

Controlling Across Complex Networks

Emerging links between networks and control

A. Clauset^a, H.G. Tanner^b, C.T. Abdallah^c, R.H. Byrne^d

^a*Santa Fe Institute, 1399 Hyde Park Rd. Santa Fe, NM 87501, USA.*

^b*Mechanical Engineering Department, MSC01 1150, 1 University of New Mexico, Albuquerque, NM 87131-0001, USA.*

^c*Electrical & Computer Engineering Department, MSC01 1100, 1 University of New Mexico, Albuquerque, NM 87131-0001, USA*

^d*Data Analysis And Exploitation, Department 5535, Sandia National Laboratories, MS 1243, PO Box 5800, Albuquerque, NM 87185-1243I, USA.*

Abstract

This paper discusses the interplay between networks and control systems. As we gain more understanding about the structure and dynamics of physical networks, their effects on the performance of closed-loop control systems, as well as the ability to control such networks, provide fertile areas of research. The paper reviews such research with special emphasis on the connectivity and delays in the information transfer across networks.

Key words: complex networks, networked control systems, delays, connectedness, navigability, efficiency

1 Introduction

Networks are a powerful metaphor for understanding the organization of systems from disciplines as diverse as biology, computer science, physics, and social science. In control systems, communication networks are becoming increasingly pervasive, forcing control engineers to expand their application domain by incorporating the communication infrastructure into their designs,

Email addresses: aaronc@santafe.edu (A. Clauset), tanner@unm.edu (H.G. Tanner), chaouki@ece.unm.edu (C.T. Abdallah), rhbyrne@sandia.gov (R.H. Byrne).

and by considering the impact of link capacity, latency, and packet loss on control systems Zhang et al. (2001); Walsh et al. (2002); Verriest and Egerstedt (2002). Insight is sought to better understand how systems can be controlled across networks, how to design distributed, multi-agent control systems, or to predict when the network's structure gives rise to undesirable network behaviors such as congestion.

Consider for example the system depicted in Figure ??, where a plant is being controlled across a network shared by various systems, computers, and communication devices. From a control perspective, the communication links of Figure ?? are a means of information exchange, which is generally assumed to be instantaneous. The impact of the network's connectivity on the closed-loop system performance is discussed in Jadbabaie et al. (2002); Moreau (2005); Ren and Beard (2005); Olfati-Saber and Murray (2004); Tanner et al. (2003b). Note that this paper focuses on controlling systems across a network, rather than controlling a network (such as congestion or access control). As such, we do not review the vast recent work in Kelly (2001) and elaborated upon in Srikant (2004); Tarbouriech et al. (2004).

The need for new paradigms for control design is evident in large-scale interconnected multi-agent systems. In this class of systems, signals need to flow quickly and efficiently, but interconnected components may not be able to store and manipulate the complete state of the system. While complexity barriers make the design of controllers for high-dimensional systems extremely difficult, the ability to reason about global network properties based on locally available information enables the design of decentralized control laws. When scaled to networked systems with hundreds of thousands of components, decentralized control laws allow us to deal with unrealistic computation, communication, and storage requirements.

In control design, a network model such as a graph is used to enable control, while, in network theory, models of network dynamics and growth are constructed to simulate physical or engineering processes. Despite the use of different analysis tools, network properties such as connectivity, efficiency, and robustness are common to both control and network theory. A question that arises is whether pervasive ideas in network theory can suggest new control design directions.

Network theory provides tools that characterize the growth and topology of distributed networks Faloutsos et al. (1999) in relation to their navigability, congestion, clustering, and robustness to failure Bollobás (1985). For some systems, such as social networks and the World Wide Web, not only do short paths exist between every pair of nodes, but such paths can be found under certain conditions using only local information Milgram (1967); Kleinberg (2000); Clauset and Moore (2003). In this article we review analysis tools used

to study complex networks Newman (2003); Dorogovtsev and Mendes (2003); Watts and Strogatz (1998); Albert and Barabási (2002); Kleinberg (1999), and discuss the possibility of using them to facilitate control design.

This article presents an overview of problems at the intersection of control theory and complex network research. After a brief introduction of the relevant aspects of complex network theory and its methodological differences and similarities to control engineering, we discuss the potential benefits of knowledge transfer between the two fields. Within a brief review of recent results in cooperative control, we show how topological network properties affect control performance and that:

- (1) Increased network connectivity does not necessarily yield robustly connected networks with respect to node failures;
- (2) The structure of sensor networks and their algebraic graph properties determine the performance of distributed estimation;
- (3) Properly interleaving communication and control can protect against the effect of delayed information.

2 Network-theoretic Issues

By considering the network as a communication service, we identify three properties that critically impact the flow of information:

- (1) Connectedness, which expresses the existence of a path between the information transmitter and the information receiver.
- (2) Navigability, quantified by the difficulty of finding a connecting path. Typically, this difficulty depends on whether the path is predetermined, or whether it is discovered in an ad hoc fashion.
- (3) Efficiency, as represented by the latency of each utilized path. This latency, usually a function of the number of hops and the individual link latencies, must be sufficiently small so that bounds on the end-to-end communication delays are respected.

All three properties affect the robustness of a network with respect to node or link failures, as well as the reliability of network protocols with respect to corruption.

