

Input-Output Performance of Linear-Quadratic Saddle-Point Algorithms with Application to Distributed Resource Allocation Problems

John W. Simpson-Porco, Bala Kameshwar Poolla, Nima Monshizadeh, and Florian Dörfler

Abstract—Saddle-point or primal-dual methods have recently attracted renewed interest as a systematic technique to design distributed algorithms which solve convex optimization problems. When implemented online for streaming data or as dynamic feedback controllers, these algorithms become subject to disturbances, and quantifying input-output performance becomes important. We analyze the input-output performance of the continuous-time saddle-point method applied to linearly constrained quadratic programs, providing explicit expressions for the saddle-point \mathcal{H}_2 norm under a relevant input-output configuration. We then proceed to derive analogous results for regularized and augmented versions of the saddle-point algorithm. We observe some rather peculiar effects – a modest amount of regularization significantly improves the transient performance, while augmentation does not necessarily offer improvement. We then propose a distributed dual version of the algorithm which overcomes some of the performance limitations imposed by augmentation. Finally, we apply our results to a resource allocation problem to compare the input-output performance of various centralized and distributed saddle-point implementations and show that distributed algorithms may perform as well as their centralized counterparts.

I. INTRODUCTION

Saddle-point methods are a class of continuous-time gradient-based algorithms for solving constrained convex optimization problems. Introduced in the early 1950s [1], [2], these algorithms are designed to seek the saddle points of the optimization problem’s Lagrangian function. These saddle points are in one-to-one correspondence with the solutions of the first-order optimality (KKT) conditions, and the algorithm therefore drives its internal state towards the global optimizer of the convex program; see [3]–[5] for convergence results.

Recently, these algorithms have attracted renewed attention in the control literature for solving *distributed* optimization problems, where agents cooperate through a communication network to solve an optimization problem with minimal or no centralized coordination. Applications of distributed optimization include utility maximization [3], congestion management in communication networks [6], and control in power systems

[7]–[12]. While most standard optimization algorithms require centralized information to compute the optimizer, saddle-point algorithms often yield distributed strategies in which agents perform state updates using only locally measured information and communication with some subset of other agents. We refer the reader to [4], [13]–[17] for control-theoretic interpretations of these algorithms.

Rather than solve the optimization problem offline, it is desirable to run these distributed algorithms online as “controllers” in feedback with system measurements and actuation, so that the optimizer can be tracked in real-time as parameters vary or when the system is subject to external disturbances. Such algorithms also offer promise for online-optimization, especially in scenarios with streaming data. However, when saddle-point methods are implemented online as distributed controllers, they become subject to disturbances arising from fluctuating parameters and noise; the precise nature of these disturbances being application dependent. Therefore, understanding the input-output properties of these algorithms becomes key for assessing dynamic performance, quantifying how well the algorithm can reject the effect of these disturbances on a chosen controlled output. Recent work in this direction includes input-to-state-stability results [18], [19], finite \mathcal{L}_2 -gain analysis [20], and the robust control framework proposed in [21], [22]. The purpose of this paper is to continue this line of research, by providing explicit analytical performance results for a restricted class of linear time-invariant saddle-point algorithms.

Contributions: The three main contributions of this paper are as follows.¹ First, in Section III we consider the effect of disturbances on the saddle-point dynamics arising from linearly constrained, convex quadratic optimization problems. We quantify the input-output performance of the method via the \mathcal{H}_2 system norm, and — for a relevant input-output configuration — derive an explicit expression for the norm as a function of the algorithm parameters. We find (Theorem 3.1) that the squared \mathcal{H}_2 norm scales linearly with the number of disturbances to both the primal and dual variable dynamics.

Second, we study two common modifications to the Lagrangian optimization paradigm: regularization and augmentation. We show that regularization strictly improves the transient

J. W. Simpson-Porco is with the Department of Electrical and Computer Engineering, University of Waterloo, ON, Canada. Email: jwsimpson@uwaterloo.ca. N. Monshizadeh is with the Engineering and Technology Institute, University of Groningen, The Netherlands. Email: n.monshizadeh@rug.nl. B. K. Poolla and F. Dörfler are with the Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH) Zürich, Switzerland. Email: {bpoolla, dorfler}@ethz.ch. The work of J. W. Simpson-Porco is supported in part by the NSERC Discovery Grant RGPIN-2017-04008. The work of B. K. Poolla and F. Dörfler is supported by ETH start-up funds and the SNF AP Energy Grant #160573.

