

Dynamic Network Visualization¹

James Moody
Ohio State University

Daniel McFarland and Skye Bender-deMoll
Stanford University

Increased interest in longitudinal social networks and the recognition that visualization fosters theoretical insight create a need for dynamic network visualizations, or network “movies.” This article confronts theoretical questions surrounding the temporal representations of social networks and technical questions about how best to link network change to changes in the graphical representation. The authors divide network movies into (1) static *flip books*, where node position remains constant but edges cumulate over time, and (2) dynamic *movies*, where nodes move as a function of changes in relations. Flip books are particularly useful in contexts where relations are sparse. For more connected networks, movies are often more appropriate. Three empirical examples demonstrate the advantages of different movie styles. A new software program for creating network movies is discussed in the appendix.

INTRODUCTION

Ranging from simple histograms to dynamic images of the birth of galaxies (Abel, Bryan, and Norman 2000), visualization tools have always been key elements in scientific advancement. The ability to see data clearly creates a capacity for building intuition that is unsurpassed by summary statistics. Wide ranges of distributional shapes (Handcock and Morris

¹ This work was partially supported by grants from the National Science Foundation (IIS-0080860) and the National Institutes for Health (DA12831, HD41877) awarded to Moody, and a Research Incentive Award provided by Stanford University’s Office of Technology and Licensing (grant 2-CDZ-108) to McFarland. Thanks to Mark Handcock, Martina Morris, Walter Powell, Doug White, the participants of James G. March’s Monday Munch at Stanford University, and participants of the Social Structure Research Group at Ohio State University. Direct correspondence about this article to James Moody, 372 Bricker Hall, 190 North Oval Mall, Ohio State University, Columbus, Ohio 43210. E-mail: Moody.77@sociology.osu.edu

1999), nonlinear relations, or spatial (geographic or social) proximity can be easily summarized with an image that helps scientists develop theory. While the basic principle of substantive clarity is key to successful data visualization, work in this field is often as much art as science. To build network visualization tools, we need to examine carefully questions about the meaning and implication of time in the formation of social networks. In this article, we hope to extend a bit more science into the art of dynamic network visualizations and further the theoretical discussion about social network dynamics.

Social network research has made extensive use of visualization since Moreno first introduced the sociogram (Brandes, Raab, and Wagner 2001; Freeman 2000*a*, 2000*b*; Freeman, Webster, and Kirke 1998). Actors are usually represented as points, and relations among actors are represented by lines, with relational direction indicated by arrows. Early sociograms were drawn by hand (Whyte 1943; Coleman 1961), and the layout was determined by the artistic and analytic eye of the author. Such early graphs were usually simple, having few relations per person or a clear hierarchical structure.

The state of the art has progressed remarkably since Whyte and Coleman, and a growing body of research has developed around various definitions for optimal network layout (Brandes et al. 2001; Freeman 2000*b*; McGrath et al. 1997). Most network images do a poor job of representing *change* in networks, and researchers make do by presenting successive snapshots of the network over time (Bearman and Everett 1993; Powell et al. 2005; Roy 1983). The problem is fundamental to the media. To effectively display the relational structure of a social network, at least two dimensions are needed to represent proximity, and that leaves no effective space (on a printed page) to represent time.² However, recent media advances allow us to use space to represent social distance and *movement* to represent change over time (Bender-deMoll and McFarland 2002; Freeman 2000*b*).

² There have been a number of creative attempts to overcome this limitation, such as producing multiple images of the network and placing them “next” to each other temporally (see, e.g., <http://www.stanford.edu/group/esrg/siliconvalley/docs/network-memo.pdf>). These approaches require a fair amount of reader training to see the difference between each plotted time frame. The less-than-intuitive results follow because these approaches are attempting to replicate a dynamic process in a static medium. Our approach is to use the dynamic nature of such networks directly, thereby producing a more readily interpretable visualization.

RESEARCH ON LONGITUDINAL SOCIAL NETWORKS

Much of the recent interest in longitudinal social networks revolves around understanding how networks develop and change, as scientists seek to build models of social processes that result in observed structures (Doreian et al. 1996; Leenders 1996; Nakao and Romney 1993; Snijders 1998; Suitor et al. 1997; Weesie and Flap 1990; Zeggelink et al. 1996). Such dynamics are important for understanding network stability, which is necessary for understanding the effect of networks on individual and group behavior. The clear importance of such questions has prompted a good deal of methodological research on network change. For example, Snijders (1996, 1998, 2001) has developed models for evaluating the extent of change in a social network, conditional on structural features of the graph.

The purpose of dynamic network visualizations is to help augment theoretical intuition provided by summary statistics and standard static visualizations. Until now, visualizations of network change have tended to take two forms. The first common visualization approach plots network summary statistics as line graphs over time. For example, Doreian et al. (1996) present change in reciprocity and transitivity for the Newcomb data (see also Gould 2002). However, such summary statistics provide information on a single dimension of a network's structure. For example, one might find that a network reaches a given equilibrium transitivity level, but since transitivity is a single average for the graph as a whole, we cannot know if this—*in itself*—means the graph is now relationally stable. The second common visualization approach is to examine separate images of the network at each point in time. Unfortunately, such images are often difficult to interpret, since it is impossible to identify the sequence linking node position in one frame to position in the next.

THEORETICAL IMPLICATIONS OF NETWORK DYNAMICS

A standing critique of social network research has focused on a “structural bias” that implicitly denies much of the dynamic nature of social relations (Emirbayer 1997; Emirbayer and Goodwin 1994). For some types of relations (such as conversations that occur in real time), one could argue that the networks are largely artificial constructions built by aggregating dead past events. The network “structure” as such only emerges from this aggregation. While we do not think this argument should be pushed too far, it raises important questions about *how* the temporal embeddedness of relations defines a dynamic social space. While discussions of meaning and temporal abstraction in themselves are not new (Abbott 1997; Bearman et al. 1999; Danto 1985), our goal is to identify a meaningful way

to capture network dynamics empirically that simultaneously allows for the importance of past relations typically captured in static network images.

Basic Terms

Much research on social networks is filled with static nouns, such as “roles,” “relations,” “obligations,” and so forth. Longitudinal research on social networks requires a different set of *process* terms, such as “ritual,” “dance,” “pulse,” “tempo,” “congealing,” or “dispersal.” The key distinction is that an apparently static network pattern often *emerges* through a set of temporal interactions, with important implications for the relational process under investigation. For example, when viewed in continuous time, networks may develop by spurts or build slowly and steadily, or they may reflect repeated ritual behaviors that mix moments of order and chaos.

The most basic dimension for dynamic relations is *relational pace*. These are questions concerning the *rate of change* in relations, with particular interest in irregularities. The pace of relation formation can thus be described with respect to levels (fast, slow), change (accelerating, decelerating), or stability (cascades, jumps and starts, etc.). Clearly, the exact meaning of such terms depends on context, since the relevant scale will vary across types of relations. Think, for example, of seconds for conversation networks, weeks for friendship networks, decades for world trade networks, and generations for kinship networks (Nadel 1955; Collins 1981).

A second specifically temporal aspect of relations revolves around the order of relations, or their *sequence*. In many circumstances, being able to explain the prevalence of given structures depends on identifying the order in which relations occur. For example, the distinguishing characteristics of Johnsen’s (1985, 1986) process agreement models for friendship formation depend on whether one *first* forms reciprocal friends, which *then* generates agreement on third parties and creates transitive ties, or if two people first agree on their admiration of a third, which draws them together in friendship. Similarly, Chase’s (1980) explanation for the development of hierarchy in social relations rests on a particular order of relationship unfolding.

A special case of the intersection of pace and sequence is found in notions of concurrency developed for disease networks (Kretzschmar and Morris 1996; Morris 1993). Two relations are concurrent if they overlap in time and share at least one person. While static views of networks focus on multiplicity (the overlap of types of relations), concurrency is effectively “temporal multiplicity,” which can dramatically complicate our

understanding of relations. Moody (2002) shows that the temporal unfolding of relations reduces reachability, which has implications for any measures on networks with relations enacted in real time. Consider, for example, the different disease implications of a relational “switching” model, where actors retain past relations until they are secure in new relations, leading to a sequence of short-duration overlaps and a multiple partner model, where people overlap for extended periods of time with multiple partners (Laumann and Youm 1999).

Finally, the richness of a relational structure further expands when we link relational timing to types of tie. When we allow simultaneous consideration of the pace and sequence of ties with variation across types of ties, we can start to tap questions of how sets of interaction sequences transform into stable relations (see Nadel 1955). For example, if theoretical discussions of local action (Leifer 1988; Padgett and Ansell 1993) are correct, then we should be able to map social interaction as a sequence of seemingly random “milling around” that quickly cascades once role positions are fixed. This type of rich network visualization provides a more holistic understanding of the network’s structure than any single-dimensional index could.

