

Topology Design for Optimal Network Coherence

Tyler Summers, Iman Shames, John Lygeros, Florian Dörfler

Abstract—We consider a network topology design problem in which an initial undirected graph underlying the network is given and the objective is to select a set of edges to add to the graph to optimize the coherence of the resulting network. We show that network coherence is a submodular function of the network topology. As a consequence, a simple greedy algorithm is guaranteed to produce near optimal edge set selections. We also show that fast rank one updates of the Laplacian pseudoinverse using generalizations of the Sherman-Morrison formula and an accelerated variant of the greedy algorithm can speed up the algorithm by several orders of magnitude in practice. These allow our algorithms to scale to network sizes far beyond those that can be handled by convex relaxation heuristics.

I. INTRODUCTION

Among the most challenging and important problems in control and optimization of dynamical networks is the design of topologies for sensing, control, and communication. One important dynamical process in a variety of networks is synchronization, and it is widely recognized that network topology properties play a key role. There is now a growing literature on synchronization and consensus and many associated applications, including power networks and robotic vehicle networks.

Recently, there has been work on quantifying the robustness of consensus dynamics to stochastic disturbances [2], [1], [25]. The concept of network coherence has been proposed to quantify variance of states around the consensus subspace [2] and is also closely related to the effective resistance of graphs [13], [10]. Several recent papers have focused on leader selection problems to optimize coherence [22], [23], [16], [7], [17], [9], [4]. However, designing network topologies to optimize coherence by choosing sets of edges has received less attention. In [29], [10] the problems of choosing edge weights for a given network topology to optimize network coherence and effective resistance are considered and shown to be convex optimization problems. It is possible to modify the algorithm in [29], [10] to obtain an convex relaxation heuristic for edge selection. However, the resulting problems can still be difficult to solve for large networks, and purely combinatorial versions of these problems appear not to have been considered.

We consider a network topology design problem in which an initial undirected graph underlying the network is given and the objective is to select a set of edges to add to the graph to optimize the coherence of the resulting network. Our main result is to show that network coherence is a submodular set function of the network topology. As a consequence, a simple greedy algorithm is guaranteed to produce near optimal edge set selections. The problem has a similar mathe-

matical structure to recently studied set function optimization problems linking submodularity to controllability [26], [27], [5] and rigidity [24]. We also show that fast rank one updates of the Laplacian pseudoinverse using generalizations of the Sherman-Morrison formula [19] and an accelerated variant of the greedy algorithm [20], [14] can speed up the algorithm by several orders of magnitude in practice. These techniques allow our algorithms to scale to network sizes far beyond those that can be handled by convex relaxation heuristics based on [29], [10]. The results are illustrated with numerical examples.

The rest of the paper unfurls as follows. Section II provides background on network coherence and submodular set functions. Section III presents our main results on topology design and algorithmic speed ups. Section IV exhibits the performance of the accelerated algorithm and illustrates the results on various types of fixed and random initial networks. Finally, Section V concludes.

II. NETWORK COHERENCE AND SUBMODULAR SET FUNCTIONS

This section reviews notions of network coherence and submodular set functions. Topology design problems for optimizing network coherence can be formulated as set function optimization problems.

A. Network coherence

Consider a network with underlying weighted undirected graph $G = (V, E, w)$ where $V = \{1, \dots, n\}$ is a set of nodes, $E \subseteq V \times V$ is a set of edges, and $w \in \mathbf{R}^{|E|}$ is a set of nonnegative weights associated with each edge. Suppose a scalar state variable is associated with each node and the network has consensus dynamics modeled by the stochastic differential equation

$$dx(t) = -Lx(t)dt + dW \quad (1)$$

where L is the weighted Laplacian matrix and dW is a vector of independent Gaussian white noise stochastic processes.

Without noise, it is well known that when the graph is connected, the states converge exponentially to a point on the consensus subspace corresponding to the average value of the initial conditions. With additive noise, the state evolution becomes a stochastic process. The expected values of the states evolve according to deterministic consensus dynamics, but the actual average value undergoes Brownian motion and the states stochastically fluctuate around this value. When the graph is connected, the state variance relative to the average converges to a steady state value. Network coherence quantifies the steady-state variance of these fluctuations

and can be considered as a measure of robustness of the consensus process to the additive noise; networks with small steady-state variance have high network coherence and can be considered to be more robust to noise than networks with low coherence.