2.1 Connectedness

Connectedness is mathematically identified with notion of percolation Bollobás (1985). Percolation theory characterizes how connected clusters in a

random graph aggregate as a function of the edge probability. Given this fixed probability p , percolation can be illustrated as a wildfire, initiated at a source node that spreads across an edge connected to the burning node with probability p . By locating the nodes reached by the process, it is possible to determine whether a path connecting a given pair of nodes exists Bollobás (1985).

In network theory, percolation is typically analyzed in two ways. The constructive approach determines the number of random edges that must be successively added to a collection of disconnected nodes before the vast majority of nodes, termed the *giant component*, are connected. In the destructive approach, edges or nodes are successively removed until the giant component vanishes and most pairs of nodes are no longer connected. Surprisingly, the appearance and disappearance of the giant component can be quite sudden, and is often a genuine phase transition Stanley (1983).

One feature of many real world networks is a power-law degree distribution, in which the probability of a randomly chosen node having k neighbors scales as $P(k) \propto k^{-\alpha}$, where α is the scaling exponent Newman (2003). The ubiquity of the power-law degree distribution motivates the study of graph models that exhibit this feature, but whose topological structure is otherwise random. A network with many redundant paths between all pairs of nodes is obviously very robust to node and edge failures. However, if a minimal fraction of nodes p_c is removed, the giant component vanishes. This disappearance *shatters* the network.

Consider a random graph with a power-law degree distribution where a node fail or is removed with probability p . The after-failure degree distribution $P'(k)$ is given by

$$P'(k) = \sum_{k_0=k}^{\infty} P(k_0) \binom{k_0}{k} (1-p)^k p^{k_0-k}, \quad (1)$$

where k_0 is the degree of a node before failure, k is its degree after failure, and p is the probability of failure. When the scaling exponent α for $P(k_0)$ is larger than 3, (1) is used in Cohen et al. (2000) to show that the critical threshold of p for maintaining the giant component of the network is $p_c \approx 0.99$. In other words, more than 99% of the nodes must fail or be removed before the network shatters. Hence, large random structures are robust to random failures. For finite-size networks, the exact value of p_c is related to the number of nodes n , and approaches 1 as the number of nodes n increases. A recent study shows, however, that the value of p_c can be significantly smaller than 1 for a specific subclass of these graphs Link et al. (2005). In such networks, removing nodes according to the probability $p_c \ll 1$, shatters the network. Therefore, random graphs with a power-law degree distribution exhibit various degrees of robustness to random failures.

Unfortunately, the random removal of nodes is not the only kind of failure

that networks can suffer. For instance, when nodes in a random graph are preferentially removed according to a specified rule, for example removing the 10% of nodes with the highest degree, the network quickly shatters Gallos et al. (2005), Guillaume et al. (2004). More subtle forms of failure, in which some fraction of nodes disobey the network communication protocols, possibly in a malicious way, are considered in the context of peer-to-peer networks Engle and Khan (2006). These Byzantine faults have been extensively studied, and continue to drive much of the research in developing secure distributed communication protocols Pease et al. (1980), Ben-Or et al. (1993).

2.2 Navigability

In a connected network, several paths may link a transmitter to a receiver. The navigability of a network is determined by how easily a connecting path can be found, as well as by how many links or edges such a path contains. The navigability problem may be solved in one of two ways:

- (1) using central authorities, in which the communication path between two nodes is determined by an external source then communicated to the network's routers, and
- (2) using decentralized techniques, in which routing decisions are made independently by network routers, possibly in an ad hoc fashion.

For a static network, namely one for which the number of nodes and the topology are fixed, a central authority is easy to construct. A decentralized navigation approach however, is called for when routers are added or removed from the network. Current standards for routing on Internet-like networks, such as the Internet protocol (IP), the open shortest path first (OSPF) Force (????a) protocol, and the border gateway protocol (BGP) Force (????b), are a mixture of both centralized and decentralized techniques. Each protocol involves an initial consensus phase among the nodes of the network that allows local connectivity information to propagate, until each router constructs its own map of the network for routing packets in the future. After the consensus phase is completed, the routers' maps remain fixed until the local topology changes sufficiently to trigger a new consensus phase.

While standard protocols, such as IP, OSPF, and BGP, perform well for networks that change only occasionally, dynamic networks pose a more challenging problem, since the overhead of reaching consensus must be balanced against the efficiency of the network as a communication medium. An alternative approach is to use decentralized or ad hoc routing strategies, where routing decisions are made on the fly based on the relative position of the current router, the packet's destination, and possibly the current local connectivity.

Navigability of the resulting network requires that short paths between source and destination nodes be easily found in a decentralized way. A network is *efficiently navigable* if the average length of a path T grows sublinearly with the number of nodes in the network, and preferably as a polylogarithm such as $O(\log^2 n)$. Theoretical work in Clauset and Moore (2003) indicates that such decentralized protocols can be developed under reasonable assumptions. A brief description of these results, along with their implications for control systems follows.

