

Structurally Observable Distributed Networks of Agents under Cost and Robustness Constraints

Stephen Kruzick, Sérgio Pequito, Soumya Kar, José M. F. Moura, and A. Pedro Aguiar

Abstract— In many problems, agents cooperate locally so that a leader or fusion center can infer the state of every agent from probing the state of only a small number of agents. Versions of this problem arise when a fusion center reconstructs an extended physical field by accessing the state of just a few of the sensors measuring the field, or a leader monitors the formation of a team of robots. Given a link cost, the paper presents a polynomial time algorithm to design a minimum cost coordinated network dynamics followed by the agents, under an observability constraint. The problem is placed in the context of structural observability and solved even when up to k agents in the coordinated network dynamics fail.

Index Terms— distributed networks, linear dynamics, structural systems, observability, minimum cost, Steiner subgraph

I. INTRODUCTION

IT is important in many distributed settings to appreciate the impact of local cooperation. Much work has been done in distributed inference [1], [2], distributed diffusion [3], and distributed optimization [4], [5], to name a few applications and references. This paper is concerned with a design question under the following different but related scenario. A group of agents share a common goal. They cooperate at a cost to update their individual state, but only a few agents share their state with a leader or fusion center. This leads to network dynamics that we will describe more formally below. The mission of the leader is to reconstruct the network state, i.e., the state of all the agents. We consider how to design the minimal cost network of (local) cooperation among agents and with which agents the leader interacts, i.e., the design of minimal cost coordinated network dynamics, subject to the constraint that the leader can reconstruct the network state—an observability constraint. Before being more specific, we illustrate with two motivating examples. The first is static, while the second is dynamic like the problem we address. References [6], [7] consider the problem of reconstructing a field from information (state) provided by only a few agents (sensors) in a network of sensors. For instance, the field may

be the contamination level of a pollutant over a wide area. It turns out that, through local cooperation (sensors updating iteratively their state with the state of their neighbors), the fusion center can, under appropriate conditions, reconstruct (through a basis pursuit type algorithm) the entire field from a snapshot of the states of a very sparse subset of the agents. A question of interest is to design the cooperation dynamics and to determine a minimal set of sensors to be probed so that a fusion center can reconstruct the field.

The second considers formation with flocks of birds or schools of fish that can be described by models in which each individual constantly monitors or senses the distance to neighboring birds or fish and reacts to ensure that it maintains a suitable position. For example, each individual might accelerate to move closer if the distance to neighbors increases or slow down if the distance to neighbors becomes dangerously close. The collective networked dynamics of birds or fish and their qualitative properties are discussed in [8], [9], [10], among others. These references provide several interesting and revealing examples. The self-organization of these adaptive networks has also been studied in [11]. Similarly, in several robotic problems like in formation control of multi-robot teams and in the coordination of groups of mobile autonomous agents, robots move by approaching or gaining distance from their neighbors in order to maintain the formation. This leads the multi-robot team to follow coordinated network dynamics. Furthermore, the robots do not explicitly communicate with each other in many of these applications. Each robot finds its own relevant information with respect to neighboring robots [12], for example, range and bearing from on-board sensors, such as cameras. However, for reasons derived from the application, the robots do not directly communicate this information to other individuals. The resulting networked dynamics have been used in the literature to obtain simplified linear distributed control or update laws achieving mission specific formation or flocking objectives [13], [14]. Finally, we assume that a leader is charged with monitoring that the robots remain in formation by probing only the state of a few agents and then computing the (global) network state. This task is trivial if the fusion center accesses every agent or node in the network. It becomes interesting when at the same time it is desirable to reduce the costs to a certain minimum, see below, including agents measuring their relative position to neighbors and the direct access to node states by the leader. It is then of interest to design the network structure of the coordinated dynamics and to determine which node states should be measured by the leader so that it can track the states of all the nodes in the network while minimizing the cost of

Stephen Kruzick (skruzick@andrew.cmu.edu), Sérgio Pequito (sergo@seas.upenn.edu), Soumya Kar (soumyyak@andrew.cmu.edu), and José M. F. Moura (moura@ece.cmu.edu) affiliate with the Department of Electrical and Computer Engineering at Carnegie Mellon University in Pittsburgh, PA, USA. A. Pedro Aguiar (pedro.aguiar@fe.up.pt) affiliates with the Department of Electrical and Computer Engineering at University of Porto in Porto, Portugal. The work of Kruzick was partially conducted with government support under and awarded by the DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a and by NSF grants CCF-1018509, CCF-1011903, and CCF-1513936. The work of Pequito was partially supported by grant SFRH/BD/33779/2009 from Fundação para a Ciência e a Tecnologia (FCT) and by the CMU-Portugal (ICTI) program.

network coordination and measurements.

The above are difficult questions. We cast problems of these types below in the context of structural observability, which entails designing the structure of a dynamical system with observable dynamics, a framework studied in [15], [16], [17]. Once the system is observable, the actual problem of continuously monitoring the state of the nodes by the remote fusion station can then be achieved by implementing a recursive observer such as, for example, of the Luenberger type [18].

Remark 1 *We emphasize that there may not be explicit communication among the nodes in the applications we envision, even though we use the consensus context to describe the coordination among agents. Accordingly, in this paper, we replace communication costs usually assumed in consensus problems by link costs. Link costs subsume costs associated with local neighbor sensing interactions (proxies for costs associated with sensing and extracting from measurements the state of neighbors), as well as costs incurred by a fusion center to learn (also, possibly by sensing) the current state (position) of a few selected agents. For simplicity, the interactions between the fusion center and nodes are described as through a backbone network.*

We model the coordinated network dynamics followed by the aggregate of all the agents by a linear system

$$\mathbf{x}(n+1) = A\mathbf{x}(n) \quad (1)$$

$$\mathbf{y}(n) = C\mathbf{x}(n). \quad (2)$$

The system dynamics matrix A captures the network graph of interactions among agents. Its nonzero entries represent interaction links between corresponding nodes or agents. Each agent or sensor node maintains a single scalar state variable $x_i(n)$ initially set to $x_i(0)$. The vector of agent or sensor states at a given iteration n is denoted by $\mathbf{x}(n)$. The nodes update their states according to the coordinated dynamics given by (1), with initial state vector $\mathbf{x}(0)$. The nodes that are monitored by the fusion center, referred to as backbone nodes, are collected in the output vector $\mathbf{y}(n)$ given by (2). Because network structure must be respected, disallowed interactions among nodes or among the fusion center and nodes restrict corresponding entries of the A and C matrices to equal zero, whereas entries corresponding to network links represent design parameters.

We use structural systems theory, a survey of which can be found in [19], to find network dynamics (A, C) requiring a minimum set of link costs while guaranteeing that the initial state can be recovered from the backbone outputs collected over time at the fusion center despite some number of sensor node failures, which are known to the central node and occur before the dynamics begin. Networks that operate according to this framework are discussed in [15]. This reference provides a necessary topological condition for the initial state to be reconstructed from the backbone outputs for a specific choice of dynamics related to the network graph structure. These networks are also studied under a structural systems context in [16] and [17], which describe types of networks where this

can be achieved. Our paper focus on designing robust networks that operate according to the above model and that minimize a specified cost function. This work extends a preliminary version of the optimal network dynamics design problem that we previously introduced in [20], by considering arbitrary costs, arbitrary backbone topologies, and sensing robustness requirements. From a technical standpoint, the methodology and combinatorial optimization tools are also more general and computationally efficient. Practical implementable dynamic systems of the optimal network structures that we obtain rely on results presented in [17]. While the optimization problem involving cost minimization under robustness constraints that appears in this paper does not have a direct comparison in the existing structural systems literature, other related works regarding structural systems may be of interest for further reading. The structural systems framework for networks appears in [15], [16], and [17]. Other optimal network design problems involving this structural systems framework with significantly differing objectives and constraints can be found in [21], [22], and [23].

The paper is organized as follows. Section II introduces key background concepts, including relevant information from graph theory, combinatorial optimization, and systems theory. Section III formally describes our network operation model. Section IV examines our minimum cost design problem and provides a solution algorithm followed by a proof of correctness and practical discussion. Finally, Section V concludes the paper.

II. BACKGROUND CONCEPTS

This section provides supporting background information. Of particular importance are graph theory definitions and concepts used to describe network structure as well as topics from combinatorial optimization relevant to the minimum cost design problem solutions. Network operation closely relates to system theory concepts, which are also introduced. Finally, results from structural system theory provide a bridge between network structure, described by graphs, and network operation, described by dynamical systems.

A. Graph Theory Concepts

This section introduces terminology and concepts regarding graphs. A **directed graph** \mathcal{G} is an ordered pair $(\mathcal{V}, \mathcal{E})$ in which \mathcal{V} denotes a set of **nodes (vertices)** and \mathcal{E} denotes a set of directed **links (edges)**. These directed links are ordered pairs (v_i, v_j) of nodes $v_i, v_j \in \mathcal{V}$. Note that **self-loops**, links formed as (v, v) for $v \in \mathcal{V}$, are not excluded from this definition. Furthermore, an **directed, weighted graph** \mathcal{G} is the ordered triple $(\mathcal{V}, \mathcal{E}, w)$ in which $w : \mathcal{E} \rightarrow \mathbb{R}^+$ assigns a cost or **weight** $w(e)$ to each link $e \in \mathcal{E}$. Any graph $\mathcal{G}_S = (\mathcal{V}_S, \mathcal{E}_S)$ with $\mathcal{V}_S \subseteq \mathcal{V}$ and $\mathcal{E}_S \subseteq \mathcal{E}$ is called a **subgraph** of \mathcal{G} . Furthermore, if $\mathcal{V}_S = \mathcal{V}$, then \mathcal{G}_S **spans** \mathcal{G} . Two graphs $\mathcal{G}_1 = (\mathcal{V}_1, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{V}_2, \mathcal{E}_2)$ are said to be **isomorphic**, written as $\mathcal{G}_1 \simeq \mathcal{G}_2$, if there is bijective function $f : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ such that $(u, v) \in \mathcal{E}_1$ if and only if $(f(u), f(v)) \in \mathcal{E}_2$.