Formally, network coherence is defined for connected graphs as the average steady-state deviation from the average value

$$\begin{aligned} \mathcal{C} &= \lim_{t \rightarrow \infty} \sum_{i=1}^n \mathbf{E} \left[\left(x_i(t) - \frac{1}{n} \sum_{j=1}^n x_j(t) \right)^2 \right] \\ &= \lim_{t \rightarrow \infty} \mathbf{E}[x(t)^T P x(t)], \end{aligned} \quad (2)$$

where $P = I - \frac{1}{n} \mathbf{1}\mathbf{1}^T$ is the projection operator onto the disagreement subspace. The coherence then relates to the system H_2 norm as

$$\mathcal{C} = \mathbf{trace} \int_0^\infty e^{-L^T t} P e^{-L t} dt. \quad (3)$$

This in turn can be shown to be related to the spectrum of the Laplacian matrix [22]

$$\mathcal{C} = \frac{1}{2} \mathbf{trace}(L^\dagger) = \frac{1}{2} \sum_{i=2}^n \frac{1}{\lambda_i(L)}, \quad (4)$$

where $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n$, i.e., the pseudoinverse trace of the Laplacian matrix is proportional to network coherence. This quantity is also proportional to the total effective resistance of a graph [13], which is known to be a monotone convex function of the edge weights for a given topology [10].

B. Submodularity

Many combinatorial problems can be formulated as *set function* optimization problems. For a given finite set $V = \{1, \dots, M\}$, a *set function* $f : 2^V \rightarrow \mathbf{R}$ assigns a real number to each subset of V . For the set function optimization problem

$$\underset{S \subseteq V, |S|=k}{\text{maximize}} \quad f(S), \quad (5)$$

the objective is to select a k -element subset of V that maximizes f . This can be solved by brute force by enumerating all possible subsets of size k , evaluating f for all of these subsets, and picking the best subset. However, the number of possible subsets grows factorially as $|V|$ increases, so the brute force approach quickly becomes infeasible even for moderate $|V|$.

Instead, there are structural properties of the set function f that facilitate optimization. In particular, *submodularity* plays similar roles in combinatorial optimization as convexity and concavity play in continuous optimization [18], [14]. It occurs often in applications [3], [12], [15]; is supported by an elegant and practically useful mathematical theory; and there are efficient methods for minimizing and approximation guarantees for maximizing submodular functions.

Definition 1 (Submodularity): A set function $f : 2^V \rightarrow \mathbf{R}$ is called submodular if for all subsets $A \subseteq B \subseteq V$ and all elements $s \notin B$, it holds that

$$f(A \cup \{s\}) - f(A) \geq f(B \cup \{s\}) - f(B), \quad (6)$$

or equivalently, if for all subsets $A, B \subseteq V$, it holds that

$$f(A) + f(B) \geq f(A \cup B) + f(A \cap B). \quad (7)$$

Intuitively, submodularity is a diminishing returns property where adding an element to a smaller set gives a larger gain than adding one to a larger set. The following definition and result from [18] makes this precise and will be used to prove submodularity of a set function associated with network coherence.

Definition 2: A set function $f : 2^V \rightarrow \mathbf{R}$ is called monotone increasing if for all subsets $A, B \subseteq V$ it holds that

$$A \subseteq B \Rightarrow f(A) \leq f(B) \quad (8)$$

and is called monotone decreasing if for all subsets $A, B \subseteq V$ it holds that

$$A \subseteq B \Rightarrow f(A) \geq f(B). \quad (9)$$

Theorem 1 ([18]): A set function $f : 2^V \rightarrow \mathbf{R}$ is submodular if and only if the derived set functions $f_a : 2^{V-\{a\}} \rightarrow \mathbf{R}$

$$f_a(X) = f(X \cup \{a\}) - f(X)$$

are monotone decreasing for all $a \in V$.

