separators

In addition to calculating the domain matrix, we locate magnetic null points in the model magnetic field. As the name implies, a null point $\boldsymbol{x}_a$, is a location at which the magnetic field vector vanishes, $\boldsymbol{B}(\boldsymbol{x}_a) = \mathbf{0}$. In the vicinity of $\boldsymbol{x}_a$, the magnetic field can be approximated by

$$\boldsymbol{B}(\boldsymbol{x}_a + \delta\boldsymbol{x}) \approx \mathrm{M}^a \cdot \delta\boldsymbol{x} \tag{6}$$

where $M^a_{ij} = \partial B_i / \partial x_j$ is the Jacobian matrix evaluated at $\boldsymbol{x}_a$. An A-type (B-type) null is one for which $\det \mathbf{M}^a$ is positive (negative). Of the eigenvectors of $\mathbf{M}^a$, the two with like-signed eigenvalues define a plane called the fan plane, and the third defines the spine direction. Field lines originating in the fan plane extend into the corona forming a fan surface, which is one of the field's separatrices. Field lines leaving parallel and anti-parallel to the spine direction are spine field lines. The intersection of a fan surface with the photospheric surface ($z = 0$) is called the *fan trace*, and mapping them provides valuable insight into the structure of the coronal domains (Longcope & Klapper 2002).

A photospheric null of either type is classified as a *prone null* if its spine field lines lie in the photosphere, and an *upright null* otherwise. A separator is a field line which originates and terminates on a null; it is the intersection of two separatrix surfaces, which form the boundaries of flux domains (see Fig. 5).

We locate null points using a Newton-Raphson root-finding algorithm. Most of the nulls lie in the photospheric plane so this is the only place we seek them. To initiate the Newton-Raphson method, we use the location of the null point associated with each pair of sources considered in isolation, as well as the midpoint of all triplets of sources of the same polarity. Initial locations of the first kind typically lead to prone nulls, while initial locations

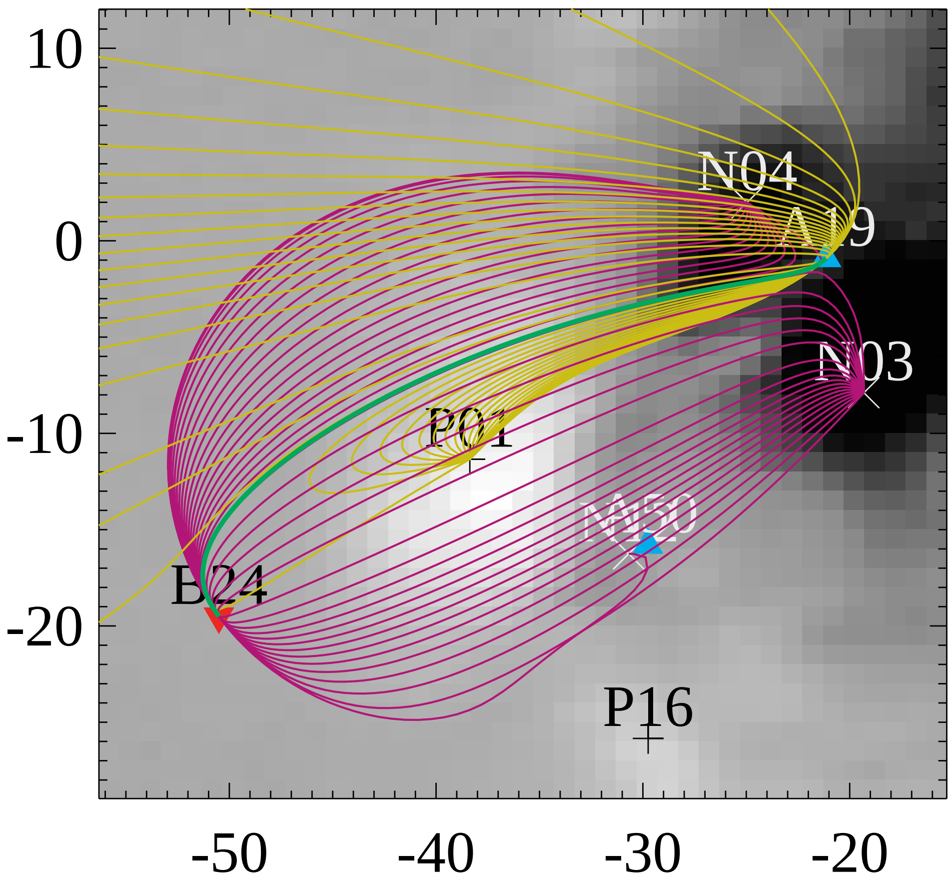

Fig. 5.— Field lines in the separatrix surfaces associated with null points `A19` (yellow) and `B24` (purple), along with the corresponding separator (green). Note how purple field lines originating in the fan of `B24` terminate on either source `N03` or `N04`. Null `A19` has spine field lines which terminate on the same pair of sources, so these nulls are linked. The fan field lines for `A19` terminate at the same sources as the spine field lines for `B24`, thus there is a candidate for a separator connecting these two nulls.

of the second kind typically lead to upright nulls. Figure 6 shows the nulls and separators, spines, and fan traces from the model field of AR 8210.

Our algorithm located $N_p = 51$ prone nulls, $N_A = 19$ of type-A and $N_B = 32$ of type-B, and $N_u = 0$ upright nulls. These values satisfy both Poincaré indices (Longcope & Klapper 2002)

$$N_p = N_+ + N_- + N_u - 2, \quad \text{and} \quad N_B - N_A = N_+ - N_-, \tag{7}$$

where $N_+ = 33$ and $N_- = 20$ are the numbers of positive and negative sources; since the model has negative net flux, infinity has been counted as a positive source. If these relations were not both satisfied we would know that some nulls had not been found.

The first step in locating the separators is to determine which nulls are linked. With the exception of separator field lines, each field line originating in a null's fan surface terminates on a source. We begin by following one of the fan traces until it terminates on a source, then initiate another field line in the fan surface in a direction at an angle θ from the fan trace. If that field line terminates on the same source as the fan trace, we increment the angle determining the initial direction by another θ and trace that field line until it terminates on a source. We continue stepping in angle through the fan surface until a field line terminates on a different source (or the other fan trace is reached).

When the two adjacent field lines terminate on different sources, there must be at least one separator field line between the two (see Fig. 5). This initial pair of angles bracketing a separator is refined to more closely bracket the separator by initiating a new field line at an angle midway between the bracketing angles. If the field line from the new direction terminates on the same source as one of the original bracketing field lines, its initial angle replaces that bracketing angle. The process of bisecting the angle is repeated until the initial direction of the separator is constrained to lie within a range of less than θ/n, where n is the number of bisections. However, it is frequently the case that more than one separator is contained in the initial bracketing range. Thus, the field line originating from the bisection of the angle may terminate on a third source, which does not match either of the bracketing field lines. In this case, it replaces the larger of the bracketing angles, and the bisection continues until *a* separator has been located. When the search resumes, the latest refinement of bracketing angle is incremented by the standard angle θ, and the procedure continues, thus separators which differ in initial angle by less than θ are generally located, although those which differ by less than θ/n are not.