Consider the network of Figure ??, which is a lattice with nodes having bidirectional local connections to their nearest neighbors, as well as a single, unidirectional nonlocal connection to specified node. The distance between nodes u and v is evaluated using the Manhattan metric or l_1 metric denoted by $d(u, v)$, and the dimensionality of the lattice is denoted by D (in this example, $D = 2$). The diamonds in Figure ?? define the set of nodes at a fixed distance from the node at the center. If each node (router) forwards packets to its neighbor with the smallest remaining distance to the packet's destination, then this decentralized routing protocol, for this particular topology, guarantees packet delivery in an average of $O(\log^2 n)$ steps Kleinberg (1999). The receiving neighbor is found as follows: first we choose a distance ℓ from the distribution $P(\ell) \sim \ell^{-D}$. The distance $e\ell$ is the distance from the destination node to all potential neighbors of the sending node. Then, out of all the nodes at distance $e\ell$ from the destination node, we choose uniformly at random a receiving node.

To verify the average number of steps required for delivery, we assume that a packet travels in phases and that a phase ends when the remaining distance is halved. Thus, there are at most $\log_2 n$ phases in a network of n nodes. If the distribution of lengths for the nonlocal links is a power law with exponent D , the packet visits a router with a non-local neighbor that is roughly half as distant from the destination after $O(\log n)$ trials. Thus, the expected routing time is $O(\log^2 n)$.

The algorithm presented in Clauset and Moore (2003) constructs the Kleinberg-routable network through a dynamic, decentralized rewiring process. The algorithm assumes that local connections are fixed. Given a source-destination pair (x, y) , a packet is routed according to the current topology, and a time threshold t is chosen uniformly from the interval $[1, d(x, y)]$. If the routing time of the packet T , exceeds the threshold t at a node z that is not the destination, x “rewires” its non-local link so that it terminates at z . In Clauset and Moore (2003), it is empirically shown that this rewiring algorithm produces the power-law link-length distribution $P(\ell) \sim \ell^{-D}$. This then guarantees fast ad-hoc routing over the entire network $T = O(\log^2 n)$, after a modest number of rewiring actions $R \sim n^{1.77}$.

With the availability of global positioning (GPS) systems that provide simple distance measurements, these results may be adapted as a routing protocol for packets on a wireless array of devices. In such cases, local links are either physical connections or low-power broadcast transmissions, and non-local links are occasional high-power broadcast transmissions or unidirectional long-range transmissions.

The development of dynamic and decentralized routing algorithms that guarantee efficient navigability under a variety of assumptions is an active topic of research in network theory. In the ad-hoc routing algorithm presented in Şimşek and Jensen (2005), packets are routed under assumptions about the connectivity of nodes with similar properties (homophily), and the assumption that higher degree nodes are likely to be closer to the target. In the model used in Şimşek and Jensen (2005), it is assumed that each node has a set of attributes, and that nodes are linked to others that are similar to themselves. Thus, a homophily-sensitive algorithm adjusts the routing based on the assumption that a node close to the destination node in their attribute space, is in fact geographically closer to the destination.

2.3 Efficiency

In network theory, *efficiency* is quantified by the cost of a network property as a function of the number of nodes n in this network. Thus in this context, efficiency is related to that of scalability, and the bounds on the related cost are expressed using asymptotic O -notation. Generally, for a property to have a small cost, it should scale sub-linearly, and ideally as a poly-logarithm $O(\log^k n)$. For example, the decentralized routing algorithm Kleinberg (1999) described in the previous section, guarantees that the average number of intermediate nodes that a message passes through is $O(\log^2 n)$. On the other hand, if a property needs to be true for the largest possible portion of the network, then it must scale as a constant fraction of the nodes $O(1)$, and ideally to be $1 - o(1)$. For connectedness, the question of efficiency boils down to determining what fraction of the network remains connected, after a fraction of the nodes is removed. In this context, Gallos et al. (2005) shows that random networks with a power-law degree distribution, are increasingly more efficient at guaranteeing connectedness under random failures as the network grows.

These definitions of efficiency are highly applicable to random graph models used in network theory. Efficiency, however, is being understood quite differently in control. Depending on optimality criteria in a given control application, efficiency may be related to the input signal strength, or the output rise and settling times. Bridging this gap, and producing network theoretic results for control performance specifications may be a fertile area of cross-

collaboration between the two fields.

3 The Control Designer's Point of View

Networked control system (NCS) applications such as teleoperation and robot formation control, require measurement and control signals to travel across communication networks. Even when the distance traveled is short (as in the case of a modern car or a smart house), a general purpose communication network introduces new issues into the feedback loop, such as time-varying delays, and the potential loss of information. While some communication applications may suffer from the same limitations, a feedback control system is especially vulnerable, not only to the unavailability of sensory information and control signals, but also to their timing. In particular, in a NCS, the issues of connectedness, navigability, and efficiency of message propagation manifest themselves as described in the following sections.

