

Interacting with networks

How does structure relate to controllability in single-leader,
consensus networks?

Magnus Egerstedt, Simone Martini, Ming Cao, Kanat Camlibel,
and Antonio Bicchi — April 10, 2012

As networked dynamical systems appear around us at an increasing rate, questions concerning how to manage and control such systems are becoming more important. Examples include multi-agent robotics, distributed sensor networks, interconnected manufacturing chains, and data networks. In response to this growth, a significant body of work has emerged focusing on how to organize such networks in order to facilitate their control and make them amenable to human interactions. In this article, we summarize these activities by connecting the network topology, that is, the layout of the interconnections in the network, to the classic notion of controllability.

In manufacturing, one of the technological bottlenecks can be found in the general assembly phase. This is the last stage of the manufacturing chain where the pieces, such as doors, locks, and cup-holders in automotive manufacturing, are assembled into a finished product. If a single worker could command and interact with a number of flexible, mobile manipulators in an effective manner, it is expected that this process could be improved significantly. For this to be

possible, means must be made available to the operator to be able to effectively influence the state of the manipulators. Similarly, the current mode of operation when piloting unmanned aerial vehicles (UAVs) is that multiple operators are required to operate a single UAV. An explicit aim is to be able to invert this many-to-one relationship so that a single operator can pilot multiple UAVs, which again calls for the operator to be able to influence the state of the system. In both of these applications, we are lacking the tools for systematically characterizing and designing useful interaction models. In this article, we take one step towards achieving such a characterization by focusing on the controllability properties of the underlying interaction network itself. By itself, controllability does not provide answers to how these interactions should be structured. It does, however, provide insights into what is possible.

At a high level of abstraction, a network can be viewed as a graph, that is, as a collection of vertices and edges. In particular, given a collection of N interconnected nodes, we let the network graph be given by $G = (V, E)$, where the vertex set is simply the set of nodes $V = \{1, \dots, N\}$, and the edge set $E \subseteq V \times V$ encodes the information flow in the network. The interpretation is that an edge exists between nodes i and j (denoted by $(i, j) \in E$) if information is flowing between these nodes. In this article, we only consider undirected graphs in the sense that $(i, j) \in E$ if and only if $(j, i) \in E$, which corresponds to a bidirectional information flow in the network. Such graphs have a convenient graphical representation, as shown in Figure 1.

Now, imagine that the nodes in the network are mobile robots that somehow coordinate their movements, akin to swarming insects or schooling fish. If one were to try to control such a swarm, one approach could be to select key individuals and then drag them around in order to induce a desired, global behavior in the robot swarm. Although other approaches can be

envisioned, this is the basic setup in this article and we will investigate how effective such a strategy might be. In particular, we will see that the organization of the underlying network structure plays a central role when addressing this issue.

The effectiveness of the interactions with a networked control system can, at least partially, be understood in terms of its controllability properties. In particular, we are interested in whether or not the system is completely controllable, that is, if it is possible to drive it from any initial configuration to any target configuration. And, questions related to controllability become meaningful only when the nodes are endowed with dynamics and if there is some way of injecting exogenous control signals into the network. We achieve this latter objective by dividing the vertex set $V = V_f \cup V_\ell$ into *follower* nodes, V_f , and *leader* nodes, V_ℓ , with the understanding that control signals can be injected only at the leader nodes. Moreover, the followers execute their own coordination strategies and the control inputs propagate through the network by virtue of the fact that these strategies take neighboring nodes' states into account. In Figure 1, the black nodes are the leader nodes while the remaining nodes are the follower nodes.

In order to define the dynamics over the network, we first need to associate a state with each of the nodes, $x_i \in \mathbb{R}^d$, $i = 1, \dots, N$, where d is the dimension of the state. These states could for example correspond to the positions of the nodes in a mobile robot network, or the processed sensor values in a sensor network. In this article, we assume that we can control the leader nodes' states directly in the sense that $x_i = u_i$, $i \in V_\ell$, where u_i is the control input at node i . This assumption can be somewhat relaxed but it helps keep the notation simple and allows us to focus directly on the connections between the network structure and its associated

controllability properties. By assuming that each of the follower nodes are executing a particular coordination strategy $\dot{x}_i = f_i(x_1, \dots, x_N)$, $i \in V_f$, where f_i is allowed only to depend on the state values associated with those nodes adjacent to vertex i in the network, one can ask whether or not it is possible to drive the follower states from any configuration to any other configuration. The answer to this question depends on the choice of interaction law as well as on the underlying network topology.