To identify the other end of the separator, it is necessary to look at the sources associated with the spine field lines of the nulls, as well as the sources associated with the bracketing field lines in the fan surface. A pair of nulls is considered to be linked if the sources associated with

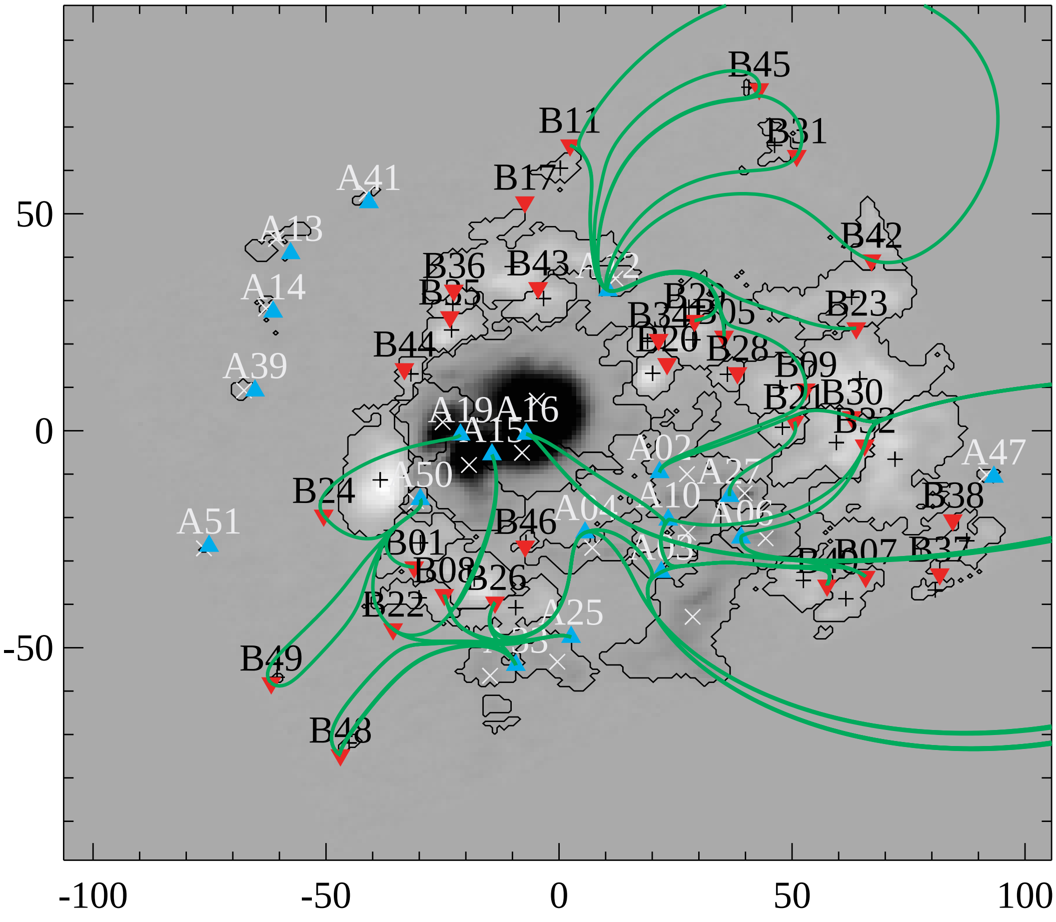

Fig. 6.— Magnetogram of NOAA AR 8210 at 17:07 UT, showing in heliographic coordinates the vertical field, the locations of the poles and nulls (blue for A-type, red for B-type), and the projection of the separators onto the photosphere (green lines).

the spines of the first null match the sources associated with the bracketing field lines of the second null, and vice versa. The process of linking nulls is straightforward, if computationally intensive when a sufficiently large number of field lines to bracket all separators are originated in each fan surface.

Once a pair of linked nulls has been identified, the existence of the separator must be confirmed. The process we employ makes use of a plane, located midway between the linked nulls, with a normal parallel to a line connecting the two nulls. The field lines bracketing the separator in the fan surface of the first null are traced to their intersections with the plane, say at positions $\boldsymbol{x}_1$, $\boldsymbol{x}_2$. The bracketing field lines from the second null are traced to their intersections with the plane, say at $\boldsymbol{y}_1$, $\boldsymbol{y}_2$. The method is deemed to have converged to the separator if the segment connecting $\boldsymbol{x}_1$ with $\boldsymbol{x}_2$ intersects the segment connecting $\boldsymbol{y}_1$ with $\boldsymbol{y}_2$. If the segments do not intersect, then the bracketing field lines are revised, and the segments are reconstructed until they intersect or a maximum number of iterations is reached, at which point the search for the separator is abandoned.

The largest limitation to the method is that the bracketing field lines may not in fact intersect the plane. Because field lines rapidly diverge from the separator, it is sometimes necessary to bracket the separator with extremely high precision in order to have a chance for the method to work. One could also apply the method with different choices of the plane, but this typically requires human intervention, which is not consistent with our goal of automating the implementation of the MCT model. Thus, the separator associated with linked pairs of nulls is not always located.

For the results presented here, we use an angle $\theta = \pi/5000$, and bisect the resulting range $n = 25$ times, so that a separator is typically determined to lie within a range in angle of 1.9×10^{-12}. With this precision, every A-type null linked to a B-type null also has the B-type null linked to the A-type null, and the associated separator was found, for a total of 59 separators. Of course, that is not a guarantee that no other links or separators are present.

The corona is divided into domains by the fan surfaces, which intersect each other along separators. The number of domains is related to the number of surfaces, edges (separators and spines), and vertices (sources and prone nulls) through a topological formula called *Euler's relation* (Longcope & Klapper 2002). A field with X separators and N_c null points in the corona, which is anchored to S different sources (including a source at infinity if necessary) will be divided into exactly

$$D \;=\; X + S - N_c - 1 \tag{8}$$

distinct domains (Beveridge & Longcope 2005).

In our model of AR 8210, with $S = N_+ + N_- = 53$ sources (including infinity), we found $X = 59$ separators and $N_c = 0$ coronal null points. Topology dictates that there be $D = 111$ domains in such a field, according to equation (8). Since our algorithm can miss separators the number we found, $X = 59$, should itself be taken as a lower bound and the number of domains would be larger by the number of missed separators. It is unlikely that a coronal null has gone unnoticed, since both Poincare relations are satisfied already.

Each separator lies at the interface of 4 different domains connecting each of its two positive spine sources to each of its two negative spine sources (Longcope et al. 2001). Counting up all the connections implied by all 59 of the separators yields a total of 101. These include all 99 connections present in the domain matrix (see Fig. 4) and two others which were apparently missed by the Monte Carlo integrator (specifically `P05`–`N01` and `P23`–`N01`). Had connections been found by the Monte Carlo integrator which were *not* implied by the spine sources of our separators we would have concluded that some separators were missed. As it stands we have no evidence of missing separators.