These are just a small number of the theoretical considerations evident when one takes time seriously in the modeling and understanding of social networks, and we expect that others will contribute much to this discussion. Our belief is that any such discussions are best grounded in empirical investigation (Abbott 2000), which will be enhanced with a set of flexible tools for displaying dynamic network data.

REPRESENTING LONGITUDINAL NETWORK DATA

To develop dynamic network images, we need to conceptualize clearly how time is encoded in social networks. We conceive of time in two analytically distinct forms: *discrete* and *continuous*. *Discrete* renditions of time consist of cross-sectional snapshots of the network. Hence, longitudinal analysis focuses on the change from one network state to another without any (explicit) reference to the sequence of changes that generate change. In such cases, a process is generally inferred from the total network change across time. Due to research costs and design, most longitudinal network studies use discrete time. *Continuous* renditions of time consist of streaming relational events or interactions recorded with exact starting and ending times. Streaming relational events consist of sequential dyadic events or interactions whose visual representation should unfold as a continuous social process. Continuous representations of time enable researchers to identify how overall network changes emerge through *or-*

dered dyadic events. But, researchers need not analyze the data in its most reductive form, as they can aggregate relational events into larger time units (such as minutes instead of seconds). Hence, it is possible to develop discrete notions of time from continuous data but not vice versa.

If researchers have many panels of network data or continuous representations of network change, they have several decisions to make before they represent the data visually. First, *what constitutes their networks?* With discrete waves of sociometric surveys, each wave becomes the network used in graphic representation. However, with continuous data it becomes more difficult to define a network's temporal boundary (Laumann, Marsden, and Prensky 1983).³ Is the network defined by 10 relational events, all the events in one minute, in 10 minutes, or in a single day? The answer to this may depend on the empirical focus of the researcher, but what is undeniable is that we cannot select each dyadic event as a network. A network consists of a pattern of social relations, and therefore we must identify chunks of time that substantively capture the nature of relational events and the character of temporary networks that arise in the focal context.⁴

For visualizing continuous network data, we characterize a network by a *time window* that spans a set of relational events (Bender-deMoll and McFarland 2002). The relational events that transpire within a win-

³ One could argue that Laumann et al.'s (1983) discussion of nominal and realist notions of networks' spatial boundaries also applies to networks' temporal boundaries. Realist notions of time will have natural boundaries that actors acknowledge and recognize (i.e., school weeks and days, class periods, etc.). Nominalist notions of time boundaries are defined by the researcher for a variety of theoretical concerns (i.e., development focus, period of historical change, etc.).

⁴ There is a conceptual tension here regarding what a network consists of and how different kinds of representations will offer different understandings of network change. This is a conceptual issue that runs to the heart of network theory. The general structuralist view is that relations form from repeated interactions and their aggregate patterns (Nadel 1955; Hinde 1971). In contrast, interactionists argue that relations and roles are established by ritual patterns of interaction that have a particular sequence, such that reordering the sequence of interactions would undermine the meaning of the ritual and relation being enacted (Goffman 1967). Work in social movements and collective action suggest that it is not so much the accrual of everyday interactions that constitutes a relation but the rare, jarring events which lock in or rewire relational patterns (Gould 1995). Finally, work on diffusion treats relations as conduits for the transferal of goods, and relational timing acts as a "switch" determining temporally "downstream" nodes from "upstream" nodes (Moody 2002; Morris and Kretzschmar 1997). Regardless of the outlook, a more fine-grained perspective may reveal change processes in more substantive detail. From the study of continuous relational events, we may learn how ordered dyadic acts can form ritual patterns emblematic of roles and relations (enriching the aggregate interaction view to include order), and how rare jarring events reverberate through relational structures so as to create network change, or how the unfolding of relations creates diffusion opportunities (Emirbayer 1997; Emirbayer and Goodwin 1994; McFarland and Bender-deMoll 2003).

dow of time are aggregated into a network. One implication of the time window is that relational events have a residue that extends beyond their occurrence, defined by the length of the time window. For example, in the classroom data discussed below we have streaming interactions specifiable to seconds. Since many interactions take hold of actors' attention for a short time after they occur (at least), we decided to make the time window equal to 2.5 minutes, which affords a relatively meaningful graph for characterizing the interaction pattern arising in that class (see fig. 1).

The "time window" defines a right open interval (from "slice start" to "slice end"). Events that fall inside the interval (arcs 2–5, node C) or ongoing events that extend past the interval's bounds (nodes A, B) are included in the network. This method is quite flexible as it deals fairly well with the aggregation of both "instantaneous" events and those having a defined duration. Successive networks are generated by incrementing the values for "slice start" and "slice end" points by a constant (delta) value.

After we define the window of time, the next decision is whether successively defined windows are to be *overlapping or nonoverlapping*.⁵ For example, if we aggregate streaming dyadic relations into 2.5 minutes of class time for defining the network, we can make the next window consist of the next 2.5 minutes (nonoverlapping) or have it partially overlap with the current one, as in a moving average (see Doreian 1980, 1986). The most incremental relational change consists of a new window that only has one tie added to the front of it and another dropped off the back.

The third set of concerns center on the *timing of arc representation*. The placement of nodes in a graph is distinct from the animation of a particular relation. One can use the network data to generate coordinates for nodes and then control the animation of when relations are shown by a different procedure. We have found that the timing of arc representation helps viewers see the network process more clearly because it clarifies the link between relational change and node placement. In essence, the animation of arcs follows a second time window, a "render window" that "slides" from one network time window to the next. Nodes and arcs are drawn and deleted as the render moves over them, and the node coordinates are adjusted incrementally between the two layouts (see fig. 2).

The process of creating a visual transition from one network's layout coordinates to another involves drawing a number of intermediate "renders" of the network while gradually adjusting node coordinates and adding or deleting nodes and arcs as they fall within the moving render window. Successive network windows may overlap, which is controlled

⁵ In SoNIA, we define the degree of overlap by a delta measure, which indicates how much time is added or subtracted from the prior window.




FIG. 1.—Construction of a “network slice” by binning node and arc events

by the relationship between the duration of the time slice window and the offset between successive windows (bin delta).

If the network data are continuous, adjusting the size of the render window in relation to the network window allows the researcher limited control over how many new relations tend to *lead or lag* node movement. For example, A’s interaction with B can *lead* nodal position such that the arc transpires and *then* the nodes shift position; or we can *lag* interaction such that nodes move closer to one another and then an arc forms (as if nodes first draw near each other and then talk). However, the actual timing of arc representation is determined by where an arc’s time coordinates land relative to those of the bin.

NETWORK VISUALIZATION

Graph Layout Principles

Good reviews of the history of network visualization can be found in the literature (Freeman 2000b). The effectiveness of network visualization technique differs by network size. Small networks can focus on detailed elements of the graph structure while larger networks can mainly capture gross topology. Visualizing networks of tens of thousands of nodes requires further abstraction yet.⁶ Our interest in this article is identifying layout principles that are useful for research in small to moderate sized networks (fewer than about 100 people).

A useful goal for most social network layouts is to represent *social*

⁶ See app. B (online only) in this article for a static example and Powell et al. (2005) for a dynamic example.

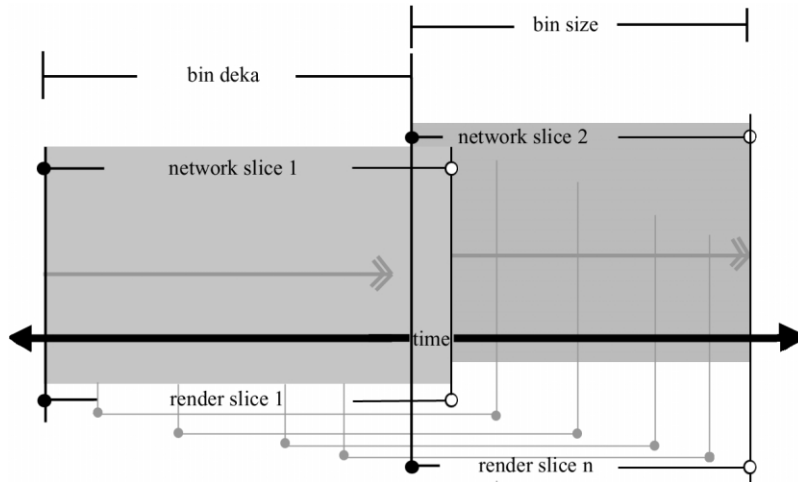


FIG. 2.—Transitioning between slices with a sliding “render window”

*distance as physical distance.*⁷ This representation allows viewers to get a spatial understanding of social relations, as nodes with many relations in common are placed close together on the printed page. An intuitive impression of the network structure then emerges from the proximities in the image. For example, one can easily see racial segregation in a network image as the clustering of races into distinct spaces on the page (Moody 2001b, p. 683). A social-distance-based representation of network structure is facilitated when edge lengths are equated to relational strengths. Various force-directed layout techniques are usually successful at this.⁸

Within the larger topology, aesthetic features of the graph help facilitate readability. All else equal, edge crossing, running edges “under” nodes that are not connected to each other, and stacking nodes on top of each other

⁷ For this article, social distance is defined graph theoretically as the length of the shortest path in the network connecting two nodes. However, alternate formulations might be interesting and advisable depending on a researcher’s theoretical framework.