Many different decentralized interaction laws and coordination strategies have been designed for networked multi-agent systems to achieve a vast array of objectives such as swarming, flocking, alignment, cohesion, rendezvous, formation maintenance, and coverage, [1], [5], [9], [12], [15]. In this article, we focus on a particular such choice, namely, on the linear agreement protocol, which has proved useful for providing cohesion in the network and has served as a starting point for a large class for other types of networked controllers. The reason for this is that the linear agreement protocol ensures that each state asymptotically approaches the stationary average of all the states in the network when the underlying graph is connected, that is, there is a path (not necessarily direct) through the network between each pair of two vertices in the graph. For an introduction to this topic, see [12], [15].

Once the leader nodes are selected and the interaction laws are decided upon, what makes different networks respond differently to control inputs becomes solely a question of the network topology, that is, on the graph structure itself. As discussed in [11], [14], [16], certain network topologies are better than others when it comes to being able to effectively control the system. This matters since the network design is typically decoupled from the control design. But, if the network structure can be explicitly designed with the aim of making the system amenable to

control, this would improve the performance of the overall system. For instance, it turns out that more interactions are not necessarily a good thing. If the network topology is given by a complete graph, where every vertex is directly connected to every other vertex, what can effectively be controlled by a single leader under the linear agreement protocol is just the centroid of the node states. In other words, this is a particularly poor choice of network topology from a controllability vantage point even though it has the largest number of edges possible. In this article, we take this observation one step further and summarize the connections between the graph topology and the controllability properties of the *controlled agreement dynamics*. It should be noted already at this point that other types of interaction dynamics and user-network interactions can be envisioned. As this type of inquiry into the controllability properties of networked control systems is still in its infancy, we here report on the results that have been obtained so far, but acknowledge that significant work remains to be done before this topic has been completely understood.

The Controlled Agreement Dynamics

Consider a network whose node states evolve according to the nearest neighbor-averaging rule known as the *consensus equation*, as defined, for example, in [12], [8], [13], [17],

$$\dot{x}_i = - \sum_{j \in \mathcal{N}_i} (x_i - x_j). \quad (1)$$

Here, \mathcal{N}_i is the set of nodes adjacent to node i in the static and undirected information exchange graph $G = (V, E)$, in the sense that $\mathcal{N}_i = \{j \in V \mid (i, j) \in E\}$. An example of using this coordination strategy is shown for 10 nodes in Figure 2. As dictated by the theory, the agents' states converge to the same value.

Assume that we would like to control this network, and we achieve this by injecting

control signals at the leader nodes, as

$$x_i = u_i, \quad i \in V_\ell, \quad (2)$$

while all the remaining follower nodes execute the coordination strategy given in (1). To be able to characterize the controllability properties of this network from a purely graph-theoretic vantage point, we first need some basic tools from algebraic graph theory. (For a comprehensive treatment of this subject, see [7].) What algebraic graph theory helps us with is to associate matrices to graphs, which is crucial in order to arrive at a formulation that is amenable to control theoretic tools.

Let Δ be the $N \times N$ *degree matrix* associated with a graph, with entries given by

$$[\Delta]_{i,j} = \begin{cases} \text{deg}(i) & \text{if } i = j \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where the degree, $\text{deg}(i) = |\mathcal{N}_i|$, is the size of the neighborhood set to node i , and where $|\cdot|$ denotes cardinality. Similarly, the *adjacency matrix* \mathcal{A} is given by

$$[\mathcal{A}]_{i,j} = \begin{cases} 1 & \text{if } (i,j) \in E \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

The final matrix, the *graph Laplacian*, needed for the discussion is given by

$$L = \Delta - \mathcal{A}. \quad (5)$$

If we index the nodes in such a way that the last M nodes are the leader nodes and the first $N - M$ nodes are the followers, we can decompose L as

$$L = - \left[\begin{array}{c|c} A & B \\ \hline B^T & \lambda \end{array} \right], \quad (6)$$

where $A = A^T$ is $(N-M) \times (N-M)$, B is $(N-M) \times M$, and λ is $M \times M$. The point behind this decomposition is that if we assume that the state values are scalars, that is, $x_i \in \mathbb{R}$, $i = 1, \dots, N$, and gather the states from all follower nodes as $x = [x_1, \dots, x_{N-M}]^T$ and the leader nodes as $u = [x_{N-M+1}, \dots, x_N]^T$, the dynamics of the controlled network can be written as

$$\dot{x} = Ax + Bu, \quad (7)$$

as shown in [14]. Note that if the states were non-scalar, the analysis still holds even though one has to decompose the system dynamics along the different dimensions of the states.