It is possible for multiple domains to link the same pair of sources. While these constitute distinct domains, and are therefore counted separately in equation (8), the domain matrix will include only a single non-zero entry between the sources. Thus we can expect a discrepancy between the number of connections, 101 in the case, and the number of domains, $D = 111$, when these are present.

Redundant domains of this type are, in fact, present in our model field. Sufficiently close to the source, the field can be considered to be radial, so each field line can be associated with a unique direction specified by the latitude and longitude given in a spherical coordinate system centered on the source. The number of domains present can be determined by considering the set of angles from which field lines are originated on one source and terminate on the second source. If this set is connected, then there is a single domain; if the set is expressed as the disjoint union of maximal connected subsets, then the number of domains is equal to the number of subsets.

Figure 7a shows the sets of angles for source `N13`, with each set shown in a different color. The boundary between each set and its neighbor is a separatrix surface, and the point at approximately 7 o'clock where the separatrices come together on the edge of the figure is a spine field line for null `A12`. There are three distinct domains (shown in yellow) connecting the positive source `P17` to the negative source `N13`. Separating these from each other are the domains `P25`–`N13` (red) and `P31`–`N13` (gray), which are completely enclosed by separatrices from null points `B45` and `B31` respectively. There are several extremely small domains which are nearly invisible in the figure (e.g., `P15`–`N13`, with a flux of about 1.5 G Mm^2, and `P08`–`N13` with a flux of about 11.4 G Mm^2, which both lie between `P27`–`N13` [orange] and `P04`–`N13` [light blue]).

The existence of the multiple domains can also be inferred by considering the photospheric footprints of the domains (Fig. 7b). All of the domains `P17`–N13 are surrounded by the separatrix from null `B42`, but their footprints are separated from each other by the footprints of the domains `P25`–`N13` and `P31`–`N13`, and the spines of nulls `B45` and `B31`. The separatrix from each of these null points intersects the separatrix from `A12` *twice*. Such double intersections create a looped pair of separators, visible in the upper right of Figure 6. We believe that such looped separators will always cause one domain to be split into two components. These two loops along with a third between `B42`–`A13` cause the source `N13` to be connected to nine different sources through *twelve* different domains; three separator loops lead to three redundant domains.

Among the 59 separators there are *ten* pairs of looped separators, including the three just discussed. It is our belief that each of these must be accompanied by a redundant domain. That would bring the number of domains to 111 in full agreement with topological constraint given by equation (8). Therefore it would seem that we have found every separator and every domain, and that our domain matrix ψ_{ij} is missing two very small entries.

4. Uncertainties for the MCT Model

The foregoing describes an implementation of the MCT at a single time. As the source positions and fluxes change so will the the domain matrix of the potential field ψ_{ij}. It is possible that the the topology will change as well, however, this can happen only as distinct events called bifurcations (Brown & Priest 2001; Maclean et al. 2005). Without reconnection the field must depart from a potential field in order to hold fixed its domain matrix. The Minimum Current Corona Model (MCC, Longcope 1996; Longcope & Klapper 2002) is a self-consistent model for such a departure from a potential field. In the MCC, the field evolves through a series of equilibrium each of which minimizes the magnetic energy subject to the fixed domain matrix. These equilibria are potential except for current sheets along each separator. The current on each separator depends on changes in all elements of the domain matrix of the potential field, $\Delta\psi_{ij}$.

In order to implement the MCC, it is important to know the accuracy and reliability of the topology and domain matrix derived from the data. The overall accuracy will be affected by errors in the magnetograms, errors in the partitioning, and errors in calculations using the model field. Of these, the errors in the partitioning turn out to be the most significant and also the most difficult to characterize. One might hope to determine whether the partitioning

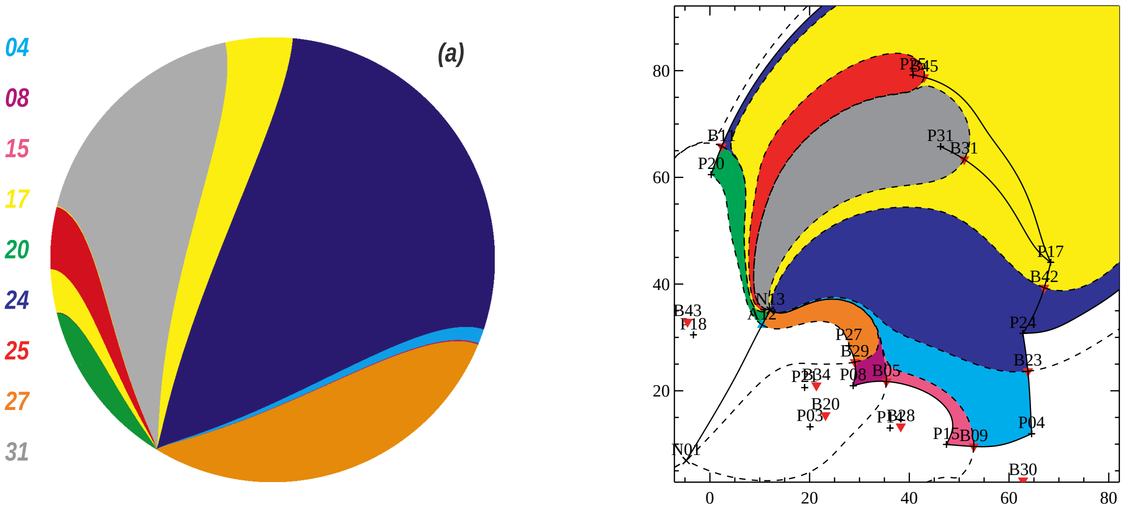

Fig. 7.— Left panel show sets of field lines which terminate on a given source (colorscale) as a function of the direction in which they were initiated from source `N13`. The colatitude and longitude specifying the direction in a spherical coordinate system centered on the source have been mapped to a polar coordinate system in which $r = \sqrt{1 - \cos(colat)}$, in order to preserve the apparent area. Thus, the photosphere is the outer edge of the plot, and the vertical direction is in the center of the plot. The list on the left of the figure gives the color of the label of each source. The sequence of colors starting on the left of the figure is green, dark blue, yellow, red, yellow, gray, yellow, dark blue, cyan, fuchsia, purple, orange; the three bands of yellow indicate that there are three distinct domains connecting source `N13` with source `P17`. Several domains are small enough to be nearly invisible in the figure. The right panel shows the spines (solid lines) and fan traces (dashed lines) for those nulls associated with source `N13`, located at $(x, y) \approx (10, 35)$. The footpoints of the multiple domains `P17`–`N13` (yellow) can be identified with the three areas of the photosphere separated by the spines and fan traces from nulls `B45` and `B31`, and bounded by the fan traces from null `B42`.

into source regions faithfully represents actual flux concentrations. It is not clear how to do this at the present time, since it is not clear what the partitions or concentrations are meant to represent. Instead we characterize only the repeatability and robustness of partitions produced. If they are repeatable and robust measurements they might ultimately prove useful regardless of what they might physically represent.