⁸ There is a difficulty here when ties are directed, since the visual distance from *i* to *j* has to equal the visual distance from *j* to *i*, even though social distance need not be symmetric. In SoNIA, we use the basic correspondence between screen distance and social distance to evaluate layout fit, using a modification of Kruskal’s stress measure,

$$Stress = \frac{\sum_{ij} (ScreenDist - GraphDist)^2}{\sum_{ij} GraphDist^2},$$

where *ScreenDist* is the Euclidean distance of the layout *x*–*y* coordinates, and *GraphDist* is the geodesic distance based on an underlying symmetric graph.

all hinder readability of the graph (Davidson and Harel 1996). While we do not explicitly control these features in the movies below, they are key elements in our evaluation of the resulting graph layouts. Color, size, and shape are useful ways to add additional information on actor or relational attributes. We have found that visualizing multiplex ties—multiple ties of different types—can be effectively represented with colors and “transparent” arcs, which allow you to look through the edge to see multiple relations on a single dyad as a blend of colors.

Graph Layout Algorithms

Force-directed or spring-embedder algorithms are among the most common automatic network layout strategies. These algorithms function substantively on an analogy, treating the collection of nodes and arcs as a system of forces, and the layout as an “equilibrium state” of the system. Generally, edges between nodes are represented as an attractive force (a “spring” pulling them together), while nodes that do not share a tie are pushed apart by some constraint to help prevent overlap. The two most common layout algorithms are Kamada Kawai (KK; Kamada and Kawai 1989) and Fruchterman Reingold (FR; Fruchterman and Reingold 1991). For FR, the underlying model roughly corresponds to one of electrostatic attraction in which the attractive force between connected nodes is balanced by a repulsive force between all nodes. For KK, it is as if all nodes are connected by springs with a resting length proportional to the shortest-path distance between them. For both KK and FR, the relations between nodes must be expressed as distances rather than adjacencies. In KK, this “dissimilarity” matrix is constructed from geodesic distances between nodes. In FR, the dissimilarity matrix is constructed directly from adjacencies.

Both of these algorithms are available in current network drawing software, such as Pajek (Batagelj and Mrvar 2001) and NetMiner (Cyram 2003). The algorithms work by iterative optimization—adjusting a node’s position by reacting to the positions of others. As such, the starting position of the network affects the outcome. The details of the optimization procedure can also affect layout. Because there are no exact calculations of global minima, layouts are subject to local convergence problems.

A second class of algorithms plots nodes using the dimensions that result from multidimensional scaling (MDS) techniques based on the geodesic distances or some alternate measure of node similarities.⁹ These layouts are available in common software packages, such as NetDraw, Krackplot, or MultiNet. Substantively, metric MDS models reduce the $n - 1$ dimen-

⁹ Technically, KK could be considered a variant of nonmetric MDS.

sions present in a network to the two dimensions that capture most of the variance in observed multidimensional distance. Because they take input distances as their starting point, these models do not, in themselves, attempt to stop nodes from stacking on top of each other, and as such they can easily generate layouts where structurally equivalent nodes occupy the same location. Many authors correct for this by adding a small amount of noise or other correction to the resulting structure. One theoretical value to the metric MDS model is the direct linkage between input relational distances and resulting display. Nonmetric MDS algorithms use one of a number of optimization techniques to find a “low-stress” (well-balanced relation between input and screen distances) layout.

Because of the necessity of giving results in terms of two-dimensional distance, both force-directed and MDS models are symmetric, using either a symmetric input distance matrix (MDS) or having symmetric forces driving/pulling nodes together in force-directed layouts. While the distinction between direction and distance is irrelevant in static layouts, it can be substantively important for dynamic layouts, as we often want nodes to follow their nominations. To accommodate direction, we introduce an algorithm that follows a simple peer influence analogy, building on a suggestion in Moody (2001*a*). A node’s position is a function of their prior position based on a self-weight and the average position of those they nominate (see fig. 3).

In figure 3, the thin black arrows indicate nominations, while the wider gray arrows indicate net force for each person’s movement. Person 2 should move down in response to the force of ties sent to persons 1 and 3, while 3 moves slightly to the right based on the net pull of 4 and 5 over the pull of 2, and slightly up as a function of the net pull of 2 and 4 over 5. Similarly, node 6 moves toward the center of nodes 1, 3, and 5, based on the strong difference in their positions. It is important to note that nodes 1, 4, and 5 should not move at all, since they send no ties elsewhere in the network.

The peer influence algorithm works by computing the weighted average of a node’s current x and y positions and that of those that node nominates. This process is iterative, adjusting each node’s position to that of others multiple times.¹⁰ As with the metric MDS layouts described above, the algorithm makes no explicit reference to resulting node positions, which can result in nodes stacked on top of each other. The result is often similar in flavor to a metric MDS plot, which should not be surprising given their

¹⁰ In practice, the number of needed iterations is quite small, as the procedure tends toward high “consensus” positions. Similarly, high self-weights generate more stable movies.




FIG. 3.—An example of the peer influence layout algorithm

common roots (under a particular set of assumptions) to a general class of eigenvalue models (Friedkin 1998).

EMERGING NETWORKS

In many settings, the substantive question is the emergence of a cumulative network structure. Relations in such settings are temporally sparse, requiring the aggregation of many periods (i.e., a large time window) to generate a meaningful image of the network structure. For example, a high school romantic network viewed on any given day will consist largely of a set of completely disconnected dyads, since only a small number of students are involved in multiple romantic relations on any given day. Viewed contemporaneously, this network is essentially structureless. However, sexually transmitted disease risk resides in the *history* of previous relations, which cumulate and thus provide pathways for disease transmission. As such, we are substantively interested in both the cumulative structure of the network and the process through which the structure unfolds.

We have found that one of the most effective ways for displaying such sparse networks is to show how the network emerges over time, by adding

nodes and relations as they appear, but placing them in the display plane based on the final aggregate structure. The dynamic element then appears by revealing the network over time and seeing the pieces grow together. As new nodes and relations enter the population, they are added to the graph. To effectively distinguish current relations from past relations, we “ghost” relations when they end (i.e., fade out).

As an example, consider the graph given in figure 4 below, which is a small section of the sociology collaboration network (Moody 2003), recording all article collaborations (among this small sample) between 1963 and 1999. Based only on the image in figure 4, it appears clear that node A is at the center of this network.

However, our impression of this network changes if we examine how it develops over time, as can be seen in movie 1. (All movies in this article may be viewed in the online copy of *AJS*.) Here we see that the relations connecting the full component only form much later in the unfolding of the network. Note also that the *pace* of relations is evident in the movie, as the structure admits to a four-year “dormant” period between 1969 and 1973 when no changes occur. One’s understanding of the betweenness of these center nodes changes once the temporal nature of the network is revealed.

As a second example, consider relations among all people in the largest connected component of an adolescent romantic network (see Bearman, Moody, and Stovel [2004] for details of this structure’s development; Moody [2002] for a discussion of dynamic diffusion implications). Figure 6 shows all relations that were active in the 18 months prior to the interview.

The most striking feature of the static network is a large connected component, linking 288 students through romantic involvement. When viewed dynamically, we see that this structure emerges quite late in the local history of the network and is never observed as a complete structure at any point in time. Moreover, a close review of the movie reveals that certain contemporaneous substructures never occur—namely four-cycles. Prior work suggests that avoiding such four-cycles helps generate the sparse treelike structure of the network (Bearman et al. 2004). In general, this flip book approach is among the simplest network movies to implement. The technique will likely be most effective when one is interested in the cumulative graph structure, or when one can meaningfully fix node position in an x - y space.

