

Topology Design for Stochastically Forced Consensus Networks

Sepideh Hassan-Moghaddam, *Student Member, IEEE*, and Mihailo R. Jovanović , *Senior Member, IEEE*

Abstract—We study an optimal control problem aimed at adding a certain number of edges to an undirected network, with a known graph Laplacian, in order to optimally enhance closed-loop performance. The performance is quantified by the steady-state variance amplification of the network with additive stochastic disturbances. To promote controller sparsity, we introduce ℓ_1 -regularization into the optimal \mathcal{H}_2 formulation and cast the design problem as a semidefinite program. We derive a Lagrange dual, provide interpretation of dual variables, and exploit structure of the optimality conditions for undirected networks to develop customized proximal gradient and Newton algorithms that are well suited for large problems. We illustrate that our algorithms can solve the problems with more than million edges in the controller graph in a few minutes, on a PC. We also exploit structure of connected resistive networks to demonstrate how additional edges can be systematically added in order to minimize the \mathcal{H}_2 norm of the closed-loop system.

Index Terms—Convex optimization, coordinate descent, effective resistance, ℓ_1 -regularization, network coherence, proximal gradient and Newton methods, semidefinite programming, sparsity-promoting control, stochastically forced networks.

I. INTRODUCTION

CONVENTIONAL optimal control of distributed systems relies on centralized implementation of control policies. In large networks of dynamical systems, centralized information processing imposes a heavy burden on individual nodes and is often infeasible. This motivates the development of distributed control strategies that require limited information exchange between the nodes to reach consensus or guarantee synchronization. Over the last decade, a vast body of literature has dealt with analysis, fundamental performance limitations, and design of distributed averaging protocols; e.g., see [1]–[8].

Optimal design of the edge weights for networks with pre-specified topology has received significant attention. In [2], the design of the fastest averaging protocol for undirected networks was cast as a semidefinite program (SDP). Two customized algorithms, based on primal barrier interior-point (IP) and

subgradient methods, were developed and the advantages of optimal weight selection over commonly used heuristics were demonstrated. Similar SDP characterization, for networks with state-dependent graph Laplacians, was provided in [3]. The allocation of symmetric edge weights that minimize the mean-square deviation from average for networks with additive stochastic disturbances was solved in [4]. A related problem, aimed at minimizing the total effective resistance of resistive networks, was addressed in [6]. In [7], the edge Laplacian was used to provide graph-theoretic characterization of the \mathcal{H}_2 and \mathcal{H}_∞ symmetric agreement protocols.

Network coherence quantifies the ability of distributed estimation and control strategies to guard against exogenous disturbances [5], [8]. The coherence is determined by the sum of reciprocals of the nonzero eigenvalues of the graph Laplacian and its scaling properties cannot be predicted by algebraic connectivity of the network. In [8], performance limitations of spatially localized consensus protocols on regular lattices were examined. It was shown that *the fundamental limitations for large-scale networks are dictated by the network topology rather than by the optimal selection of the edge weights*. Moreover, epidemic spread in networks is strongly influenced by their topology [9]–[11]. Thus, optimal topology design represents an important challenge. It is precisely this problem, for undirected consensus networks, that we address in the paper.

More specifically, we study an optimal control problem aimed at achieving a *desired tradeoff* between the *network performance* and *communication requirements* in the distributed controller. Our goal is to add a certain number of edges to a given undirected network in order to optimally enhance the closed-loop performance. One of our key contributions is the formulation of topology design as an optimal control problem that admits convex characterization and is amenable to the development of efficient optimization algorithms. In our formulation, the plant network can contain disconnected components and optimal topology of the controller network is an integral part of the design. In general, this problem is NP-hard [12] and it amounts to an intractable combinatorial search. Several references have examined convex relaxations or greedy algorithms to design topology that optimizes algebraic connectivity [13] or network coherence [14]–[17].

We tap on recent developments regarding sparse representations in conjunction with regularization penalties on the level of communication in a distributed controller. This allows us to

Manuscript received October 17, 2016; revised January 27, 2017; accepted February 20, 2017. Date of publication February 24, 2017; date of current version September 17, 2018. This work was supported in part by the 3M Graduate Fellowship, in part by the UMN Informatics Institute Transdisciplinary Faculty Fellowship, and in part by the National Science Foundation under award ECCS-1407958. Recommended by Associate Editor S. Martinez

The authors are with the Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 USA (e-mail: hassanmo@usc.edu; mihailo@usc.edu).

Digital Object Identifier 10.1109/TCNS.2017.2674962

formulate convex optimization problems that exploit the underlying structure and are amenable to the development of efficient optimization algorithms. To avoid combinatorial complexity, we approach optimal topology design using a sparsity-promoting optimal control framework introduced in [18] and [19]. Performance is captured by the \mathcal{H}_2 norm of the closed-loop network and ℓ_1 -regularization is introduced to promote controller sparsity. While this problem is in general nonconvex [19], for undirected networks we show that it admits a convex characterization with a nondifferentiable objective function and a positive definite constraint. This problem can be transformed into an SDP and, for small size networks, the optimal solution can be computed using standard IP method solvers, e.g., SeDuMi [20] and SDPT3 [21].

To enable design of large networks, we pay particular attention to the computational aspects of the edge-addition problem. We derive a Lagrange dual of the optimal control problem, provide interpretation of dual variables, and develop efficient proximal algorithms. Furthermore, building on preliminary work [22], we specialize our algorithms to the problem of growing connected resistive networks described in [13] and [6]. In this, the plant graph is connected and inequality constraints amount to nonnegativity of controller edge weights. This allows us to simplify optimality conditions and further improve computational efficiency of our customized algorithms.

Proximal gradient algorithms [23] and their accelerated variants [24] have recently found use in distributed optimization, statistics, machine learning, image and signal processing. They can be interpreted as generalization of standard gradient projection to problems with nonsmooth and extended real-value objective functions. When the proximal operator is easy to evaluate, these algorithms are simple yet extremely efficient.

For networks that can contain disconnected components and nonpositive edge weights, we show that the proximal gradient algorithm iteratively updates the controller graph Laplacian via convenient use of the soft-thresholding operator. This extends the iterative shrinkage thresholding algorithm (ISTA) to optimal topology design of undirected networks. In contrast to the ℓ_1 -regularized least squares, however, the step-size has to be selected to guarantee positivity of the second smallest eigenvalue of the closed-loop graph Laplacian. We combine the Barzilai–Borwein (BB) step-size initialization with backtracking to achieve this goal and enhance the rate of convergence. The biggest computational challenge comes from evaluation of the objective function and its gradient. We exploit problem structure to speed up computations and save memory. Finally, for the problem of growing connected resistive networks, the proximal algorithm simplifies to gradient projection which additionally improves the efficiency.

We also develop a customized algorithm based on the proximal Newton method. In contrast to the proximal gradient, this method sequentially employs the second-order Taylor series approximation of the smooth part of the objective function; e.g., see [25]. We use cyclic coordinate descent over the set of active variables to efficiently compute the Newton direction by consecutive minimization with respect to individual coordinates. Similar approach has been recently utilized in a number of

applications, including sparse inverse covariance estimation in graphical models [26].

Both of our customized proximal algorithms significantly outperform a primal-dual IP method developed in [22]. It is worth noting that the latter is significantly faster than the general-purpose solvers. While the customized IP algorithm of [22] with a simple diagonal preconditioner can solve the problems with hundreds of thousands of edges in the controller graph in several hours, on a PC, the customized algorithms based on proximal gradient and Newton methods can solve the problems with millions of edges in several minutes. Furthermore, they are considerably faster than the greedy algorithm with efficient rank-one updates developed in [17].

Our presentation is organized as follows. In Section II, we formulate the problem of optimal topology design for undirected networks subject to additive stochastic disturbances. In Section III, we derive a Lagrange dual of the sparsity-promoting optimal control problem, provide interpretation of dual variables, and construct dual feasible variables from the primal ones. In Section IV, we develop customized algorithms based on the proximal gradient and Newton methods. In Section V, we achieve additional speedup by specializing our algorithms to the problem of growing connected resistive networks. In Section VI, we use computational experiments to design optimal topology of a controller graph for benchmark problems and demonstrate efficiency of our algorithms. In Section VII, we provide a brief overview of the paper.

II. PROBLEM FORMULATION

We consider undirected consensus networks with n nodes

$$\dot{\psi} = -L_p \psi + u + d \quad (1)$$

where d and u are the exogenous disturbance and the control input, respectively, ψ is the state of the network, and L_p is a symmetric $n \times n$ matrix that represents graph Laplacian of the open-loop system, i.e., plant. Such networks arise in applications ranging from load balancing to power systems to opinion formation to control of multiagent systems. The goal is to improve performance of a consensus algorithm in the presence of stochastic disturbances by adding a certain number of edges (from a given set of candidate edges). We formulate this problem as a feedback design problem with

$$u = -L_x \psi$$

where the symmetric feedback-gain matrix L_x is required to have the Laplacian structure. This implies that each node in (1) forms control action using a weighted sum of the differences between its own state and the states of other nodes and that information is processed in a symmetric fashion. Since a nonzero ij th element of L_x corresponds to an edge between the nodes i and j , the communication structure in the controller graph is determined by the sparsity pattern of the matrix L_x .