¹A preliminary version of these results with application to power system control appeared in the conference article [23]. In contrast to the conference article, this paper reports the proofs of Theorems 3.1 and 3.4, studies the effect of regularization (Theorem 3.2), studies both centralized and distributed dual saddle-point approaches (Proposition 4.1 and Corollary 4.2), and studies the application of these results to resource allocation problems (Section V).

\mathcal{H}_2 performance of saddle-point algorithms (Theorem 3.2). However, this improvement in performance is not usually monotone in the regularization parameter, and the system norm may achieve its global minimum at some finite regularization parameter value. For augmented Lagrangian saddle-point methods, we derive an explicit expression for the \mathcal{H}_2 norm (Theorem 3.4); the results show that augmentation may either improve or deteriorate the \mathcal{H}_2 performance. For cases when standard augmentation deteriorates performance, we propose augmented dual distributed saddle-point algorithm which strictly improves performance (Section IV-B).

Third and finally, in Section V we apply our results to resource allocation problems, comparing and contrasting the input-output performance of the different algorithms we have considered. The results show that distributed implementations can perform equally well as centralized implementations, but that significant performance differences can appear between the algorithms once augmentation is considered.

As background, in Section II we review saddle-point algorithms for the relevant class of optimization problems, then recall the basic facts about the \mathcal{H}_2 norm as a measure of input-output system performance.

Notation: The $n \times n$ identity matrix is I_n , $\mathbf{0}$ is a matrix of zeros of appropriate dimension, while $\mathbf{1}_n$ (resp. $\mathbf{0}_n$) are n -vectors of all ones (resp. zeros). If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is differentiable, then $\frac{\partial f}{\partial x} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is its gradient. For $A \in \mathbb{R}^{n \times n}$, A^\top is its transpose and $\text{Tr}(A) = \sum_{i=1}^n A_{ii}$ is its trace. If $S \in \mathbb{R}^{r \times n}$ has full row-rank, then $SS^\dagger = I_r$ where $S^\dagger = S^\top(SS^\top)^{-1}$ is the Moore-Penrose pseudoinverse of S . For a positive semidefinite matrix $Q \succeq \mathbf{0}$, $Q^{\frac{1}{2}}$ is its square root. The symbol \otimes denotes the Kronecker product.

Graphs and graph matrices: A graph is a pair $\mathcal{G} = (\mathcal{V}, \mathcal{E}_u)$, where \mathcal{V} is the set of vertices (nodes) and \mathcal{E}_u is the set of undirected edges (unordered pairs of nodes). The set of neighbours of node $i \in \mathcal{V}$ are denoted by $\mathcal{N}(i)$. If a label $e \in \{1, \dots, |\mathcal{E}_u|\}$ and an arbitrary orientation is assigned to each edge, we can define a corresponding directed edge set $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ with elements $e \sim (i, j) \in \mathcal{E}$. The *node-edge incidence matrix* $E \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{E}|}$ is defined component-wise as $E_{ke} = 1$ if node k is the source node of edge e and as $E_{ke} = -1$ if node k is the sink node of edge e , with all other elements being zero. If the graph is connected, then $\ker(E^\top) = \text{Im}(\mathbf{1}_{|\mathcal{V}|})$. A graph is a *tree* (or *acyclic*) if it contains no cycles, and in this case $\ker(E) = \{\mathbf{0}_{|\mathcal{E}|}\}$.