3.1 Connectivity, dropped packets, and lost links

From the perspective of control design for networked control systems, connectedness (or connectivity) expresses the ability of two systems to communicate information and actuation signals over the network connecting them. Connectivity is therefore related to the existence of a network path from any node u to any other node v . In recent studies linking the dynamics of the networked systems to the connectivity properties of the network, certain graph algebraic properties of the latter seem to be pervasive. In Jadbabaie et al. (2002), the dynamics of the networked system is formally related to the Laplacian matrix of the graph representing the network of interconnections between the system components. The researchers in Jadbabaie et al. (2002) established algebraic conditions for the matrices related to the graph, to guarantee that all interconnected subsystems asymptotically reach consensus over a quantity of interest. The consensus is reached by when each subsystem replaces the value of its quantity of interest by the average value of its network's neighbors. For this consensus update algorithm to be asymptotically stable, i.e. for all individual quantities of interest to asymptotically converge to the same value, the communication network should be connected. In algebraic graph theoretic terms, connectivity is quantified by means of the second smallest Laplacian eigenvalue, also known as the *algebraic connectivity* of the graph. In Tanner et al. (2003a) it was shown that if connectivity were permanently lost, stability can no longer be guaranteed. If connectivity is however regained across a sequence of compact intervals $[t_i, t_{i+1})$, reference Jadbabaie et al. (2002) demonstrates that consensus stability may still be reached.

A network may not have a constant topology when communication links are dynamically established and lost Tanner et al. (2003b), Jadbabaie et al. (2002), Olfati-Saber and Murray (2004). Physical ad-hoc networks are typically modeled by nearest-neighbor type graphs, where nodes are distributed uniformly at random over a certain area, and are assumed connected if nodes are within a certain distance r_0 from each other. Thus nodes u , and v , are connected if $|u - v| < r_0$, where $|u - v|$ denotes the Euclidean distance between them. The question of whether such an ad-hoc network is connected or not does not have a deterministic answer, especially when the number of nodes grows very large. Results in this area are typically asymptotic and probabilistic in nature. Whether the network is connected is thus given with a certain probability, which usually relates to the minimum degree of the nodes in the network, as exemplified in Xue and Kumar (2004), or to the minimum communication range r_0 Bettstetter (2002), Santi and Blough (2003). In Xue and Kumar (2004) it is shown that if each node is connected to less than $0.074 \log n$ other nodes, the network is disconnected with probability one, as the total number of nodes n increases. If, on the other hand, each node has more than $5.1774 \log n$ neighbors, the network is asymptotically connected with probability one when n tends to infinity. In Bettstetter (2002) it is shown that if the network is required to be connected with probability p , the transmission range r_0 must satisfy $r_0 \geq \sqrt{\frac{-\ln(1-p^{1/n})}{\pi\rho}}$, where ρ is the node density in nodes per unit area.

In networks where information flows in a unidirectional manner, directed graphs are used to capture the network topology. For directed graphs we differentiate between strong and weak connectivity, with the former property guaranteeing that a message originating from one node can reach any other node, following paths in the graph that respect the orientation of all edges. The existence of a (directed) spanning tree over the union of the graphs that describe the evolution of the network over time Ren and Beard (2005) however, may be sufficient to ensure asymptotic consensus in the network, provided that the graph switching frequency is bounded, on average. This condition is definitely weaker than strong connectivity, though still stronger than weak connectivity (expressed again by the second smallest eigenvalue) for which edge orientation is irrelevant. The gap between these conditions seems to be the missing piece in a uniform characterization of stability in terms of network topology. Of course, another approach for ensuring stability is to restrict the dynamics, as described in Moreau (2005), Angeli and Bliman (2005).

In most practical models connectivity is binary, that is, two nodes are either connected at a particular time instant, or disconnected. In the former case the second smallest eigenvalue is positive, in the latter it is zero. In order to capture the quality of a communication link, or the cost of broadcasting information from one node to another, weighted graph models may be used. The edge

weights quantify the energy required for a message to be sent over an edge (u, v) , usually expressed as $|u - v|^e$, where $e \geq 2$ is a constant. Weighted graphs are not as well understood as their unweighted counterparts, but connectivity analysis using the second smallest eigenvalue of the (weighted) Laplacian can be extended to this situation as well.

The effect of network topology and connectivity on the performance of cooperative localization algorithms is pointed out in Hidaka et al. (2005), in which a genetic algorithm is used. The genetic algorithm selects network topologies that result in smaller traces of the covariance matrix for the extended Kalman filter (EKF) constructed for the whole networked system. The analysis in Hidaka et al. (2005) suggests that increased connectivity may be beneficial for localization accuracy. Intuitively, “the more sensor links between robotic nodes, the better.” Define the *sensor graph* as one in which nodes are mapped to mobile robots and environment landmarks, and where directed edges denote relative position measurements. While a genetic algorithm favors complete sensor graphs, other approaches may suggest “cheaper” solutions. In the special case where a landmark’s location is accurately known, the expression of the upper right submatrix $P_{rr\infty}$ in the steady-state value for the EKF covariance matrix Hidaka et al. (2005), Mourikis and Roumeliotis (2005) contains the eigenvalues of a minor of the sensor graph Laplacian weighted by the variances of the relative distance measurements. Specifically,

$$P_{rr\infty} = Q_o^{1/2} U \text{diag} \left\{ \frac{1}{2} + \left(\frac{1}{4} + \frac{1}{\lambda_i} \right)^{\frac{1}{2}} \right\} U^T Q_o^{1/2}, \quad (2)$$

where Q_o is a diagonal matrix with entries that depends on the characteristics of the mobile sensors and their speed, U , is the matrix of eigenvectors, and λ_i is the i th eigenvalue of the matrix

$$C = Q_o^{1/2} H_o^T R_o^{-1} H_o Q_o^{1/2}, \quad (3)$$

in which R_o is a diagonal matrix of the noise covariance, and H_o relates to the incidence matrix of the sensor graph. Suppose for example that the location of a single landmark is accurately known and three robots can measure distances and bearings to each other or the landmark. In this case, H_o contains a block of zeros that eliminates the graph node corresponding to the landmark, and $H_o^T R_o^{-1} H_o$ turns out to be a minor of a weighted Laplacian.