Upon closing the loop, we obtain

$$\dot{\psi} = -(L_p + L_x) \psi + d. \quad (2a)$$

For a given L_p , our objective is to design the topology for L_x and the corresponding edge weights x in order to achieve the desired tradeoff between controller sparsity and network performance. The performance is quantified by the steady-state variance amplification of the stochastically forced network, from the white-in-time input d to the performance output ζ ,

$$\zeta := \begin{bmatrix} Q^{1/2} \\ 0 \end{bmatrix} \psi + \begin{bmatrix} 0 \\ R^{1/2} \end{bmatrix} u = \begin{bmatrix} Q^{1/2} \\ -R^{1/2} L_x \end{bmatrix} \psi \quad (2b)$$

which penalizes deviation from consensus and control effort. Here, $Q = Q^T \succeq 0$ and $R = R^T \succ 0$ are the state and control weights in the standard quadratic performance index.

The interesting features of this problem come from structural restrictions on the Laplacian matrices L_p and L_x . Both of them are symmetric and are restricted to having an eigenvalue at zero with the corresponding eigenvector of all ones,

$$L_p \mathbb{1} = 0, \quad L_x \mathbb{1} = 0. \quad (3)$$

Since each node uses relative information exchange with its neighbors to update its state, in the presence of white noise, the average mode $\bar{\psi}(t) := (1/n) \mathbb{1}^T \psi(t)$ experiences a random walk and its variance increases linearly with time. To make the average mode unobservable from the performance output ζ , the matrix Q is also restricted to having an eigenvalue at zero associated with the vector of all ones, $Q \mathbb{1} = 0$. Furthermore, to guarantee observability of the remaining eigenvalues of L_p , we consider state weights that are positive definite on the orthogonal complement of the subspace spanned by the vector of all ones, $Q + (1/n) \mathbb{1} \mathbb{1}^T \succ 0$; e.g., $Q = I - (1/n) \mathbb{1} \mathbb{1}^T$ penalizes mean-square deviation from the network average.

In what follows, we express L_x as

$$L_x := \sum_{l=1}^m x_l \xi_l \xi_l^T = E \operatorname{diag}(x) E^T \quad (4)$$

where E is the incidence matrix of the controller graph L_x , m is the number of edges in L_x , and $\operatorname{diag}(x)$ is a diagonal matrix containing the vector of the edge weights $x \in \mathbb{R}^m$. The matrix E is given and it determines the set of candidate edges in controller network. This set can contain all possible edges in the network or it can only include edges that are not in the plant network. Many other options are possible as long as the union of the sets of edges in the plant and controller networks yields a connected graph. We note that the size of the set of candidate edges in controller network influences computational complexity of our algorithms.

It is desired to select a subset of edges in order to balance the closed-loop performance with the number of added edges. Vectors $\xi_l \in \mathbb{R}^n$ determine the columns of E and they signify the connection with weight x_l between nodes i and j : the i th and j th entries of ξ_l are 1 and -1 and all other entries are equal to 0. Thus, L_x given by (4) satisfies structural requirements on the controller graph Laplacian in (3) by construction.

To achieve consensus in the absence of disturbances, the closed-loop network has to be connected [1]. Equivalently, the second smallest eigenvalue of the closed-loop graph Laplacian, $L := L_p + L_x$, has to be positive, i.e., L has to be positive

definite on $\mathbb{1}^\perp$. This amounts to positive definiteness of the “strengthened” graph Laplacian of the closed-loop network

$$\begin{aligned} G &:= L_p + L_x + (1/n) \mathbb{1} \mathbb{1}^T \\ &= G_p + E \operatorname{diag}(x) E^T \succ 0 \end{aligned} \quad (5a)$$

where

$$G_p := L_p + (1/n) \mathbb{1} \mathbb{1}^T. \quad (5b)$$

Structural restrictions (3) on the Laplacian matrices introduce an additional constraint on the matrix G ,

$$G \mathbb{1} = \mathbb{1}. \quad (5c)$$

A. Design of Optimal Sparse Topology

Let d be a white stochastic disturbance with zero-mean and unit variance,

$$\mathbf{E}(d(t)) = 0, \quad \mathbf{E}(d(t_1) d^T(t_2)) = I \delta(t_1 - t_2)$$

where \mathbf{E} is the expectation operator. The square of the \mathcal{H}_2 norm of the transfer function from d to ζ ,

$$\|H\|_2^2 = \lim_{t \rightarrow \infty} \mathbf{E}(\psi^T(t) (Q + L_x R L_x) \psi(t))$$

quantifies the steady-state variance amplification of closed-loop system (2). As noted earlier, the network average $\bar{\psi}(t)$ corresponds to the zero eigenvalue of the graph Laplacian and it is not observable from the performance output ζ . Thus, the \mathcal{H}_2 norm is equivalently given by

$$\begin{aligned} \|H\|_2^2 &= \lim_{t \rightarrow \infty} \mathbf{E}(\tilde{\psi}^T(t) (Q + L_x R L_x) \tilde{\psi}(t)) \\ &= \operatorname{trace}(P (Q + L_x R L_x)) = \langle P, Q + L_x R L_x \rangle \end{aligned}$$

where $\tilde{\psi}(t)$ is the vector of deviations of the states of individual nodes from $\bar{\psi}(t)$,

$$\tilde{\psi}(t) := \psi(t) - \mathbb{1} \bar{\psi}(t) = (I - (1/n) \mathbb{1} \mathbb{1}^T) \psi(t)$$

and P is the steady-state covariance matrix of $\tilde{\psi}$,

$$P := \lim_{t \rightarrow \infty} \mathbf{E}(\tilde{\psi}(t) \tilde{\psi}^T(t)).$$

The above measure of the amplification of stochastic disturbances is determined by $\|H\|_2^2 = (1/2)J(x)$, where

$$J(x) := \left\langle (G_p + E \operatorname{diag}(x) E^T)^{-1}, Q + L_x R L_x \right\rangle. \quad (6)$$

It can be shown that J can be expressed as

$$\begin{aligned} J(x) &= \left\langle (G_p + E \operatorname{diag}(x) E^T)^{-1}, Q_p \right\rangle + \\ &\quad \operatorname{diag}(E^T R E)^T x - \langle R, L_p \rangle - 1 \end{aligned} \quad (7)$$

with

$$Q_p := Q + (1/n) \mathbb{1} \mathbb{1}^T + L_p R L_p.$$

Note that the last two terms in (7) do not depend on the optimization variable x and that the term $L_p R L_p$ in Q_p has an interesting interpretation: it determines a state-weight that guarantees inverse optimality (in LQR sense) of $u = -L_p \psi$ for a system with no coupling between the nodes, $\dot{\psi} = u + d$.

We formulate the design of a controller graph that provides an optimal tradeoff between the \mathcal{H}_2 performance of the closed-loop network and the controller sparsity as

$$\begin{aligned} & \underset{x}{\text{minimize}} && J(x) + \gamma \|x\|_1 \\ & \text{subject to} && G_p + E \text{diag}(x) E^T \succ 0 \end{aligned} \quad (\text{SP})$$

where $J(x)$ and G_p are given by (7) and (5b), respectively. The ℓ_1 norm of x , $\|x\|_1 := \sum_{l=1}^m |x_l|$, is introduced as a convex proxy for promoting sparsity. In (SP), the vector of the edge weights $x \in \mathbb{R}^m$ is the optimization variable; the problem data are the positive regularization parameter γ , the state and control weights Q and R , the plant graph Laplacian L_p , and the incidence matrix of the controller graph E .

The sparsity-promoting optimal control problem (SP) is a constrained optimization problem with a convex nondifferentiable objective function [14] and a positive definite inequality constraint. This implies convexity of (SP). Positive definiteness of the strengthened graph Laplacian G guarantees stability of the closed-loop network (2a) on the subspace $\mathbb{1}^\perp$, and thereby consensus in the absence of disturbances [1].

The consensus can be achieved even if some edge weights are negative [2], [4]. By expressing x as a difference between two nonnegative vectors x_+ and x_- , (SP) can be written as

$$\begin{aligned} & \underset{x_+, x_-}{\text{minimize}} && \left\langle (G_p + E \text{diag}(x_+ - x_-) E^T)^{-1}, Q_p \right\rangle + \\ & && (\gamma \mathbb{1} + c)^T x_+ + (\gamma \mathbb{1} - c)^T x_- \\ & \text{subject to} && G_p + E \text{diag}(x_+ - x_-) E^T \succ 0 \\ & && x_+ \geq 0, x_- \geq 0 \end{aligned} \quad (8)$$

where $c := \text{diag}(E^T R E)$. By utilizing the Schur complement, (8) can be cast to an SDP, and solved via standard IP method algorithms for small size networks.