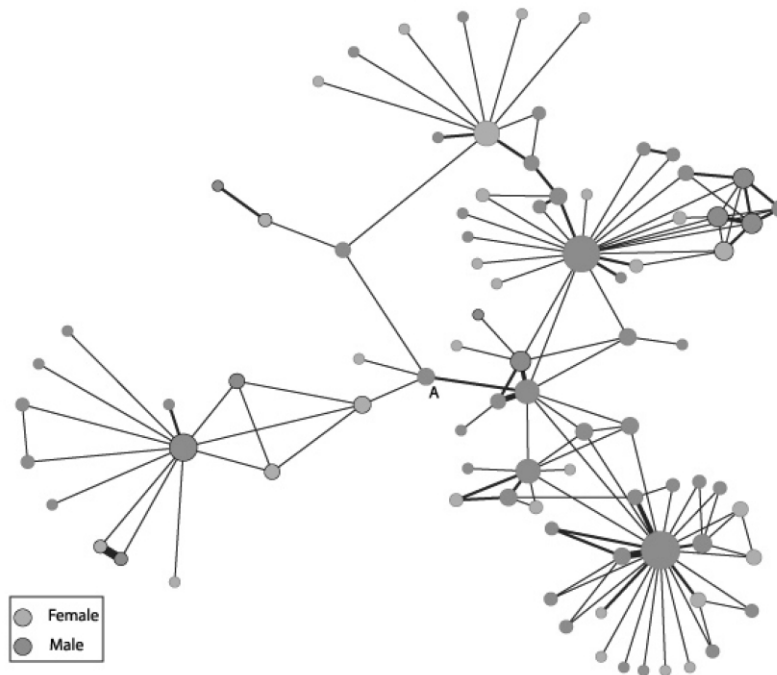



FIG. 4.—Sociology coauthor selection; nodes colored by sex, size proportional to degree

DYNAMIC NETWORK MOVIES

Perhaps the theoretically most promising advance for dynamic network visualization lies in using node movement to map changes in the underlying network structure. Substantively we are often interested in endogenous network processes. Examples include identifying how conversations build in classrooms to transform an ordered teaching environment into a cacophonous menagerie, the development of powerful positions through business exchange, or the endogenous emergence of social structure from actor-based interaction rules. In all of these cases, we suspect that the structure at time t influences the structure at time $t + 1$ in systematic ways, and being able to visualize the transformation of the structure can help identify the mechanisms through which such changes occur (McFarland and Bender-deMoll 2003).

Requirements for Meaningful Dynamic Layouts

All of the aesthetic requirements for static graphs apply to dynamic graphs. In addition, however, a number of factors are unique to dynamic




MOVIE 1.—Online

representations that need to be addressed. An animation or interpolation technique is needed to create meaningful movement between temporally adjacent network slices. Most useful for this is a sinusoidal animation technique that gradually interpolates the position of a node from one resting position to the next. This interpolation is what helps the eye follow changes in the graph structure over time.

Given the ability to animate change in the network structure, one must then assure that the layout at time $t + 1$ is linked to the layout at time t to avoid meaningless movement in the graph. While simple on its face, the separate application of standard layout algorithms to each time slice will rarely give a satisfactory result. Instead, as network layouts usually have no inherent coordinate axes, the entire graph tends to “rotate” and “flip” in the display space. A partial solution to the problem of spurious movement rests in developing an adequate starting position or “anchor” for the network that results in a meaningful orientation for the graph. The anchor choice is not theoretically neutral, as it will affect the resulting layout. Below we identify multiple anchor possibilities and discuss the implication of each.

In addition to providing meaningful movement linked to relational change, a number of additional features of nodes and relations should be temporally variable as well. A successful layout should be able to accommodate changes in relational strength, type, and valence (expressed as edge attributes), as well as changes in node attributes (expressed as shape, color, and size of nodes), so that temporal changes in attributes will also be apparent. Finally, while static graphs are usually agnostic with respect to relational direction, in a dynamic setting we might want to distinguish movement based on tie direction.



MOVIE 2.—Online

Anchoring Temporal Networks

The first layout possibility is really not an anchor but the use of *random starting positions*, which we term “random anchoring.” Points are initially distributed randomly throughout the x - y plane, and the algorithm uses these initial starting positions. In general, a random anchoring is effective only for determining the initial (time 1) layout, and using a new random set of starting coordinates at each observed time point usually results in meaningless motion. However, one can test (qualitatively at least) the effect of starting position on eventual layout by running the same figure layout multiple times with new starting values.

Alternatively, one can use a *constant fixed position* as the starting anchor for all times, resulting in a “fixed anchor.” The most common such anchor might be a simple circle. A second common starting point would be based on some function of the overall graph structure. For example, one might build a cumulative graph as the sum of all interaction over the observation period and generate initial layout positions from this aggregate graph. Or, one might use a meaningful nonrelational distance metric for the starting conditions, such as the latitude and longitude of cities in a trade network or the seating chart in a classroom network.

In any such fixed anchor case, movement from time 1 to time 2 reflects differences in the structure of relations between time 1 and time 2, since starting positions are “held constant” across graph observations, which greatly reduces superfluous movement. Substantively, using the circle as a fixed start point can result in systematic distortion in the overall display fit (if, e.g., nodes that are often connected to each other are placed on opposite sides of the starting circle), though *movement* will still be consistently related to graph changes. Using an aggregate graph layout position can sometimes lend a bit more consistency, as each observation window starts “near” the bulk of the structure observed for the whole

period. In practice, we have found these types of starting values most effective for the metric MDS and peer influence algorithms.

Finally, an obvious anchor for a network layout at time t is the graph layout at time $t - 1$. We term this a “chaining anchor,” and it fits well, substantively, with our intuitive notion of a network movie, as we are literally plotting the change in node position from time t to time $t + 1$. Substantively, this model uses only the information embedded in prior positions for graph layout. As we show below, the chaining anchor seems to work best with force-directed layout algorithms, particularly when network change is small, as starting from $t - 1$ coordinates helps the algorithm find an optima at t that is geometrically close. But when used with the metric MDS or peer influence algorithms, we often see a network “inertia” effect, where nodes quickly converge on a very small portion of the total display space, resulting in a largely unreadable movie.

MOVIE EXEMPLARS

Data

In the examples that follow, we use data from three sources chosen to reflect a wide range of potential applications, moving from theoretical a priori simulations, to a well-known classical dynamic network, and finally to complex, multiple relations, streaming interaction data within classrooms. Below we first describe the three data sources and then describe the dynamic movies for each.

Simulating Social Balance

Social balance theory encapsulates the folk notion that “a friend of a friend is a friend,” suggesting that people avoid friendships where their friends are not friends with each other and form friendships when others relink the pair. Primarily a theory of relational change, social balance theory is a clear candidate for exploration with dynamic tools. The essential prediction of much of the work on balance theory is that through a relational adjustment process, ordered social structure can emerge entirely endogenously (Davis 1963, 1970; Davis and Leinhardt 1972; Doreian et al. 1996; Johnsen 1985). Part of the power of this model rests in the understanding that global implications follow from local relational action. The simulation starts with a simple random network of 45 actors who each nominate (on average) four other people. At each of 200 iterations, five randomly chosen nodes evaluate their local network with respect to transitivity, intransitivity, and reciprocity; nominations are changed if doing so increases the comfort of the actors’ overall position with respect

to these characteristics. Actors favor relations that are transitive, avoid those that are intransitive, seek to reciprocate nominations, and avoid making long-term asymmetric nominations (Gould 2002). In addition, actors have a small preference for current ties over changing ties and do not have perfect vision, but instead evaluate the returns to changes in their local network with a small amount of random error.¹¹ This model is useful for demonstrating the endogenous emergence of order from randomness and demonstrates how seemingly stable summary statistics on one network dimension can mask significant structural change on other dimensions, highlighting the holistic-view payoff to this technique. For the purposes of the movie, we sample the network at every other iteration, resulting in 100 discrete images of the network.

Newcomb Fraternity

Newcomb's fraternity is among the best-known dynamic data sets in common usage. We use a version of this data that comes standard with UCINET (Borgatti et al. 1999). The data consists of each student providing friendship rankings for every other student in the fraternity. Following Doreian et al. (1996), we recode the original rank data so that the top four positive ties are retained as "friends," but use the original rank data for summary transitivity measures. These data provide a real-world complement to questions of social balance raised in the simulation above.

McFarland Classrooms

A third data set stems from McFarland's repeated observations of social interactions in over 150 high school classrooms during the 1996–97 school year. We show dynamic network representations of social interaction from two of these classes below. The first class (class 173) is an accelerated trigonometry class at a magnet high school. It is composed of tenth graders who are tightly controlled by an authoritarian male teacher. The second class (class 182) is an honors algebra 2 class at the same magnet high school. It is composed of mostly tenth graders (light gray), but also a few eleventh graders (dark gray), and it is taught by a progressive but bumbling male teacher (McFarland 2001).

The data on classroom interactions consists of streaming observations of conversation turns. The conversation turns were recorded as pairs of senders and receivers and for types of content. Speakers were viewed as directing their communication in one of two fashions: (1) indirect sound-

¹¹ The simulation is implemented in SAS IML (interactive matrix language) and is available from the first author.

ings, such as lectures (where a teacher addresses all students); and (2) direct interactions that are focused on particular others. Each type of directional speech is viewed as having different forms of network reception—indirect speech reaches all bystanders as passive hearers, and directed speech reaches the focal person as an active coauthor of the conversation (Goffman 1981).

Two types of interaction are found to prevail in classroom contexts: task and sociable (McFarland 2003). Task interactions are those behaviors that pertain to the ongoing teacher-prescribed task (content is academic). In contrast, sociable interactions concern everyday concerns of adolescents' social lives, such as parties, dating, social outings, plans, and so on. While the content is the key distinction, it is often the case that these speech acts are distinguishable in style as well, where sociable behaviors are more playlike, fast paced, and free than the more constrained academic styles of speech during lessons (Cazden 1988). Last, observations also recorded when task and sociable forms of interaction were laminated with evaluative meaning. Such evaluations were seen as being either positive or negative—either giving praise or attempting a reprimand (Ridgeway and Johnson 1990).¹²

RESULTS

Example 1: Dynamic Order from Randomness

Our first example is of the simulated balance process. Figure 5 presents summary statistics over time, including the reciprocity and transitivity rates and the proportion of arcs that change from iteration to iteration.