It is possible to estimate the detection threshold for the vertical field by looking at measurements of the field in quiet sun, but it is much more difficult to estimate the uncertainty in the measured vertical field at each pixel. We make the assumption here that the percent error in the vertical field is constant for vertical fields stronger than the threshold, $\delta B_z(x,y) = \beta B_z(x,y),\ B_z > 3\sigma_{\rm detection}$, and that the noise is uncorrelated. With these assumptions, the uncertainty in the flux in a fixed region of $N_{\rm pix}$ pixels each mapping to an area $A_{\rm pix}$ on the photosphere is

$$\delta\Phi \;\;=\;\; \beta A_{\rm pix}\Big[\sum_{i=1}^{N_{\rm pix}} B_z^2(x_i,y_i)\Big]^{1/2} . \qquad (9)$$

The term "uncertainty" here means the degree by which the same underlying distribution of flux would lead to different source regions or the same regions with different fluxes or locations. Since active region fields evolve on time scales of hours, the high time cadence IVM observations of AR 8210 can be used as an ensemble of realizations of the same, or similar, underlying flux distributions. We therefore use the partitions of the 20 magnetograms to infer the inherent variability or uncertainty in the partitioning algorithm. This is compared to our estimate of the uncertainty (eq. [9]) to determine the value of β.

Each region appears in different partitions with a different flux and a different centroid position. We assume there to be an actual value of each quantity which varies slowly enough over the one hour epoch to be approximated by a linear function of time. We therefore take the residuals from a linear least-squares fit to be the uncertainty in the partitioning algorithm. Figure 8 shows the result for one region (`N13`) whose centroid is moving at $v_x = 417\,{\rm m\,s^{-1}}$ and whose net flux is increasing at $\dot{\Phi} = 8.2 \times 10^{15}\,{\rm Mx\,s^{-1}}$. On top of these trends there is a scatter of $\delta x \simeq 0.26\,{\rm Mm}$ and $\delta\Phi = 750\,{\rm G\,Mm^2}$, which we attribute to uncertainty (indicated by the error bars).

We repeat this procedure for all 52 regions to give a flux uncertainty, $\delta\Phi$, for each one. Figure 9 shows the results, plotted as a function of our estimate of the uncertainty from equation (9). The best fit linear relation of the form given in equation (9) has a slope $A_{\rm pix}\beta = 0.66$, which implies a percent uncertainty of $\beta \approx 1.0$. Clearly, this is not a physically reasonable result, as for the IVM we expect the percent error to be no worse than a few tens of percent. Thus, we conclude that the dominant source of error is indeed the partitioning,

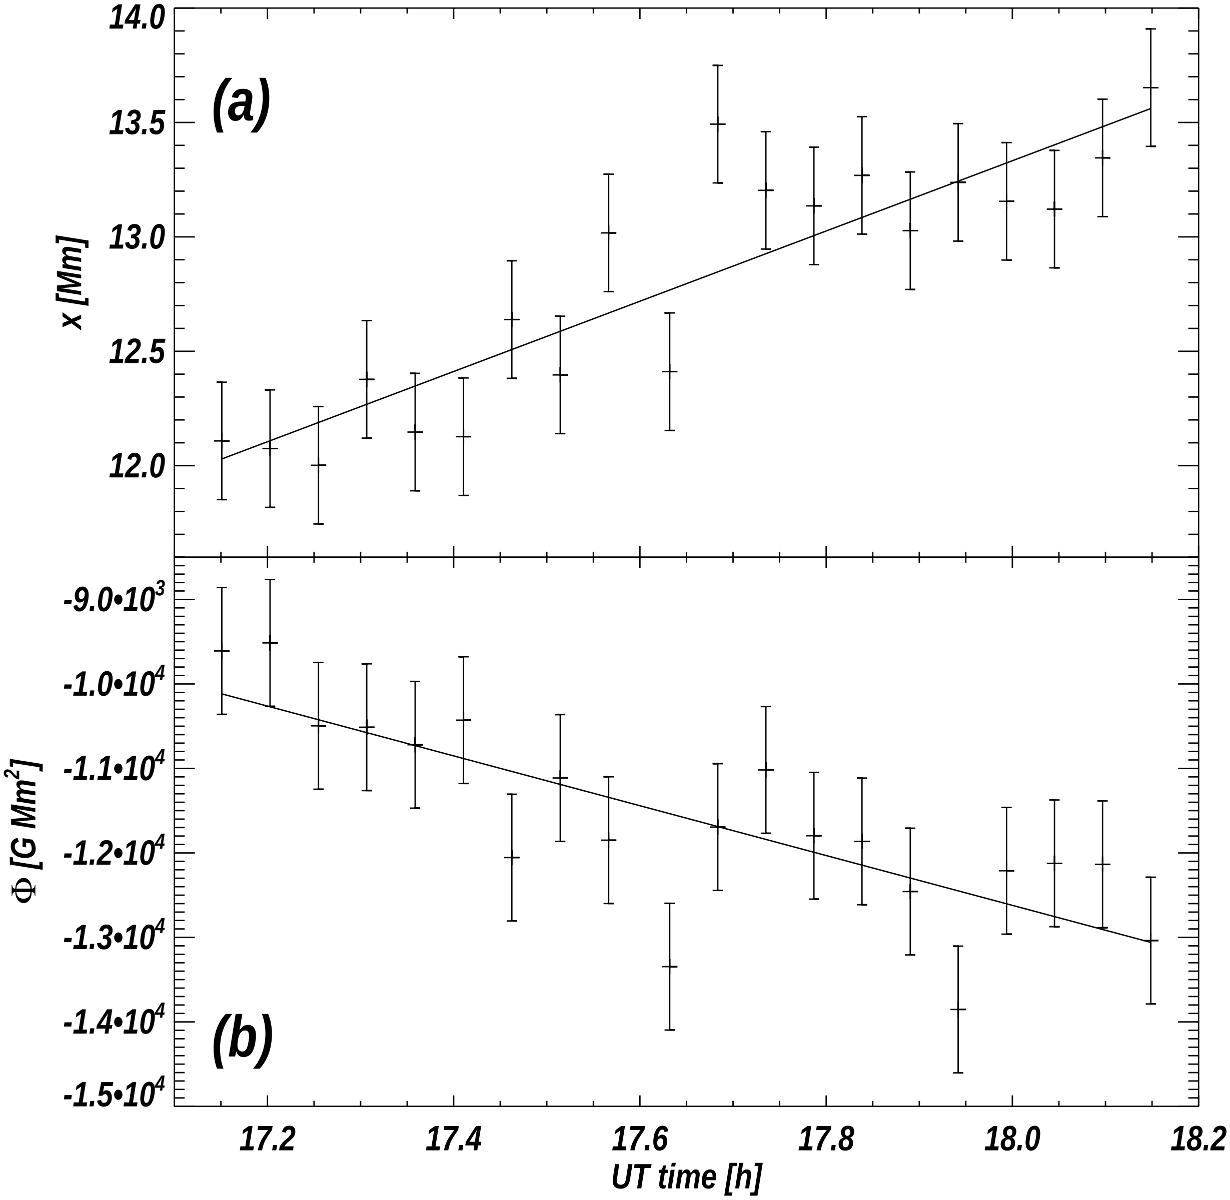

Fig. 8.— For source `N13`, the x-position (a) and strength (b) as a function of time. The error bars are estimated from the residuals from a linear least-squares fit. The slope to the position indicates that this source is moving while the slope to the strength indicates that it is also emerging.

rather than the variation in the field strength in individual pixels. It is interesting, however, how well the functional form of equation (9) agrees with result from the residuals from the timeseries. Since there is no particular reason to expect this functional form for the uncertainty, we also tried fitting $\delta\Phi$, as determined by the residuals, to products of powers of the regions' flux, perimeter and area. The most compelling scaling found was $\delta\Phi \sim 4.2\Phi^{0.60}$ (fluxes expressed in units of $\mathrm{G\,Mm}^2$). However, the agreement with the residuals was only marginally better than that of equation (9). Such an empirical scaling law may allow us, in future, to estimate the uncertainties in the flux from a single magnetogram, instead of requiring a time series.