A sequence $(v_1, v_2), (v_2, v_3), \dots, (v_{k-1}, v_k)$ of directed links in which the **head** (destination) node of the previous link is the

tail (origin) node of the subsequent link constitutes a **directed path**. Provided $v_i \neq v_j$ for all $i \neq j$, it comprises an **directed elementary path**. When $v_1 = v_k$ but all other nodes are distinct, the path forms a **directed cycle**. Two directed paths are **internally node-disjoint** if they share no nodes apart from the start node and end node. Likewise, two directed paths are **link-disjoint** if they share no links. The **directed local node-connectivity** $\kappa_{\mathcal{G}}(u, v)$ from node u to node v gives the minimum number of nodes that must be removed from the graph \mathcal{G} such that there is no directed path from u to v , equal to the number of internally node-disjoint directed paths from u to v . The **directed local link-connectivity** $\lambda_{\mathcal{G}}(u, v)$ from node u to node v gives the minimum number of links that must be removed from the graph such that there is no directed path from u to v , equal to the number of link-disjoint directed paths from u to v .

Several notions of connectedness exist for directed graphs. In this paper, a specially labeled **root** node $r \in \mathcal{V}$ is given, and a graph is **r-rooted connected** if there is an elementary directed path from each node $v \in \mathcal{V}$ to r . An **arborescence**, also known as a directed rooted tree, is a directed graph in which there exists exactly one elementary directed path from each node to r . A collection of disjoint arborescences with root nodes collected into a set R is called a **branching**. A directed graph is **r-rooted k-node-connected** if removing fewer than k nodes leaves at least one elementary directed path from each node to r , that is $\kappa_{\mathcal{G}}(v, r) \geq k$ for all $v \in \mathcal{V}$. Similarly, a directed graph is **r-rooted k-link-connected** if removing fewer than k links leaves at least one elementary directed path from each node to r , that is $\lambda_{\mathcal{G}}(v, r) \geq k$ for all $v \in \mathcal{V}$.

B. Optimal Connectivity Problems

Network design problems often involve finding optimal subgraph structures within a graph of possible network connections with the restriction that some notion of connectivity be ensured. The following discussion introduces problems involving optimality objectives and connectedness requirements that will prove useful in solving the problem presented in Section IV.

Consider identification of the least costly set of directed links that connects all the nodes of a network to a root node. The **minimum spanning arborescence**, also known as a minimum directed rooted spanning tree, of a directed, weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$ rooted at a node $r \in \mathcal{V}$ formalizes this concept. The problem consists of finding a subgraph of minimum cost $\text{MStA}(\mathcal{G}, r)$ that has exactly one directed path from each node $v \in \mathcal{V}$ to r . Note that the solution must be an r -rooted arborescence and that \mathcal{G} must be r -rooted connected for a solution to exist. The minimum spanning arborescence can be computed by the Chu-Liu/Edmonds Algorithm, which can be implemented with complexity $O(|\mathcal{E}| + |\mathcal{V}| \log |\mathcal{V}|)$ [24]. Generalizations for higher connectivity, such as the minimum r -rooted k -node-connected spanning subgraph and the minimum r -rooted k -link-connected spanning subgraph are similarly defined. Unlike the corresponding undirected counterparts, these directed, rooted problems are solvable in polynomial time using a maximum weight matroid intersection formulation [25], [26].

In a problem related to the minimum spanning arborescence, consider identification of the least costly set of directed links that connect a required subset of the network nodes to a root node while possibly involving other nodes. The general **minimum Steiner arborescence** problem, also known as the minimum directed rooted Steiner tree, formalizes this concept. Let a directed, weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$ with root node $r \in \mathcal{V}$ be given along with a set $\mathcal{S} \subseteq \mathcal{V}$ of **terminal nodes**. The problem consists of finding a minimum cost subgraph $\text{MStA}(\mathcal{G}, \mathcal{S}, r)$ of \mathcal{G} that has exactly one directed path from each node $v \in \mathcal{S}$ to r . Note that non-terminal nodes in $\mathcal{V} \setminus \mathcal{S}$ may or may not be used, and that a path must exist from each $v \in \mathcal{V}$ to r for a solution rooted at r to exist. The general minimum Steiner arborescence problem is NP-hard, so no polynomial time solution algorithm is known. However, polynomial time approximation algorithms with nontrivial performance guarantee ratios exist and may be used to obtain approximate solutions [27].

A generalization to greater connectivity requirements, the **minimum r-rooted k-node-connected Steiner subgraph** problem, provides the minimum cost connectivity problem of greatest utility to this paper. For a directed, weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$, terminal node set \mathcal{S} , and root node r , the problem consists of finding a minimum cost subgraph $\text{MStNCS}(\mathcal{G}, \mathcal{S}, r, k)$ that has k internally node-disjoint directed paths from each node $v \in \mathcal{S}$ to r . Note that non-terminal nodes in $\mathcal{V} \setminus \mathcal{S}$ do not necessarily have k node-disjoint paths to r . It is clear that at least k internally node-disjoint paths must exist from each $v \in \mathcal{S}$ to r for a solution rooted at r to exist. The problem is, in general, at least as computationally difficult as the minimum Steiner arborescence problem. However, in the particular restricted case in which all links with positive weight, called **augmenting links**, originate in \mathcal{S} , the problem is solvable in polynomial time using a submodular flow algorithm [26].

For networks that engage in data forwarding, the problem of finding the path between two nodes with the least costly total link sum often finds relevance. The concept of each node connecting to a root node over the least costly path possible is formalized by the **shortest path spanning arborescence**, which will also contribute to solving the design problem in Section IV. Let a directed, weighted, and r -rooted connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$ with nonnegative link weights w and a root node $r \in \mathcal{V}$ be given. The problem consists of finding a path from each $v \in \mathcal{V}$ to the root node r that has the minimum possible link sum. The solution set of links is not necessarily unique. However, at least one possible set of links forms a spanning tree called the shortest path spanning arborescence $\text{SPSpA}(\mathcal{G}, r)$ rooted at r . The shortest path lengths, using which the arborescence may be built starting at the root, can be computed using methods such as Dijkstra's algorithm, which can be implemented with complexity $O(|\mathcal{E}| + |\mathcal{V}| \log |\mathcal{V}|)$ [28].

C. System Theory Concepts

The weighted networks discussed in this paper follow the coordinated dynamics given by the discrete time dynamical system (1) and (2). The concept of system observability plays a key role in understanding the communication scheme

under consideration. A system is said to be **observable** if the initial state $\mathbf{x}(0)$ can be uniquely determined from the collected outputs over a finite time period with knowledge of the system matrices and input values. For discrete linear, time-invariant systems with zero-input such as in (1)-(2), the output at time n is $\mathbf{y}(n) = CA^n\mathbf{x}(0)$. The collected outputs $\mathbf{y} = [\mathbf{y}^\top(0), \dots, \mathbf{y}^\top(N-1)]^\top$ over N iterations are described by a linear transformation of the initial state

$$\mathbf{y} = O_{(A,C)}\mathbf{x}(0) \quad (3)$$

where $O_{(A,C)}$ is the **observability matrix**

$$O_{(A,C)} = [(CA^0)^\top \quad \dots \quad (CA^{N-1})^\top]^\top \quad (4)$$

and N is the number of state variables. Observability of a system can be inferred from the observability matrix as stated in Theorem 1.

Theorem 1 (Observability [29]) *The N state linear, time-invariant system described by matrices (A, C) is observable if and only if $\text{rank}(O_{(A,C)}) = N$.*

In contexts involving an interval $[0, T-1]$ other than the first N iterations, the following notations apply. The collected outputs $\mathbf{y}_{[0, T-1]} = [\mathbf{y}^\top(0), \dots, \mathbf{y}^\top(T-1)]^\top$ with no input over T iterations are described by a linear transformation of the initial state

$$\mathbf{y}_{[0, T-1]} = O_{(A,C), [0, T-1]}\mathbf{x}(0) \quad (5)$$

such that

$$O_{(A,C), [0, T-1]} = [(CA^0)^\top \quad \dots \quad (CA^{T-1})^\top]^\top \quad (6)$$

with $[0, T-1]$ explicitly written for specificity.

A well known dual to the concept of observability is the concept of controllability for systems driven by inputs. Under appropriate conditions on the system matrices, these concepts are dual and results for observability hold with corresponding adaptation to controllability. This paper is primarily concerned with ensuring that the network systems designed are observable. Duality allows the techniques developed here to apply in situations requiring the design of controllable coordinated network systems. In the subsequent text, we consider only observability and will not address controllability.

D. Structural System Theory

Because the dynamical system that describes network coordination must respect local network connections as represented in a directed graph, analysis of how network structure affects system properties provides design insights. Denote by (\tilde{A}, \tilde{C}) a pair of structural matrices composed of entries that are zero or one, with $\tilde{A} \in \{0, 1\}^{N \times N}$ and $\tilde{C} \in \{0, 1\}^{M \times N}$. Structural system theory examines the general system properties of all dynamic matrix pairs (A, C) that respect the structure (\tilde{A}, \tilde{C}) in the following sense. An entry of (A, C) is zero if the corresponding entry of (\tilde{A}, \tilde{C}) is zero, while an entry of (A, C) is an arbitrary parameter if the corresponding entry of (\tilde{A}, \tilde{C}) is one. In a straightforward way, (\tilde{A}, \tilde{C}) may be summarized by a directed graph $\mathcal{D}(\tilde{A}, \tilde{C})$. Nodes of the graph for each of the N system state variables have connections described by

\tilde{A} , with a directed link from state j to state i if and only if $\tilde{A}_{ij} = 1$. Nodes of the graph for each of the M system output variables have connections to states as described by \tilde{C} , with a directed link from state j to output i if and only if $\tilde{C}_{ij} = 1$.

Specifically, conditions on the network structure that ensure recoverability of the initial state are desired. Several results concerning structural systems are discussed in the survey paper [19], including conditions for structural observability. The pair (\tilde{A}, \tilde{C}) is said to be **structurally observable** if there is an observable pair (A, C) that respects the structure (\tilde{A}, \tilde{C}) [30]. Additionally, if (\tilde{A}, \tilde{C}) is structurally observable, nearly all realizations that respect the structure over suitable fields, such as \mathbb{R} or \mathbb{C} , are observable in the sense that the set of feasible unobservable realizations must have measure zero [30].