Balance theory predicts that people should adjust their relations until they reach largely consonant friendship groups, at which point the network should stabilize around a pattern where everyone's ties are (largely) transitive. The simulation summary evident in figure 5 suggests this is essentially the case, as transitivity rises steadily for the first 100 iterations or so, then largely stabilizes after that. Reciprocity quickly rises early in the simulation, then falls briefly (around iteration 20), then continues to

¹² This coding method was piloted in two prior studies by McFarland (2001) and had a high degree of accuracy in more teacher-centered classroom lessons due to the turn-taking sequential nature of discourse. In more chaotic classrooms, simultaneous turns at talk often prevented the observer from acquiring perfect accuracy. However, some record was made of when such diminished accuracy was acquired. The two class periods used in this paper were considered to have a high degree of accuracy in their coding. For the movies included below, we use 2.5-minute moving windows that overlap in incremental shifts of 0.5 minutes. We believe this method best captures the fluidity of interaction patterns but diminishes the amount of artificial fluctuation across frames so that some continuity is had.

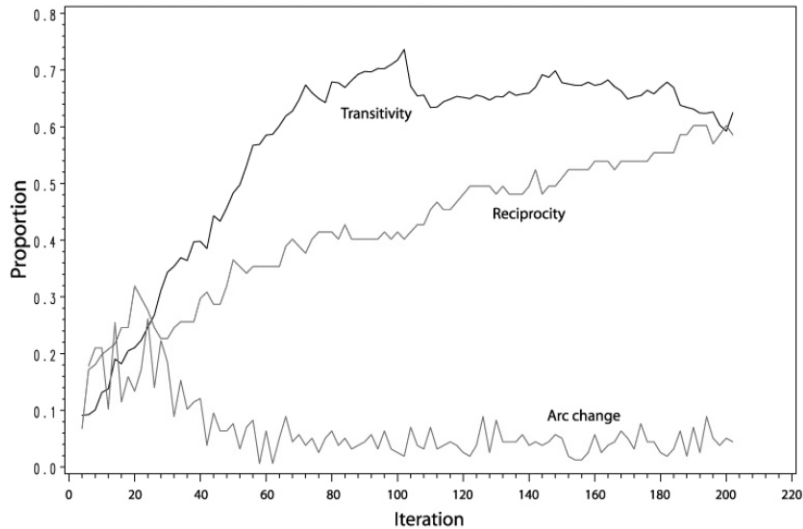



FIG. 5.—Reciprocity and transitivity in a dynamic balance simulation

rise steadily throughout the simulation.¹³ Movie 3 presents a dynamic visualization of this simulation. In this figure, blue ties represent asymmetric nominations and green ties represent symmetric relations.

The movie shows the emergence of structure out of randomness. The initial stages of the movie show very little structure. Instead, actors are simply “milling around”—as they adjust relations in response to the initial random tie allocation. After about 50 iterations, two patterns emerge. First, a set of initial “protogroups” forms around four-person cliques. These small protogroups form the structural kernels that future relations orient around as time progresses. Second, a small number of “stars”—people receiving a relatively large number of asymmetric ties—become evident. Star status proves somewhat unstable, however, as the search for reciprocity means that nodes only nominate stars for a short time. However, as a counterforce, the push for transitivity means that connected pairs will nominate thirds in concert, resulting in a continuous existence of stars in the setting. About halfway through the simulation, it becomes

¹³ The network movie provides a nice explanation for this pattern. The drop in reciprocity seems to occur because the initial rise in reciprocity leads to a large number of intransitive triads. As the number of intransitive triads increases, the transitivity-seeking pressure seems to overpower reciprocity seeking, and adjustments are made to build transitive ties, and then within the transitive sets, reciprocity increases. This pattern is consistent with theoretical and empirical expectations on real networks (Sorensen and Hallinan 1976).



MOVIE 3.—Online

clear that the network is no longer a purely random collection of nodes, but instead a loose federation of groups weakly connected through asymmetric ties. At this point, the group-group structure becomes interesting, as high-degree actors change asymmetric ties between groups, changing the relative links of groups to the remaining network. Here we see the component structure change, as groups break away from each other into smaller collections.

Throughout the remainder of the movie, the structure never crystallizes, but instead “bubbles” as nodes form links between different groups. Nodes make nominations to others in the setting that are then built on by others in response to the changing conditions. The resulting temporal network story, then, becomes one of changes building on change, and a group-level dance that emerges as some group members form ties to others. That some of these changes are, surely, the result of the random component built into the simulation is less interesting than the fact that such random events can have systematic effects in the structuring of others’ responses.

While not intended to model any particular real setting, the image one gets from the resulting movie might well mirror the development of conversation groups in a cocktail party, as sets of people mingle and eventually find those they wish to speak with on a more engaging basis. The simulation also fits well with recent work on the systematic effects of a small number of random ties (Newman 2000; Watts and Strogatz 1998).

Technically, we consider this a successful visualization because almost all of the movement evident in the graph results directly from interpretable changes in underlying network structure. The visualization used the KK layout algorithm and a chain-based anchor. (For alternative movies using different layout algorithms, see app. B).

Example 2: Multiple Visions of Newcomb's Fraternity

The substantive story usually told about Newcomb's fraternity is one of structural convergence as a group of new college students meet and form friendships (Doreian et al. 1996; Newcomb 1961; White et al. 1976). This description of the data rests on evaluations of summary measures (or aggregate blockmodels) of the network. However, the "convergence" story of the Newcomb data is not uncontroversial since up to 20% of the ties change even in the last few weeks of observation. Finding local fluidity and global stability (as measured in blocks or summary statistics) suggests that change moves through structurally equivalent actors. Figure 6 plots the reciprocity and transitivity rate over the 15 observations.

Figure 6 plots the change in reciprocity and transitivity, using the coding scheme described in Doreian et al. (1996). The figure suggests, and many have interpreted these data as showing, that the network reached a particular level of reciprocity early, while transitivity increased throughout the 15-week observation period. However, when combined with the arc change information, one could argue that reciprocity never converges in any meaningful sense but instead fluctuates substantially over the entire observation period. Similarly, transitivity increases steadily, and the change is more than you would expect by chance (Doreian et al. 1996, p.124), but the degree of change is not dramatic (moving from 0.75 to 0.80).¹⁴ How does our image of the Newcomb data differ when we view changes in the entire network simultaneously? In the movies that follow, reciprocated nominations are colored green and given a width equal to the average of the two arcs. Asymmetric ties are blue, with width equal to the strength of the nomination.

Movie 4 presents a KK algorithmic chain-anchor layout version of the Newcomb fraternity network, matching the version used in the balance simulation above. The first impression of this movie is that the network seems to change a great deal over time. This is partially due to the fact that, unlike the simulation example, the Newcomb data were collected in discrete waves, so there are large numbers of edge changes between subsequent networks. Close examination suggests that many nodes are moving even when their relational ties do not change (though the ties of the people they are tied to change, and thus the whole system should change shape). Some of this movement, however, seems excessive and hard to follow. Consider instead movie 5, which uses the FR layout algorithm.

Movement in movie 5 is more subdued, and the pattern of interaction is somewhat easier to follow. For example, one can see that nodes 1, 6,

¹⁴ The transitivity measure used here is based on a sorted rank model.





FIG. 6.—Reciprocity, transitivity, and change in Newcomb's fraternity

8, and 13 remain strongly connected to each other throughout the observation period, occupying a small cluster at the right of the graph. Nodes 7, 12, and 4 form a cluster early in the group's history, but node 4 then breaks with this group at about week 8, instead nominating nodes 17 and 2. In general, there is still a great deal of movement in this graph, and the large number of asymmetric ties suggests that we might gain some insight by using a layout method that accounts for this asymmetry.

Movie 6 uses the peer-influence layout algorithm, with a cumulative graph anchor (the starting position is based on a KK layout of the average positive tie value between every pair cumulated over the entire observation period). Here we immediately see a quick break between those embedded in largely symmetric relations and those hanging on to the edge of the structure. Nodes 10 and 15, for example, quickly emerge as nodes on the edge of the social structure. While they nominate each other symmetrically early in the observation period, they lose interest in each other by the end. Neither node receives top-five nominations from any other node in the network. Their nominations to others seem to dance around the graph, never resting for long on a single person. The group-structure dynamics also become clearer, as node 17 seems to be a popular actor bridging the cluster formed around nodes 1, 6, 8, and 13 and a smaller (and less stable) cluster surrounding nodes 7 and 4.

Hence, the movies suggest that the overall structure does not converge on a single form, and that the process of change is heterogeneous, with

Dynamic Network Visualization



MOVIE 4.—Online


some actors forming stable relations while others dance between friends throughout the observation period. This insight suggests disaggregating the network change statistics evident in figure 6, which we do in figure 7.