As an example, returning to the graph in Figure 1, the corresponding system dynamics become

$$\dot{x} = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & -3 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -3 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -2 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & -4 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -2 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & -1 \end{bmatrix} x + \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} u. \quad (8)$$

What we would like to know is what the controllability properties associated with this system are. In particular, we would like to avoid the standard rank tests and instead obtain

characterizations of what the network topology should look like in order to render the system completely controllable. The reason for this is that if we had a clear understanding of the usefulness of different network structures, this would help guide our design choices when constructing the underlying information exchange network. There are various approaches to obtain connections between network structure and controllability, and we start with the most general and then focus in on methods for analyzing special classes of graphs.

Controllability Through External Equitable Partitions

One interesting fact about the controlled agreement dynamics is that the followers tend to cluster together due to the cohesion provided by the consensus equation. This clustering effect can actually be exploited when analyzing the network's controllability properties. We thus start with a discussion of how such clusters can be obtained.

By a *partition* of the graph $G = (V, E)$ we understand a grouping (clustering) of nodes into cells, that is, a map $\pi : V \rightarrow \{C_1, \dots, C_K\}$, where we say that $\pi(i)$ denotes the *cell* that node i is mapped to, and we use $\text{range}(\pi)$ to denote the *codomain* to which π maps, that is, $\text{range}(\pi) = \{C_1, \dots, C_K\}$. Similarly, the operation $\pi^{-1}(C_i) = \{j \in V \mid \pi(j) = C_i\}$ returns the set of nodes that are mapped to cell C_i . An example of such a node partition is given in Figure 3.

But, we are not interested in arbitrary clusters. Instead, we want to partition the nodes into cells in such a way that all nodes inside a cell have the same number of neighbors in adjacent cells. To this end, the *node-to-cell degree* $\text{deg}_\pi(i, C_j)$ characterizes the number of neighbors that

node i has in cell C_j under the partition π ,

$$\deg_{\pi}(i, C_j) = |\{k \in V \mid \pi(k) = C_j \text{ and } (i, k) \in E\}|. \quad (9)$$

A partition π is said to be *equitable* if all nodes in a cell have the same node-to-cell degree to all cells, that is, if, for all $C_i, C_j \in \text{range}(\pi)$, $\deg_{\pi}(k, C_j) = \deg_{\pi}(\ell, C_j)$, for all $k, \ell \in \pi^{-1}(C_i)$.

This is almost the construction one needs in order to obtain an initial characterization of the controllability properties of the network. However, what we need to do is produce partitions that are equitable between cells in the sense that all agents in a given cell have the same number of neighbors in adjacent cells, but where we do not care about the structure *inside* the cells themselves. This leads to the notion of an *external equitable partition* (EEP), and we say that a partition π is an *EEP* if, for all $C_i, C_j \in \text{range}(\pi)$, where $i \neq j$,

$$\deg_{\pi}(k, C_j) = \deg_{\pi}(\ell, C_j), \text{ for all } k, \ell \in \pi^{-1}(C_i). \quad (10)$$

A Necessary Condition for Single-Leader Networks

One key objective when trying to understand controllability of networked systems is to enable users to interact with such networks. As a first step, one can start by analyzing leader-follower networks with a single leader, which thus corresponds to a sole operator interacting with the network. Hence, we assume that we have a single leader acting as the leader node, and we are particularly interested in EEPs that place this leader node in a singleton cell, that is, in partitions where $\pi^{-1}(\pi(N)) = \{N\}$, and we refer to such EEPs as *leader-invariant*. Moreover, we say that a leader-invariant EEP is *maximal* if its codomain has the smallest cardinality, that is, if it contains the fewest possible cells, and we let π^* denote this maximal, leader-invariant EEP. We note that given a graph G and a single leader, π^* always exists uniquely [2], [6], [7]. The

maximal, equitable partition (and as a consequence, π^* as well) can be computed in polynomial time (polynomial in the size of the graph) and different algorithms have been given to this end, [3], [6]. Examples of the construction of π^* are shown in Figure 4, which allow us to state the following key result from [11].