A result critical to this work, Theorem 2 shows that structural observability of (\tilde{A}, \tilde{C}) is equivalent to the existence of an output cactus patch, a specific type of graph defined below, that spans $\mathcal{D}(\tilde{A}, \tilde{C})$ [19]. **Output cacti** are defined recursively. Consider a collection of nodes labeled either as state nodes or as output nodes. An **output stem** graph consists of an output node and a single directed elementary path potentially containing several state nodes rooted at the output node, with all links directed toward the output node. All output stems are defined to be output cacti. Furthermore, any existing output cactus to which a directed cycle of state nodes has been attached via a directed link from one node in the cycle to any node of the output cactus is also an output cactus. Finally, a union of node-disjoint output cacti is an **output cactus patch**.

Theorem 2 (Structural Observability [30][31]) *The structural system matrix pair (\tilde{A}, \tilde{C}) is structurally observable if and only if the structural system graph $\mathcal{D}(\tilde{A}, \tilde{C})$ is spanned by an output cactus patch.*

III. FORMAL NETWORK DESCRIPTION

To formally describe the operation of the networks under consideration, we distinguish between two closely related network structures, namely, the physical network \mathcal{G}_P describing link connections and the underlying dynamic system network \mathcal{G}_D describing the relationships among sensor states and network outputs. The physical network is composed of three types of nodes: a set \mathcal{X} of **sensor nodes**, a set \mathcal{Q} of **backbone nodes**, and a **fusion center node** $\mathcal{Z} = \{z\}$. In this network, sensor nodes contain data and form a local state dynamics. Backbone nodes accept outputs from the sensor nodes and route them to the fusion center node, which aggregates data for state observation tasks. Network connectivity is described by a directed, weighted graph $\mathcal{G}_P = (\mathcal{V}_P, \mathcal{E}_P, w_P)$ where $\mathcal{V}_P = \mathcal{X} \cup \mathcal{Q} \cup \mathcal{Z}$ is the set of all nodes, \mathcal{E}_P is the set of directed feasible links, and $w_P : \mathcal{E}_P \rightarrow \mathbb{R}^+$ is the weight function describing individual directed link cost. Self-links are assumed to always be possible at zero-cost $w_P(x, x) = 0$ for all $x \in \mathcal{X}$. Figure 1 provides an example of such a network, showing the sensor nodes (black), backbone nodes (green), and fusion center node (red). Conceptually, this network divides into two subnetworks, the sensing subnetwork and the backbone subnetwork.




Fig. 1: Physical connectivity graph \mathcal{G}_P showing feasible links among sensor nodes \mathcal{X} (black), backbone nodes \mathcal{Q} (green), and central fusion node \mathcal{Z} (red).




Fig. 2: Dynamic system graph \mathcal{G}_D showing feasible computation links among state nodes \mathcal{X} (black), output nodes \mathcal{Y} (blue), and central fusion node \mathcal{Z} (red).

The **sensor subnetwork** $\mathcal{G}_S = (\mathcal{V}_S, \mathcal{E}_S, w)$ consists of all directed links between two sensor nodes or from a sensor node to a backbone node. Thus, $\mathcal{V}_S = \mathcal{X} \cup \mathcal{Q}$ and $\mathcal{E}_S = \mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Q}}$ where $\mathcal{E}_{\mathcal{X}\mathcal{X}} = \mathcal{E}_P \cap (\mathcal{X} \times \mathcal{X})$ and $\mathcal{E}_{\mathcal{X}\mathcal{Q}} = \mathcal{E}_P \cap (\mathcal{X} \times \mathcal{Q})$. Each sensor $x_i \in \mathcal{X}$ makes a single initial scalar measurement α_i from field \mathbb{F} and maintains a single scalar variable of state $\mathbf{x}_i(n)$ over discrete time, with initial state $\mathbf{x}_i(0) = \alpha_i$. At every iteration, the network updates the state variable of each sensor node using local data according to a linear combination as described by a matrix $A \in \mathbb{F}^{N \times N}$ such that $\mathbf{x}(n+1) = A\mathbf{x}(n)$ where $N = |\mathcal{X}|$. Each row of A indicates which neighboring states are used to compute the updated value of a given state and the coefficients of the linear combination. In this way, each sensor need only know the values in the corresponding row of A rather than the entire network topology and dynamics. When $A_{ij} \neq 0$, sensor node x_j is linked to x_i , incurring cost. Otherwise, no link exists between the two nodes.

The **backbone subnetwork** $\mathcal{G}_B = (\mathcal{V}_B, \mathcal{E}_B, w)$ consists of all directed links between two backbone nodes or from a backbone node to the fusion center. Thus, $\mathcal{V}_B = \mathcal{Q} \cup \mathcal{Z}$ and $\mathcal{E}_B = \mathcal{E}_{\mathcal{Q}\mathcal{Q}} \cup \mathcal{E}_{\mathcal{Q}\mathcal{Z}}$ where $\mathcal{E}_{\mathcal{Q}\mathcal{Q}} = \mathcal{E}_P \cap (\mathcal{Q} \times \mathcal{Q})$ and $\mathcal{E}_{\mathcal{Q}\mathcal{Z}} = \mathcal{E}_P \cap (\mathcal{Q} \times \mathcal{Z})$. Some of the states are output (sensed) through the network backbone nodes in \mathcal{Q} and available at the central node z as the vector $\mathbf{y}(n)$. It is assumed that the backbone subnetwork operates much faster than the sensor subnetwork iterations, such that $\mathbf{y}(n)$ is available to z without delay. For link cost efficiency, the backbone network uses the least costly path to the central node. These outputs can be described by a matrix $C \in \mathbb{F}^{M \times N}$ such that $\mathbf{y}(n) = C\mathbf{x}(n)$ where $M = |\mathcal{E}_{\mathcal{X}\mathcal{Q}}|$ is the number of feasible connections between the sensors and backbone nodes. Each row i of C either is composed of all zeros, indicating an output is not made over the corresponding connection and incurring no cost, or has a single nonzero entry in column j , indicating state x_j is output over the connection and incurring cost.

Consider the system dynamics network \mathcal{G}_D that describes the relationships among sensor states, backbone outputs, and fusion center in the physical network. The sensor nodes \mathcal{X} of the physical network correspond to the **state nodes**, and

maintain the same links $\mathcal{E}_{\mathcal{X}\mathcal{X}}$ with $w_D(x_1, x_2) = w_P(x_1, x_2)$ for $x_1, x_2 \in \mathcal{X}$. Additionally, the fusion center node $\mathcal{Z} = \{z\}$ appears in this network. Because the backbone nodes may output the states from multiple sensor nodes, backbone nodes do not correspond to only one output. Rather, each link in $\mathcal{E}_{\mathcal{X}\mathcal{Q}}$ from a sensor node to a backbone node represents a potential output of the system. Thus, let a set \mathcal{Y} of **output nodes** be indexed by $\mathcal{E}_{\mathcal{X}\mathcal{Q}}$ such that $y_{(x,q)} \in \mathcal{Y}$ corresponds to $(x, q) \in \mathcal{E}_{\mathcal{X}\mathcal{Q}}$. Each output node $y_{(x,q)}$ will be linked to the corresponding sensor node x with weight $w_D(x, y_{(x,q)}) = w_P(x, q)$ and to the central node with weight $w_D(y_{(x,q)}, z) = w_{\text{sp}}(q, z; \mathcal{E}_B)$ where $w_{\text{sp}}(q, z; \mathcal{E}_B)$ is the weight of the shortest path from q to z over the backbone links \mathcal{E}_B . Hence, $\mathcal{E}_{\mathcal{X}\mathcal{Y}} = \{(x, y_{(x,q)}) | (x, q) \in \mathcal{E}_{\mathcal{X}\mathcal{Q}}\}$ and $\mathcal{E}_{\mathcal{Y}\mathcal{Z}} = \{(y_{(x,q)}, z) | (x, q) \in \mathcal{E}_{\mathcal{X}\mathcal{Q}}\}$. Note that state nodes may connect to multiple output nodes, but each output node corresponds to only one state node. Also, note that output nodes do not have links to each other, so the output nodes together with the fusion center form a star topology subnetwork. Thus, the graph that describes the state and output update dependencies is $\mathcal{G}_D = (\mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}, \mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Y}} \cup \mathcal{E}_{\mathcal{Y}\mathcal{Z}}, w_D)$. As with the physical network, a link only contributes cost to the network operation if it is required for updating states and outputs. The nodes, links, and weights of this network are entirely derived from the physical network \mathcal{G}_P , and Figure 2 provides an example of such a modification. The figure shows the organization into state nodes (black), output nodes (blue), and the central node (red).

The network state $\mathbf{x}(n)$ and the network output $\mathbf{y}(n)$ behave according to the dynamical system (1)-(2) with $\mathbf{x}(0) = \alpha$ being the initial state. Observability of the system must be ensured when designing the matrix pair (A, C) in (1)-(2), which translates to the full rank condition on $O_{(A,C)}$. Furthermore, the fusion center must be aware of the entire network dynamics so that the observability matrix is known and state observation tasks may be performed.

Additionally, each node must only access the state from its neighbors in \mathcal{G}_P to update its state or output values but does not necessarily use information from every neighbor. Let $\mathcal{G}_P(A, C) = (\mathcal{V}_P, \mathcal{E}_P(A, C), w_P)$ be the subgraph of \mathcal{G}_P defined by the set $\mathcal{E}_P(A, C)$ of physical links required for the computation of (1)-(2) and outputs along the least costly backbone path. Note that $\mathcal{E}_P(A, C) = \mathcal{E}_{\mathcal{X}\mathcal{X}}(A) \cup \mathcal{E}_{\mathcal{X}\mathcal{Q}}(C) \cup \mathcal{E}_B(C)$ with links in $\mathcal{E}_{\mathcal{X}\mathcal{X}}(A)$ used to compute (1), links in $\mathcal{E}_{\mathcal{X}\mathcal{Q}}(C)$ used in (2), and links in $\mathcal{E}_B(C)$ used in the shortest path from every used backbone node q directly connected to some sensor x by $(x, q) \in \mathcal{E}_{\mathcal{X}\mathcal{Q}}(C)$. For feasibility of network operations, the subset constraints $\mathcal{E}_{\mathcal{X}\mathcal{X}}(A) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{X}}$, $\mathcal{E}_{\mathcal{X}\mathcal{Q}}(C) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{Q}}$, and $\mathcal{E}_B(C) \subseteq \mathcal{E}_B$ must hold.