The errors in centroid location, on the other hand, are all less than one pixel (they average $\delta x \sim 320$ km) and do not appear to scale with any property of the region.

For the domain fluxes, there are two sources of uncertainty: one is the statistical error associated with the Monte-Carlo method; the second is the uncertainty due to the uncertainties in the strengths and positions of the sources, although the positions have so little uncertainty their contribution is likely to be negligible. We estimate the statistical error from the width of the peak in the likelihood function [eq. (3)]. The inverse square root of its relative curvature gives a width

$$\delta\psi_{ij} = \frac{\psi_c}{3}\sqrt{m_{ij} + m_{ji}}. \tag{10}$$

By choosing ψ_c sufficiently small, we can make this uncertainty smaller than that due to the uncertainties in the source strengths. It is difficult to deduce from first principles the latter uncertainty so we estimate it from our data. We begin with the sources for the reference magnetogram. The uncertainties in the position and strength of each source are multiplied by normally distributed random numbers, and the results are added to the location and strength of that source, to produce a new set of sources. The connectivity calculation is repeated for 100 different sets of sources, to produce a distribution of flux for each domain. The standard deviation of that distribution is taken to be the uncertainty in the domain flux $\delta\psi_{ij}$.

Having estimated the uncertainty from the data, we looked for an empirical relationship between the uncertainty in a domain flux and the domain flux itself. We found excellent agreement for

$$\delta\psi_{ij} = 3.2\psi_{ij}^{0.33} \tag{11}$$

From this relation, we note that the percent error in the domain flux will be less than 10% for $\psi_{ij} \gtrsim 200\,\mathrm{G\,Mm}^2$, however, the absolute error continues to grow slowly as the domain

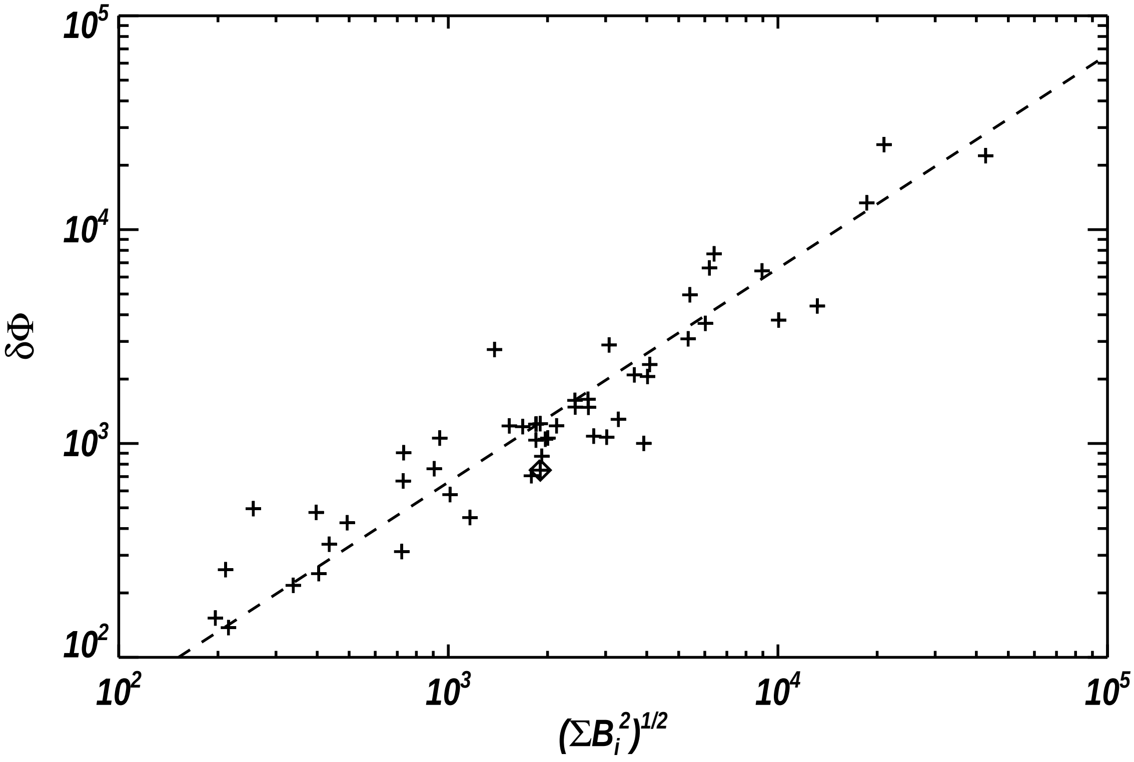

Fig. 9.— Uncertainty estimates for the flux of all 52 sources, calculated from the variations in the time series, as a function of the uncertainty based on an assumption of constant percent error in the vertical field measurements. The dashed line is the best fit line, $\delta\Phi = 0.66(\Sigma B_i^2)^{1/2}$. Source `N13` is highlighted by a diamond symbol.

flux increases. This will certainly have implications for determining changes in the domain matrix, $\Delta\psi_{ij}$, necessary for estimating the separator currents in the MCC model.

For the null points, we follow the same procedure as for the sources: the uncertainty in the location of each null is taken to be the residuals from the linear least-squares fit to the time series. A null at one time is identified as being the same null at a different time if the two have spines which end on the same sources, and are within 15 Mm of each other. In some instances, a null present at one time will be replaced by a null in almost the same location, but with spines which terminate on different sources; these nulls are considered to be distinct. This same behavior is evident in the null points in the noise added magnetograms: nearly cospatial null points can have different sources associated with their spines. Typically, at least one of these sources will be of small magnitude, and the corresponding change in the magnetic topology will involve small flux domains.

The uncertainty in the number of null points is estimated from the standard deviation of the number of nulls found in the noise-added reference magnetograms. While our null finding scheme certainly does not always find all of the photospheric null points, the variation in the number of nulls due to changes in the locations and strengths of the sources is likely to be higher.