The networked system in (7) is completely controllable only if G is connected and π^ is trivial, that is, $\pi^{*-1}(\pi^*(i)) = \{i\}$, for all $i \in V$.*

This result allows us to obtain necessary conditions for controllability purely in terms of the network's graph topology, that is, it does not rely on any rank tests. Examples of this topological condition for controllability are given in Figure 5.

One particularly intriguing aspect of letting the interaction dynamics be given by the consensus equation (1) is that it provides cohesion in the network. A consequence of that, as shown in [11], is that the difference between states within cells in $\text{range}(\pi^*)$ is uncontrollable. Moreover, if G is connected, these differences decay asymptotically due to the fact that A in (6) is negative definite if the graph is connected. In other words,

$$\lim_{t \rightarrow \infty} (x_k(t) - x_\ell(t)) = 0, \text{ for all } k, \ell \in \pi^{*-1}(C_i). \quad (11)$$

What this tells us is that no matter what the control input is, inside cells, the state values will inevitably converge to the same value.

An example of this effect is shown in Figure 6. In that figure, six follower agents are running the consensus equation (1), while the leader agent's state is given by a harmonic function. As can be seen, agents 2, 3, and 4 end up with the same state value since they share the same cell in the maximal, leader-invariant EEP. Similarly, agents 5 and 6 end up with the same value

while agent 1 belongs to a singleton cell. What is at play here is that nodes inside the same cells are symmetric with respect to the leader. And said symmetries are obstructions to controllability. A surprising consequence of this is discussed in [10], where the electrical power grid was found to be more symmetric (and hence less controllable) than biological or social networks.

But, we can do even better than this in that we can characterize an upper bound on what the dimension of the controllable subspace is, as shown in [4]. In fact, let (A, B) be given in (7) and let Γ be the corresponding controllability matrix. Then

$$\text{rank}(\Gamma) \leq |\text{range}(\pi^*)| - 1. \quad (12)$$

We note that since this result is given in terms of an inequality instead of an equality, we have only necessary conditions for controllability rather than a, as of yet elusive, necessary and sufficient condition. One instantiation where this inequality is indeed an equality is when π^* is also a distance partition, as shown in [18]. What this means is that when all nodes that are at the same distance from the leader (counting hops through the graph) also occupy the same cell under π^* , we have that $\text{rank}(\Gamma) = |\text{range}(\pi^*)| - 1$. These types of situations will be discussed in subsequent sections.

Quotient Graph Dynamics

One question one can ask now is if it is possible to give the part of the network that we can in fact control a graph-theoretic interpretation, that is, if there is a network structure associated with the controllable subspace. In order to answer this question, we need to introduce the notion of a *quotient graph*. Given a graph G together with an EEP π , the *quotient graph*

$G \setminus \pi = (V_\pi, E_\pi, w_\pi)$ is the weighted and directed graph whose node set is $V_\pi = \text{range}(\pi)$, the edge set is the set of ordered pairs such that $(C_i, C_j) \in E_\pi$ if and only if edges connect nodes in cells C_i and C_j , and the weight between cells is given by the cell-to-cell degree, that is, the number of edges connecting nodes in cells C_i and C_j . An example is shown in Figure 7.

As $V_{\pi^*} = \text{range}(\pi^*)$ and, within cells, state values converge to the same value. we expect be able to endow the quotient graph with a dynamics that is somehow related to the original system. As the difference between state values inside a cell in the EEP vanishes asymptotically, what we can in fact have some hope of controlling is the average inside a cell. For this, we let ξ_i be the average state value of a cell $C_i \in \text{range}(\pi^*)$,

$$\xi_i = \frac{1}{|\pi^{*-1}(C_i)|} \sum_{j \in \pi^{*-1}(C_i)} x_j, \quad (13)$$

which allows us to state a result involving the quotient graph dynamics, found in [4].