Maximization of monotone increasing submodular functions is NP-hard, but a greedy heuristic can be used to obtain a solution that is provably close to the optimal solution [21]. The greedy algorithm for (5) starts with an empty set, $S_0 \leftarrow \emptyset$, computes the gain $\Delta(a | S_i) = f(S_i \cup \{a\}) - f(S_i)$ for all elements $a \in V \setminus S_i$ and adds any element with the highest gain:

$$S_{i+1} \leftarrow S_i \cup \{\arg \max_a \Delta(a | S_i) \mid a \in V \setminus S_i\}.$$

The algorithm terminates after k iterations.

Performance of the greedy algorithm is guaranteed by a well known bound [21]:

Theorem 2 ([21]): Let f^* be the optimal value of the set function optimization problem (5), and let $f(S_{\text{greedy}})$ be the value associated with the subset S_{greedy} obtained from applying the greedy algorithm on (5). If f is submodular and monotone increasing, then

$$\frac{f^* - f(S_{\text{greedy}})}{f^* - f(\emptyset)} \leq \left(\frac{k-1}{k} \right)^k \leq \frac{1}{e} \approx 0.37. \quad (10)$$

This means that the greedy algorithm is guaranteed to produce a subset whose function value is within a constant factor of the value of the optimal subset. This is the best any polynomial time algorithm can achieve [8], assuming $P \neq NP$. Note that this is a worst-case bound; the greedy algorithm often performs much better than the bound in practice, which we will verify for the considered set of problems.

III. OPTIMAL TOPOLOGY DESIGN FOR NETWORK COHERENCE

Consider the problem of choosing a subset \mathcal{E} of k edges, each with a given weight, to add to a given weighted undirected graph $G = (V, E)$ to maximize the network coherence of the resulting graph $G_{\mathcal{E}}$, which can be formulated as a set function optimization problem:

$$\underset{\mathcal{E} \subseteq V \times V \setminus E}{\text{minimize}} \quad \text{trace}(L_{\mathcal{E}}^{\dagger}), \quad (11)$$

where $L_{\mathcal{E}}$ is the resulting Laplacian. In the first subsection, we will assume that the given graph is connected so that the rank of the Laplacian remains constant as edges are added. In the third subsection, we will relax this assumption and present a modified algorithm for constructing trees with good network coherence.

A. Network coherence is a submodular function of network topology

To prove that network coherence is a submodular function of the network topology, the key structure is additivity of the Laplacian matrix in the edges. Specifically, let $M \in \mathbf{R}^{|E| \times |V|}$ denote the weighted incidence matrix of a graph $G = (V, E, w)$ which has a row m_e^T for each edge $e = (i, j) \in E$ with elements (for $i > j$) $M_{ev} = w_e$ if $v = i$, $M_{ev} = -w_e$ if $v = j$, and $M_{ev} = 0$ otherwise. The Laplacian associated with any edge set E can be written

$$L_E = M^T M = \sum_{e=1}^{|E|} m_e m_e^T = \sum_{e=1}^{|E|} L_e, \quad (12)$$

which implies $L_{E_1 \cup E_2} = L_{E_1} + L_{E_2}$ for any disjoint pair of edge sets E_1 and E_2 . We have the following result; the proof has almost identical structure to the proof for the inverse of the controllability Gramian in [26] and the proof for the pseudoinverse of the rigidity Gramian in [24].

Theorem 3: Let $G = (V, E, w_E)$ be a given connected weighted graph, let $\mathcal{E} \subseteq V \times V \setminus E$ with weights $w_{\mathcal{E}}$, and let $L_{\mathcal{E}}$ be the weighted graph Laplacian matrix associated with the edge set $E \cup \mathcal{E}$. Then the set function $f : V \times V \setminus E \rightarrow \mathbf{R}$ defined by $f(\mathcal{E}) = -\text{trace}(L_{\mathcal{E}}^{\dagger})$ is submodular.