II. SADDLE-POINT METHODS AND \mathcal{H}_2 PERFORMANCE

A. Review of Saddle-Point Method

We consider the constrained quadratic optimization problem

$$\begin{aligned} \underset{x \in \mathbb{R}^{n_x}}{\text{minimize}} \quad & J(x) := \frac{1}{2}x^\top Qx + x^\top c \\ \text{subject to} \quad & Sx = W_b b, \end{aligned} \quad (1)$$

where $x \in \mathbb{R}^{n_x}$, $Q = Q^\top \succ \mathbf{0}$ is positive definite, $c \in \mathbb{R}^{n_x}$ and $b \in \mathbb{R}^{n_b}$ are parameter vectors, and $S \in \mathbb{R}^{n_r \times n_x}$, $W_b \in \mathbb{R}^{n_r \times n_b}$ with $n_r < n_x$. We make the blanket assumption that S and W_b have full row rank, which simply means that the

constraints $Sx = W_b b$ are not redundant. While the right-hand side $W_b b$ of the constraints is apparently over-parameterized, this formulation is natural when considering particular problem instances. In the resource allocation problem of Section V, $W_b = [1 \ 1 \ \dots \ 1]$ and b is a vector of demands; the product $W_b b$ is simply the total demand.

The problem (1) describes only a subclass of the optimization problems to which saddle-point algorithms are applicable; more generally one considers strictly convex costs and convex inequality constraints as well. We restrict our attention to (1), as this case will allow LTI system analysis techniques to be applied, and represents a large enough class of problems to yield some general insights. Intuitively, the performance of the saddle-point algorithm on (1) should indicate a “best” case performance for the general case, as the objective $J(x)$ is smooth and strongly convex, and (1) is free of hard inequality constraints. See Section VI for further discussion.

Under these assumptions, the convex problem (1) has a finite optimum, the equality constraints are strictly feasible, and (1) may be equivalently studied through its Lagrange dual with zero duality gap [24]. The Lagrangian $L : \mathbb{R}^{n_x} \times \mathbb{R}^{n_r} \rightarrow \mathbb{R}$ of the problem (1) is

$$L(x, \nu) = \frac{1}{2}x^\top Qx + c^\top x + \nu^\top (Sx - W_b b), \quad (2)$$

where $\nu \in \mathbb{R}^{n_r}$ is a vector of Lagrange multipliers. By strong duality, the KKT conditions

$$\begin{aligned} \frac{\partial L}{\partial x}(x, \nu) = \mathbf{0}_{n_x} & \iff \mathbf{0}_{n_x} = Qx + S^\top \nu + c, \\ \frac{\partial L}{\partial \nu}(x, \nu) = \mathbf{0}_{n_r} & \iff \mathbf{0}_{n_r} = Sx - W_b b, \end{aligned} \quad (3)$$

are necessary and sufficient for optimality. From these linear equations one can quickly compute the unique global optimizer (x^*, ν^*) to be

$$\begin{bmatrix} x^* \\ \nu^* \end{bmatrix} = \begin{bmatrix} -Q^{-1}(S^\top \nu^* + c) \\ -(SQ^{-1}S^\top)^{-1}(W_b b + SQ^{-1}c) \end{bmatrix}. \quad (4)$$

While (4) is the exact solution to the optimization problem (1), its evaluation requires centralized knowledge of the matrices S, Q, W_b and the vectors c and b . If any of these parameters change or evolve over time, the optimizer should be recomputed. In many multi-agent system applications, the cost matrix Q is diagonal or block-diagonal and $J(x) = \sum_i \frac{q_i}{2} x_i^2 + c_i x_i$ is therefore a sum of local costs. Finally, the constraints encoded in S are often sparse, mirroring the topology of an interaction or communication network between agents. These factors motivate the solution of (1) in an online distributed fashion, where agents in the network communicate and cooperate to calculate the global optimizer.

A simple continuous-time algorithm to seek the optimizer is the *saddle-point* or *primal-dual* method [3]–[5], [25], [26]

$$\tau_x \dot{x} = -\frac{\partial}{\partial x} L(x, \nu), \quad \tau_\nu \dot{\nu} = +\frac{\partial}{\partial \nu} L(x, \nu),$$

which here reduces to the affine dynamical system

$$\tau_x \dot{x} = -Qx - S^\top \nu - c \quad (5a)$$

$$\tau_\nu \dot{\nu} = Sx - W_b b, \quad (5b)$$

where $\tau_x, \tau_\nu \succ \mathbb{0}$ are positive definite diagonal matrices of time-constants. By construction, the equilibrium points of (5) are in one-to-one correspondence with the solutions of the KKT conditions (3), and the system is internally exponentially stable [5].