Figure 7 presents the results of a cluster analysis on the sequence of network change scores (calculated as the correlation between each person's time t and time $t - 1$ nomination vector, to retain the full range of rank data). A three-cluster solution is instructive. Two groups follow a simple convergence story, with their nominations getting progressively more stable as time passes. The first of these groups (group 1 in fig. 7) has seven members, including the cluster at the right of the movie (1, 6, 13, 8), and presents a gradual convergence of nomination patterns, while the second (group 2, with six members) does not converge on stable nomination patterns until week 5. Finally, group 3 (with four members, including the hanger-on nodes 10 and 15) never seems to settle on a particular nomination pattern, but changes nominations steadily over the observation period.

The substantive advantage of the dynamic movie, in this case, is to identify a level of internal heterogeneity to the network evolution that had never been noticed before, given the focus on global network summary statistics. This type of exploratory interaction between visualization and analysis is one of the best reasons to push progress on dynamic network visualizations.

Example 3: Ritual and Rebellion in the Classroom

Our final example presents a network representation of the streaming interaction data in two classrooms observed by McFarland. Each interaction observed is represented by a directed arc. Statements that were directed “to all” appear as a “star” of links. Direct interactions were weighted as 1 (thick lines), and indirect interactions were weighted as



MOVIE 5.—Online


$1/n$ (where n = class size). Ties are coded as “social,” “task,” and “sanction” to control the red, green, and blue color values of the ties. “Task” interactions appear black, “social” interactions as blue, “praise” as green, and “sanction” as red. The networks are shown in slices of 2.5 minutes duration, so it is possible that multiple ties will exist between a given pair. In all figures, the teacher is represented as a yellow node. Gender is coded in the shape of the nodes (square = male, circle = female) and age in color (tenth graders are light gray, eleventh graders are dark gray).

Class 173: Development of an Interaction Ritual

This example presents an orderly classroom, and movie 7 demonstrates this order. To start, figure 8 presents the summary graph scores for class 173 comparable to those seen in the figures above. The bottom axis reveals half-minute increments of class time so that a value of 70 indicates the thirty-fifth minute of a class period. Viewing the figure, we see an immediate difference with the previous friendship networks: transitivity seems much less important in the conversation networks examined here. On the other hand, reciprocity tends to be quite strong throughout the observation period, though it admits to a high variance over time. Note that none of the summary graph statistics converge toward a particular value, suggesting that the interesting story in this setting deals with the shape of change in the network rather than with the shape of the resultant structure. Relational change remains relatively constant until 35 minutes into the class period (see unit 70) when the level of interaction drops significantly.

The movie for this class shows the transition across two primary activity structures: recitation and group quizzes. Each type of segment calls upon students to organize their behaviors in different sequential and relational patterns. In the opening phase of the class, the teacher does some maintenance, collecting papers and making announcements as students slowly

Dynamic Network Visualization



MOVIE 6.—Online

discard their sociable routines (minutes 0–2). Then the teacher settles the class into a recitation routine concerning homework problems (minutes 3–28). Once the recitation routine is established, a clear ritual structure emerges where the teacher broadcasts questions to the class, and individual students sequentially reply to those questions as they move to the center of the network image (Mehan 1979; Cazden 1988). At around minute 26 (see unit 52) the teacher makes a joke (sociable broadcast) and describes a humorous problem for the students to solve. Here, the task retains the same structure of recitation but has a “blip” of bracketed social activity. In the final phase of the class, there is a clear shift to dyads and triads. Here the teacher asks students to get into assigned pairs and triads to work quietly on their quiz. This final segment shows lower rates of behavior that arise in an intermittent, localized nature. Near the end of class, the teacher makes a final announcement, and then the students break off and leave.

This movie illustrates how interaction dynamics can be represented in a network form that captures both the sequential and relational nature of ritual interaction. Moreover, it illustrates how the coordination and mobilization of students through these routines is accomplished (i.e., through some sanctions and sociable activity, in this case). What we observe are two distinct kinds of relational “dances” and clear switches across them.

Class 182: Ignoring authority

Moving from order to disorder, we now turn to class 182, with the now-familiar summary scores presented in figure 9. As with the previous classroom, we find no general pattern toward convergence in these values. The amount of change and transitivity are generally higher in this classroom. As will become clear, the higher transitivity levels follow from the generally higher level of sociable activities.




FIG. 7.—Nomination stabilization, Newcomb's fraternity

Movie 8 shows the students in honors algebra 2 move through essentially two phases of classroom activity. In the first phase, the teacher lectures on and then invites discussion on student test performance (minutes 0–15). The students ignore and deride this effort, so we see a great deal of sociable interaction and conflict. About halfway through this shortened class period (to make time for an assembly), the teacher shifts into discussing and modeling homework problems with calculators. This shift offers little change in form, as no one really listens and the sociable routines continue. In the final 10 minutes of class, the teacher allows students to work alone in their groups, and here their sociable interactions can persist unabated (minutes 16–35). Throughout the class period, students stay involved in their social affairs, and the teacher is forced to interrupt them, but still fails to secure their attention.

The point of this movie is to show what disorder and lack of control look like in dynamic form. While the movie shows that different patterns of relations arise with each shift in activity, it does not show as clear a shift as observed in the more controlled class above. This makes sense since the class never really pays attention to the teacher's prescribed scripts for behavior. What we see is a high degree of social cliquing among students that never abates in spite of the teacher's sanctions efforts. It is in the later phases of unabated social activity that we see the most stability in network form.

In comparison to the other movies of discrete time, the movies of class-





FIG. 8.—Reciprocity, transitivity, and change, class 173

room interaction and continuous time seem to be qualitatively different. This may be more of an artifact of the examples and slices of time we chose than of the methods. The simulation on balance and Newcomb's acquaintance story both show a progression in affective ties toward some end state and uses of discrete notions of time. In classroom interactions, the ties are represented in minutes of continuous class time, and since the movie is specific to a single class period, there really is no story of progressive equilibrium in affective ties. Instead, the aim is to mobilize and coordinate collective action in various formats (lecture, recitation, group work, student presentations, etc.) and types of ties (task, social, etc.). Hence, states of equilibrium are specific to tasks like lecture, recitation, and group quizzes. While the shifts in rates of network change somewhat align with changing ritual patterns, it is the movies that offer the most meaningful evidence of change in relational sequence and form.

By using longer time slices (of say, class periods) and relating them over the course of the school year, we may acquire something akin to the pattern of tie formation in the Newcomb study or the balance simulation. A stable pattern of behavior may eventually form in a classroom and act like a central tendency (norm) adopted across tasks. The stability of this pattern may even depend on how close ties reach an equilibrium state of balance and transitivity. Future work will explore this empirical question further.



MOVIE 7.—Online

SUMMARY AND CONCLUSION

This article has presented a first step toward developing a new class of dynamic network visualizations. As with most methodological developments in social networks, the use and development of dynamic network movies is grounded in the substantive need present in empirical research. To build intuition about network dynamics effectively, a fine-grained tool for visualizing relational change is necessary. We have focused on two types of dynamic network visualizations: network flip books and dynamic movies. The flip books are a combination of fixed node layout and dynamic social relations, where nodes remain in a constant position and arcs fill in the holes among these nodes. They are particularly effective at showing how sparse relational structures emerge from temporally disaggregated social interaction. In contrast, fully dynamic network movies allow nodes to move as a function of relational change. We have examined three substantive examples, and in each case the movie makes it possible to observe directly a relational feature that would be lost if we focused only on summary statistics of the complete network.

The theoretical promise of network visualization rests in helping scientists see their data. We hope that this type of visualization will spark theoretical development as people are able to engage their data in new ways. Abbott (2000) makes a similar point, when he argues, “Important general theory always grows out of extensive empirical work; every great sociological theorist has been a datahead” (p. 299). Social network movies allow a limited form of data abstraction and exploration. We are neither embedded directly in the relations themselves (which we could never do due to the constraints of time and place) nor forced to limit our observations to one-dimensional summary statistics that filter out much of the interesting temporal and relational variation in the data. Instead, like much recent work in geography and meteorology, dynamic maps provide a combined synthesis of information, allowing one to view the relevant abstract features of a given interaction system.





FIG. 9.—Reciprocity, transitivity, and change, class 182

Technically, we have described the use and performance of a number of particular layout algorithms for building movies. More detail on constructing such movies can be found in the software description given in appendix A (Bender-deMoll and McFarland 2002). In many ways these algorithms are somewhat crude, so additional and ongoing research is necessary. While we think it is much too early in the development of this form of data visualization to make strong claims for a superior method, a few summary statements of our experience are in order. In general, layout methods should seek to minimize movement such that any change can be directly related to a particular relational event. As with static graph visualization procedures, simplicity is the goal. While judicious use of colors, shapes, and sizes can help, too much of any of these elements leads to a confusing visual cacophony.