For the separators, we also make use of the noise-added reference magnetograms: the separators are located for each random number seed, then matched to the separators in the data. A separator is judged to match if it connects the same pair of nulls. If more than one separator connects a pair of nulls, then the separator in the noise-added calculation is associated with the separator which most closely matches in angle. Once the separators have been matched, the standard deviation of the distribution of the noise-added separators is used to estimate the uncertainty in any given property of the separator. For example, one may wish to estimate the uncertainty in the flux enclosed between the separator and the photosphere.

Figure 10(a) shows the position of null `A12` as a function of time, with error bars determined as described above. One of the spine field lines from this null terminates on source `N13`, whose position was shown in Figure 8(a). The high correlation between the position of the source and the position of the null is not surprising, since the location of `A12` is determined in part by the location and strength of `N13`. On average, the uncertainty in the position of the null points is slightly larger than the uncertainty in the position of the sources, but is still less than one pixel.

Also shown in Figure 10(b) is the flux enclosed by one of the separators from null `A12`. The error bars in this case are determined from the noise added reference magnetograms. The

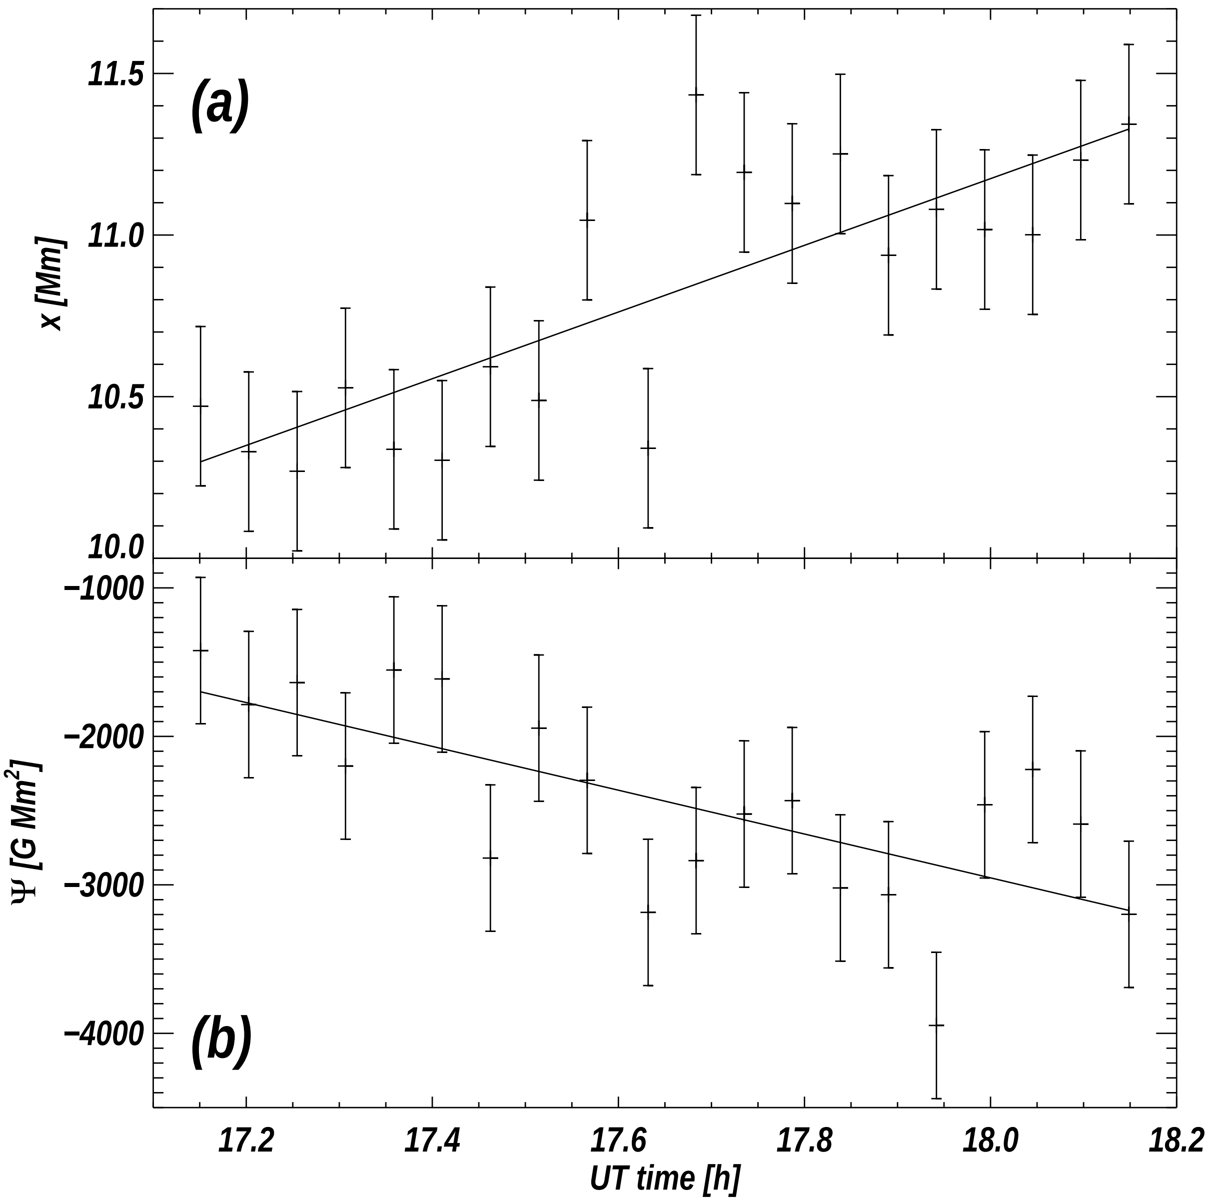

Fig. 10.— For null `A12`, the x-position as a function of time (a) and the flux enclosed between the separator connecting `A12` with `B23` and the photosphere (b). The error bars are estimated from the residuals from a linear least-squares fit. Note the high correlation between the position of the null and the position of the spine source `N13` shown in Fig. 8(a). The correlation for the flux is much weaker, due to the more sensitive dependence of the location of the separator on the sources.

percent uncertainty in the separator parameters is generally much larger than in the source, null and domain flux parameters. For the example shown in the figure, the uncertainty is larger than 10% in a flux with a typical magnitude $\Psi \gtrsim 10^3\,\mathrm{G\,Mm}^2$. In comparison, we found the percent error in the domain flux to be less than 10% for $\psi \gtrsim 200\,\mathrm{G\,Mm}^2$.

5. Discussion

We have outlined an automated way to implement a Magnetic Charge Topology model for real observational data. Given the complexity of the manipulations performed to obtain a result that primarily reflects real changes in the solar photosphere, we must concern ourselves with whether we have succeeded. One comparison that can be made is of the horizontal velocities. There have been several studies made on the velocities of AR 8210 (Nindos et al. 2003; Welsch et al. 2004; Longcope 2004) which differ in detail, but have several common features. In particular, the central sunspot is seen to rotate in a clockwise direction, while the opposite polarity plage beside it rotates in the reverse direction. In addition, there is a patch of negative polarity close to the spot that is moving very rapidly. All of these features can also be found in our MCT analysis of AR 8210, as shown in Figure 11.