Hence, these link feasibility and system observability constraints must be applied when designing the network dynamics (A, C) . Furthermore, minimization of the sum total cost

$$F(A, C) = \sum_{e \in \mathcal{E}_P(A, C)} w_P(e) \quad (7)$$

of all used network physical links is desired, including sensor state updates, sensor output production, and backbone activity.

Section IV formulates a robust version of the minimum cost network design problem with this objective function and set of constraints.

IV. NETWORK DESIGN PROBLEM

The problem examined in this paper concerns optimal design of networks operating according to the description in Section III. As equation (7) shows, the underlying system dynamics determines the physical link cost of operating such networks. A less complete formulation of this problem appeared in [20], which provides an efficient solution algorithm based on minimum spanning tree methods for the simpler and much more fragile case in which only one backbone node exists and no node failures may occur. In contrast, this paper significantly generalizes this previous work by formulating the optimal design problem in the context of nontrivial backbone subnetworks and sensor node failures.

Therefore, this section examines minimization of the objective function in (7) with respect to the dynamic matrices (A, C) subject to the aforementioned system observability and link-connectivity constraints under added robustness requirements arising from sensor failures. Specifically, consider optimal design of a link feasible network where system observability must be guaranteed, while any subset of sensor nodes $\mathcal{U} \subset \mathcal{X}$ of size $|\mathcal{U}| \leq k$ less than the robustness design parameter k experiences total failure before coordination begins. In this way, the original states of nodes that did not fail could still be recovered from the collected outputs. Sensor node failure results in the elimination of associated rows and columns in the network dynamics matrices. Note that the row of A , columns of A , and columns of C are indexed by the sensor nodes \mathcal{X} , while the rows of C are indexed by $\mathcal{E}_{\mathcal{X}\mathcal{Q}}$. Denote by $A_{[\mathcal{X}\setminus\mathcal{U}, \mathcal{X}\setminus\mathcal{U}]}$ and $C_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X}\setminus\mathcal{U}]}$ submatrices of A and C , respectively, in which rows and columns corresponding to \mathcal{U} have been removed. This yields the optimization problem appearing below.

$$\underset{A, C}{\operatorname{argmin}} F(A, C) \quad (8)$$

$$\text{s. t.} \quad \mathcal{E}_{\mathcal{X}\mathcal{X}}(A) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{X}} \quad (9)$$

$$\mathcal{E}_{\mathcal{X}\mathcal{Q}}(C) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{Q}} \quad (10)$$

$$\mathcal{E}_B(C) \subseteq \mathcal{E}_B \quad (11)$$

$$(A_{[\mathcal{X}\setminus\mathcal{U}, \mathcal{X}\setminus\mathcal{U}]}, C_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X}\setminus\mathcal{U}]}) \quad (12)$$

observable

for all $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| \leq k$

The optimization described by (8)-(12) appears seemingly difficult due to the large number of observability constraints that must be satisfied. Recall that Theorem 1 states the equivalence between observability of (A, C) and the condition $\operatorname{rank}(O_{(A, C)}) = N$. Thus, the constraint given in (12) translates to $\operatorname{rank}(O_{(A_{[\mathcal{X}\setminus\mathcal{U}, \mathcal{X}\setminus\mathcal{U}]}, C_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X}\setminus\mathcal{U}]})}) = N - |\mathcal{U}|$ for all $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| \leq k$. Because there are potentially very many ways to form subsets of \mathcal{X} of size at most k , a large number of rank constraints must be satisfied. As such, it is not immediately clear that this problem is efficiently solvable. Furthermore, it is not obvious how to check that a solution

exists for a given k or to determine the largest k for which a solution exists. Later in the paper, after additional discussion of this problem, Remark 2 addresses existence of the solution in terms of flow problems. However, the optimization problem must first be related more closely to the combinatorial structure of the graph.

A structural systems approach renders the problem tractable by decoupling the problem of finding the optimal structure for the dynamic matrices from the problem of finding an observable instantiation of the optimal structure. Note that the objective function in (8) only depends on which links were used, as do the link feasibility constraints (9)-(12). Thus, it depends on only the zero-nonzero structure (\tilde{A}, \tilde{C}) of (A, C) and can be restated in structural terms. By the definition of structural observability, if (\tilde{A}, \tilde{C}) is structurally observable then an observable instantiation (A, C) respecting that structure must exist and may be found subsequently. Therefore, the final constraint can be rewritten to guarantee structural observability under limited node failures. The problem appears below with this reformulation.

$$\underset{\tilde{A}, \tilde{C}}{\operatorname{argmin}} F(\tilde{A}, \tilde{C}) \quad (13)$$

$$\text{s. t.} \quad \mathcal{E}_{\mathcal{X}\mathcal{X}}(\tilde{A}) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{X}} \quad (14)$$

$$\mathcal{E}_{\mathcal{X}\mathcal{Q}}(\tilde{C}) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{Q}} \quad (15)$$

$$\mathcal{E}_B(\tilde{C}) \subseteq \mathcal{E}_B \quad (16)$$

$$(\tilde{A}_{[\mathcal{X}\setminus\mathcal{U}, \mathcal{X}\setminus\mathcal{U}]}, \tilde{C}_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X}\setminus\mathcal{U}]}) \quad (17)$$

structurally observable

for all $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| \leq k$

Replacement of the observability constraint with a structural observability constraint may not appear, at first, to suggest a solution. However, by appealing to Theorem 2, the structural observability constraint can be replaced by the equivalent condition that the associated directed graph be spanned by an output cactus patch after any set of at most k sensor node deletions are applied. Hence, the original analytic rank constraint has been transformed to a combinatorial structural constraint in the problem below.

$$\underset{\tilde{A}, \tilde{C}}{\operatorname{argmin}} F(\tilde{A}, \tilde{C}) \quad (18)$$

$$\text{s. t.} \quad \mathcal{E}_{\mathcal{X}\mathcal{X}}(\tilde{A}) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{X}} \quad (19)$$

$$\mathcal{E}_{\mathcal{X}\mathcal{Q}}(\tilde{C}) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{Q}} \quad (20)$$

$$\mathcal{E}_B(\tilde{C}) \subseteq \mathcal{E}_B \quad (21)$$

$$D(\tilde{A}_{[\mathcal{X}\setminus\mathcal{U}, \mathcal{X}\setminus\mathcal{U}]}, \tilde{C}_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X}\setminus\mathcal{U}]}) \quad (22)$$

spanned by output cactus patch

for all $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| \leq k$

This optimization problem is solved by Algorithm 1, which takes as input the graph \mathcal{G}_P of all possible physical network links. Through six steps, illustrated in Figure 3, it produces as output an optimal dynamics structure $(\tilde{A}^*, \tilde{C}^*)$ for the operation of the network. This also defines the subgraph $\mathcal{G}_P(\tilde{A}^*, \tilde{C}^*)$ of \mathcal{G}_P that is necessary for the operation of the network, assuming backbone forwarding occurs along the least

Algorithm 1

Input : Physical network communication cost graph $\mathcal{G}_P = (\mathcal{X} \cup \mathcal{Q} \cup \mathcal{Z}, \mathcal{E}_S \cup \mathcal{E}_B, w_P)$ with $\mathcal{E}_S = \mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Q}}$ and $\mathcal{E}_B = \mathcal{E}_{\mathcal{Q}\mathcal{Q}} \cup \mathcal{E}_{\mathcal{Q}\mathcal{Z}}$
Required degree of robustness k

Step 1: Compute the least costly path $\text{sp}(q, z; \mathcal{E}_B)$ from each $q \in \mathcal{Q}$ to the central node z over the backbone subnetwork links \mathcal{E}_B . This may be accomplished by computing the shortest path spanning arborescence of $\mathcal{G}_B = (\mathcal{Q} \cup \mathcal{Z}, \mathcal{E}_B, w_P)$ rooted at z [28]. Let the weight of the shortest path from q to z be $w_{\text{sp}}(q, z; \mathcal{E}_B)$.

Step 2: Generate the dynamic system cost graph $\mathcal{G}_D = (\mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}, \mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Y}} \cup \mathcal{E}_{\mathcal{Y}\mathcal{Z}}, w_D)$ from the physical links cost graph \mathcal{G}_P by removing the backbone nodes \mathcal{Q} and adding the output nodes \mathcal{Y} as described in Section III. For each $(x_1, x_2) \in \mathcal{E}_{\mathcal{X}\mathcal{X}}$, $w_D(x_1, x_2) = w_P(x_1, x_2)$. The link $(x, y_{(x,q)}) \in \mathcal{E}_{\mathcal{X}\mathcal{Y}}$ if and only if $(x, q) \in \mathcal{E}_{\mathcal{X}\mathcal{Q}}$, with $w_D(x, y_{(x,q)}) = w_P(x, q)$. For each $y_{(x,q)} \in \mathcal{Y}$ and $\mathcal{Z} = \{z\}$, $(y_{(x,q)}, z) \in \mathcal{E}_{\mathcal{Y}\mathcal{Z}}$ and $w_D(y_{(x,q)}, z) = w_{\text{sp}}(q, z; \mathcal{E}_B)$.

Step 3: Modify \mathcal{G}_D to produce the graph $\mathcal{G}'_D = (\mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}, \mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Y}} \cup \mathcal{E}_{\mathcal{Y}\mathcal{Z}}, w'_D)$. For each $(x_1, x_2) \in \mathcal{E}_{\mathcal{X}\mathcal{X}}$, $w'_D(x_1, x_2) = w_P(x_1, x_2)$. For each $(x, y_{(x,q)}) \in \mathcal{E}_{\mathcal{X}\mathcal{Y}}$ and $\mathcal{Z} = \{z\}$, $w'_D(x, y_{(x,q)}) = w_P(x, q) + w_{\text{sp}}(q, z; \mathcal{E}_B)$ and $w'_D(y_{(x,q)}, z) = 0$. Note that all augmenting links of \mathcal{G}'_D have tail in the sensor nodes \mathcal{X} .