A problem that arises in this cases how to select a new observation (add a new edge in the graph) that can most improve the accuracy of position estimates. In view of (2) and (3), and using eigenvalue interlacing theorems, it can be shown that the trace of $P_{rr\infty}$ is related to the nonzero eigenvalues of the

(weighted) sensor graph Laplacian L_w as follows

$$\frac{(n-1)^2}{\sqrt{\text{trace}\{L_w\}}} \leq \sum_{i=1}^{n-1} \frac{1}{\sqrt{\lambda_i}} \leq \frac{n-1}{\sqrt{\lambda_2(L_w)}}.$$

Thus, forming a complete sensing graph minimizes the localization error, but comes at a cost of obtaining and processing the maximum number of observations.

Network connectivity appears to be a catalyst since nothing useful can happen without it. Messages cannot reach their destination, consensus among the network nodes over a certain quantity cannot be achieved using only nearest-neighbor communicated information, and estimation errors may grow unbounded. Nonetheless, in Byrne et al. (2005) it is shown that high algebraic connectivity does not necessarily imply high robustness in terms of maintaining connectedness in the presence of randomly failing links. In other words, while the network connectivity is certainly improved as the second smallest eigenvalue increases thus decreasing the diameter of the network, the network remains vulnerable to targeted attacks at edges. In particular, there may be few nodes or links that guarantee the connectivity of the network, and the removal of as few as one or two such nodes may shatter the network.

Consider a network represented by a graph G . The algebraic connectivity of G , $\lambda_2(G)$, satisfies (Fiedler's inequality Fiedler (1973))

$$\lambda_2(G) \leq \nu(G) \leq \eta(G), \quad (4)$$

where $\nu(G)$ measures the nodes connectivity, and $\eta(G)$ denotes the edges connectivity. While increasing the algebraic connectivity increases the lower bound on node-connectivity $\nu(G)$, it was shown in Byrne et al. (2005), that for circular and mesh lattice graphs, an increase in algebraic connectivity often corresponds to a decrease in node-connectivity $\nu(G)$ and edge-connectivity $\eta(G)$.

In fact, let us consider first the *small-world network* introduced in Watts and Strogatz (1998). This network is based on an n -nodes one-dimensional lattice on a ring where each node is connected to its k nearest neighbors. In Watts and Strogatz (1998) it is shown that the random rewiring of nodes according to a small probability p greatly reduces the characteristic path length, and results in a small-world network. Figure ?? shows the effects of random rewiring for a network with 20 nodes, and $k = 4$. In Olfati-Saber (2005), the author shows that this random rewiring also results in a large increase in algebraic connectivity for ring lattices, and the author then concludes that the network becomes more robust to node and link failures.

For certain types of networks however, large increases in algebraic connectiv-

ity often correspond to a decrease in node-connectivity and edge-connectivity. For example, let us start with a ring lattice of $n = 100$ nodes, and $k = 4$ edges per node (Figure ??), then rewire each edge at random with a probability p . As p increases from 0 to 0.9, the graph algebraic connectivity increases sharply, and the mean path length of the network decreases. However, the node-connectivity and edge-connectivity of the network decrease as the probability p increases. Similar results were obtained for a regular mesh lattice of 100 nodes (Figure ??), with each node having a communication radius $R = 1$. As p becomes larger, the edges connecting nearest neighbors are increasingly rewired, and link nodes in remote locations are directly connected. As shown here however, this does not necessarily improve the node or edge connectivities.

In a system where nodes are redundant or dispensable, improving algebraic connectivity does indeed improve the overall robustness of the network in terms of link failures, by reducing the characteristic path length. In systems where each node is critical however, node-connectivity and edge-connectivity are the truly important parameters for assessing robustness of connectivity with respect to randomly failing links. Evaluating node-connectivity or edge-connectivity for large networks is unfortunately much more costly than computing the graph's algebraic connectivity.

3.2 Navigability: path lengths and hops

When designing controllers for a networked system, it is typically assumed that paths exist between arbitrary node pairs of the communication network. The problem of determining these paths is usually ignored, or assumed solved by the routers that direct the flow of information through the network.

Standard routing protocols make use of assumptions that may not be generally favorable to control system design. For instance, Ethernet is a broadcast protocol, and thus only a limited number of participants can communicate over a given portion of a network. The open shortest path first (OSPF), as well as the border gateway (BGP) protocols, are susceptible to the propagation of corrupt or maliciously faulty information Nordstrom and Dovrolis (2004).

To provide the most basic packet delivery service, such as on the Internet at the IP-level, protocols like Ethernet, OSPF and BGP, combined with the commodity network hardware do well enough when most nodes are connected, the network is navigable, paths are relatively short, and service is fairly reliable. Such is the case on the Internet, where the average hop-count at the IP-level is at most a few dozens, despite there being potentially billions of routable IP addresses. Notably however, deviations from ideal conditions result in several

serious interruptions in global Internet service. For applications such as sensor networks or ad-hoc networks among mobile devices (such as cell phones), all of these issues are active areas of research in both the control systems and network theory communities.