Given a connected network, G , with a single leader node, whose node dynamics are given in (7). Let π^* be the maximal, leader-invariant EEP associated with this network, with $G \setminus \pi^*$ being the corresponding quotient graph. We now chose to associate a dynamics with the quotient graph as

$$\dot{\xi}_i = - \sum_{C_j \in \mathcal{N}_{\pi^*, C_i}} w_{i,j} (\xi_i - \xi_j), \quad (14)$$

for all i such that $\pi^{*-1}(C_i) \neq \{N\}$, that is, cell i does not contain the input node, and let

$$\xi_i = u, \quad (15)$$

if $\pi^{*-1}(C_i) = \{N\}$. This choice of dynamics is consistent with the original dynamics in the sense that the dynamics (14-15), describing the evolution of ξ_i , satisfy

$$\xi_i(t) = \frac{1}{|\pi^{*-1}(C_i)|} \sum_{j \in \pi^{*-1}(C_i)} x_j(t) \quad (16)$$

as long as

$$\xi_i(0) = \frac{1}{|\pi^{\star-1}(C_i)|} \sum_{j \in \pi^{\star-1}(C_i)} x_j(0). \quad (17)$$

What this result tells us is that given a network, what we can control is in fact another smaller network, given by the quotient graph. The equivalent dynamics over the quotient graph is given in terms of the average state values inside cells in the EEP. As the differences between state values inside the cells vanish asymptotically, it describes the behavior of the actual states in the original system as t approaches infinity.

The reason why it is beneficial to be able to view the controllable subspace as a network is that this vantage point allows control designers to focus directly on smaller structures with a physical interpretation. It also allows for the network design to be done in such a way that the desired quotient graphs are obtained. An example is shown in Figure 8, in which different edges are removed from the graph in order to produce different quotient graphs.

What we have arrived at, thus far, is a necessary condition for controllability based solely on a characterization of the network topology. There are stronger conditions for specialized classes of graphs, whose eigenstructure can be more clearly established. In the next section, we investigate two such classes, namely, chain graphs and multi-chain graphs.

Chain and Multi-Chain Graphs

We now move on to networks that exhibit a rather specialized structure, yet are quite common in different application domains. In particular, we consider systems consisting of $n > 0$ followers, labeled by $1, \dots, n$, and one leader, labeled by $n + 1$. In view of the system dynamics

(7), we know there is a one-to-one correspondence between the system matrix A and its associated graph $G(A)$, which is the unique graph with $-A$ as its upper left-most block in the Laplacian, as per Equation (6). As we will take advantage of how the graph's structure affects its spectral properties, we, for simplicity, call the spectrum of A the *spectrum of the graph* $G(A)$.

We let a *chain graph* with $n + 1$ vertices to be the graph for which one can label its vertices in such a way that the edge set contains exactly the edge $(n + 1, 1)$ and the edges $(i, i - 1), (i - 1, i), 1 < i \leq n$. What this entails is simply a graph laid out as a chain, where each node (except the end nodes) has two neighbors, and the leader node being an end node connected to node n , as shown in Figure 9. Such chain structures can for instance be found in platooning autonomous vehicles and manufacturing chains. Their widespread use is, in part, why such structures deserve special attention.

We call n the *length* of the chain and we immediately note that there are a number of interesting relationships between the spectra of two chain graphs if their lengths satisfy certain relationships. To be more specific, if λ is an eigenvalue of $A(G_1)$ with associated eigenvector v ($A(G_1)v = \lambda v$), where G_1 is the chain graph with $n + 1$ vertices, then λ is also an eigenvalue of $A(G_2)$, where G_2 is the chain graph with $k(2n + 1) + n + 1$ vertices, for any positive integer $k = 1, 2, \dots$. What this means is that one can understand the spectral properties of longer chains through collections of shorter chain graphs. An additional important fact about such single-leader chain graphs is that *they are always completely controllable*, that is $\text{rank}(\Gamma) = n$, as shown in [2].

The chain construction can be generalized to other types of structures that take the form of the union of several chains. We say that a graph G with $n+1$ vertices is an *m-chain graph*, $m > 1$,

if one can label its vertices in such a way that there exist integers $1 \leq k_1 < k_2 < \dots < k_{m-1} < n$ such that its edge set is the union of the edge set $\{(n+1, 1), (n+1, k_1+1), \dots, (n+1, k_{m-1}+1)\}$ and the edge set $\{(i-1, i), (i, i-1), 1 < i \leq n \text{ and } i \neq 1, k_1+1, \dots, k_{m-1}+1\}$. A typical m -chain graph is shown in Figure 10. These types of interconnection structures can be found among transportation networks and flexible manufacturing lines. But, more importantly, they serve as generators of examples that highlight that the external equitable partition results are indeed only sufficient and not necessary for controllability.