Proof: Denote the set of potential edge choices by $\mathcal{E}_c = V \times V \setminus E$. Take any $e \in \mathcal{E}_c$ and consider the derived set functions $f_e : 2^{\mathcal{E}_c \setminus \{e\}} \mapsto \mathbf{R}$ given by

$$\begin{aligned} f_e(\mathcal{E}) &= -\text{trace}(L_{\mathcal{E} \cup \{e\}}^{\dagger}) + \text{trace}(L_{\mathcal{E}}^{\dagger}) \\ &= -\text{trace}((L_{\mathcal{E}} + L_e)^{\dagger}) + \text{trace}(L_{\mathcal{E}}^{\dagger}). \end{aligned}$$

Take any $\mathcal{E}_1 \subseteq \mathcal{E}_2 \subseteq \mathcal{E}_c \setminus \{e\}$. By the additivity property of the Laplacian, it is clear that $\mathcal{E}_1 \subseteq \mathcal{E}_2 \Rightarrow L_{\mathcal{E}_1} \preceq L_{\mathcal{E}_2}$. Now define $L(t) = L_{\mathcal{E}_1} + t(L_{\mathcal{E}_2} - L_{\mathcal{E}_1})$ for $t \in [0, 1]$. Obviously, $L(0) = L_{\mathcal{E}_1}$ and $L(1) = L_{\mathcal{E}_2}$. Now define

$$\hat{f}_e(L(t)) = -\text{trace}((L(t) + L_e)^{\dagger}) + \text{trace}(L(t)^{\dagger}).$$

Note that $\hat{f}_e(L(0)) = f_e(\mathcal{E}_1)$ and $\hat{f}_e(L(1)) = f_e(\mathcal{E}_2)$. We have

$$\begin{aligned} \frac{d}{dt} \hat{f}_e(L(t)) &= \frac{d}{dt} [-\text{trace}((L(t) + L_e)^{\dagger}) + \text{trace}(L(t)^{\dagger})] \\ &= \text{trace}[(L(t) + L_e)^{\dagger}(L_{\mathcal{E}_2} - L_{\mathcal{E}_1})(L(t) + L_e)^{\dagger}] \\ &\quad - \text{trace}[L(t)^{\dagger}(L_{\mathcal{E}_2} - L_{\mathcal{E}_1})L(t)^{\dagger}] \\ &= \text{trace}\left[\left((L(t) + L_e)^{\dagger,2} - L(t)^{\dagger,2}\right)(L_{\mathcal{E}_2} - L_{\mathcal{E}_1})\right] \leq 0. \end{aligned}$$

To obtain the second equality we used the matrix derivative formula $\frac{d}{dt} \text{trace}(L(t)^{\dagger}) = \text{trace}(L(t)^{\dagger} \frac{d}{dt} L(t) L(t)^{\dagger})$ which holds whenever $L(t)$ has constant rank for all t [11], which we have here since the given graph is connected and thus $\text{rank}(L_{\mathcal{E}}) = n - 1, \forall \mathcal{E} \subseteq \mathcal{E}_c$. To obtain the third equality we used the cyclic property of trace. Since $(L(t) + L_e)^{\dagger,2} - L(t)^{\dagger,2} \preceq 0$ and $L_{\mathcal{E}_2} - L_{\mathcal{E}_1} \succeq 0$, the last inequality holds because the trace of the product of a positive and negative semidefinite matrix is non-positive. Since

$$\hat{f}_e(L(1)) = \hat{f}_e(L(0)) + \int_0^1 \frac{d}{dt} \hat{f}_e(L(t)) dt,$$

it follows that $\hat{f}_e(L(1)) = f_e(\mathcal{E}_2) \leq \hat{f}_e(L(0)) = f_e(\mathcal{E}_1)$. Thus, f_e is monotone decreasing, and f is submodular by Theorem 1.

Finally, it can be seen from additivity of the Laplacian that f is monotone increasing, which just means that adding an edge to the graph cannot decrease its coherence. ■

As a consequence, the greedy algorithm is guaranteed to produce a near optimal edge set selection. If the given graph is not connected, the Laplacian changes rank as edges are added. This means that $L(t)$ in the proof does not have constant rank, so \hat{f}_e is not differentiable, and the proof breaks down.

B. Accelerated greedy algorithm and fast rank-one updates

For a sparse connected network, the number of possible edges to be added scales quadratically with the number of nodes. So for the standard greedy algorithm, the marginal gain function may need to be evaluated many times.