Step 4: Find the minimum z -rooted $(k+1)$ -node-connected Steiner subgraph of \mathcal{G}'_D with terminal nodes \mathcal{X} using a suitable algorithm, such as that provided in [26]. The modification suggested in Remark 3 may also be applied. Refer to this graph as $\mathcal{T}^* = (\mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}, \mathcal{E}_{\mathcal{T}^*}, w'_D)$.

Step 5: Construct a graph $\mathcal{P}^* = (\mathcal{X} \cup \mathcal{Y}, \mathcal{E}_{\mathcal{P}^*}, w'_D)$ by adding a zero-cost self-loop to each sensor node in the graph \mathcal{T}^* and by removing z . Hence, $\mathcal{E}_{\mathcal{P}^*} = (\mathcal{E}_{\mathcal{T}^*} \setminus \mathcal{E}_{\mathcal{Y}\mathcal{Z}}) \cup \{(x, x) | x \in \mathcal{X}\}$.

Step 6: Form $(\tilde{A}^*, \tilde{C}^*)$ such that $\mathcal{P}^* \simeq \mathcal{D}(\tilde{A}^*, \tilde{C}^*)$.
Set $\tilde{A}^*_{ij} = 1$ if and only if $(x_j, x_i) \in \mathcal{E}_{\mathcal{P}^*}$.
Set $\tilde{C}^*_{kj} = 1$ if and only if $(x_j, q_k) \in \mathcal{E}_{\mathcal{P}^*}$.

Output: Optimal dynamics structure $(\tilde{A}^*, \tilde{C}^*)$
Corresponding physical graph $\mathcal{G}_P(\tilde{A}^*, \tilde{C}^*)$

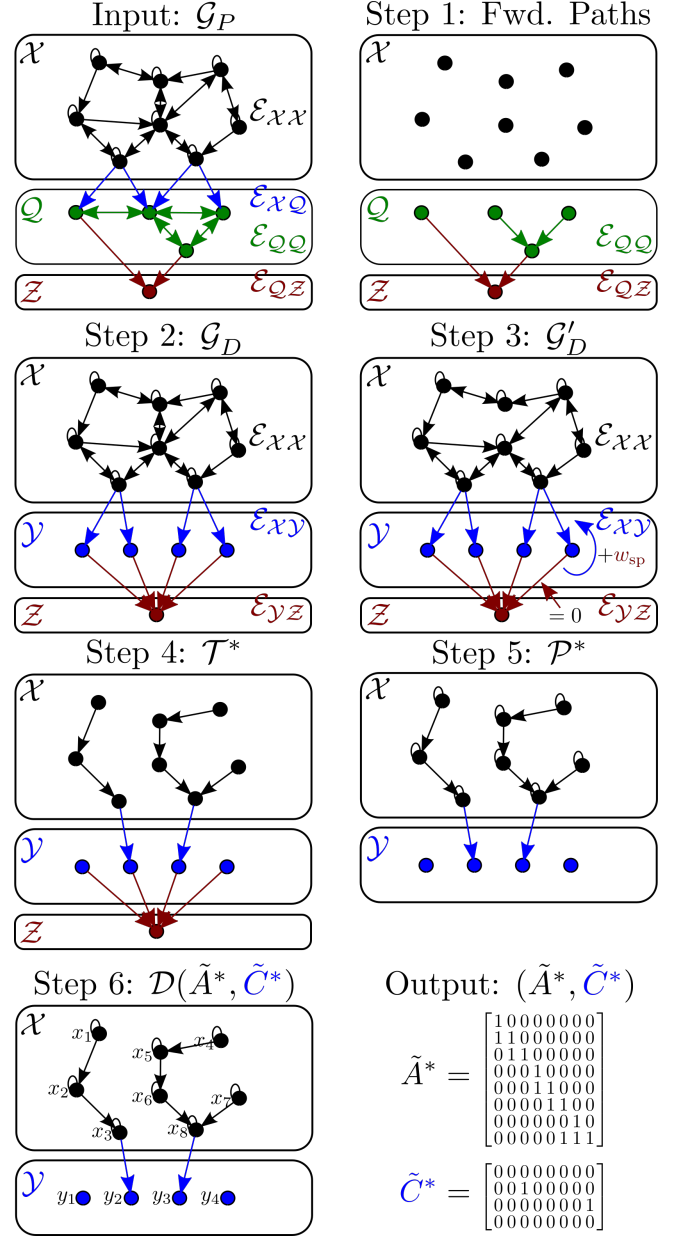


Fig. 3: The illustration demonstrates the steps of Algorithm 1 for an example input graph \mathcal{G}_P and $k = 0$, resulting in a \mathcal{Y} -rooted branching with added loops.

costly path. Practical discussion of how to find an observable instantiation of this structure appears later in this section.

Algorithm 1 begins by first finding the cost of the fusion center accessing a single output in the physical links cost graph \mathcal{G}_P from each backbone node by computing the shortest path spanning arborescence of the backbone subnetwork in Step 1. Subsequently, Step 2 generates the dynamic system computation cost graph \mathcal{G}_D , which it accomplishes by removing the backbone nodes and creating an output node $y_{(x,q)}$ for every $(x, q) \in \mathcal{E}_{\mathcal{X}\mathcal{Q}}$. This output node connects to only one sensor node x with cost $w_D(x, y_{(x,q)}) = w_P(x, q)$ and the fusion center z with cost $w_D(y_{(x,q)}, z) = w_{\text{sp}}(q, z; \mathcal{E}_B)$ obtained

from the shortest path spanning arborescence found in Step 1. In order to frame the optimization problem as an efficiently solvable (by submodular flow) case of the Steiner subgraph problem with \mathcal{X} as the terminal node set, all augmenting links must have tail in the terminal node set. Therefore, Step 3 produces the modified graph \mathcal{G}'_D in which $w'_D(x, y_{(x,q)}) = w_P(x, q) + w_{\text{sp}}(q, z; \mathcal{E}_B)$ and $w'_D(y_{(x,q)}, z) = 0$. Note that any minimal z -rooted Steiner subgraph of \mathcal{G}'_D with terminal nodes \mathcal{X} either include both $(x, y_{(x,q)})$ and $(y_{(x,q)}, z)$ or includes neither $(x, y_{(x,q)})$ nor $(y_{(x,q)}, z)$. Hence, this modification does not affect the total cost of any minimal solution. Step 4 performs the optimization step, which is a minimum z -rooted

$(k+1)$ -node-connected Steiner subgraph computation for \mathcal{G}'_D with terminal nodes \mathcal{X} . Because all augmenting links have tail in \mathcal{X} due to the modification in Step 3, the problem is solvable in polynomial time using submodular flows [26]. In fact, it can be solved through a maximum weighted matroid intersection algorithm [26] by the method in Remark 3. Step 5 constructs a graph that is guaranteed to be spanned by an output cactus patch through the addition of zero-cost self loops to the sensor nodes. Finally, Step 6 interprets the graph from Step 5 as a structural system, allowing an optimal structurally observable network dynamics matrix pair to be output from the algorithm. Theorem 3 and the corresponding proof show that the supplied algorithm solves the minimum cost design problem for any input physical links cost network \mathcal{G}_P and robustness requirement k such that a solution exists.

Theorem 3 (Main Result) *For the physical links cost network \mathcal{G}_P and robustness requirement k the structural dynamics pair $(\tilde{A}^*, \tilde{C}^*)$ output by Algorithm 1 is a solution to the optimization problem in (13)-(17), provided that for each $x \in \mathcal{X}$ there is a set of at least $k+1$ internally node-disjoint directed paths in \mathcal{G}_P each beginning at x and ending in \mathcal{Q} and that there is a directed path from each $q \in \mathcal{Q}$ to z .*

Proof We first demonstrate the feasibility and existence of the solution $(\tilde{A}^*, \tilde{C}^*)$. Subsequently, we show by contradiction that no other feasible solution incurring lesser cost exists. Hence, we conclude the solution found is optimal.

Note that the shortest path spanning arborescence of the backbone subnetwork computed in Step 1 must exist because there is a directed path from each $q \in \mathcal{Q}$ to z . There are $k+1$ internally node-disjoint directed paths in \mathcal{G}_P beginning at x and ending in \mathcal{Q} , each of which contains a distinct link $(x_i, q_i) \in \mathcal{E}_{\mathcal{X}\mathcal{Q}}$. Thus, the graph \mathcal{G}_D constructed in Step 2 has $k+1$ internally node-disjoint directed paths beginning at x and ending in a distinct $y_i = y_{(x_i, q_i)} \in \mathcal{Y}$, which can be constructed from the paths in \mathcal{G}_P by substituting (x_i, y_i) for the final link. Furthermore, because each $y_i \in \mathcal{Y}$ connects directly to z , this implies that $\kappa_{\mathcal{G}_D}(x, z) \geq k+1$ for all $x \in \mathcal{X}$. Since construction of \mathcal{G}'_D in Step 3 does not alter connectivity, it also follows that $\kappa_{\mathcal{G}'_D}(x, z) \geq k+1$ for all $x \in \mathcal{X}$. Hence, the minimum z -rooted $(k+1)$ -node-connected Steiner subgraph of \mathcal{G}'_D with terminal nodes \mathcal{X} exists and \mathcal{T}^* can be found in Step 4. With deletion of any failing node set $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| \leq k$, there is at least one path from each node in $\mathcal{X} \setminus \mathcal{U}$ to z by z -rooted $(k+1)$ -node-connectedness. These paths form a spanning tree \mathcal{T}^*_U for the subgraph of \mathcal{T}^* in which \mathcal{U} has been excluded. Because zero-cost self-loops are always permitted at every $x \in \mathcal{X}$, \mathcal{P}^* may be formed in Step 5, and the subgraph of \mathcal{P}^* that excludes \mathcal{U} is spanned by \mathcal{T}^*_U with z removed. Let \mathcal{P}^*_U be the graph formed from \mathcal{T}^*_U by addition of zero cost self-loops to all $x \in \mathcal{X} \setminus \mathcal{U}$ and with z removed. Note that \mathcal{P}^*_U satisfies the recursive definition of an output cactus patch, with output nodes \mathcal{Y} , and that \mathcal{P}^*_U spans the subgraph of \mathcal{P}^* with \mathcal{U} excluded. By the construction in Step 6, $\mathcal{D}(\tilde{A}^*, \tilde{C}^*)$ is isomorphic to \mathcal{P}^* , and, consequently, $\mathcal{D}(\tilde{A}^*_{[\mathcal{X} \setminus \mathcal{U}, \mathcal{X} \setminus \mathcal{U}]}, \tilde{C}^*_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X} \setminus \mathcal{U}]})$ is spanned by an output cactus patch isomorphic to \mathcal{P}^*_U (by the same isomorphism). Hence, by Theorem 2, $(\tilde{A}^*_{[\mathcal{X} \setminus \mathcal{U}, \mathcal{X} \setminus \mathcal{U}]}, \tilde{C}^*_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X} \setminus \mathcal{U}]})$ is structurally ob-

servable for all $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| < k$. The other constraints, $\mathcal{E}_{\mathcal{X}\mathcal{X}}(\tilde{A}^*) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{X}}$, $\mathcal{E}_{\mathcal{X}\mathcal{Q}}(\tilde{C}^*) \subseteq \mathcal{E}_{\mathcal{X}\mathcal{Q}}$, and $\mathcal{E}_B(\tilde{C}^*) \subseteq \mathcal{E}_B$, are also satisfied through subgraph constructions. Therefore, $(\tilde{A}^*, \tilde{C}^*)$ is a feasible solution.