Even if determining paths from source to destination is not an issue, the lengths of such paths matter, especially when information is processed as it propagates through the nodes of the network. One such example is the case of leader-follower control architectures. When the leading vehicle in a platoon suddenly decelerates, the more vehicles are between a follower and the leader, the faster this follower must decelerate. Depending on the size of the platoon and the dynamics of the vehicles, there comes a point where actuators reach their physical limits, control signals saturate, and collisions between vehicles occur.

String stability Swaroop (2002) is a theoretical framework that addresses this issue by treating propagating destabilizing information as a disturbance. By appropriate control design, these disturbance signals are attenuated as they propagate through the string of interconnected systems, and stability is preserved. Mesh stability Pant et al. (2002) generalizes this idea to multiple (physical) dimensions.

When the propagated information is not regulated in terms of its effect on the receiving nodes, it is shown in Tanner et al. (2004) that the network distance of a follower from the source of the signals (the leader), has an adverse effect on the ability of the follower to track its desired position in the formation. In such cases, routing the information signals through shorter paths improves stability Tanner et al. (2004). Thus, in network control system design, two options seem to be available: either regulate the system dynamics so that it can cope with information traveling over long paths, or make sure that short paths (up to a certain length) can be found. Regarding the latter, ad-hoc routing algorithms that improve the navigability of the network are needed.

3.3 Efficiency: capacity, link quality, and delays

As far as control design is concerned, a communication channel is merely a medium for obtaining or sending information (measurement signals, or control commands). From this perspective, what seems to be important is: (i) how much information can be carried, and (ii) how fast can it be transferred.

The first question is related to the channel's capacity as studied in Information theory, and results linking information theory to control have recently been reported Wong and Brockett (1997, 1999); Nair and Evans (2000); Ballieul (2002); Brockett and Liberzon (2000). While information theory models the

communication channel as an information transmitting medium that corrupts portions of the signal, the main issue for control-based applications are the delays (as well as corruption) suffered by the signals as they are carried across the channel. In the case of noiseless channels, a necessary condition for asymptotic observability and stabilizability for linear, time-invariant, discrete-time systems, is that the rate of communication R (which must be less than the capacity C of the channel) is bounded below as $R > \sum_{\lambda_u(A)} \max\{0, \log|\lambda_u(A)|\}$, where $\lambda_u(A)$ are the unstable eigenvalues of the system matrix A . In some cases, this condition is also sufficient. Similar results hold in the case of noisy channels, as described in Tatikonda and Elia (2005). Article Martins and Dahleh (2005) investigates the fundamental limitation of performance for networked feedback systems, in which the feedback loop is comprised of a discrete-time, linear, time-invariant plant, a channel, as well as an encoder and a decoder. The disturbance rejection ability is found to be bounded from below by $\sum_{\lambda_u(A)} \max\{0, \log|\lambda_u(A)|\} - C$. This particular result shows that the excess capacity $C - \sum_{\lambda_u(A)} \max\{0, \log|\lambda_u(A)|\}$ is all that is available for disturbance rejection.

The speed at which information travels from source to destination is usually measured by a "communication delay," the time elapsed between transmission and reception. Depending on where the network is included in the feedback loop of the network control system, such communication delays can cause actuation delays, measurement delays, or both. It is generally recognized that the delays degrade the performance of control systems. It is natural to expect, therefore, that communication delays will adversely impact the performance of a networked control system, possibly even causing instabilities.

Initial investigation seemed to support this claim. In Olfati-Saber and Murray (2004), stability analysis in the frequency domain suggests the existence of an upper limit in the (uniform) communication delays that a continuous, nearest neighbor interconnected system can tolerate before becoming unstable. However, more recent analysis of state space, discrete-time models of interconnected systems, have led to different conclusions: in some (not so special) cases, arbitrary (but bounded) communication delays may be tolerated at the expense of convergence speed. Moreau Moreau (2005) was among the first to address the consensus problem in the presence of time delays, giving convexity conditions on the set of admissible control inputs that ensure asymptotic velocity synchronization. In Angeli and Bliman (2005), the approach of Moreau (2005) is extended, showing that if the agents dynamics are appropriately restricted, stability can still be maintained. A different approach in Tanner and Christodoulakis (2005) focuses on the communication protocol, and shows that velocity synchronization in a connected group of autonomous mobile agents, may still be achieved when the agent controllers use delayed information, regardless of the size of this delay, if control and communication are properly interleaved. In Morse (2006) the composition properties of graphs are used to

show that under certain assumptions on the communication topology, delays have no effect on the stability of the system. In fact, in a somehow counter-intuitive situation, it turns out that longer delays (if used judiciously) can improve the stability of some systems Abdallah et al. (1993).

4 Conclusion

Any conceptual links between networked control systems, cooperative control, and complex networks through graph theoretic analysis, provide opportunities for control theory to reach out and exploit the arsenal available in complex network research and computer science. This article offers such a suggestion by highlighting the recently revealed power of randomized algorithms in routing, network design, resource allocation, and game theory.