1) *Reweighted ℓ_1 Norm:* An alternative proxy for promoting sparsity is given by the weighted ℓ_1 norm [27], $\|w \circ x\|_1 := \sum_{l=1}^m w_l |x_l|$ where \circ denotes element-wise product. The vector of nonnegative weights $w \in \mathbb{R}^m$ can be selected to provide better approximation of nonconvex cardinality function than the ℓ_1 norm. An effective heuristic for weight selection is given by the iterative reweighted algorithm [27], with w_l inversely proportional to the magnitude of x_l in the previous iteration,

$$w_l^+ = 1/(|x_l| + \varepsilon). \quad (9)$$

This puts larger emphasis on smaller optimization variables, where a small positive parameter ε ensures that w_l^+ is well defined. If the weighted ℓ_1 norm is used in (SP), the vector of all ones $\mathbb{1}$ should be replaced by the vector w in (8).

B. Structured Optimal Control Problem: Debiasing Step

After the structure of the controller graph Laplacian L_x has been designed, we fix the structure of L_x and optimize the corresponding edge weights. This “polishing” or “debiasing” step is used to improve the performance relative to the solution of the regularized optimal control problem (SP); see [28, Sec. 6.3.2] for additional information. The structured optimal

control problem is obtained by eliminating the columns from the incidence matrix E that correspond to zero elements in the vector of the optimal edge weights x^* resulting from (SP). This yields a new incidence matrix \hat{E} and leads to

$$\begin{aligned} & \underset{x}{\text{minimize}} && \left\langle (G_p + \hat{E} \text{diag}(x) \hat{E}^T)^{-1}, Q_p \right\rangle + \\ & && \text{diag}(\hat{E}^T R \hat{E})^T x \\ & \text{subject to} && G_p + \hat{E} \text{diag}(x) \hat{E}^T \succ 0. \end{aligned}$$

Alternatively, this optimization problem is obtained by setting $\gamma = 0$ in (SP) and by replacing the incidence matrix E with \hat{E} . The solution provides the optimal vector of the edge weights x for the controller graph Laplacian with the desired structure.

C. Gradient and Hessian of $J(x)$

We next summarize the first- and second-order derivatives of the objective function J , given by (7), with respect to the vector of the edge weights x . The second-order Taylor series approximation of $J(x)$ around $\bar{x} \in \mathbb{R}^m$ is given by

$$J(\bar{x} + \tilde{x}) \approx J(\bar{x}) + \nabla J(\bar{x})^T \tilde{x} + \frac{1}{2} \tilde{x}^T \nabla^2 J(\bar{x}) \tilde{x}.$$

For related developments, we refer the reader to [6].

Proposition 1: The gradient and the Hessian of J at $\bar{x} \in \mathbb{R}^m$ are determined by

$$\begin{aligned} \nabla J(\bar{x}) &= -\text{diag}(E^T(Y(\bar{x}) - R)E) \\ \nabla^2 J(\bar{x}) &= H_1(\bar{x}) \circ H_2(\bar{x}) \end{aligned}$$

where

$$\begin{aligned} Y(\bar{x}) &:= (G_p + E D_{\bar{x}} E^T)^{-1} Q_p (G_p + E D_{\bar{x}} E^T)^{-1} \\ H_1(\bar{x}) &:= E^T Y(\bar{x}) E \\ H_2(\bar{x}) &:= E^T (G_p + E D_{\bar{x}} E^T)^{-1} E \\ D_{\bar{x}} &:= \text{diag}(\bar{x}). \end{aligned}$$

III. DUAL PROBLEM

Herein, we study the Lagrange dual of the sparsity-promoting optimal control problem (8), provide interpretation of dual variables, and construct dual feasible variables from primal feasible variables. Since minimization of the Lagrangian associated with (8) does not lead to an explicit expression for the dual function, we introduce an auxiliary variable G and find the dual of

$$\begin{aligned} & \underset{G, x_{\pm}}{\text{minimize}} && \langle G^{-1}, Q_p \rangle + (\gamma \mathbb{1} + c)^T x_+ + (\gamma \mathbb{1} - c)^T x_- \\ & \text{subject to} && G - G_p - E \text{diag}(x_+ - x_-) E^T = 0 \\ & && G \succ 0, x_+ \geq 0, x_- \geq 0. \end{aligned} \quad (\text{P})$$

In (P), G represents the “strengthened” graph Laplacian of the closed-loop network and the equality constraint comes from (5a). As we show next, the Lagrange dual of the primal optimization problem (P) admits an explicit characterization.

Proposition 2: The Lagrange dual of the primal optimization problem (P) is given by

$$\begin{aligned} & \underset{Y}{\text{maximize}} \quad 2 \text{trace} \left((Q_p^{1/2} Y Q_p^{1/2})^{1/2} \right) - \langle Y, G_p \rangle \\ & \text{subject to} \quad \| \text{diag} (E^T (Y - R) E) \|_\infty \leq \gamma \\ & \quad Y \succ 0, Y \mathbb{1} = \mathbb{1} \end{aligned} \quad (\text{D})$$

where $Y = Y^T \in \mathbb{R}^{n \times n}$ is the dual variable associated with the equality constraint in (P). The duality gap is

$$\eta = y_+^T x_+ + y_-^T x_- = \mathbb{1}^T (y_+ \circ x_+ + y_- \circ x_-) \quad (10)$$

where

$$y_+ = \gamma \mathbb{1} - \text{diag} (E^T (Y - R) E) \geq 0 \quad (11a)$$

$$y_- = \gamma \mathbb{1} + \text{diag} (E^T (Y - R) E) \geq 0. \quad (11b)$$

are the Lagrange multipliers associated with element-wise inequality constraints in (P).

Proof: The Lagrangian of (P) is given by

$$\begin{aligned} \mathcal{L} = & \langle G^{-1}, Q_p \rangle + \langle Y, G \rangle - \langle Y, G_p \rangle + \\ & (\gamma \mathbb{1} - \text{diag} (E^T (Y - R) E) - y_+)^T x_+ + \\ & (\gamma \mathbb{1} + \text{diag} (E^T (Y - R) E) - y_-)^T x_-. \end{aligned} \quad (12)$$

Note that no Lagrange multiplier is assigned to the positive definite constraint on G in \mathcal{L} . Instead, we determine conditions on Y and y_\pm that guarantee $G \succ 0$.

Minimizing \mathcal{L} with respect to G yields

$$G^{-1} Q_p G^{-1} = Y \quad (13a)$$

or, equivalently,

$$G = Q_p^{1/2} \left(Q_p^{1/2} Y Q_p^{1/2} \right)^{-1/2} Q_p^{1/2}. \quad (13b)$$

Positive definiteness of G and Q_p implies $Y \succ 0$. Furthermore, since $Q_p \mathbb{1} = \mathbb{1}$, from (5c) and (13a) we have

$$Y \mathbb{1} = \mathbb{1}.$$

Similarly, minimization with respect to x_+ and x_- leads to (11a) and (11b). Thus, nonnegativity of y_+ and y_- amounts to

$$-\gamma \mathbb{1} \leq \text{diag} (E^T (Y - R) E) \leq \gamma \mathbb{1}$$

or, equivalently,

$$\| \text{diag} (E^T (Y - R) E) \|_\infty \leq \gamma.$$

Substitution of (13) and (11) into (12) eliminates y_+ and y_- from the dual problem. We can thus represent the dual function, $\inf_{G, x_\pm} \mathcal{L}(G, x_\pm; Y, y_\pm)$, as

$$2 \text{trace} \left((Q_p^{1/2} Y Q_p^{1/2})^{1/2} \right) - \langle Y, G_p \rangle$$

which allows us to bring the dual of (P) to (D). ■

Any dual feasible Y can be used to obtain a lower bound on the optimal value of the primal problem (P). Furthermore, the difference between the objective functions of the primal (evaluated at the primal feasible (G, x_\pm)) and dual (evaluated at the dual feasible Y) problems yields expression (10) for the

duality gap η , where y_+ and y_- are given by (11a) and (11b). The duality gap can be used to estimate distance to optimality.