Assume by way of contradiction that a feasible solution $(\tilde{A}^\dagger, \tilde{C}^\dagger)$ of lesser cost than $(\tilde{A}^*, \tilde{C}^*)$ with respect to the objective function F exists. That is,

$$F(\tilde{A}^\dagger, \tilde{C}^\dagger) < F(\tilde{A}^*, \tilde{C}^*). \quad (23)$$

Because $(\tilde{A}^\dagger_{[\mathcal{X} \setminus \mathcal{U}, \mathcal{X} \setminus \mathcal{U}]}, \tilde{C}^\dagger_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X} \setminus \mathcal{U}]})$ must be structurally observable for all $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| \leq k$, $\mathcal{D}(\tilde{A}^\dagger_{[\mathcal{X} \setminus \mathcal{U}, \mathcal{X} \setminus \mathcal{U}]}, \tilde{C}^\dagger_{[\mathcal{E}_{\mathcal{X}\mathcal{Q}}, \mathcal{X} \setminus \mathcal{U}]})$ must have a minimum cost subgraph $\mathcal{P}^\dagger(\mathcal{X} \cup \mathcal{Y}, \mathcal{E}_{\mathcal{P}^\dagger}, w'_D)$, taking into account sensor to output link costs, which is spanned by an output cactus patch when any such node set \mathcal{U} is removed. Note that

$$W(\mathcal{P}^\dagger) \leq W(\mathcal{D}(\tilde{A}^\dagger, \tilde{C}^\dagger)) = F(\tilde{A}^\dagger, \tilde{C}^\dagger) \quad (24)$$

where $W(\cdot)$ gives the total weight of a graph. Furthermore there are at least $k+1$ internally node-disjoint paths beginning at x and ending in \mathcal{Y} in \mathcal{P}^\dagger for all $x \in \mathcal{X}$ because any graph disconnected from all outputs by k sensor node failures could not have that property. Reversing the process in Step 5, construct $\mathcal{T}^\dagger(\mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}, \mathcal{E}_{\mathcal{T}^\dagger}, w'_D)$ where $\mathcal{E}_{\mathcal{T}^\dagger} = (\mathcal{E}_{\mathcal{P}^\dagger} \setminus \{(x, x) | x \in \mathcal{X}\}) \cup \{(y, z) | (x, y) \in \mathcal{E}_{\mathcal{P}^\dagger} \text{ for some } x \in \mathcal{X}\}$. Because the self-loops have zero cost and output links connecting to z have weight 0 with respect to w'_D , it follows that

$$W(\mathcal{T}^\dagger) = W(\mathcal{P}^\dagger). \quad (25)$$

Noting that $\kappa_{\mathcal{T}^\dagger}(x, z) \geq k+1$ for all $x \in \mathcal{X}$, it is clear that \mathcal{T}^\dagger is a z -rooted $(k+1)$ -node-connected Steiner subgraph of \mathcal{G}'_D with terminal nodes \mathcal{X} . By (23)-(25)

$$W(\mathcal{T}^\dagger) \leq F(\tilde{A}^\dagger, \tilde{C}^\dagger). \quad (26)$$

Similarly,

$$W(\mathcal{T}^*) = W(\mathcal{P}^*), \quad (27)$$

and, this time with equality,

$$W(\mathcal{P}^*) = W(\mathcal{D}(\tilde{A}^*, \tilde{C}^*)) = F(\tilde{A}^*, \tilde{C}^*). \quad (28)$$

Hence, by (27)-(28),

$$W(\mathcal{T}^*) = F(\tilde{A}^*, \tilde{C}^*), \quad (29)$$

so it follows from (23) that

$$W(\mathcal{T}^\dagger) < W(\mathcal{T}^*). \quad (30)$$

This contradicts the fact that \mathcal{T}^* is a minimum z -rooted $(k+1)$ -node-connected Steiner subgraph for \mathcal{G}'_D with terminal nodes \mathcal{X} . Thus, $(\tilde{A}^*, \tilde{C}^*)$ is a minimum cost solution to the optimization problem. ■

Remark 2 (Existence Verification and Maximum Robustness) *The condition that for each $x \in \mathcal{X}$ there is a set of at least $k+1$ internally node-disjoint directed paths in \mathcal{G}_P each beginning at x and ending in \mathcal{Q} may be efficiently verified through $N = |\mathcal{X}|$ maximum $(\{x_s\}, \mathcal{Q})$ -flow computations (with restricted vertex capacity) where the source node is $x_s = x$, non-source nodes in \mathcal{X} have capacity 1, and links in $\mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Q}}$ have capacity 1. This process can also be used to find the largest value of k*

for which the solution exists. Among all the nodes $x \in \mathcal{X}$ find the one with smallest node-capacitated flow to \mathcal{Q} . The largest admissible value of k is 1 less than that flow value.

Before applying Algorithm 1 to obtain an optimal dynamics structure $(\tilde{A}^*, \tilde{C}^*)$ that solves the problem in (13) – (17), it would be prudent to ensure that a solution providing robustness factor k exists for the input physical links cost graph \mathcal{G}_P . In order to accomplish this, it must be verified that for each $x \in \mathcal{X}$ there is a set of at least $k + 1$ internally node-disjoint paths beginning at x and ending in \mathcal{Q} , in which case Theorem 3 guarantees that a solution exists. This condition can be tested by computing the maximum flow from each x_s , treated as the only source, to \mathcal{Q} , treated as a set of sinks, where all links $\mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Q}}$ have capacity 1 and the nodes $x \in \mathcal{X} \setminus \{x_s\}$ also have capacity 1. Through node splitting techniques [26], the node capacities may be converted to link capacities. By connecting all sinks to a common sink by links of infinite capacity, the result can then be formulated as a standard maximum flow problem, which can be solved in polynomial time by a number of algorithms, such as the Ford-Fulkerson algorithm [28]. Because each node can only be used once due to the node flow capacity, the desired condition holds if the maximum $(\{x_s\}, \mathcal{Q})$ -flow is at least $k+1$ for all $x_s \in \mathcal{X}$.

Remark 3 (Implementation and Complexity) *The optimization in Step 4 may be simplified to a minimum z -rooted k -node-connected spanning subgraph computation by inserting an additional k zero-cost parallel (duplicate) links from each $y \in \mathcal{Y}$ to z in \mathcal{G}'_D . This computation can be computed in polynomial time through maximum weighted matroid intersection methods detailed in [26].*

While the optimization in Step 4 could be solved as a Steiner subgraph problem because all augmenting links have tail in the terminal nodes \mathcal{X} , it is simpler to convert the problem to a z -rooted k -node-connected spanning subgraph computation. All nodes in \mathcal{X} have $k + 1$ internally node-disjoint paths to \mathcal{Q} assuming a solution exists. Thus, each $x \in \mathcal{X}$ has $k + 1$ internally node-disjoint paths to \mathcal{Y} , each ending at a different node. Nodes in \mathcal{Y} have only one direct connection to z . Therefore, addition of k zero-cost parallel links from each $y \in \mathcal{Y}$ to z makes \mathcal{G}'_D z -rooted k -node-connected as a whole. By computing the z -rooted k -node-connected spanning subgraph and discarding redundant zero-cost edges from \mathcal{Y} to \mathcal{Z} , the solution is obtained.

Note that Step 4 is, by far, the most costly step and subsumes the complexity of the other steps. The z -rooted k -node-connected spanning subgraph can be computed in polynomial time as a maximum weighted common independent set of two matroids as described in [26]. Matroids provide rules that define which subsets of a set are independent. These rules must satisfy several axioms that will not be detailed in this paper. For two matroids over edge set E , the basic weighted matroid intersection algorithm involves $O(|S|^2|E|)$ calls to the two matroid independence oracles, functions that determine membership of a set in the matroid [28], where $|S|$ is the maximum size of a common independent set. In this case $|S|$, the maximum size of a common independent set,

is $|S| = (k + 1)(|\mathcal{X}| + |\mathcal{Y}|)$. Links from Y to z are always included in the solution, so they can be trivially included in the final set, making $|S|$ is $O(k|\mathcal{X}|)$ for the purpose of counting calls to the matroid oracles. The set is $E = \mathcal{E}_{\mathcal{X}\mathcal{X}} \cup \mathcal{E}_{\mathcal{X}\mathcal{Y}} \cup \mathcal{E}_{\mathcal{Y}\mathcal{Z}}$, but links from Y to z are always included in the final solution, so $|E|$ is $O(|\mathcal{E}_{\mathcal{X}\mathcal{X}}| + |\mathcal{E}_{\mathcal{X}\mathcal{Y}}|)$ for purposes of counting. For this problem, one matroid oracle has relatively low complexity, while the other has complexity $O(|V||S|^2)$ [26] where $|V|$ is $O(|\mathcal{X}| + |\mathcal{Y}|)$. Thus the total complexity of the algorithm is $O(|V||E||S|^4)$ or $O(k^4|\mathcal{X}|^4(|\mathcal{X}| + |\mathcal{Y}|)(|\mathcal{E}_{\mathcal{X}\mathcal{X}}| + |\mathcal{E}_{\mathcal{X}\mathcal{Y}}|))$. It is claimed in [26] that the matroid oracle can be computed in $O(|V|^3)$, which would reduce this to $O(|V|^3|E||S|^2)$ or $O(k^2|\mathcal{X}|^2(|\mathcal{X}| + |\mathcal{Y}|)^3(|\mathcal{E}_{\mathcal{X}\mathcal{X}}| + |\mathcal{E}_{\mathcal{X}\mathcal{Y}}|))$. Improved matroid intersection algorithms also exist.