In fact, using the relationships between the spectra of chain graphs, one can show that the spectrum of an m -chain graph has the following property. If G is an m -chain graph and the length of each chain i , $1 \leq i \leq m$, is $3l_i + 1$ for some $l_i \geq 0$, then $A(G)$ has -1 as an eigenvalue whose geometric multiplicity is at least m . But, the system (7) is not completely controllable if A has an eigenvalue whose geometric multiplicity is greater than one, as shown in [2]. As a result, a topological test for complete controllability is thus to check if G is an m -chain graph and the length of each chain i is $3k_i + 1$ for some $k_i \geq 0$, then the system is not controllable. One can also compute the EEPs of multi-chain graphs. In fact, if the lengths of the chains of an m -chain graph G are different, then its maximal leader-invariant EEP is trivial [2].

It turns out that some m -chain graphs can be augmented by adding edges connecting different chains. The augmentation can be carried out in such a way that the augmented graph still has a trivial maximal leader-invariant EEP and is at the same time uncontrollable. We show two examples of this construction in Figures 11 for such uncontrollable augmented multi-chain graphs.

Conclusions

To be able to infer controllability properties directly from the network structure is useful since it allows the network designer to build networks that satisfy desired controllability properties. This is important since we typically want to be able to command and control networks in an efficient manner. In this article, we discuss this issue and collect some of the key results that have emerged in this area during the last five years. Necessary conditions for controllability are given in terms of the networks' maximal, leader-invariant EEPs. These conditions are quite general and can be extended in a straightforward manner beyond the single-leader case, as is done in [14]. Unfortunately, these conditions are not sufficient and the quest for such a graph-based necessary and sufficient condition remains an open issue. However, for certain classes of systems, we have obtained a more complete characterization, and in this article we report on two such classes, namely, chain and multi-chain graphs.




Figure 1: A graphical representation of a network graph. The circles are nodes in the network and the edges between nodes encode that information can flow between adjacent nodes. In the figure, the leader nodes (nodes 12, 13, 14) are given in black, while the remaining nodes (nodes 1 to 11) are follower nodes.




Figure 2: Running the consensus equation (1). Ten agents are executing the coordination protocol in (1) and their states converge to the same value.




Figure 3: A partition of the node set into cells. The partition has four cells C_1, \dots, C_4 and each vertex belongs to exactly one cell.

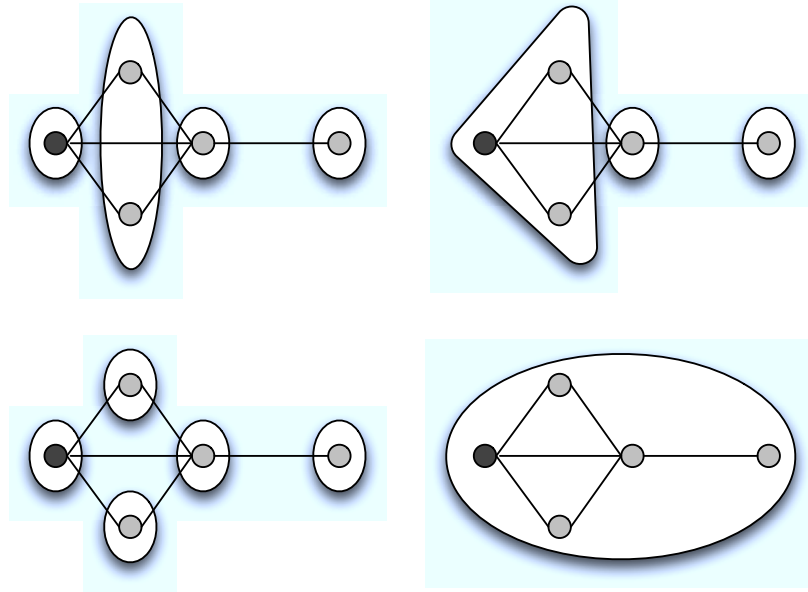


Figure 4: A graph with four possible EEPs. The leader-node (black node) is in a singleton cell in the two left-most figures and, as such, they correspond to leader-invariant EEPs. Of these two leader-invariant EEPs, the top-left partition has the fewest number of cells and that partition is thus maximal. We note that this maximal partition is not trivial since one cell contains two nodes.




Figure 5: Networks (a), (b) are not completely controllable, as their partitions π^* are not trivial. The partitions π^* associated with networks (c), (d), (e) are indeed trivial, but we cannot directly conclude anything definitive about their controllability properties since the topological condition is only necessary. Indeed, (c) is completely controllable, while (d) and (e) are not completely controllable.