Two techniques can be used to significantly speed up the greedy algorithm. First, an accelerated form of the greedy algorithm can be used to reduce the number of times that the marginal gain function is evaluated by exploiting submodularity of the set function [20]. In particular, at each iteration an element is selected that maximizes the marginal benefit of the function given previously chosen elements. The key observation is that submodularity implies $\Delta(s|S_{i+1}) \leq \Delta(s|S_i)$; i.e., the marginal benefits of each element can never increase between algorithm iterations. In the accelerated greedy algorithm, after the first iteration, a list of marginal benefits from the previous iteration sorted in decreasing order is maintained. The marginal benefits for the next iteration are then updated starting from the top of this list. If during this process an element remains at the top of the list after this update, submodularity guarantees that this element has maximal marginal gain, and the algorithm can move to the next iteration without needing to compute the marginal gain

for a potentially very large number of elements. Otherwise, the list is resorted and the process continues. Although the worst case complexity of this accelerated variant is the same as the naive greedy algorithm, speedups of multiple orders of magnitude have been observed in practice [14].

Second, the individual marginal gain function calls can be cheaply performed as rank-one updates using a generalized Sherman-Morrison formula. Note that computing the marginal gain requires computing the trace of the pseudoinverse of a matrix following a rank one update. Although the standard Sherman-Morrison formula does not hold in general for updating the pseudoinverse of a matrix (as opposed to the inverse) [19], we show that for the case of the Laplacian pseudoinverse, one can use the standard formula with the inverse replaced by the pseudoinverse.

Lemma 1: For any connected weighted graph $G = (V, E, w)$ with weighted Laplacian matrix L_E and any edge $e \in V \times V \setminus E$ with given weight w_e and associated weighted incidence matrix row m_e , we have

$$L_{E \cup \{e\}}^\dagger = (L_E + m_e m_e^T)^\dagger = L_E^\dagger - \frac{1}{\beta} L_E^\dagger m_e m_e^T L_E^\dagger, \quad (13)$$

where $\beta = 1 + m_e^T L_E^\dagger m_e$, and correspondingly,

$$\text{trace}(L_{E \cup \{e\}}^\dagger) = \text{trace}(L_E^\dagger) - \frac{1}{\beta} \|L_E^\dagger m_e\|^2. \quad (14)$$

Proof: From Theorem 3 in [19], it holds that

$$(L_E + m_e m_e^T)^\dagger = L_E^\dagger + \frac{1}{\beta} u k^T L_E^\dagger - \frac{\beta}{\sigma} p q,$$

where $u = (I - L_E L_E^\dagger) m_e$, $k = L_E^\dagger m_e$, $\beta = 1 + m_e^T L_E^\dagger m_e$, $p = \frac{\|k\|^2}{\beta} u + k$, $q = \left(\frac{\|u\|^2}{\beta} k^T L_E^\dagger + k^T\right)$, and $\sigma = \|k\|^2 \|u\|^2 + \beta^2$. But since $I - L_E L_E^\dagger = \frac{1}{n} \mathbf{1}\mathbf{1}^T$, it follows that for any incidence matrix row, $(I - L_E L_E^\dagger) m_e = u = 0$, so the expression collapses immediately to (13). Finally, (14) follows from the linearity and cyclic properties of trace. ■ This means that the full Laplacian pseudoinverse needs to be computed from scratch only once at the beginning for the given connected graph, requiring $\mathcal{O}(n^3)$ operations. Then within each iteration, (14) can be used to evaluate the marginal gain for each edge in $\mathcal{O}(n)$ operations since m_e has only two non-zero entries. Finally, after the optimizing edge has been found, the Laplacian pseudoinverse can be updated using (13) in $\mathcal{O}(n^2)$ operations.

C. Constructing Tree Graphs with Optimal Coherence and Nonidentical Edge Weights

We now relax the assumption that the given graph is connected and consider the problem of constructing a tree graph with optimal network coherence. It is known that when the edge weights are identical, a star graph is the tree with optimal network coherence; see, e.g., [6]. However, when non-identical edge weights are given, it is not obvious how to find the tree with optimal coherence. Given a node set $V = \{1, \dots, n\}$ and weights $w_e \geq 0$ associated with each possible edge $e \in V \times V$, the goal is to find an edge set

\mathcal{E} such that the undirected graph $\mathcal{G} = (V, \mathcal{E})$ is a connected tree with minimum $\text{trace}(L_{\mathcal{E}}^\dagger)$. This can be expressed as

$$\begin{aligned} & \underset{\mathcal{E} \subset V \times V}{\text{minimize}} && \text{trace}(L_{\mathcal{E}}^\dagger) \\ & \text{subject to} && |\mathcal{E}| = n - 1, \quad \text{rank}(L_{\mathcal{E}}) = n - 1. \end{aligned} \quad (15)$$

A modified greedy algorithm detailed in Algorithm 1 can be used as a heuristic for this problem. The only difference is

Algorithm 1 Finding a tree with small $\text{trace}(L_{\mathcal{E}}^\dagger)$.