Both of these rapidly evolving areas are potentially of interest for solar energetic events. Chromospheric images, courtesy of the solar MURI effort, from the U. Hawai‘i Mees CCD Imaging Spectrograph (“MCCD”, Penn et al. 1991) and the high-resolution Hα camera at Big Bear Solar Observatory (Denker et al. 1999) show that numerous homologous events of various sizes on 1998 May 01 between 17:00 UT and 23:30 UT produced line-center brightenings primarily in the east and south quadrants of the active region. In Figure 12, we show a sample image from the MCCD, taken at 18:06:44 UT. The northernmost extension of the large brightening is in the vicinity of null `B24`, and the brightening extends toward nulls `A50` and `B01`, `B08`, `B22`. In addition, the brightening near the sunspot is in the vicinity of null `A19`. While there are certainly other parts of the region containing more separators, this area is of interest because of the *combination* of the existence of separators and the relative motion of the positive and negative polarity areas, which may induce currents to flow along the separators.

There were also, as expected, numerous small-scale short-lived brightenings during the course of the day that occurred over the entire active region including in the area near source `N13` and the nearby null `A12`. Hα images from Kanzelhohe Observatory on 1998 May 02, on the other hand, show brightenings related to an event beginning at 13:35 UT concentrated in the north and west quadrants, again near `A12`. This area is also of particular interest because of the complex topology, as evidenced by the multiple domains and the

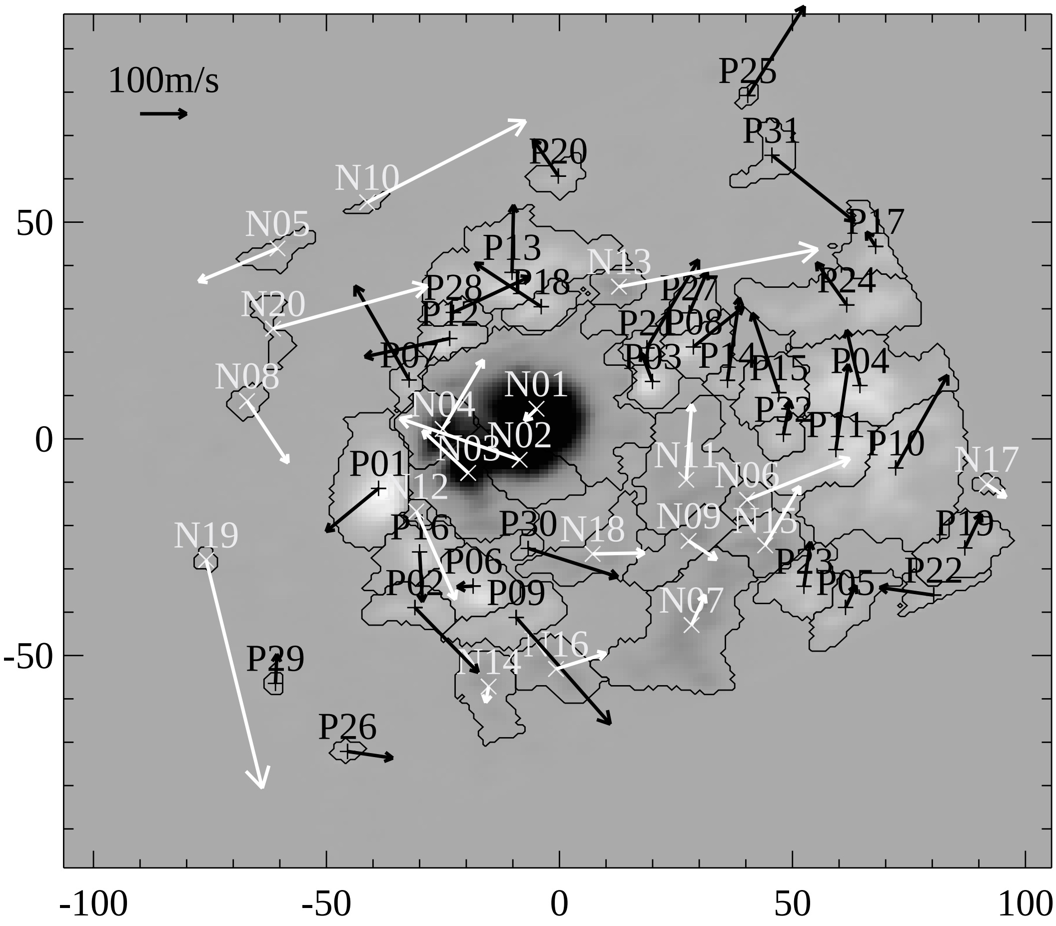

Fig. 11.— Horizontal velocity for each of the sources. We are able to reproduce the clockwise rotation of the large negative spot, the counter-rotation of the surrounding plage, and the rapid motion of source `N13`, as determined by local correlation tracking methods.