Strong duality follows from Slater's theorem [28], i.e., convexity of the primal problem (P) and strict feasibility of the constraints in (P). This implies that at optimality, the duality gap η for the primal problem (P) and the dual problem (D) is zero. Furthermore, if (G^*, x_\pm^*) are optimal points of (P), then $Y^* = (G^*)^{-1} Q_p (G^*)^{-1}$ is the optimal point of (D). Similarly, if Y^* is the optimal point of (D),

$$G^* = Q_p^{1/2} \left(Q_p^{1/2} Y^* Q_p^{1/2} \right)^{-1/2} Q_p^{1/2}$$

is the optimal point of (P). The optimal vector of the edge weights x^* is determined by the nonzero off-diagonal elements of the controller graph Laplacian, $L_x^* = G^* - G_p$.

A. Interpretation of Dual Variables

For electrical networks, the dual variables have appealing interpretations. Let $\iota \in \mathbb{R}^n$ be a random current injected into the resistor network satisfying

$$\mathbb{1}^T \iota = 0, \mathbf{E}(\iota) = 0, \mathbf{E}(\iota \iota^T) = Q + L_p R L_p.$$

The vector of voltages $\vartheta \in \mathbb{R}^m$ across the edges of the network is then given by $\vartheta = E^T G^{-1} \iota$. Furthermore, since

$$\mathbf{E}(\vartheta \vartheta^T) = E^T G^{-1} \mathbf{E}(\iota \iota^T) G^{-1} E = E^T Y E$$

the dual variable Y is related to the covariance matrix of voltages across the edges. Moreover, (11) implies that y_+ and y_- quantify the deviations between variances of edge voltages from their respective upper and lower bounds.

Remark 1: For a primal feasible x , Y resulting from (13a) with G given by (5a) may not be dual feasible. Let

$$\hat{Y} := \beta Y + \frac{1 - \beta}{n} \mathbb{1} \mathbb{1}^T \quad (14a)$$

and let the control weight be $R = r I$ with $r > 0$. If

$$\beta \leq \frac{\gamma + 2r}{\| \text{diag} (E^T (Y - R) E) \|_\infty + 2r} \quad (14b)$$

then \hat{Y} satisfies the inequality constraint in (D) and it is thus dual feasible.

IV. CUSTOMIZED ALGORITHMS

We next exploit the structure of the sparsity-promoting optimal control problem (SP) and develop customized algorithms based on the proximal gradient and Newton methods. The proximal gradient algorithm is a first-order method that uses a simple quadratic approximation of J in (SP). This yields an explicit update of the vector of the edge weights via application of the soft-thresholding operator. In the proximal Newton method a sequential quadratic approximation of the smooth part of the objective function in (SP) is used and the search direction is efficiently computed via cyclic coordinate descent over the set of active variables.

A. Proximal Gradient Method

We next use the proximal gradient method to solve (SP). A simple quadratic approximation of $J(x)$ around the current iterate x^k ,

$$J(x) \approx J(x^k) + \nabla J(x^k)^T (x - x^k) + \frac{1}{2\alpha_k} \|x - x^k\|_2^2$$

is substituted to (SP) to obtain

$$x^{k+1} = \arg \min_x g(x) + \frac{1}{2\alpha_k} \|x - (x^k - \alpha_k \nabla J(x^k))\|_2^2.$$

Here, α_k is the step-size and the update is determined by the proximal operator of the function $\alpha_k g$,

$$x^{k+1} = \text{prox}_{\alpha_k g}(x^k - \alpha_k \nabla J(x^k)).$$

In particular, for $g(x) = \gamma \|x\|_1$, we have

$$x^{k+1} = \mathcal{S}_{\gamma\alpha_k}(x^k - \alpha_k \nabla J(x^k))$$

where $\mathcal{S}_\kappa(y) = \text{sign}(y) \max(|y| - \kappa, 0)$ is the soft-thresholding function.

The proximal gradient algorithm converges with rate $O(1/k)$ if $\alpha_k < 1/L$, where L is the Lipschitz constant of ∇J [23], [24]. It can be shown that ∇J is Lipschitz continuous but, since it is challenging to explicitly determine L , we adjust α_k via backtracking. To provide a better estimate of L , we initialize α_k using the Barzilai-Browne (BB) method which provides an effective heuristic for approximating the Hessian of the function J via the scaled version of the identity [29], $(1/\alpha_k)I$. At the k th iteration, the initial BB step-size $\alpha_{k,0}$,

$$\alpha_{k,0} := \frac{\|x^k - x^{k-1}\|_2^2}{(x^{k-1} - x^k)^T (\nabla J(x^{k-1}) - \nabla J(x^k))} \quad (15)$$

is adjusted via backtracking until the inequality constraint in (SP) is satisfied and

$$\begin{aligned} J(x^{k+1}) &\leq J(x^k) + \nabla J(x^k)^T (x^{k+1} - x^k) \\ &\quad + \frac{1}{2\alpha_k} \|x^{k+1} - x^k\|_2^2. \end{aligned}$$

Since J is continuously differentiable with Lipschitz continuous gradient, this inequality holds for any $\alpha_k < 1/L$ and the algorithm converges sublinearly [24]. This condition guarantees that objective function decreases at every iteration. Our numerical experiments in Section VI suggest that BB step-size initialization significantly enhances the rate of convergence.

Remark 2: The biggest computational challenge comes from evaluation of the objective function and its gradient. Since the inverse of the strengthened graph Laplacian G has to be computed, with direct computations these evaluations take $O(n^3)$ and $O(nm^2)$ flops, respectively. However, by exploiting the problem structure, ∇J can be computed more efficiently. The main cost arises in the computation of $\text{diag}(E^T Y E)$. We instead compute it using $\text{sum}(E^T \circ (Y E))$ which takes $O(n^2 m)$ operations. Here, $\text{sum}(A)$ is a vector that contains summation of each row of the matrix A in its entries. For networks with $m \gg n$ this leads to significant speed up. Moreover, in contrast

to direct computation, we do not need to store the $m \times m$ matrix $E^T Y E$. Only formation of the columns is required, which offers memory saving.

B. Proximal Newton Method

In contrast to the proximal gradient algorithm, the proximal Newton method benefits from second-order Taylor series expansion of the smooth part of the objective function in (SP). Herein, we employ cyclic coordinate descent over the set of active variables to efficiently compute the Newton direction.

By approximating the smooth part of the objective function J in (SP) with the second-order Taylor series expansion around the current iterate \bar{x} ,

$$J(\bar{x} + \tilde{x}) \approx J(\bar{x}) + \nabla J(\bar{x})^T \tilde{x} + \frac{1}{2} \tilde{x}^T \nabla^2 J(\bar{x}) \tilde{x}$$

the problem (SP) becomes

$$\begin{aligned} \underset{\tilde{x}}{\text{minimize}} \quad & \nabla J(\bar{x})^T \tilde{x} + \frac{1}{2} \tilde{x}^T \nabla^2 J(\bar{x}) \tilde{x} + \gamma \|\bar{x} + \tilde{x}\|_1 \\ \text{subject to} \quad & G_p + E \text{diag}(\bar{x} + \tilde{x}) E^T \succ 0. \end{aligned} \quad (16)$$

Let \tilde{x} denote the current iterate approximating the Newton direction. By perturbing \tilde{x} in the direction of the i th standard basis vector e_i in \mathbb{R}^m , the objective function in (16) becomes

$$\begin{aligned} \nabla J(\bar{x})^T (\tilde{x} + \delta_i e_i) + \frac{1}{2} (\tilde{x} + \delta_i e_i)^T \nabla^2 J(\bar{x}) (\tilde{x} + \delta_i e_i) \\ + \gamma |\bar{x}_i + \tilde{x}_i + \delta_i|. \end{aligned}$$

Elimination of constant terms allows us to bring (16) into

$$\underset{\delta_i}{\text{minimize}} \quad \frac{1}{2} a_i \delta_i^2 + b_i \delta_i + \gamma |c_i + \delta_i| \quad (17)$$

where the optimization variable is the scalar δ_i and $(a_i, b_i, c_i, \bar{x}_i, \tilde{x}_i)$ are the problem data with

$$\begin{aligned} a_i &:= e_i^T \nabla^2 J(\bar{x}) e_i \\ b_i &:= (\nabla^2 J(\bar{x}) e_i)^T \tilde{x} + e_i^T \nabla J(\bar{x}) \\ c_i &:= \bar{x}_i + \tilde{x}_i. \end{aligned}$$

The explicit solution to (17) is given by

$$\delta_i = -c_i + \mathcal{S}_{\gamma/a_i}(c_i - b_i/a_i).$$

After the Newton direction \tilde{x} has been computed, we determine the step-size α via backtracking. This guarantees positive definiteness of the strengthened graph Laplacian and sufficient decrease of the objective function. We use generalization of Armijo rule [30] to find an appropriate step-size α such that $G_p + E \text{diag}(\bar{x} + \alpha \tilde{x}) E^T$ is positive definite matrix and

$$\begin{aligned} J(\bar{x} + \alpha \tilde{x}) + \gamma \|\bar{x} + \alpha \tilde{x}\|_1 &\leq J(\bar{x}) + \gamma \|\bar{x}\|_1 + \\ \alpha \sigma (\nabla J(\bar{x})^T \tilde{x} + \gamma \|\bar{x} + \tilde{x}\|_1 - \gamma \|\bar{x}\|_1). \end{aligned}$$

Remark 3: The parameter a_i in (17) is determined by the i th diagonal element of the Hessian $\nabla^2 J(\bar{x})$. On the other hand, the i th column of $\nabla^2 J(\bar{x})$ and the i th element of the gradient vector $\nabla J(\bar{x})$ enter into the expression for b_i . All of these can be obtained directly from $\nabla^2 J(\bar{x})$ and $\nabla J(\bar{x})$ and forming them

does not require any multiplication. Computation of a single vector inner product between the i th column of the Hessian and \tilde{x} is required in b_i , which typically takes $O(m)$ operations. To avoid direct multiplication, in each iteration after finding δ_i , we update the vector $\nabla^2 J(\tilde{x})^T \tilde{x}$ using the correction term $\delta_i(E^T Y E_i) \circ ((G^{-1} E_i)^T E)^T$ and take its i th element to form b_i . Here, E_i is the i th column of the incidence matrix of the controller graph. This also avoids the need to store the Hessian of J , which is an $m \times m$ matrix, thereby leading to a significant memory saving.