Mechanism design, as recently developed in the area of computer science, seeks to allow selfish individuals to interact in a networked environment in such a way that no outcome is particularly disadvantageous to any of the nodes. Such approaches yield results for routing, network design, and resource allocation Tardos (2004), and seem directly applicable to open networked control systems in which the control engineer must ensure that corrupt or misbehaving nodes do not negatively affect the functionality of the system.

With this brief review of a small selection of intriguing ideas, we hope to establish further links between network theory, physics, and control systems.

5 Acknowledgment

The work of the first author is supported by NSF grant no. PHY-0200909; the second author is supported in part by the NSF Career grant no. 0447898; the third author is supported in part by NSF grant no. CNS 0626380 under the FIND initiative; the work of the last author was supported by the United States Department of Energy under Contract DE-AC04-94AL8500.

References

- Abdallah, C., Dorato, P., Benitez-Read, J., Byrne, R., 1993. Delayed positive feedback can stabilize oscillatory systems. In: Proc. of the American Control Conference. pp. 3106–3107.
- Albert, R., Barabási, A.-L., 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics* 74, 47–97.

- Angeli, D., Bliman, P.-A., 2005. Extension of a result by Moreau on stability of leaderless multi-agent systems. In: Proc. of the 44th IEEE Conference on Decision and Control, ECC-CDC05. pp. 759–764.
- Ballieul, J., 2002. Feedback designs in information-based control. In: Pasik-Duncan, B. (Ed.), Stochastic Theory and Control – Proceedings of Workshop held at the University of Kansas. Lecture Notes in Control and Information Sciences. Springer-Verlag, New York, pp. 35–57.
- Ben-Or, M., Canetti, R., Goldreich, O., 1993. Asynchronous secure computation. In: Proc. 25th ACM Symposium on Theory of Computing. pp. 52–61.
- Bettstetter, C., 2002. On the minimum node degree and connectivity of a wireless multihop network. In: Proc. 3rd ACM Int. Symp. on Mobile Ad Hoc Net. and Comp. (MobiHoc02). pp. 80–91.
- Bollobás, B., 1985. Random graphs. Academic Press.
- Brockett, R., Liberzon, D., 2000. Quantized feedback stabilization of linear systems. IEEE Transactions on Automatic Control 45 (7), 1279–1289.
- Byrne, R., Feddema, J., Abdallah, C., 2005. Algebraic connectivity and graph robustness. Tech. rep., Tech. Rep., Sandia National Laboratories.
- Clauset, A., Moore, C., 2003. How do networks become navigable? e-print arXiv:cond-mat/0309415.
URL <http://arxiv.org/pdf/cond-mat/0309415>
- Cohen, R., Erez, K., ben Avraham, D., Havlin, S., 2000. Resilience of the Internet to random breakdowns. Physical Review Letters 85 (21), 4626–4628.
- Şimşek, O., Jensen, D., 2005. Decentralized search in networks using homophily and degree disparity. In: Proc. 19th International Joint Conference on Artificial Intelligence. pp. 304–310.
- Dorogovtsev, S. N., Mendes, J. F. F., 2003. Evolution of Networks: From Biological Nets to the Internet and WWW. Oxford University Press.
- Engle, M., Khan, J., 2006. Vulnerabilities of p2p systems and a critical look at their solutions.
URL <http://www.medianet.kent.edu/techreports/TR2006-11-01-p2pvuln-EK.pdf>
- Faloutsos, M., Faloutsos, P., Faloutsos, C., 1999. On power-law relationships of the Internet topology. Computer Communications Review 29, 251–262.
- Fiedler, M., 1973. Algebraic connectivity of graphs. Czechoslovak Mathematical Journal 23 (98), 298–305.
- Force, I. E. T., ???a. Rfc 1247 (open shortest path first).
<http://www.ietf.org/rfc/rfc1247.txt>.
- Force, I. E. T., ???b. Rfc 1771 (border gateway protocol).
<http://www.ietf.org/rfc/rfc1771.txt>.
- Gallos, L. K., Cohen, R., Argyrakis, P., Bunde, A., Havlin, S., 2005. Stability and topology of scale-free networks under attack. Physical Review Letters 94 (18), 188701(1–4).
- Guillaume, J.-L., Latapy, M., Magnien, C., 2004. Comparison of failures and attacks on random and scale-free networks. In: in Lecture Notes in Computer Science, Proceedings of the 8th International Conference On Principles Of