Require: V, w_e for each $e \in V \times V$

- 1: $\mathcal{E} \leftarrow \{\text{argmax}_{e \in V \times V} w_e\}$
 - 2: $\bar{\mathcal{V}} \leftarrow \{i, j \mid (i, j) \in \mathcal{E}\}$
 - 3: **while** $|\mathcal{E}| \leq n - 1$ **do**
 - 4: $\bar{\mathcal{E}} \leftarrow \{(i, j) \mid (i, j) \in V \times V \setminus \mathcal{E}, \{i, j\} \cap \bar{\mathcal{V}} \neq \emptyset\}$
 - 5: $e = \text{argmin}_{e \in \bar{\mathcal{E}}} \text{trace}(L_{\mathcal{E} \cup \{e\}}^\dagger)$
 - 6: $\mathcal{E} \leftarrow \mathcal{E} \cup e$
 - 7: $\bar{\mathcal{V}} \leftarrow \bar{\mathcal{V}} \cup \{i, j \mid e = (i, j)\}$
 - 8: **end while**
-

that the feasible edge set described in line 4 prevents cycles from forming, which would involve redundant edges in terms of connectivity. As in Lemma 1, the value of $\text{trace}(L_{\mathcal{E} \cup \{e\}}^\dagger)$ can be calculated more efficiently. From Theorem 1 of [19]

$$\begin{aligned} \text{trace}(L_{\mathcal{E} \cup \{e\}}^\dagger) &= \text{trace}(L_{\mathcal{E}} + m_e m_e^T)^\dagger \\ &= \text{trace}(L_{\mathcal{E}}^\dagger) - \frac{\text{trace}\left(L_{\mathcal{E}}^\dagger m_e m_e^T (I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger)\right)}{\|(I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e\|^2} \\ &\quad - \frac{\text{trace}\left((I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e m_e^T L_{\mathcal{E}}^\dagger\right)}{\|(I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e\|^2} + \frac{1 + m_e^T L_{\mathcal{E}}^\dagger m_e}{\|(I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e\|^2}. \end{aligned} \quad (16)$$

Since $L_{\mathcal{E}}^\dagger L_{\mathcal{E}} L_{\mathcal{E}}^\dagger = L_{\mathcal{E}}^\dagger$, we have $\text{trace}\left((I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e m_e^T L_{\mathcal{E}}^\dagger\right) = \text{trace}\left(L_{\mathcal{E}}^\dagger (I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e m_e^T\right) = 0$. Hence,

$$\text{trace}(L_{\mathcal{E} \cup \{e\}}^\dagger) = \text{trace}(L_{\mathcal{E}}^\dagger) + \frac{1 + m_e^T L_{\mathcal{E}}^\dagger m_e}{\|(I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e\|^2}. \quad (17)$$

Thus, the marginal gain computations in line 5 of Algorithm 1 can be written as

$$e = \underset{e \in \bar{\mathcal{E}}}{\text{argmin}} \frac{1 + m_e^T L_{\mathcal{E}}^\dagger m_e}{\|(I - L_{\mathcal{E}} L_{\mathcal{E}}^\dagger) m_e\|^2}.$$

Note that Algorithm 1 can be used to add a new node or a set of new vertices to an existing graph such that resulting graph with a small $\text{trace}(L_{\mathcal{E}}^\dagger)$ as well.

IV. ILLUSTRATIVE NUMERICAL EXAMPLES

In this section, we illustrate the results with numerical examples.




Fig. 1. Computation times for the naive and fast versions of the greedy algorithm. The algorithms were applied to compute a set of n edges to add to n -node Erdős-Rényi random networks with the edge probability chosen to be slightly above $\ln n/n$ to ensure connectivity of the generated base graph. For the data at 120 nodes, the fast algorithm exhibits a factor of 350 speed-up.