With respect to particular graph layout algorithms, we have invested the most time and had the most general success with the KK layout algorithm. It is best at preventing node stacking, making it possible to disentangle relational patterns at any given moment in the movie. KK suffers, however, from superfluous local movement generated by changes in the behaviors of *others*. Because KK pushes away from all nodes simultaneously, and due to details of its optimization procedure, nodes tend to “float” around a space, filling in vacancies left by nodes that are drawn into new spaces. This is evident in the Newcomb fraternity movie and



MOVIE 8.—Online

can be seen as well in the movement of students around the teacher in our classroom movies.

When asymmetric ties form a significant part of the relational story, the peer-influence-based model seems to work well, though we find it very sensitive to initial conditions and anchor points. When used with a constant graph layout anchor, such as in the Newcomb example, movement is often usefully limited to the most direct changes in graph structure. When applied to our classrooms, for example, students tend not to cycle around the initial circle, but remain fixed until they engage in a moving interaction (see app. B, available in the online version of this article, for examples).

The most obvious directions for extending network movies is to move from an exploratory data analysis stage to a confirmatory analytic modeling stage (McFarland and Bender-deMoll 2003). A good deal of headway could be made, for example, by linking network movies to statistical models of network change (Snijders 1998, 1990, 1996), or to models that generate positional confidence intervals for nodes (Hoff et al. 2001). This linkage would help us judge the degree of change in a network. This linkage will be complicated, however, because much of the promise of exploratory analysis is to identify elements from the data that are not customarily included in global models. We suspect, however, that a closer linkage between data and visualization will help us build better statistical models.

APPENDIX A

Introduction to SoNIA Software

SoNIA (Social Network Image Animator) is a Java-based software pack-

age for visualizing networks that change over time.¹⁵ SoNIA is being developed as an open-source freeware project and is available with limited documentation (see <http://sonia.stanford.edu/>). Work on SoNIA began initially out of a desire for a set of visualization tools specifically tailored for dealing with network data that include explicit time information and push the limits of the more traditional single-matrix approach to network data collection. Although time is dealt with to some degree in existing network packages such as Pajek, it seemed useful to focus attention on algorithms and techniques aimed at creating dynamic visualizations or animated “movies” of network change. Also, implementations of the most common layout algorithms rest on assumptions seldom discussed in the network literature. SoNIA attempts to make some of these decisions more explicit, provides some limited criteria for assessing the accuracy of a given layout, and aims to produce layouts with a high degree of comparability and replicability. In addition, SoNIA serves a useful function as a “browser” for time-based network data, and its modular construction allows it to serve as test bed for the development and comparison of network visualization algorithms.

SoNIA has the ability to read in a somewhat limited version of Pajek’s arc-list-based “time network” format, and its own arc-list/spreadsheet style format (.son files) in which tab-delimited columns control node and arc attributes. Once data are loaded into SoNIA, the user has the ability to choose a region of time to examine the duration and offset of the “slices” used to bin the data and the layout algorithm to be employed. Currently implemented are versions of Kamada Kawai, Furchterman-Riengold, Moody’s Peer Influence, metric MDS, file coordinates, and circle. Plans are to include GEM, Davidson-Harel, nonmetric MDS, and several other multidimensional projection techniques. For each layout, SoNIA uses the user’s bin criteria to divide up arc events in regions of time and aggregate them to form network matrices to be fed into the layout algorithm. For most algorithms, these matrices are converted to dissimilarities and symmetrized for use in distance computation. Several options are provided for layout preprocessing (node starting positions, desired scaling, treatment of isolates), postprocessing (rescaling, centering, etc.), as well as parameters specific to each algorithm. For algorithms that use an optimization technique, options are provided to control the “cooling schedule” and stopping times, and feedback is provided about algorithm convergence. Several algorithms (KK and FR) attempt to create layouts in which

¹⁵ SoNIA is currently under development by Dan McFarland and Skye Bender-deMoll. The software, source files, installation instructions, and additional documentation are available at <http://sonia.stanford.edu/>. Comments, suggestions, and programming assistance would be greatly appreciated; please contact skyebend@stanford.edu.

screen distances are proportional to values in the network matrix. For most real-world networks, it is impossible to create a zero-distortion representation. However, the degree of distortion can be estimated using a modified version of Kruskal's stress statistic to compare the matrix and screen distances. This value and a "Shepard's plot" showing the distribution of distortions can be displayed.

The animation procedure for transitioning between layouts is similar to "tweening," in which a user-specified number of frames are displayed showing gradual interpolation of nodes' positions between successive layouts. The transition takes place in "network time," meaning that for real-valued continuous data, node and arc additions and deletions between the networks are shown in sequence as they occur. In addition, the time coordinates for a "render slice" can be entered manually, allowing the user to "browse" the data by displaying any arbitrarily sized region of time. Node and arc attributes such as color, size, shape, labeling, and position can be controlled from the input file (or to a limited degree in the program) and are permitted to vary over time. SoNIA takes advantage of some of Java's 2D capabilities, which means graphics can be anti-aliased, and arcs can have some degree of transparency (useful in situations where multiplex ties exist). SoNIA has the ability to save out QuickTime movie files of the network animation, a log file describing the sequence of operations and parameter settings used to create the layout, and a series of "slice" matrices for analysis in other SNA programs.

Because SoNIA is coded in Java, it is cross-platform compatible, and there are many possibilities for extending its capabilities and integrating it with other packages. SoNIA's internal structure and Java's modular object-oriented design make it relatively easy to add additional layout techniques and explore modifications. In addition, source files are available for inspection if the details of algorithm behavior are not clear from documentation. However, a great deal of additional work is needed, particularly in the area of extending the design to facilitate communication with other open-source projects and network software.

REFERENCES

- Abbott, Andrew. 1997. "Of Time and Space: The Contemporary Relevance of the Chicago School." *Social Forces* 75:1149-82.
- . 2000. "Reflections on the Future of Sociology." *Contemporary Sociology* 29: 296-300.
- Abel, Tom, Greg L. Bryan, and Michael L. Norman. 2000. "The Formation and Fragmentation of Primordial Molecular Clouds." *Astrophysical Journal* 540:39-44.
- Batagelj, Vladimir, and Andrej Mrvar. 2001. *PAJEK*, ver. 91. <http://vlado.fmf.uni-lj.si/pub/networks/pajek/>
- Bearman, Peter S., and K. Everett. 1993. "The Structure of Social Protest." *Social Networks* 15:171-200.

Dynamic Network Visualization

- Bearman, Peter S., Robert Faris, and James Moody. 1999. "Blocking the Future: New Solutions for Old Problems in Historical Social Science." *Social Science History* 23: 501–33.
- Bearman, Peter S., James Moody, and Katherine Stovel. 2004. "Chains of Affection." *American Journal of Sociology* 110:44–91.
- Bender-deMoll, Skye, and Daniel A. McFarland. 2002. *SoNIA* (Social Network Image Animator). See <http://sonia.stanford.edu/>
- Borgatti, S., M. G. Everett, and L. C. Freeman. 1999. *UCINET V for Windows: Software for Social Network Analysis*, ver. 5.2.0.1. Natick, Mass.: Analytic Technologies.
- Brandes, Ulrik, Jorg Raab, and Dorothea Wagner. 2001. "Exploratory Network Visualization: Simultaneous Display of Actor Status and Connections." *Journal of Social Structure* 2 (4). <http://www.cmu.edu/joss/content/articles/volume2/BradesRaabWagner.html>
- Cazden, Courtney B. 1988. *Classroom Discourse: The Language of Teaching and Learning*. Portsmouth, N.H.: Heineman.
- Chase, Ivan D. 1980. "Social Process and Hierarchy Formation in Small Groups: A Comparative Perspective." *American Sociological Review* 45:905–24.
- Coleman, James S. 1961. *The Adolescent Society*. New York: Free Press.
- Collins, Randall. 1981. "On the Microfoundations of Macrosociology." *American Journal of Sociology* 68 (5): 984–1014.
- Cyram. 2004. *NetMiner II*, ver. 2.5.0a. Seoul: Cyram.
- Danto, Arthur C. 1985. *Narration and Knowledge*. New York: Columbia University Press.
- Davidson, R., and D. Harel. 1996. "Drawing Graphics Nicely Using Simulated Annealing." *ACM Transactions on Graphics* 15:301–31.
- Davis, James A. 1963. "Structural Balance, Mechanical Solidarity, and Interpersonal Relations." *American Journal of Sociology* 68:444–62.
- . 1970. "Clustering and Hierarchy in Interpersonal Relations: Testing Two Graph Theoretical Models on 742 Sociomatrices." *American Sociological Review* 35: 843–51.
- Davis, James A., and Samuel Leinhardt. 1972. "The Structure of Positive Relations in Small Groups." Pp. 218–51 in *Sociological Theories in Progress*, vol. 2. Edited by Joseph Berger, Morris Zelditch, and Benedict Anderson. Boston: Houghton Mifflin.
- Doreian, Patric, Roman Kapuscinski, David Krackhardt, and Janusz Szczypula. 1996. "A Brief History of Balance through Time." *Journal of Mathematical Sociology* 21 (1–2): 113–31.
- Doreian, Patrick. 1980. "On the Evolution of Group and Network Structure." *Social Networks* 2:235–32.
- . 1986. "On the Evolution of Group and Network Structure II: Structures within Structures." *Social Networks* 8:22–64.
- Emirbayer, Mustafa. 1997. "Manifesto for Relational Sociology." *American Journal of Sociology* 103:281–317.
- Emirbayer, Mustafa, and Jeff Goodwin. 1994. "Network Analysis, Culture, and the Problem of Agency." *American Journal of Sociology* 99:1411–54.
- Freeman, Linton C. 2000a. "Visualizing Social Groups." In *Proceedings of the Section on Statistical Graphics*. American Statistical Association, <http://moreno.ss.uci.edu/groups.pdf>
- . 2000b. "Visualizing Social Networks." *Journal of Social Structure* 1 (1). <http://www.cmu.edu/joss/content/articles/volume1/Freeman.html>
- Freeman, Linton C., Cynthia M. Webster, and Deirdre M. Kirke. 1998. "Exploring Social Structure Using Dynamic Three-Dimensional Color Images." *Social Networks* 20:109–18.