- Distributed Systems OPODIS'04. pp. 186–196.
- Hidaka, Y., Mourikis, A., Roumeliotis, S., 2005. Optimal formations for cooperative localization of mobile robots. In: Proc. of the International Conference on Robotics and Automation. pp. 4137–4142.
- Jadbabaie, A., Lin, J., Morse, A. S., 2002. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control* 48 (6), 988–1001.
- Kelly, F., 2001. Mathematical modelling of the internet. In: Engquist, B., Schmid, W. (Eds.), *Mathematics Unlimited - 2001 and Beyond*. Springer-Verlag, pp. 685–702.
- Kleinberg, J., 1999. The small-world phenomenon: An algorithmic perspective. Tech. Rep. 99-1776, Computer Science Department, Cornell University.
- Kleinberg, J., 2000. Navigation in a small world. *Nature* 406 (6798), 845.
- Link, H., LaViolette, R. A., Saia, J., Lane, T., 2005. Parameters affecting the resilience of scale-free networks to random failures. e-print arXiv:cs.NI/0511012.
URL <http://arxiv.org/pdf/cs.NI/0511012>
- Martins, N., Dahleh, M., 2005. Fundamental limitations of performance in the presence of finite capacity feedback. In: Proc. of the American Control Conference. pp. 79–86.
- Milgram, S., 1967. The small world problem. *Psychology Today* 2, 60–67.
- Moreau, L., 2005. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control* 50 (2), 169–182.
- Morse, A., 2006. *Lecture Notes on Logically Switched Dynamical Systems*. Springer-Verlag, (in print).
- Mourikis, A., Roumeliotis, S., 2005. Performance bounds for cooperative simultaneous localization and mapping. In: Thrun, S., Sukhatme, G., Schaal, S., Brock, O. (Eds.), *Robotics: Science and Systems I*. MIT Press, pp. 73–80.
- Nair, G., Evans, R., 2000. Stabilization with data-rate-limited feedback: Tightest attainable bounds. *Systems and Control Letters* 41, 49–76.
- Newman, M. E. J., 2003. The structure and function of complex networks. *SIAM Review* 45, 167–256.
- Nordstrom, O., Dovrolis, C., 2004. Beware of BGP attacks. *SIGCOMM Comput. Commun. Rev.* 34 (2), 1–8.
- Olfati-Saber, R., 2005. Ultra-fast consensus in small-world networks. In: Proc. of the American Control Conference. pp. 2371–2378.
- Olfati-Saber, R., Murray, R. M., 2004. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control* 49 (9), 1520–1533.
- Pant, A., Seiler, P., Hedrick, K., 2002. Mesh stability of look-ahead interconnected systems. *IEEE Transactions on Automatic Control* 47, 403–407.
- Pease, M., Shostak, R., Lamport, L., 1980. Reaching agreement in the presence of faults. *Journal of the ACM* 27 (22), 228–234.
- Ren, W., Beard, R., 2005. Consensus seeking in multiagent systems under dynamically changing interaction topologies. *IEEE Transactions on Automatic*

- Control 50 (5), 655–661.
- Santi, P., Blough, D., 2003. The critical transmitting range for connectivity in sparse wireless ad hoc networks. *IEEE Transactions on Mobile Computing* 2 (1), 25–39.
- Srikant, R., 2004. *The Mathematics of Internet Congestion Control (Systems and Control: Foundations & Applications)*. Birkhuser.
- Stanley, H. E., 1983. *Introduction to Phase Transition and Critical Phenomena*. Oxford University Press.
- Swaroop, D., 2002. A note about the stability of a string of LTI systems. *ASME Journal of Dynamic Systems, Measurement and Control* 124, 472–475.
- Tanner, H. G., Christodoulakis, D., 2005. Discrete time flocking with time delays. In: *Proc. of the 44th IEEE Conference on Decision and Control*. pp. 4945–4950.
- Tanner, H. G., Jadbabaie, A., Pappas, G. J., 2003a. Stable flocking of mobile agents, Part I: Fixed topology. In: *Proc. of the 44th IEEE Conference on Decision and Control*. pp. 2010–2015.
- Tanner, H. G., Jadbabaie, A., Pappas, G. J., 2003b. Stable flocking of mobile agents, Part II: Dynamic topology. In: *Proc. of the 44th IEEE Conference on Decision and Control*. pp. 2016–2011.
- Tanner, H. G., Pappas, G. J., Kumar, V., 2004. Leader-to-formation stability. *IEEE Transactions on Robotics and Automation* 20 (3), 433–455.
- Tarbouriech, S., Abdallah, C. T., Chiasson, J., December 2004. *Advances in Communication Control Networks (Lecture Notes in Control and Information Sciences)*, 1st Edition. Springer.
- Tardos, E., 2004. Network games. In: *Proc. ACM Symposium on Theory of Computing*. pp. 341–342.
- Tatikonda, S., Elia, N., 2005. Communication requirements for networked systems. In: S. Tarbouriech, C.T. Abdallah, J. C. (Ed.), *Advances in Communication Control Networks*. Lecture Notes in Control and Information Sciences. Springer-Verlag, New York, pp. 303–326.
- Verriest, E., Egerstedt, M., 2002. Control with delayed and limited information: A first look. In: *Proc. of the IEEE Conference on Decision and Control*. pp. 1231–1236.
- Walsh, G., Ye, H., Bushnell, L., 2002. Stability analysis of networked control systems. *IEEE Transactions on Control Systems Technology* 10 (3), 438–445.
- Watts, D. J., Strogatz, S. H., 1998. Collective dynamics of 'small-world' networks. *Nature* 393 (6684), 440–442.
- Wong, W., Brockett, R., 1997. Systems with finite communication bandwidth constraints I: State estimation problems. *IEEE Transactions on Automatic Control* 42 (9), 1294–1299.
- Wong, W., Brockett, R., 1999. Systems with finite communication bandwidth constraints II: Stabilization with limited information feedback. *IEEE Transactions on Automatic Control* 44 (5), 1049–1053.
- Xue, F., Kumar, P., 2004. The number of neighbors needed for connectivity

of wireless networks. *Wireless Networks* 10 (2), 169–181.

Zhang, W., Branicky, M., Phillips, S., 2001. Stability of networked control systems. *IEEE Control Systems Magazine* 21 (1), 84–99.