Remark 4: Active set strategy is an effective means for determining the directions that do not need to be updated in the coordinate descent algorithm. At each outer iteration, we classify the variable as either active or inactive based on the values of \tilde{x}_i and the i th component of the gradient vector $\nabla J(\tilde{x})$. For $g(x) = \gamma \|x\|_1$, the i th search direction is inactive if

$$\tilde{x}_i = 0 \quad \text{and} \quad |e_i^T \nabla J(\tilde{x})| < \gamma - \epsilon$$

and it is active otherwise. Here, $\epsilon > 0$ is a small number (e.g., $\epsilon = 0.0001\gamma$). The Newton direction is then obtained by solving the optimization problem over the set of active variables. This significantly improves algorithmic efficiency for large values of the regularization parameter γ .

1) *Convergence Analysis:* In (SP), $J(x)$ is smooth for $G_p + E \text{diag}(x) E^T \succ 0$ and the nonsmooth part is given by the ℓ_1 norm of x . The objective function of the form $J(x) + g(x)$ was studied in [26], where J is smooth over the positive definite cone and g is a separable nondifferentiable function. Superlinear (i.e., quadratic) convergence rate of the quadratic approximation method for (SP) is implied from [26, Th. 16].

2) *Stopping Criteria:* The norms of the primal and dual residuals r_p and r_d^\pm as well as the duality gap η are used as stopping criteria. In contrast to the stopping criteria available in the literature, this choice enables fair comparison of the algorithms. We use (14) to construct a dual feasible \hat{Y} and obtain y_\pm from (11) and (10) to compute the duality gap η , and

$$r_p(x, x_\pm) := x - x_+ + x_-$$

$$r_d^+(x, y_+) := \gamma \mathbb{1} - \text{diag}(E^T(\hat{Y} - R)E) - y_+$$

$$r_d^-(x, y_-) := \gamma \mathbb{1} + \text{diag}(E^T(\hat{Y} - R)E) - y_-$$

to determine the primal and dual residuals.

3) *Comparison of Algorithms:* Table I compares and contrasts features of our customized proximal algorithms and the algorithm based on the primal-dual IP method developed in [22].

V. GROWING CONNECTED RESISTIVE NETWORKS

The problem of optimal topology design for stochastically forced networks has many interesting variations. An important class is given by resistive networks in which all edge weights are nonnegative, $x \geq 0$. Here, we study the problem of growing connected resistive networks; e.g., see [13]. In this, the plant graph is connected and there are *no joint edges* between the plant and the controller graphs. Our objective is to enhance the closed-loop performance by adding a small number of edges.

As we show below, inequality constraints in this case amount to nonnegativity of controller edge weights. This simplifies optimality conditions and enables further improvement of the computational efficiency of our customized algorithms.

The restriction on connected plant graphs implies positive definiteness of the strengthened graph Laplacian of the plant, $G_p = L_p + (1/n) \mathbb{1} \mathbb{1}^T \succ 0$. Thus, $G_p + E \text{diag}(x) E^T$ is always positive definite for connected resistive networks and (SP) simplifies to

$$\underset{x}{\text{minimize}} \quad f(x) + g(x) \quad (18)$$

where

$$f(x) := J(x) + \gamma \mathbb{1}^T x$$

and $g(x)$ is the indicator function for the nonnegative orthant,

$$g(x) := I_+(x) = \begin{cases} 0, & x \geq 0 \\ +\infty, & \text{otherwise.} \end{cases}$$

As in Section III, in order to determine the Lagrange dual of the optimization problem (18), we introduce an additional optimization variable G and rewrite (18) as

$$\underset{G, x}{\text{minimize}} \quad \langle G^{-1}, Q_p \rangle + (\gamma \mathbb{1} + \text{diag}(E^T R E))^T x$$

$$\text{subject to} \quad G - G_p - E \text{diag}(x) E^T = 0$$

$$x \geq 0. \quad (\text{P1})$$

Proposition 3: The Lagrange dual of the primal optimization problem (P1) is given by

$$\underset{Y}{\text{maximize}} \quad 2 \text{trace} \left((Q_p^{1/2} Y Q_p^{1/2})^{1/2} \right) - \langle Y, G_p \rangle$$

$$\text{subject to} \quad \text{diag}(E^T(Y - R)E) \leq \gamma \mathbb{1}$$

$$Y \succ 0, Y \mathbb{1} = \mathbb{1} \quad (\text{D1})$$

where Y is the dual variable associated with the equality constraint in (P1). The duality gap is

$$\eta = y^T x = \mathbb{1}^T (y \circ x) \quad (19)$$

where

$$y := \gamma \mathbb{1} - \text{diag}(E^T(Y - R)E) \geq 0 \quad (20)$$

represents the dual variable associated with the nonnegativity constraint on the vector of the edge weights x .

Remark 5: For connected resistive networks with the control weight $R = r I$, \hat{Y} given by (14a) is dual feasible if

$$\beta \leq \frac{\gamma + 2r}{\max(\text{diag}(E^T(Y - R)E)) + 2r}. \quad (21)$$

A. Proximal Gradient Method

Using a simple quadratic approximation of the smooth part of the objective function f around the current iterate x^k

$$f(x) \approx f(x^k) + \nabla f(x^k)^T (x - x^k) + \frac{1}{2\alpha_k} \|x - x^k\|_2^2$$

TABLE I
COMPARISON OF OUR CUSTOMIZED PROXIMAL ALGORITHMS WITH THE PRIMAL DUAL IP METHOD OF [22]

| Algorithm | Primal-dual IP method | Proximal gradient | Proximal Newton |
|---------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| Order | 2nd | 1st | 2nd |
| Search direction | PCG | explicit update | coordinate descent |
| Speed-up strategy | PCG with preconditioner | BB step-size initialization | active set strategy |
| Memory | no storage of $m \times m$ matrices | no storage of $m \times m$ matrices | no storage of $m \times m$ matrices |
| Most expensive part | $O(m^3)$ | $O(n^2 m)$ | $O(m^2)$ |
| Convergence rate | super-linear | linear | super-linear (quadratic) |

the optimal solution of (18) is determined by the proximal operator of the function $g(x) = I_+(x)$,

$$x^{k+1} = (x^k - \alpha_k \nabla f(x^k))_+$$

where $(\cdot)_+$ is the projection on the nonnegative orthant. Thus, the action of the proximal operator is given by the projected gradient.

As in Section IV-A, we initialize α_k using the BB heuristics but we skip the backtracking step here and employ a nonmonotone BB scheme [31], [32]. The effectiveness of this strategy has been established on quadratic problems [29], [31], but its convergence in general is hard to prove. In Section VI, we demonstrate efficiency of this approach.