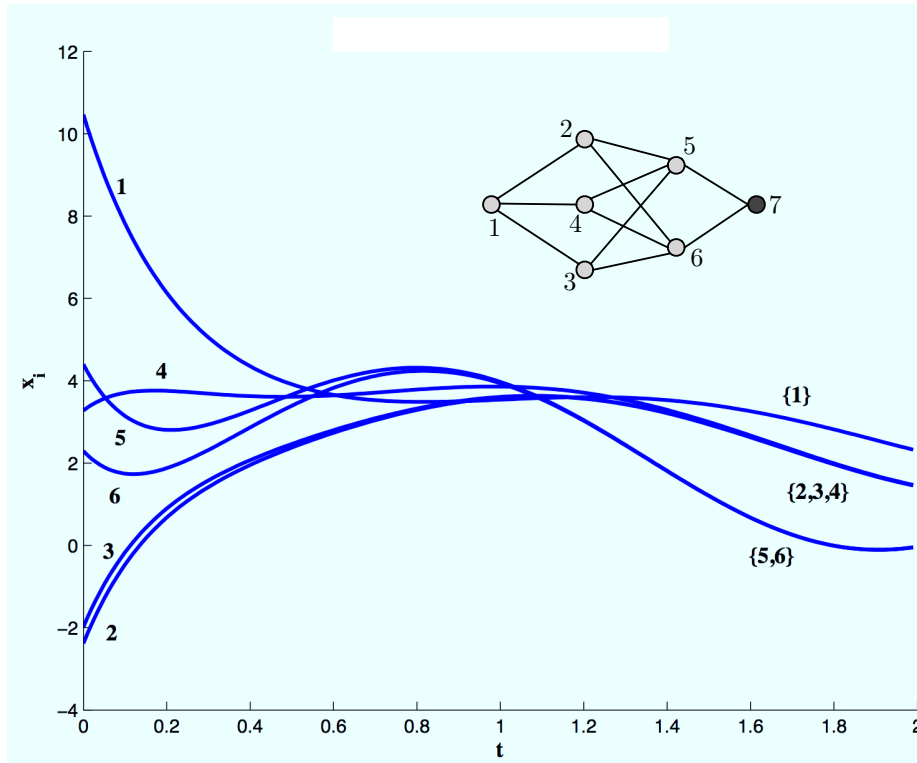


Figure 6: Asymptotically stable uncontrollable part of the dynamics. The uncontrollable part is given by the differences between state values inside the same cell in the maximal, leader-invariant EEP.




Figure 7: A graph G with an EEP π (left) and the resulting weighted and directed quotient graph $G \setminus \pi$ (right). For this quotient graph, we have $w_\pi(C_i, C_j) \neq w_\pi(C_j, C_i)$, that is, the edge weights are different along different directions.




Figure 8: An original graph (left) together with two graphs obtained through the removal of edges. As a result, the corresponding minimal, leader-invariant EEPs (leader node in black) lead to different quotient graphs (middle and right).




Figure 9: Chain graph. Control signals are injected at one of the boundary nodes and are propagated through the network.






Figure 10: m -chain graph.



(a)



(b)

Figure 11: Examples of augmented two-chain graphs that both have trivial maximal leader-invariant EEPs yet are not completely controllable.

References

- [1] F. Bullo, J. Cortes, and S. Martinez. *Distributed Control of Robotic Networks: A Mathematical Approach to Motion Coordination Algorithms*, Princeton University Press, 2009.
- [2] M. Cao, S. Zhang, and M. K. Camlibel. A Class of Uncontrollable Diffusively Coupled Multi-Agent Systems with Multi-Chain Topologies. Submitted to *IEEE Transactions on Automatic Control*, 2011.
- [3] D.G. Comeil and C.C. Cottleib. An efficient algorithm for graph isomorphism. *J. Assoc. Comput. Mach.*, Vol. 17, pp. 51-64, 1970.
- [4] M. Egerstedt. Controllability of Networked Systems. *Mathematical Theory of Networks and Systems*, pp. 57-61, Budapest, Hungary, 2010.
- [5] A. Fax and R. M. Murray. Graph Laplacian and Stabilization of Vehicle Formations. *Proceedings of the 15th IFAC World Congress*, pp. 283-288, 2002.
- [6] C.D. Godsil. Compact graphs and equitable partitions. *Linear Algebra and its Applications*, Vol. 255, pp. 259-266, 1997.
- [7] C. Godsil and G. Royle. *Algebraic Graph Theory*, Springer, New York, 2001.
- [8] A. Jadbabaie, J. Lin, and A.S. Morse. Coordination of Groups of Mobile Autonomous Agents using Nearest Neighbor Rules, *IEEE Transactions on Automatic Control*, Vol. 48, No. 6, pp. 988-1001, 2003.
- [9] Z. Lin, M. Broucke, and B. Francis. Local Control Strategies for Groups of Mobile Autonomous Agents. *IEEE Transactions on Automatic Control*, Vol. 49, No. 4, pp. 622-629, 2004.
- [10] Y.Y. Liu, J.J. Slotine, and A.L. Barabasi. Controllability of complex networks. *Nature*, Vol.