Remark 4 (Special Case) *For the case in which $k = 0$, the optimization in Step 4 reduces to a minimum z -rooted spanning arborescence of \mathcal{G}'_D with terminal nodes \mathcal{X} .*

The special case $k = 0$ of the minimum cost network design problem, which requires no robustness to node failures, was originally proposed and solved in [20] when symmetric cost structure is considered.

Under these conditions, the Steiner subgraph optimization step reduces to a minimum z -rooted Steiner arborescence problem on \mathcal{G}'_D with terminal nodes \mathcal{X} . Because each $y \in \mathcal{Y}$ may be automatically connected directly to z by the non-augmenting link (y, z) , this can be further reduced to a minimum spanning tree arborescence, leading to Remark 4. Thus, this case may be computed with greater simplicity using Edmond's algorithm [20].

Remark 5 (Observable Instantiations) *An observable instantiation (A, C) of an optimal structurally observable solution $(\tilde{A}^*, \tilde{C}^*)$ output by Algorithm 1 over a finite field \mathbb{F}_{p^n} can be found randomly with high probability for fields of large order p^n . As p^n grows without bound, the probability of an observable random instantiation over \mathbb{F}_{p^n} approaches 1 [17].*

Although the output of the algorithm gives an optimal structure $(\tilde{A}^*, \tilde{C}^*)$ of the system dynamics that satisfies the constraints, an observable instantiation (A, C) of the dynamics must be obtained for any implementation. If the field in which the system operates is $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$, such instantiations of the structure are guaranteed to exist and the set of unobservable realizations over those fields has zero measure. Consequently, a random instantiation of the structure from a suitable distribution would be almost surely observable with probability one.

While interesting, this approach suffers from the possibility of producing systems with poorly conditioned observability matrices that are full rank yet present numerical problems. Additionally, physical devices cannot truly operate with general real or complex numbers. These problems may both be avoided through the use of finite fields $\mathbb{F} = \mathbb{F}_{p^n}$ where p is prime and n is a positive integer. It has been shown in [17] that for an output tree with self-loops attached to the N state nodes, an observable instantiation is guaranteed to exist over \mathbb{F}_{p^n} if $p^n \geq N$. In fact, Corollary 1 of [17] shows that this instantiation can be produced by assigning a distinct field

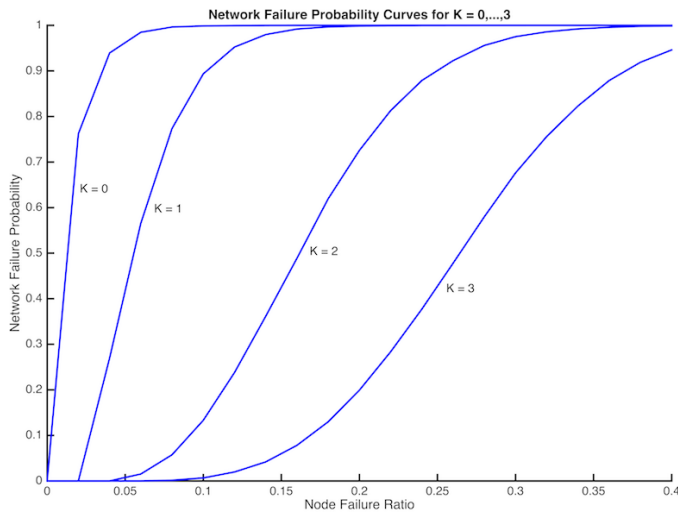


Fig. 4: The plot shows simulated network failure probability plotted against the fraction of sensor node failures for designed robustness levels $k = 0, \dots, 3$ in networks with $|\mathcal{X}| = 50$. Results show a gentler slope for higher values of k , even beyond the robustness guarantee.




Fig. 5: This illustrative example shows a small network on the unit square designed with $|\mathcal{X}| = 30$ randomly placed sensor nodes (black), and $|\mathcal{Q}| = 4$ backbone nodes (green), central node Z (red), and robustness level $k = 2$, and distance squared costs with maximum radius $R_{lim} = .4$. The colors allow comparison to Figure 1.

element to each self-loop and the multiplicative identity to each link between distinct nodes. Instead, if elements are chosen randomly at each sensor, repetition can be avoided and observability achieved with high probability if the field is sufficiently large, as shown in Theorem 5 of [17], approaching probability 1 as p^n increases without bound. This provides a more advantageous approach to practical implementation and is summarized by Remark 5.

Remark 6 (Empirical Robustness) *While Algorithm 1 guarantees robustness to failure of any node subset $\mathcal{U} \subset \mathcal{X}$ with $|\mathcal{U}| \leq k$, the solution may be robust in practice to a larger number of randomly chosen node failures. Here, as*

in the rest of the paper, robustness refers to the fact that at least one spanning output cactus remains intact after k node failures. That is, no surviving nodes are disconnected and, thus, unobservable. For more than k node failures, the idea of robustness can be relaxed from this guarantee to examine the failure probability.

In practice, a network designed by Algorithm 1 may be robust to more sensor node failures than the guaranteed degree of robustness k . Although the algorithm guarantees that removal of any subset of failing sensors $\mathcal{U} \subseteq \mathcal{X}$ with $|\mathcal{U}| \leq k$ does not break the structural observability of the remaining network, it does not necessarily follow that the network fails if $|\mathcal{U}| = \ell > k$. For instance, the structure of the sensor subnetwork for $k = 0$ is a Q rooted branching with loops, and failure of any leaf node does not affect the remainder of the network. This effect becomes more pronounced for higher values of k . With ℓ node failures selected uniformly at random, the probability of network failure cannot be easily computed. However, it is possible to empirically simulate it for networks constructed from pseudo-randomly placed nodes with pseudo-random node failures.

Figure 4 plots the empirical probability of network failure for pseudo-randomly generated node locations against the node failure ratio $\ell/|\mathcal{X}|$ for designed degree of robustness $k = 0, \dots, 3$. This simulation used $|\mathcal{X}| = 50$ uniformly distributed sensor nodes, $|\mathcal{Y}| = 3$ backbone nodes, and distance squared link cost. Figure 5 represents the graph of a smaller network designed in such a way, with the additional constraint on the communication radius. For each randomly generated network and failure ratio value, ℓ sensor nodes were randomly selected to fail. The network fails if any surviving node no longer has a directed path to a backbone node. The expected network failure probability was computed for each node failure ratio over 100 random graphs each with 1000 random sets of failing nodes. As seen in Figure 4, the resulting networks can have low failure probability even beyond the robustness guarantee, where failure occurs with probability zero. The plot also demonstrates that when designing for higher guaranteed robustness level, the failure probability grows much more slowly. For instance, when designing for robustness level $k = 3$, the failure probability is only approximately 0.2 when 10 nodes out of 50 sensor nodes fail.

V. CONCLUSIONS

The distributed networks of agents or sensors examined in this paper model situations in which nodes update their state variables using values obtained by sensing state information from nodes within their local neighborhoods. The state of some nodes is directly sensed by nodes in a network backbone, which are directly accessed by a central fusion center. The fusion center computes the (global) network state from the data it obtains. This type of scenario may arise in robotics formation control applications or in field inversion problems, for instance. By duality of observability and controllability, the same technique may be used to drive the sensor states to some desired target.

This paper posed the problem of finding minimum cost network update dynamics with respect to a link cost objective function under an observability constraint. The polynomial time algorithm proposed finds the optimal network that is structurally observable when any k sensor nodes are removed by using combinatorial optimization algorithms, guaranteeing the existence of observable dynamics that can be found separately in a straightforward way by random instantiation over sufficiently large finite fields. In practice, networks may survive more node failures than the guaranteed amount, especially as k increases. This paper improves the results presented in [20], extending consideration to robustness requirements and designable network backbone topologies. Future efforts could address a number of interesting network design problems, such as finding optimal networks that are simultaneously both structurally observable and structurally controllable. Alternate problem formulations could decrease network iterations by attempting to reduce the observability index. Additional work could also focus on implementing the network design process in a fully distributed manner within the provided framework, using known distributed algorithms for computing minimum spanning arborescences in the simple $k = 0$ case.