B. Proximal Newton Method

We next adjust the customized algorithm based on proximal Newton method for growing connected resistive networks. We approximate the smooth part of the objective function f in (18) using the second-order Taylor series expansion around the current iterate \bar{x} ,

$$f(\bar{x} + \tilde{x}) \approx f(\bar{x}) + \nabla f(\bar{x})^T \tilde{x} + \frac{1}{2} \tilde{x}^T \nabla^2 f(\bar{x}) \tilde{x}$$

and rewrite (18) as

$$\begin{aligned} \underset{\tilde{x}}{\text{minimize}} \quad & \nabla f(\bar{x})^T \tilde{x} + \frac{1}{2} \tilde{x}^T \nabla^2 f(\bar{x}) \tilde{x} \\ \text{subject to} \quad & \bar{x} + \tilde{x} \geq 0. \end{aligned} \quad (22)$$

By perturbing \bar{x} in the direction of the i th standard basis vector e_i in \mathbb{R}^m , $\bar{x} + \delta_i e_i$, the objective function in (22) becomes

$$\nabla f(\bar{x})^T (\bar{x} + \delta_i e_i) + \frac{1}{2} (\bar{x} + \delta_i e_i)^T \nabla^2 f(\bar{x}) (\bar{x} + \delta_i e_i).$$

Elimination of constant terms allows us to bring (22) into

$$\begin{aligned} \underset{\delta_i}{\text{minimize}} \quad & \frac{1}{2} a_i \delta_i^2 + b_i \delta_i \\ \text{subject to} \quad & \bar{x}_i + \tilde{x}_i + \delta_i \geq 0. \end{aligned} \quad (23)$$

The optimization variable is the scalar δ_i and a_i , b_i , \bar{x}_i , and \tilde{x}_i are the problem data with

$$\begin{aligned} a_i &:= e_i^T \nabla^2 f(\bar{x}) e_i \\ b_i &:= (\nabla^2 f(\bar{x}) e_i)^T \tilde{x} + e_i^T \nabla f(\bar{x}). \end{aligned}$$

The explicit solution to (23) is given by

$$\delta_i = \begin{cases} -b_i/a_i, & \bar{x}_i + \tilde{x}_i - b_i/a_i \geq 0 \\ -(\bar{x}_i + \tilde{x}_i), & \text{otherwise.} \end{cases}$$

After the Newton direction \tilde{x} has been computed, we determine the step-size α via backtracking. This guarantees positivity of the updated vector of the edge weights, $\bar{x} + \alpha \tilde{x}$, and sufficient decrease of the objective function, $f(\bar{x} + \alpha \tilde{x}) \leq f(\bar{x}) + \alpha \sigma \nabla f(\bar{x})^T \tilde{x}$.

Remark 6: As in Section IV-B, we use an active set strategy to identify the directions that do not need to be updated in the coordinate descent algorithm. For $g(x) = I_+(x)$, the i th search direction is inactive if $\{\bar{x}_i = 0 \text{ and } e_i^T \nabla f(\bar{x}) \geq 0\}$ and it is active otherwise.

1) *Stopping Criteria:* The norm of the dual residual r_d and the duality gap η are used as stopping criteria. The dual variable y is obtained from (20), where Y is given by (14a) and β satisfies (21). At each iteration, η is evaluated using (19) and the dual residual is determined by

$$r_d(x, y) := \gamma \mathbb{I} - \text{diag}(E^T(Y(x) - R)E) - y.$$

VI. COMPUTATIONAL EXPERIMENTS

We next provide examples and evaluate performance of our customized algorithms. Algorithm proxBB represents proximal gradient method with BB step-size initialization and proxN identifies proximal Newton method in which the search direction is found via coordinate descent. Performance is compared with the PCG-based primal-dual IP method of [22] and the greedy algorithm of [17]. We have implemented all algorithms in MATLAB and executed tests on a 3.4 GHz Core(TM) i7-3770 Intel(R) machine with 16 GB RAM.

In all examples, we set $R = I$ and choose the state weight that penalizes the mean-square deviation from the network average, $Q = I - (1/n) \mathbb{1} \mathbb{1}^T$. The absolute value of the dual residual r_d and the duality gap η are used as stopping criteria. We set the tolerances for r_d and η to 10^{-3} and 10^{-4} , respectively. Finally, for connected plant networks

$$\gamma_{\max} := \|\text{diag}(E^T G_p^{-1} Q G_p^{-1} E)\|_{\infty}$$

identifies the value of the regularization parameter γ for which all edge weights in the controller graph are equal to zero.

Additional information about our computational experiments, along with MATLAB source codes, can be found at: www.ece.umn.edu/~mihailo/software/graphsp/

TABLE II
COMPARISON OF ALGORITHMS (SOLVE TIMES IN SECONDS/NUMBER OF ITERATIONS) FOR THE PROBLEM OF GROWING CONNECTED RESISTIVE ERDŐS-RÉNYI NETWORKS WITH DIFFERENT NUMBER OF NODES n , EDGE PROBABILITY $1.05 \log(n)/n$, AND $\gamma = 0.8 \gamma_{\max}$

| number of nodes | $n = 300$ | $n = 700$ | $n = 1000$ | $n = 1300$ | $n = 1500$ |
|-----------------|---------------|----------------|----------------|----------------|-------------------|
| number of edges | $m = 43\,986$ | $m = 242\,249$ | $m = 495\,879$ | $m = 839\,487$ | $m = 1\,118\,541$ |
| IP (PCG) | 16.499/8 | 394.256/13 | 1 014.282/13 | 15 948.164/13 | 179 352.208/14 |
| proxBB | 1.279/11 | 15.353/11 | 55.944/13 | 157.305/16 | 239.567/16 |
| proxN | 1.078/4 | 11.992/4 | 34.759/4 | 82.488/4 | 124.307/4 |

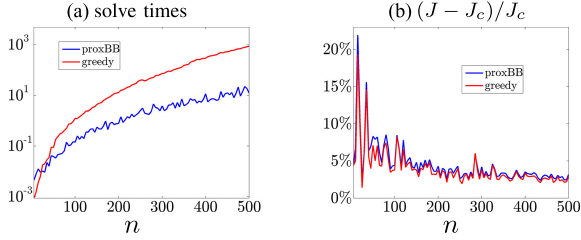


Fig. 1. (a) Solve times (in seconds); and (b) performance degradation (in percents) of proximal gradient and greedy algorithms relative to the optimal centralized controller.

A. Performance Comparison

In what follows, the incidence matrix of the controller graph is selected to satisfy the following requirements. First, in the absence of the sparsity-promoting term, the closed-loop network is given by a complete graph. Second, there are no joint edges between the plant and the controller graphs.

We first solve the problem (P1) for growing connected resistive Erdős-Rényi networks with different number of nodes. The generator of the plant dynamics is given by an undirected unweighted graph with edge probability $1.05 \log(n)/n$. Table II compares our customized algorithms in terms of speed and the number of iterations. Even for small networks, proximal methods are significantly faster than the IP method and proxN takes smaller number of iterations and converges quicker than proxBB. For a larger network (with 1500 nodes and 1 118 541 edges in the controller graph), it takes about 50 hours for the PCG-based IP method to solve the problem. In contrast, proxN and proxBB converge in about 2 and 4 minutes, respectively.

Figure 1 compares our proximal gradient algorithm with the fast greedy algorithm of [17]. We solve problem (P1) for Erdős-Rényi networks with different number of nodes ($n = 5$ to 500) and $\gamma = 0.4 \gamma_{\max}$. After proxBB identifies the edges in the controller graph, we use the greedy method to select the same number of edges. Finally, we polish the identified edge weights for both methods. Figure 1(a) shows the solve times (in seconds) versus the number of nodes. As the number of nodes increases the proximal algorithm significantly outperforms the fast greedy method. Relative to the optimal centralized controller, both methods yield similar performance degradation of the closed-loop network; see Fig. 1(b).

B. Large-Scale Facebook Network

To evaluate effectiveness of our algorithms on large networks, we solve the problem of growing a network of friendships. In

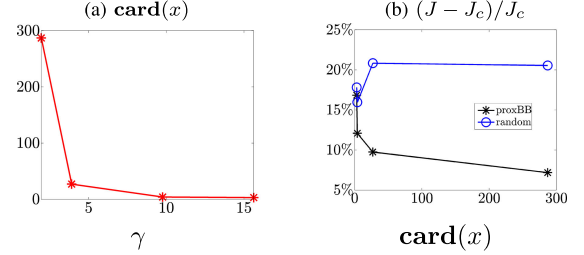


Fig. 2. (a) Sparsity level; and (b) optimal tradeoff curves resulting from the application of proximal gradient algorithm and a heuristic strategy for the Facebook network.

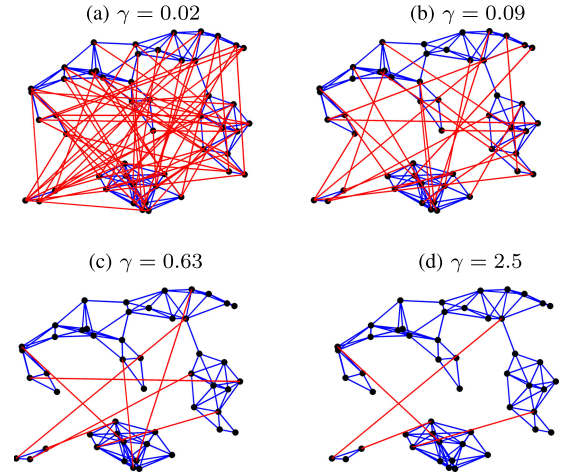


Fig. 3. Topologies of the plant (blue lines) and controller graphs (red lines) for an unweighted random network with three disconnected subgraphs.

such social networks, nodes denote people and edges denote friendships. There is an edge between two nodes if two people are friends. The network is obtained by examining social network of ten users (the so-called ego nodes); all other nodes are friends to at least one of these ego nodes [33]. The resulting network is undirected and unweighted with 4039 nodes and 88 234 edges; the data are available at <http://snap.stanford.edu/data/>. Our objective is to improve performance by adding a small number of extra edges. We assume that people can only form friendships with friends of their friends. This restricts the number of potential edges in the controller graph to 1 358 067.