A. Naive vs. fast greedy algorithm

We first compare the performance of the naive greedy algorithm with that of the modified greedy algorithm using the improvements described in Section III-B. We applied both algorithms to compute a set of n edges to add to n -node Erdős-Rényi random networks with the edge probability chosen to be slightly above $\ln(n)/n$ to ensure connectivity of the generated base graph. Note that even for the smallest networks considered here, e.g. choosing 20 edges to add from around 150, brute force computation is not feasible. However, the greedy algorithm is guaranteed by Theorem 3 to produce a network topology with suboptimality-guarantees. Figure 1 shows computation times for the two algorithms for various network sizes on a laptop with a 1.7 GHz Intel Core i7 processor. One can see a substantial increase in computation time for the naive greedy algorithm around 100 nodes (corresponding to about 4500 possible edges to add). This is also roughly where convex relaxation heuristics based on [29], [10] and using general purpose semidefinite programming solvers begin to have difficulties. The fast greedy algorithm displays significantly better scaling properties: for the data at 120 nodes, the fast algorithm exhibits a factor of 350 speed-up. Our techniques (using unoptimized Python code) were able to near-optimally modify networks with up to 1,000 nodes, with nearly half a million decision variables associated possible edges, in a few hours, which is far beyond the capabilities of current state-of-the-art general purpose semidefinite programming solvers.

B. Experiments with cycles and random graphs

We then examined the qualitative behavior of the algorithm for cycles and other types of random graphs. Figure 2 shows the results of applying to greedy algorithm to add various numbers of edges to a cycle on 50 nodes. The initial added




Fig. 2. Adding various numbers of edges to a cycle graph on 50 nodes.

edges tend to be long distance links, reminiscent of Watts-Strogatz small world graphs [28], but with the link distances intentionally chosen by the algorithm to optimize coherence. When many edges are added, the resulting graph tends to be nearly regular, indicating that small world regular graphs have near optimal coherence. A similar story emerges for Erdős-Rényi random graphs. Figure 3 shows the result of applying the greedy algorithm. In small added edge sets, the added edges tend to connect distant low-degree vertices, and in large added edge sets, the result tends to a regular graph with small-world-like long distance connections. Finally, we also applied the greedy algorithm to Barabasi-Albert scale free networks, in which a preferential attachment mechanism leads to power law degree distributions. Figure 4 shows a set of 10 edges added to a scale free tree on 100 nodes. We observe that some of the added edges tend to connect highly connected hubs together, while others make low-degree long-distance connections.

V. SUMMARY AND CONCLUSIONS

In summary, we considered a network topology design problem in which the objective is to select a set of edges to add to a given graph to optimize the coherence of the resulting network. We showed that network coherence is a submodular function of the network topology, which means that a simple greedy algorithm can be used to select a near optimal edge subset. A modified fast greedy algorithm was developed using Sherman-Morrison pseudoinverse updates and exploiting the submodularity property and provides a computational speed-up of several orders of magnitude,




Fig. 3. Adding 10 edges to an Erdős-Rényi random graph on 50 nodes (left); adding 35 edges to an Erdős-Rényi random graph on 20 nodes (right). In small added edge sets, the added edges tend to connect distant low-degree vertices, and in large added edge sets, the result converges to a regular graph with small-world-like long distance connections.




Fig. 4. A set of 10 edges added to a Barabasi-Albert scale free network with power law degree distribution. Some of the added edges tend to connect highly connected nodes, while others make low-degree long-distance connections.

allowing scaling to network sizes far beyond the capabilities of current state of the art semidefinite programming solvers.

Our current work is applying our algorithms to study various network coherence metrics for networks with second-order dynamics and associated power network models and using the algorithms for design of wide-area controllers.