American Journal of Sociology

- Friedkin, Noah E. 1998. *A Structural Theory of Social Influence*. Cambridge: Cambridge University Press.
- Fruchterman, Thomas J. J., and Edward Reingold. 1991. "Graph Drawing by Force-Directed Placement." *Software—Practice and Experience* 21:1129–64.
- Goffman, Erving. 1967. *Interaction Ritual: Essays on Face-to-Face Behavior*. Garden City, N.Y.: Anchor Books.
- . 1981. *Forms of Talk*. Philadelphia: University of Pennsylvania Press.
- Gould, Roger. 1995. *Insurgent Identities: Class, Community, and Protest in Paris from 1848 to the Commune*. Chicago: University of Chicago Press.
- . 2002. "The Origins of Status Hierarchies: A Formal Theory and Empirical Test." *American Journal of Sociology* 107:1143–78.
- Handcock, Mark, and Martina Morris. 1999. *Relative Distribution Methods in the Social Sciences*. New York: Springer.
- Hinde, R. A. 1976. "Interactions, Relationships and Social Structure." *Man* 11 (1): 1–17.
- Hoff, Peter D., Adrian E. Raftery, and Mark S. Handcock. 2001. "Latent Space Approaches to Social Network Analysis." Working Paper no. 17. University of Washington, Center for Statistics and the Social Science.
- Johnsen, Eugene C. 1985. "Network Macrostructure Models for the Davis-Leinhardt Set of Empirical Sociomatrices." *Social Networks* 7:203–24.
- . 1986. "Structure and Process: Agreement Models for Friendship Formation." *Social Networks* 8:257–306.
- Kamada, Tomihisa, and Satoru Kawai. 1989. "An Algorithm for Drawing General Undirected Graphs." *Information Processing Letters* 31:7–15.
- Kretzschmar, Mirjam, and Martina Morris. 1996. "Measures of Concurrency in Networks and the Spread of Infectious Disease." *Mathematical Biosciences* 133: 165–95.
- Laumann, Edward O., Peter Marsden, and David Prensky. 1983. "The Boundary Specification Problem in Network Analysis." In *Applied Network Analysis: A Methodological Introduction*, edited by Ron S. Burt and M. J. Minor. London: Sage Publications.
- Laumann, Edward O., and Yoosik Youm. 1999. "Race/Ethnic Group Differences in the Prevalence of Sexually Transmitted Diseases in the United States: A Network Explanation." *Sexually Transmitted Diseases* 266:250–61.
- Leenders, Roger Th. A. J. 1996. "Longitudinal Behavior of Network Structure and Actor Attributes: Modeling Interdependence of Contagion and Selection." Pp. 165–84 in *Evolution of Social Networks*, edited by P. Doreian and Frans N. Stokman. New York: Gordon & Breach.
- Leifer, Eric M. 1988. "Interaction Preludes to Role Setting: Exploratory Local Action." *American Sociological Review* 53:865–78.
- McFarland, Daniel. 2001. "Student Resistance: How Formal and Informal Organization of Classrooms Facilitates Student Defiance." *American Journal of Sociology* 107 (3): 612–78.
- . 2003. "Why Work When You Can Play? Dynamics of Formal and Informal Organization in Classrooms." In *The Social Organization of Schooling: A Tribute to Charles E. Bidwell*, edited by Barbara Schneider and Larry Hedges. Thousand Oaks, Calif.: Sage Publications.
- McFarland, Daniel, and Skye Bender-deMoll. 2003. "Classroom Structuration: How Interaction Patterns Get Reproduced and Transformed." Working Paper. Stanford University.
- McGrath, Cathleen, Jim Blythe, and David Krackgardt. 1997. "The Effect of Spatial Arrangement on Judgments and Errors in Interpreting Graphs." *Social Networks* 19:223–42.
- Mehan, Hugh. 1979. *Learning Lessons: Social Organization in the Classroom*. Cambridge, Mass.: Harvard University Press.

Dynamic Network Visualization

- Moody, James. 2001a. "Peer Influence Groups: Identifying Dense Clusters in Large Networks." *Social Networks* 23:261–83.
- . 2001b. "Race, School Integration, and Friendship Segregation in America." *American Journal of Sociology* 107 (3): 679–716.
- . 2002. "The Importance of Relationship Timing for Diffusion." *Social Forces* 81:25–56.
- . 2003. "The Structure of a Social Science Collaboration Network: Disciplinary Cohesion from 1963 to 1999." *American Sociological Review* 69:213
- Morris, Martina. 1993. "Epidemiology and Social Networks: Modeling Structured Diffusion." *Sociological Methods and Research* 22:99–126.
- Morris, Martina, and M. Kretzschmar. 1997. "Concurrent Partnerships and the Spread of HIV." *AIDS* 11:641–48.
- Nadel, S. F. 1955. *The Theory of Social Structure*. London: Cohen & West.
- Nakao, K., and A. K. Romney. 1993. "Longitudinal Approach to Subgroup Formation: A Re-Analysis of Newcomb's Fraternity Data." *Social Networks* 15:109–31.
- Newcomb, Theodore M. 1961. *The Acquaintance Process*. New York: Holt, Rinehart & Winston.
- Newman, M. E. J. 2000. "The Structure and Function of Complex Networks." *SIAM Review* 45:167–256.
- Padgett, John F., and Christopher K. Ansell. 1993. "Robust Action and the Rise of the Medici, 1400–1434." *American Journal of Sociology* 98:1259–1319.
- Powell, Walter W., Douglas R. White, and Kenneth W. Koput. 2005. "Network Dynamics and Field Evolution: The Growth of Interorganizational Collaboration in the Life Sciences." *American Journal of Sociology* 110 (4). In this issue.
- Ridgeway, Cecilia, and Cathryn Johnson. 1990. "What Is the Relationship between Socioemotional Behavior and Status in Task Groups?" *American Journal of Sociology* 95 (5) :1189–1212.
- Roy, William G. 1983. "The Unfolding of the Interlocking Directorate Structure of the United States." *American Sociological Review* 48:248–57.
- Snijders, Tom A. B. 1990. "Testing for Change in a Digraph at Two Time Points." *Social Networks* 12:163–74.
- . 1996. "Stochastic Actor-Oriented Models for Network Change." *Journal of Mathematical Sociology* 21 (1–2): 149–72.
- . 1998. "Models for Longitudinal Social Network Data." Manuscript. University of Groningen.
- . 2001. "The Statistical Evaluation of Social Network Dynamics." *Sociological Methodology* 361–95.
- Sørensen, Aage B., and Maureen T. Hallinan. 1976. "A Stochastic Model for Change in Group Structure." *Social Science Research* 5:43–61.
- Suitor, J. J., Barry Wellman, and David L. Morgan. 1997. "It's About Time: How, Why and When Networks Change." *Social Networks* 19 (1): 1–7.
- Watts, Duncan J., and Steven H. Strogatz. 1998. "Collective Dynamics of 'Small-World' Networks." *Nature* 393:440–42.
- Weesie, Jeroen, and Henk Flap. 1990. *Social Networks through Time*. Utrecht: ISOR.
- White, Harrison C., Scott A. Boorman, and Ronald L. Breiger. 1976. "Social Structure from Multiple Networks I." *American Journal of Sociology* 81:730–80.
- Whyte, William Foote. 1943. *Street Corner Society: The Social Structure of an Italian Slum*. Chicago: University of Chicago Press.
- Zeggelink, Evelien P. H., Frans N. Stokman, and Gerhard G. Van De Bunt. 1996. "The Emergence of Groups in the Evolution of Friendship Networks." *Journal of Mathematical Sociology* 21:29–55.