To avoid memory issues, we have implemented our algorithms in C++. For $\gamma = c \gamma_{\max}$ with $c = \{0.1, 0.2, 0.5, 0.8\}$ and $\gamma_{\max} = 19.525$, the proximal gradient algorithm computes the solution in about 10, 2.6, 0.87, and 0.43 hours, respectively.

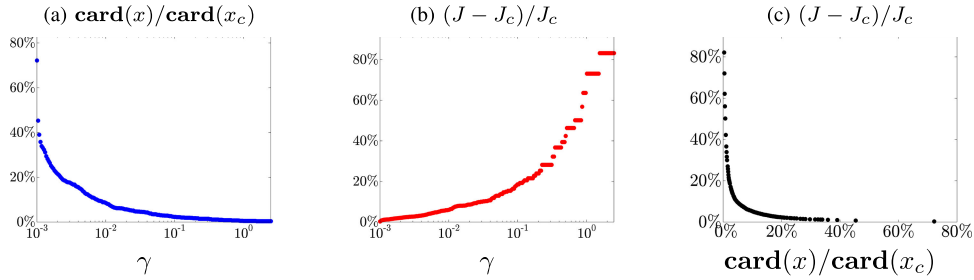


Fig. 4. (a) Sparsity level; (b) performance degradation; and (c) the optimal tradeoff curve between the performance degradation and the sparsity level of optimal sparse x compared to the optimal centralized vector of the edge weights x_c . The results are obtained for unweighted random disconnected plant network with topology shown in Fig. 3.

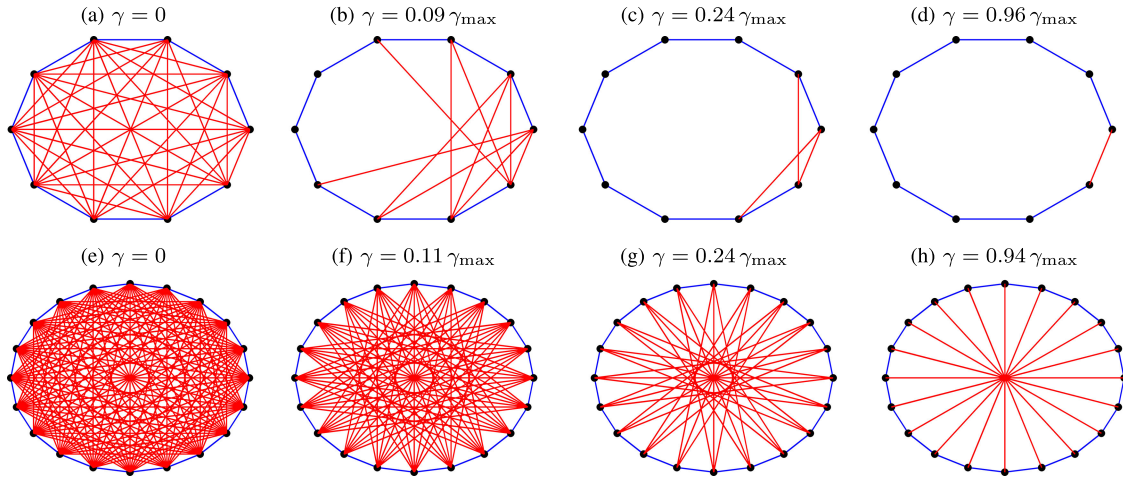


Fig. 5. Problems of growing unweighted path (top row) and ring (bottom row) networks. Blue lines identify edges in the plant graph, and red lines identify edges in the controller graph.

After designing the topology of the controller graph, we optimize the resulting edge weights via polishing.

Figure 2(a) shows that the number of nonzero elements in the vector x decreases as γ increases and Fig. 2(b) illustrates that the \mathcal{H}_2 performance deteriorates as the number of nonzero elements in x decreases. In particular, for $\gamma = 0.8 \gamma_{\max}$, the identified sparse controller has only three nonzero elements (it uses only 0.0002% of the potential edges). Relative to the optimal centralized controller, this controller degrades performance by 16.842%, $(J - J_c)/J_c = 16.842\%$.

In all of our experiments, the added links with the largest edge weights connect either the ego nodes to each other or three nonego nodes to the ego nodes. Thus, our method recognizes significance of the ego nodes and identifies nonego nodes that play an important role in improving performance.

We compare performance of the identified controller to a heuristic strategy that is described next. The controller graph contains 16 potential edges between ego nodes. If the number of edges identified by our method is smaller than 16, we randomly select the desired number of edges between ego nodes. Otherwise, we connect all ego nodes and select the remaining edges in the controller graph randomly. We then use polishing to find the optimal edge weights. The performance of resulting random controller graphs are averaged over ten trials and the

performance loss relative to the optimal centralized controller is displayed in Fig. 2(b). We see that our algorithm always performs better than the heuristic strategy. On the other hand, the heuristic strategy outperforms the strategy that adds edges randomly (without paying attention to ego nodes). Unlike our method, the heuristic strategy does not necessarily improve the performance by increasing the number of added edges. In fact, the performance deteriorates as the number of edges in the controller graph increases from 4 to 27; see Fig. 2(b).

C. Random Disconnected Network

The plant graph (blue lines) in Fig. 3 contains 50 randomly distributed nodes in a region of 10×10 units. Two nodes are neighbors if their Euclidean distance is not greater than 2 units. We examine the problem of adding edges to a plant graph which is not connected and solve the sparsity-promoting optimal control problem (SP) for controller graph with $m = 1094$ potential edges. This is done for 200 logarithmically spaced values of $\gamma \in [10^{-3}, 2.5]$ using the path-following iterative reweighted algorithm as a proxy for inducing sparsity [27]. As indicated by (9), we set the weights to be inversely proportional to the magnitude of the solution x to (SP) at the previous value of γ . We choose $\varepsilon = 10^{-3}$ in (9) and initialize weights for $\gamma = 10^{-3}$

using the solution to (SP) with $\gamma = 0$ (i.e., the optimal centralized vector of the edge weights). Topology design is followed by the polishing step that computes the optimal edge weights; see Section II-B.

As illustrated in Fig. 3, larger values of γ yield sparser controller graphs (red lines). In contrast to all other examples, the plant graph is not connected and the optimal solution is obtained using the algorithms of Section IV. Note that greedy method [17] cannot be used here. Since the plant graph has three disconnected subgraphs, at least two edges in the controller are needed to make the closed-loop network connected.

Figure 4 shows that the number of nonzero elements in the vector of the edge weights x decreases and that the closed-loop performance deteriorates as γ increases. In particular, Fig. 4(c) illustrates the optimal tradeoff curve between the \mathcal{H}_2 performance loss (relative to the optimal centralized controller) and the sparsity of the vector x . For $\gamma = 2.5$, only four edges are added. Relative to the optimal centralized vector of the controller edge weights x_c , the identified sparse controller in this case uses only 0.37% of the edges, and achieves a performance loss of 82.13%, i.e., $\text{card}(x)/\text{card}(x_c) = 0.37\%$ and $(J - J_c)/J_c = 82.13\%$. Here, x_c is the solution to (SP) with $\gamma = 0$ and the pattern of nonzero elements of x is obtained by solving (SP) with $\gamma = 2.5$ via the path-following iterative reweighted algorithm.

D. Path and Ring Networks

For path networks, our computational experiments show that for a large enough value of the sparsity-promoting parameter γ a single edge, which generates the longest cycle, is added; see Fig. 5, top row. This is in agreement with [15] where it was proved that the longest cycle is most beneficial for improving the \mathcal{H}_2 performance of tree networks. Similar observations are made for the spatially-invariant ring network with nearest neighbor interactions. For large values of γ , each node establishes a link to the node that is farthest away in the network; see Fig. 5, bottom row. This is in agreement with recent theoretical developments [34] where perturbation analysis was used to identify optimal weak links in edge-transitive consensus networks. Thus, for these regular networks and large enough values of the regularization parameter, our approach indeed provides the globally optimal solution to the original nonconvex cardinality minimization problem.