REFERENCES

- [1] B. Bamieh and D.F. Gayme. The price of synchrony: Resistive losses due to phase synchronization in power networks. In *American Control Conference (ACC)*, 2013, pages 5815–5820. IEEE, 2013.
- [2] B. Bamieh, M.R. Jovanovic, P. Mitra, and S. Patterson. Coherence in large-scale networks: Dimension-dependent limitations of local feedback. *IEEE Transactions on Automatic Control*, 57(9):2235–2249, 2012.
- [3] Y.Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. In *Proceedings of the Eighth IEEE International Conference on Computer Vision*, volume 1, pages 105–112. IEEE, 2001.
- [4] L. Bushnell, A. Clark, and R. Poovendran. A supermodular optimization framework for leader selection under link noise in linear multi-agent systems. *IEEE Transactions on Automatic Control*, 59(2):283–296, 2014.
- [5] F.L. Cortesi, T.H. Summers, and J. Lygeros. Submodularity of energy related controllability metrics. In *to appear, IEEE Conference on Decision and Control, Los Angeles, CA*, 2014.

- [6] W. Ellens, F.M. Spieksma, P. Van Mieghem, A. Jamakovic, and R.E. Koopij. Effective graph resistance. *Linear algebra and its applications*, 435(10):2491–2506, 2011.
- [7] M. Fardad, F. Lin, and M.R. Jovanovic. Algorithms for leader selection in large dynamical networks: Noise-free leaders. In *IEEE Conference on Decision and Control*, pages 7188–7193. IEEE, 2011.
- [8] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [9] K. Fitch and N.E. Leonard. Information centrality and optimal leader selection in noisy networks. In *IEEE Conference on Decision and Control*, pages 7510–7515. IEEE, 2013.
- [10] A. Ghosh, S. Boyd, and A. Saberi. Minimizing effective resistance of a graph. *SIAM Review*, 50(1):37–66, 2008.
- [11] G.H. Golub and V. Pereyra. The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate. *SIAM Journal on numerical analysis*, 10(2):413–432, 1973.
- [12] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146. ACM, 2003.
- [13] D.J. Klein and M. Randić. Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–95, 1993.
- [14] A. Krause and D. Golovin. Submodular function maximization. *Tractability: Practical Approaches to Hard Problems*, 3, 2012.
- [15] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *The Journal of Machine Learning Research*, 9:235–284, 2008.
- [16] F. Lin, M. Fardad, and M.R. Jovanovic. Algorithms for leader selection in large dynamical networks: Noise-corrupted leaders. In *IEEE Conference on Decision and Control*, pages 2932–2937. IEEE, 2011.
- [17] F. Lin, M. Fardad, and M.R. Jovanovic. Algorithms for leader selection in stochastically forced consensus networks. *IEEE Transactions on Automatic Control*, 59(7):1789–1802, 2013.
- [18] L. Lovász. Submodular functions and convexity. *Mathematical programming: the state of the art*, pages 235–257, 1983.
- [19] C.D. Meyer, Jr. Generalized inversion of modified matrices. *SIAM Journal on Applied Mathematics*, 24(3):315–323, 1973.
- [20] M. Minoux. Accelerated greedy algorithms for maximizing submodular set functions. In *Optimization Techniques*, pages 234–243. Springer, 1978.
- [21] G.L. Nemhauser, L.A. Wolsey, and M.L. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294, 1978.
- [22] S. Patterson and B. Bamieh. Leader selection for optimal network coherence. In *IEEE Conference on Decision and Control*, pages 2692–2697. IEEE, 2010.
- [23] S. Patterson and B. Bamieh. Network coherence in fractal graphs. In *IEEE Conference on Decision and Control*, pages 6445–6450. IEEE, 2011.
- [24] I. Shames and T.H. Summers. Rigid network design via submodular set function optimization. *submitted to IEEE Transactions on Network Science and Engineering*, 2014.
- [25] M. Siami and N. Motee. Graph-theoretic bounds on disturbance propagation in interconnected linear dynamical networks. *arXiv preprint arXiv:1403.1494*, 2014.
- [26] T.H. Summers, F. Cortesi, and J. Lygeros. On submodularity and controllability in complex dynamical networks. *arXiv preprint arXiv:1404.7665*, 2014.
- [27] T.H. Summers and J. Lygeros. Optimal sensor and actuator placement in complex dynamical networks. In *IFAC World Congress, Cape Town, South Africa*, pages 3784–3789, 2014.
- [28] D.J. Watts and S.H. Strogatz. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442, 1998.
- [29] L. Xiao, S. Boyd, and S.-J. Kim. Distributed average consensus with least-mean-square deviation. *Journal of Parallel and Distributed Computing*, 67(1):33–46, 2007.