REFERENCES

- [1] P. Braca, S. Marano, and V. Matta, "Enforcing consensus while monitoring the environment in wireless sensor networks," *IEEE Transactions on Signal Processing*, vol. 56, no. 7, pp. 3375–3380, 2008.
- [2] S. Kar, J. M. F. Moura, and K. Ramanan, "Distributed parameter estimation in sensor networks: nonlinear observation models and imperfect communication," *IEEE Transactions on Information Theory*, vol. 58, no. 6, pp. 3575–3605, June 2012.
- [3] C. G. Lopes and A. H. Sayed, "Diffusion least-mean squares over adaptive networks: Formulation and performance analysis," *IEEE Transactions on Signal Processing*, vol. 56, no. 7, pp. 3122–3136, July 2008.
- [4] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, Jan. 2009.
- [5] D. Jakovetic, J. Xavier, and J. M. F. Moura, "Fast distributed gradient methods," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1131–1146, May 2014.
- [6] A. Schmidt and J. M. F. Moura, "Distributed field reconstruction with model-robust basis pursuit," *37th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2012)*, pp. 2673–2676, Mar. 2012.
- [7] A. Schmidt, J. Harley, and J. M. F. Moura, "Compressed sensing radar surveillance networks," *7th IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM)*, pp. 209–212, June 2012.
- [8] C. Reynolds, "Flocks, herds, and schools: A distributed behavioral model," *ACM SIGGRAPH'87 Conference Proceedings*, vol. 21, no. 4, pp. 25–34, July 1987.
- [9] T. Vicsek, A. Czirók, E. Ben-Jacob, and O. Shochet, "Novel type of phase transition in a system of self-driven particles," *Physical Review Letters*, vol. 75, pp. 1226–1229, 1995.
- [10] F. Cucker and S. Smale, "Emergent behavior in flocks," *IEEE Transactions on Automatic Control*, vol. 52, no. 5, pp. 852–862, May 2007.
- [11] S. Tu and A. Sayed, "Mobile adaptive networks," *IEEE Journal on Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 649–664, Aug. 2011.
- [12] N. Moshtagh, N. Michael, A. Jadbabaie, and K. Daniilidis, "Vision-based, distributed control laws for motion coordination of non-holonomic robots," *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 851–860, 2009.
- [13] M. Cao, C. Yu, and B. Anderson, "Formation control using range-only measurements," *Automatica*, vol. 47, no. 4, pp. 776–781, 2011.
- [14] Z. Sun and B. Anderson, "Rigid formation control systems modeled by double integrators: System dynamics and convergence analysis," *2015 5th Australian Control Conference (AUCC)*, pp. 241–246, 2015.
- [15] M. Ji and M. Egerstedt, "Observability and estimation in distributed sensor networks," *Proceedings of the 46th IEEE Conference on Decision and Control*, pp. 4221–4226, Dec. 2007.
- [16] S. Sundaram and C. N. Hadjicostis, "Distributed calculation of linear functions in noisy networks via linear iterations," *Proc. of the 47th IEEE Conference on Decision and Control*, pp. 5462–5467, Dec. 2008.
- [17] S. Sundaram and C. N. Hadjicostis, "Structural controllability and observability of linear systems over finite fields with applications to multi-agent systems," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 60–73, Jan. 2013.
- [18] G. Ellis, *Observers in Control Systems: A Practical Guide*. Academic Press, 2000.
- [19] J. M. Dion, C. Commault, and J. van der Woude, "Generic properties and control of linear structured systems: a survey," *Automatica*, vol. 39, no. 7, pp. 1125–1144, Mar. 2003.
- [20] S. Pequito, S. Kruzick, S. Kar, J. M. F. Moura, and A. P. Aguiar, "Optimal design of distributed sensor networks for field reconstruction," *21st European Signal Processing Conference*, pp. 1–5, Sept. 2013.
- [21] S. Pequito, S. Kar, and A. P. Aguiar, "Minimum cost input/output selection for large scale linear time-invariant systems: a structured systems approach," *Automatica*, vol. 68, pp. 384–391, 2016.
- [22] S. Pequito, S. Kar, and A. P. Aguiar, "On the complexity of the constrained input selection problem for structural linear systems," *Automatica*, vol. 62, pp. 193–199, 2015.
- [23] S. Pequito, S. Kar, and A. P. Aguiar, "A framework for structural input/output and control configuration selection in large-scale systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 303–318, Feb. 2016.
- [24] H. Gabow, Z. Galil, T. Spencer, and R. Tarjan, "Efficient algorithms for finding minimum spanning trees in undirected and directed graphs," *Combinatorica*, vol. 6, no. 2, pp. 109–122, 1986.
- [25] J. Edmonds, "Edge-disjoint branchings," *Combinatorial Algorithms*, pp. 91–96, Academic Press, 1973.
- [26] A. Frank, "Rooted k -connections in digraphs," *Discrete Applied Mathematics*, vol. 157, no. 6, pp. 1242–1254, 2009.
- [27] L. Zosin and S. Khuller, "On directed Steiner trees," *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 59–63, June 2002.
- [28] W. Cook, W. Cunningham, W. Pulleyblank, and A. Schrijver, *Combinatorial Optimization*. Wiley-Interscience Series in Discrete Mathematics and Optimization, 1998, pp. 9–42.
- [29] C. T. Chen, *Linear System Theory and Design*, 3rd ed. Oxford University Press, 1999, pp.143–158.
- [30] C. T. Lin, "Structural controllability," *IEEE Transactions on Automatic Control*, vol. 19, no. 3, pp. 201–208, June 1974.
- [31] R. Shields and J. Pearson, "Structural controllability of multi-input linear systems," *IEEE Transactions on Automatic Control*, vol. 21, no. 2, April 1976.



Stephen Kruzick (S'17) currently pursues graduate studies as a PhD student in the Electrical and Computer Engineering department at Carnegie Mellon University (CMU) in Pittsburgh, Pennsylvania, USA. He received a BS degree in Electrical and Computer Engineering and a BA degree in Mathematics from Rice University in Houston, Texas, USA, in 2011. His research interests include graph signal processing, network stochastic processes, spectral statistics of random matrices, distributed algorithm design, and multi-agent systems. Stephen Kruzick has been

awarded the National Defense Science and Engineering Graduate (NDSEG) Fellowship and the Benjamin Garver Lamme/Westinghouse Graduate Fellowship as well as a National Science Foundation Graduate Research Fellowship Program (NSF-GRFP) honorable mention.



Sérgio Pequito (S'09,M'14) is a postdoctoral researcher in general robotics, automation, sensing & perception laboratory (GRASP lab) at University of Pennsylvania. He obtained his Ph.D. in Electrical and Computer Engineering from Carnegie Mellon University and Instituto Superior Tecnico, through the CMU-Portugal program, in 2014. Previously, he received his B.Sc. and M.Sc. in Applied Mathematics from the Instituto Superior Tecnico in 2007 and 2009, respectively. Pequito's research consists in understanding the global qualitative behavior of

large scale systems from their structural or parametric descriptions and provide a rigorous framework for the design, analysis, optimization and control of large scale (real-world) systems. Currently, his interests span to neuroscience, where control theoretic tools can be leveraged to develop new analysis tools for brain dynamics that, ultimately, will lead to new diagnostics and treatments of neural disorders. Pequito was awarded with the best student paper finalist in the 48th IEEE Conference on Decision and Control (2009). Also, Pequito received the ECE Outstanding Teaching Assistant Award from the Electrical and Computer Engineering Department at Carnegie Mellon University, and the Carnegie Mellon Graduate Teaching Award (university-wide) honorable mention, both in 2012. In addition, Pequito's was the 2016 EECI European PhD Award on Control for Complex and Heterogeneous Systems (finalist), and received the 2016 O. Hugo Schuck Award in the Theory Category.



Soumya Kar (S05, M10) received a B.Tech. in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, India, in May 2005 and a Ph.D. in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 2010. From June 2010 to May 2011, he was with the Electrical Engineering Department, Princeton University, Princeton, NJ, USA, as a Postdoctoral Research Associate. He is currently an Associate Professor of Electrical and Computer Engineering at Carnegie Mellon University,

Pittsburgh, PA, USA. His research interests include decision-making in large-scale networked systems, stochastic systems, multi-agent systems and data science, with applications to cyber-physical systems and smart energy systems. He has published extensively in these topics with more than 140 articles in journals and conference proceedings and holds multiple patents. Recent recognition of his work includes the 2016 O. Hugo Schuck Best Paper Award from the American Automatic Control Council, the 2016 Deans Early Career Fellowship from CIT, Carnegie Mellon, and the 2011 A.G. Milnes Award for best PhD thesis in Electrical and Computer Engineering, Carnegie Mellon University.



José M. F. Moura (S'71-M'75-SM'90-F'94) received the engenheiro electrotécnico degree from Instituto Superior Técnico (IST), Lisbon, Portugal, and the M.Sc., E.E., and D.Sc. degrees in EECS from the Massachusetts Institute of Technology (MIT), Cambridge, MA.

He is the Philip L. and Marsha Dowd University Professor at Carnegie Mellon University (CMU). He was on the faculty at IST and has held visiting faculty appointments at MIT and New York University (NYU). He founded and directs a large education

and research program between CMU and Portugal, www.icti.cmu.edu.

His research interests are on data science, graph signal processing, and statistical and algebraic signal and image processing. He has published over 550 papers and holds thirteen patents issued by the US Patent Office. The technology of two of his patents (co-inventor A. Kavčić) are in about three billion disk drive read channel chips of 60 % of all computers sold in the last 13 years worldwide and was, in 2016, the subject of the largest university verdict/settlement in the information technologies area.

Dr. Moura was the IEEE Technical Activities Vice-President (2016) and member of the IEEE Board of Directors. He served in several other capacities including IEEE Division IX Director, member of several IEEE Boards, President of the IEEE Signal Processing Society (SPS), Editor in Chief for the IEEE Transactions in Signal Processing, interim Editor in Chief for the IEEE Signal Processing Letters.

Dr. Moura has received several awards, including the Technical Achievement Award and the Society Award from the IEEE Signal Processing. In 2016, he received the CMU College of Engineering Distinguished Professor of Engineering Award. He is a Fellow of the IEEE, a Fellow of the American Association for the Advancement of Science (AAAS), a corresponding member of the Academy of Sciences of Portugal, Fellow of the US National Academy of Inventors, and a member of the US National Academy of Engineering.



A. Pedro Aguiar received the Licenciatura, M.S. and Ph.D. in electrical and computer engineering from the Instituto Superior Tecnico (IST), Technical University of Lisbon, Portugal in 1994, 1998 and 2002, respectively. Currently, Dr. Aguiar holds an Associate Professor position with the Department of Electrical and Computer Engineering (DEEC), Faculty of Engineering, University of Porto (FEUP). From 2002 to 2005, he was a post-doctoral researcher at the Center for Control, Dynamical-Systems, and Computation at the University of

California, Santa Barbara (UCSB). From 2005 to 2012, he was a senior researcher with the Institute for Systems and Robotics at IST (ISR/IST), and an invited assistant professor with the Department of Electrical and Computer Engineering, IST.

His research interests include modeling, control, navigation, and guidance of autonomous robotic vehicles, nonlinear control, switched and hybrid systems, tracking, path-following, performance limitations, nonlinear observers, the integration of machine vision with feedback control, networked control, and coordinated/cooperative control of multiple autonomous robotic vehicles.