473, pp. 167173, 2011.

- [11] S. Martini, M. Egerstedt, and A. Bicchi. Controllability Decompositions of Networked Systems Through Quotient Graphs. *IEEE Conference on Decision and Control*, pp. 2717-2722, Cancun, Mexico, Dec. 2008.
- [12] M. Mesbahi and M. Egerstedt. *Graph Theoretic Methods for Multiagent Networks*, Princeton University Press, Princeton, NJ, 2010.
- [13] R. Olfati-Saber and R.M. Murray. Consensus Protocols for Networks of Dynamic Agents, *American Control Conference*, pp. 951-956, Denver, CO, June 2003.
- [14] A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt. Controllability of Multi-Agent Systems from a Graph-Theoretic Perspective. *SIAM Journal on Control and Optimization*, Vol. 48, No. 1, pp. 162-186, Feb. 2009.
- [15] W. Ren and R.W. Beard. *Distributed Consensus in Multi-vehicle Cooperative Control*, Communications and Control Engineering Series, Springer-Verlag, London, 2008.
- [16] H. G. Tanner. On the Controllability of Nearest Neighbor Interconnections. *Proceedings of the IEEE Conference on Decision and Control*, pp. 2467-2472, Dec. 2004.
- [17] H. Tanner, A. Jadbabaie, and G.J. Pappas. Flocking in Fixed and Switching Networks, *IEEE Transactions on Automatic Control*, Vol. 52, No. 5, pp. 863-868, 2007.
- [18] S. Zhang, K. Camlibel, and M. Cao. Controllability of Diffusively-Coupled Multi-Agent Systems with General and Distance Regular Coupling Topologies. *IEEE Conference on Decision and Control*, Orlando, FL, Dec. 2011.

Author Information

Magnus Egerstedt is a Professor in Electrical and Computer Engineering at the Georgia Institute of Technology. His research interests include hybrid and networked control, with applications in motion planning, control, and coordination of mobile robots. He is the director of the Georgia Robotics and Intelligent Systems Laboratory (GRITS Lab), a Fellow of IEEE, and serves as Editor for Electronic Publications for the IEEE Control Systems Society.

Simone Martini received the master degree in Computer Science Engineering from the University of Pisa in 2008 with a thesis developed at the School of Electrical and Computer Engineering at the Georgia Institute of Technology. Currently, he is a Ph.D. student in Robotics Automation and Bioengineering at the Interdepartmental Research Center “E.Piaggio” of the University of Pisa. His main research involves distributed control, and coordination of mobile robots.

Ming Cao is an Assistant Professor at the University of Groningen, the Netherlands. He received the Bachelor degree in 1999 and the Master degree in 2002 from Tsinghua University, China, and the PhD degree in 2007 from Yale University, USA, all in electrical engineering. From September 2007 to August 2008, he was a Research Associate at Princeton University, USA. His main research interest is in autonomous agents and multi-agent system.

Kanat Camlibel received the B.Sc. and M.Sc. degrees in control and computer engineering from the Istanbul Technical University, in 1991 and 1994, respectively, and the Ph.D. degree from Tilburg University, in 2001. He currently holds an assistant professor position at the University of Groningen. His main research is focused on hybrid systems and multi-agent systems.

Antonio Bicchi is a Professor of Systems Theory and Robotics at the University of Pisa, the

Director of the Research Center “E. Piaggio” in Pisa , and a Senior Scientist at the Italian Institute of Technology, Genoa, Italy. His main research interests are in the theory and control of nonlinear systems, in particular hybrid systems; dynamics, kinematics and control of complex mechanical systems; haptics and dexterous manipulation. He is a Fellow of IEEE.