VII. CONCLUDING REMARKS

We have examined the problem of optimal topology design of the corresponding edge weights for undirected consensus networks. Our approach uses convex optimization to balance performance of stochastically forced networks with the number of edges in the distributed controller. For ℓ_1 -regularized minimum variance optimal control problem, we have derived a Lagrange dual and exploited structure of the optimality conditions for undirected networks to develop customized algorithms that are well suited for large problems. These are based on the proximal gradient and the proximal Newton methods. The proximal gradient algorithm is a first-order method that updates the controller graph Laplacian via the use of the soft-thresholding

operator. In the proximal Newton method, sequential quadratic approximation of the smooth part of the objective function is employed and the Newton direction is computed using cyclic coordinate descent over the set of active variables. Examples are provided to demonstrate utility of our algorithms. We have shown that proximal algorithms can solve the problems with millions of edges in the controller graph in several minutes, on a PC. Furthermore, we have specialized our algorithm to the problem of growing connected resistive networks. In this, the plant graph is connected and there are no joint edges between the plant and the controller graphs. We have exploited structure of such networks and demonstrated how additional edges can be systematically added in a computationally efficient manner.

ACKNOWLEDGMENT

The authors would like to thank J. W. Nichols for his feedback on earlier versions of this manuscript, T. H. Summers for useful discussion, and M. Sanjabi for his help with C++ implementation.

REFERENCES

- [1] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton, NJ, USA: Princeton Univ. Press, 2010.
- [2] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 53, pp. 65–78, 2004.
- [3] Y. Kim and M. Mesbahi, "On maximizing the second smallest eigenvalue of a state-dependent graph Laplacian," *IEEE Trans. Autom. Control*, vol. 51, no. 1, pp. 116–120, Jan. 2006.
- [4] L. Xiao, S. Boyd, and S.-J. Kim, "Distributed average consensus with least-mean-square deviation," *J. Parallel Distrib. Comput.*, vol. 67, no. 1, pp. 33–46, 2007.
- [5] P. Barooah and J. P. Hespanha, "Estimation on graphs from relative measurements: Distributed algorithms and fundamental limits," *IEEE Control Syst. Mag.*, vol. 27, no. 4, pp. 57–74, Aug. 2007.
- [6] A. Ghosh, S. Boyd, and A. Saberi, "Minimizing effective resistance of a graph," *SIAM Rev.*, vol. 50, no. 1, pp. 37–66, 2008.
- [7] D. Zelazo and M. Mesbahi, "Edge agreement: Graph-theoretic performance bounds and passivity analysis," *IEEE Trans. Autom. Control*, vol. 56, no. 3, pp. 544–555, Mar. 2011.
- [8] B. Bamieh, M. R. Jovanović, P. Mitra, and S. Patterson, "Coherence in large-scale networks: Dimension dependent limitations of local feedback," *IEEE Trans. Autom. Control*, vol. 57, no. 9, pp. 2235–2249, Sep. 2012.
- [9] Y. Wan, S. Roy, and A. Saberi, "Designing spatially heterogeneous strategies for control of virus spread," *IET Syst. Biol.*, vol. 2, no. 4, pp. 184–201, 2008.
- [10] V. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. Pappas, "Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks," in *Proc. 52nd IEEE Conf. Decis. Control*, 2013, pp. 7486–7491.
- [11] E. Ramírez-Llanos and S. Martínez, "A distributed nonlinear dynamics for virus spread control," *Automatica*, vol. 76, pp. 41–48, 2017.
- [12] M. Siami and N. Motee, "Tractable approximation algorithms for the np-hard problem of growing linear consensus networks," in *Proc. Amer. Control Conf.*, 2016, pp. 6429–6434.
- [13] A. Ghosh and S. Boyd, "Growing well-connected graphs," in *Proc. 45th IEEE Conf. Decis. Control*, 2006, pp. 6605–6611.
- [14] F. Lin, M. Fardad, and M. R. Jovanović, "Identification of sparse communication graphs in consensus networks," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput.*, 2012, pp. 85–89.
- [15] D. Zelazo, S. Schuler, and F. Allgöwer, "Performance and design of cycles in consensus networks," *Syst. Control Lett.*, vol. 62, no. 1, pp. 85–96, 2013.
- [16] M. Fardad, F. Lin, and M. R. Jovanović, "Design of optimal sparse interconnection graphs for synchronization of oscillator networks," *IEEE Trans. Autom. Control*, vol. 59, no. 9, pp. 2457–2462, Sep. 2014.
- [17] T. H. Summers, I. Shames, J. Lygeros, and F. Dörfler, "Topology design for optimal network coherence," in *Proc. Eur. Control Conf.*, 2015, pp. 575–580.

- [18] M. Fardad, F. Lin, and M. R. Jovanović, "Sparsity-promoting optimal control for a class of distributed systems," in *Proc. Amer. Control Conf.*, 2011, pp. 2050–2055.
- [19] F. Lin, M. Fardad, and M. R. Jovanović, "Design of optimal sparse feedback gains via the alternating direction method of multipliers," *IEEE Trans. Autom. Control*, vol. 58, no. 9, pp. 2426–2431, Sep. 2013.
- [20] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optim. Meth. Softw.*, vol. 11, no. 1–4, pp. 625–653, 1999.
- [21] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, "SDPT3 – A MATLAB software package for semidefinite programming, Version 1.3," *Optim. Meth. Softw.*, vol. 11, no. 1–4, pp. 545–581, 1999.
- [22] S. Hassan-Moghaddam and M. R. Jovanović, "An interior point method for growing connected resistive networks," in *Proc. Amer. Control Conf.*, Chicago, IL, USA, 2015, pp. 1223–1228.
- [23] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 123–231, 2013.
- [24] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [25] J. D. Lee, Y. Sun, and M. A. Saunders, "Proximal Newton-type methods for minimizing composite functions," *SIAM J. Optim.*, vol. 24, no. 3, pp. 1420–1443, 2014.
- [26] C.-J. Hsieh, M. A. Sustik, I. S. Dhillon, and P. Ravikumar, "QUIC: Quadratic approximation for sparse inverse covariance estimation," *J. Mach. Learn. Res.*, vol. 15, pp. 2911–2947, 2014.
- [27] E. J. Candès, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, pp. 877–905, 2008.
- [28] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [29] J. Barzilai and J. M. Borwein, "Two-point step size gradient methods," *IMA J. Numer. Anal.*, vol. 8, no. 1, pp. 141–148, 1988.
- [30] P. Tseng and S. Yun, "A coordinate gradient descent method for nonsmooth separable minimization," *Math. Program.*, vol. 117, no. 1/2, pp. 387–423, 2009.
- [31] Y.-H. Dai and R. Fletcher, "Projected Barzilai-Borwein methods for large-scale box-constrained quadratic programming," *Numerische Math.*, vol. 100, no. 1, pp. 21–47, 2005.
- [32] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. Signal Process.*, vol. 57, no. 7, pp. 2479–2493, Jul. 2009.
- [33] J. J. McAuley and J. Leskovec, "Learning to discover social circles in ego networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 539–547.
- [34] M. Fardad, X. Zhang, F. Lin, and M. R. Jovanović, "On the properties of optimal weak links in consensus networks," in *Proc. 53rd IEEE Conf. Decis. Control*, 2014, pp. 2124–2129.



Sepideh Hassan-Moghaddam (S'13) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2013, and the M.S. degree in electrical and computer engineering from the University of Minnesota, Minneapolis, MN, USA, in 2016. She is currently working toward the Ph.D. degree in the Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA.

Her primary research interests include optimization, inference, and control of large-scale networks.



Mihailo R. Jovanović (S'00–M'05–SM'13) received the Dipl. Ing. and M.S. degrees from the University of Belgrade, Serbia, in 1995 and 1998, respectively, and the Ph.D. degree from the University of California (UC), Santa Barbara, in 2004. He is a Professor in the Ming Hsieh Department of Electrical Engineering and the Founding Director of the Center for Systems and Control at the University of Southern California, Los Angeles, CA, USA. He was with the faculty in the Department of Electrical and Computer Engineering at the University of Minnesota, Minneapolis, MN, USA, from 2004 until 2017, and has held visiting positions with Stanford University and the Institute for Mathematics and its Applications. He currently serves as the Chair of the APS External Affairs Committee, an Associate Editor of the *SIAM Journal on Control and Optimization*, and had served as a Program Vice-Chair of the 55th IEEE Conference on Decision and Control and an Associate Editor of the IEEE Control Systems Society Conference Editorial Board from 2006 until 2010. His research focuses on the design of controller architectures, dynamics and control of fluid flows, and fundamental performance limitations in the design of large dynamic networks.

Prof. Jovanović received a CAREER Award from the National Science Foundation in 2007, the George S. Axelby Outstanding Paper Award from the IEEE Control Systems Society in 2013, and the Distinguished Alumni Award from the Department of Mechanical Engineering at UC Santa Barbara in 2014. Papers of his students were finalists for the Best Student Paper Award at the American Control Conference in 2007 and 2014.