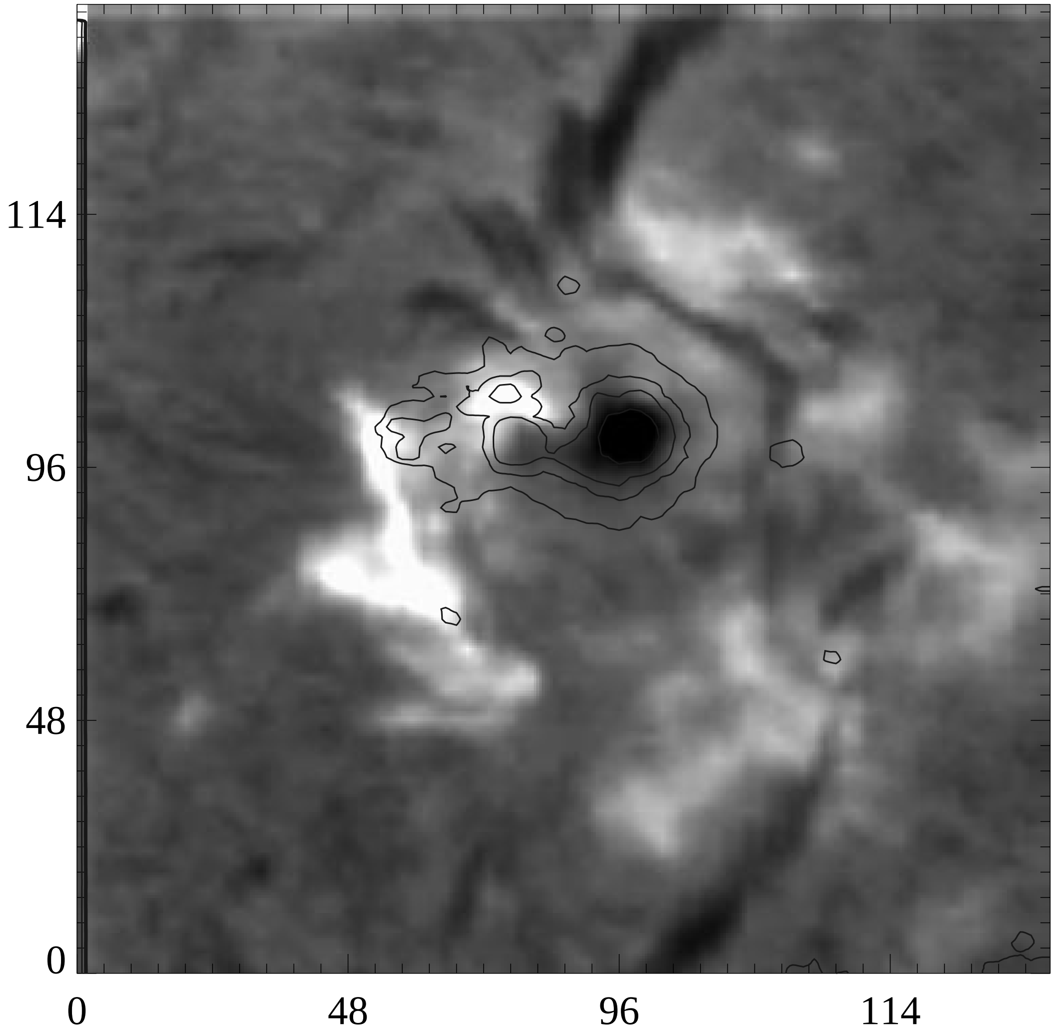

Fig. 12.— Image of line-center Hα from the MCCD on 1998 May 01 at 18:06:44 UT. Contours are continuum intensity, with the first contour roughly corresponding to the penumbral boundary, and approximately enclosing the partitions associated with sources `P01`, `N01`, `N02`, `N03` and `N04` in AR 8210. Terrestrial north/west is up/right respectively, and the tickmarks indicate the scale in arcseconds. Features to note include the brightenings in the south-east quadrant, and the extended active-region filament in the north-west quadrant which becomes active with flare activity on 1998 May 02.

large number of separators associated with null `A12`, and the rapid motion of source `N13`. Of course, the present analysis does nothing to predict why the events in the south and east quadrants occur before the event in the northwest.

Now that we have a technique for automatically implementing an MCT model, there are a number of possible applications. In addition, the partitioning scheme can be used in other contexts, for example, with more sophisticated extrapolation techniques. Perhaps the most direct application of the technique is to repeat the analysis of Longcope & Silva (1998) in attempting to match the locations of flare brightenings to separators. Such a comparison can now be done for a large number of active regions, to quantify how often there is a match. Further, the complexity of the coronal field can be quantified, and used as a predictor for flares and eruptive events, in the same manner as various previous studies (e.g., Leka & Barnes 2003b; Falconer et al. 2002) have used measures describing the photospheric field. Such a study is being undertaken by Barnes & Leka (2005, in preparation).

This work was carried out at Colorado Research Associates Division of NorthWest Research Associates, Inc. with data from the U. Hawai‘i Mees Solar Observatory. The authors wish to thank the anonymous referee for helpful comments and for such a quick response, and J. Li at the U. Hawai‘i/Institute for Astronomy for providing the MCCD data. DWL thanks Colin Beveridge for helpful discussion on the relationship between separator loops and redundant domains. Funding is gratefully acknowledged from the Air Force Office of Scientific Research under contracts F49620-00-C-0004, F49620-02-C-0191, and F49620-03-C-0019, and from NASA under grant NAG5-10489.

REFERENCES

Bagalá, L. G., Mandrini, C. H., Rovira, M. G., Démoulin, P., & Hénoux, C. H. 1995, Sol. Phys., 161, 103

Baum, P. J. & Bratenahl, A. 1980, Sol. Phys., 67, 245

Beveridge, C. & Longcope, D. W. 2005, Sol. Phys., in press

Brown, D. S. & Priest, E. R. 2001, A&A, 367, 339

D'Agostini, G. 2003, Rep. Prog. Phys., 66, 1383

Démoulin, P., Henoux, J. C., & Mandrini, C. H. 1994, A&A, 285, 1023

Démoulin, P., van Driel-Gesztelyi, L., Schmieder, B., Hénoux, J. C., Csepura, G., & Hagyard, M. J. 1993, A&A, 271, 292

Denker, C., Johannesson, A., Marquette, W., Goode, P. R., Wang, H., & Zirin, H. 1999, Sol. Phys., 184, 87

Falconer, D. A., Moore, R. L., & Gary, G. A. 2002, ApJ, 569, 1016

Gorbachev, V. S. & Somov, B. V. 1988, Sol. Phys., 117, 77

—. 1989, Soviet Ast., 33, 57

Greene, J. M. 1988, J. Geophys. Res., 93, 8583

Kankelborg, C. C. & Longcope, D. W. 1999, Sol. Phys., 190, 59

LaBonte, B., Mickey, D. L., & Leka, K. D. 1999, Sol. Phys., 189, 1

Lau, Y.-T. 1993, Sol. Phys., 148, 301

Lau, Y.-T. & Finn, J. M. 1990, ApJ, 350, 672

Leka, K. D. & Barnes, G. 2003a, ApJ, 595, 1277

—. 2003b, ApJ, 595, 1296

Leka, K. D. & Skumanich, A. 1999, Sol. Phys., 188, 3

Longcope, D. W. 1996, Sol. Phys., 169, 91

—. 2004, ApJ, 612, 1181

Longcope, D. W., Kankelborg, C. C., Nelson, J. L., & Pevtsov, A. A. 2001, ApJ, 553, 429

Longcope, D. W. & Klapper, I. 2002, ApJ, 579, 468

Longcope, D. W. & Silva, A. V. R. 1998, Sol. Phys., 179, 349

Maclean, R., Beveridge, C., Longcope, D., Brown, D., & Priest, E. 2005, Proc. Roy. Soc., in press

Mandrini, C. H., Démoulin, P., Hénoux, J. C., & Machado, M. E. 1991, A&A, 250, 541

Mickey, D. L., Canfield, R., LaBonte, B. J., Leka, K. D., Waterson, M. F., & Weber, H. M. 1996, Sol. Phys., 168, 229

Nakagawa, Y. & Raadu, M. A. 1972, Sol. Phys., 25, 127

Nindos, A., Zhang, J., & Zhang, H. 2003, ApJ, 594, 1033

Parnell, C. E., Priest, E. R., & Golub, L. 1994, Sol. Phys., 151, 57

Penn, M. J., Mickey, D. L., Canfield, R. C., & LaBonte, B. J. 1991, Sol. Phys., 135, 163

Priest, E. R. & Forbes, T. G. 1989, Sol. Phys., 119, 211

Schrijver, C. J., Hagenaar, H. J., & Title, A. M. 1997a, ApJ, 475, 328

Schrijver, C. J., Title, A. M., Van Ballegooijen, A. A., Hagenaar, H. J., & Shine, R. A. 1997b, ApJ, 487, 424

van Driel-Gesztelyi, L., Hofmann, A., Demoulin, P., Schmieder, B., & Csepura, G. 1994, Sol. Phys., 149, 309

Welsch, B. T., Fisher, G. H., Abbett, W. P., & Regnier, S. 2004, ApJ, 610, 1148

This preprint was prepared with the AAS LaTeX macros v5.2.