Lemma 2.1 (Global Convergence to Optimizer): The unique equilibrium point (x^*, ν^*) given in (4) of the saddle-point dynamics (5) is globally exponentially stable, with exponential convergence rate $\lambda_{\min}(Q)/\tau_{\max}$ where $\tau_{\max} = \max(\max_{i \in \{1, \dots, n\}} \tau_{x,ii}, \max_{i \in \{1, \dots, r\}} \tau_{\nu,ii})$.

The proof of Lemma 2.1 follows by using the Lyapunov candidate $V(x, \nu) = (x - x^*)^\top (x - x^*) + (\nu - \nu^*)^\top \tau_\nu (\nu - \nu^*)$ and an invariance argument; the convergence rate follows from a simple variant of [25, Theorem 3.6]. With stability settled, in what follows we will focus exclusively on quantifying transient input-output performance of (5) in the presence of exogenous disturbances.

B. System Performance in the \mathcal{H}_2 Norm

Consider the linear time-invariant system

$$\begin{aligned} \dot{x} &= Ax + B\eta \\ z &= Cx, \end{aligned} \quad (6)$$

where η is the disturbance input signal and z is the performance output. With $x(0) = \mathbb{0}$, we denote the linear operator from η to z by G . If (6) is input-output stable, its \mathcal{H}_2 norm $\|G\|_{\mathcal{H}_2}$ is defined as

$$\|G\|_{\mathcal{H}_2}^2 := \frac{1}{2\pi} \int_{-\infty}^{\infty} \text{Tr}(G(-j\omega)^\top G(j\omega)) d\omega,$$

where $G(j\omega) = C(j\omega I - A)^{-1}B$ is the frequency response of (6). Another interpretation of $\|G\|_{\mathcal{H}_2}^2$ is as the steady-state variance of the output

$$\|G\|_{\mathcal{H}_2}^2 = \lim_{t \rightarrow \infty} \mathbb{E}[z^\top(t)z(t)], \quad (7)$$

when each component of $\eta(t)$ is stochastic white noise with unit covariance (i.e., $\mathbb{E}[\eta(t)\eta(t')^\top] = \delta(t - t')I$). Therefore, $\|G\|_{\mathcal{H}_2}$ measures how much the output varies in steady-state under stochastic disturbances. If the state matrix A is Hurwitz, then the \mathcal{H}_2 norm is finite, and can be computed as [27, Ch.4]

$$\|G\|_{\mathcal{H}_2}^2 = \text{Tr}(B^\top XB), \quad (8)$$

where the observability Gramian $X = X^\top \succeq \mathbb{0}$ is the unique solution to the Lyapunov equation

$$A^\top X + XA + C^\top C = \mathbb{0}. \quad (9)$$

If the pair (C, A) is observable, then X is positive-definite.

III. \mathcal{H}_2 PERFORMANCE OF SADDLE-POINT METHODS

We now subject the saddle-point dynamics (5) to disturbances in both the primal and dual equations. Specifically, we assume the vectors b and c are subject to disturbances $\eta_b \in \mathbb{R}^{n_b}$ and $\eta_c \in \mathbb{R}^{n_c}$, and make the substitutions $b \mapsto b + t_b \eta_b$ and $c \mapsto c + t_c \eta_c$ in the saddle point dynamics (5). The scalar parameters $t_b, t_c \geq 0$ parameterize the strength of the

disturbances, and will help us keep track of which terms in the resulting norm expressions arise from which disturbances. As an example, when we study distributed resource allocation problems in Section V, b will have the interpretation of a vector of demands for some resource, and η_b will therefore model a fluctuation or disturbance to this nominal demand.

After translating the nominal equilibrium point (4) of the system to the origin², we obtain the LTI system

$$\begin{bmatrix} \tau_x \dot{x} \\ \tau_\nu \dot{\nu} \end{bmatrix} = \begin{bmatrix} -Q & -S^\top \\ S & \mathbb{0} \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix} - \begin{bmatrix} t_c I_{n_x} & \mathbb{0} \\ \mathbb{0} & t_b W_b \end{bmatrix} \begin{bmatrix} \eta_c \\ \eta_b \end{bmatrix}, \quad (10a)$$

$$z = \begin{bmatrix} C_1 & \mathbb{0} \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix}, \quad (10b)$$

where $C_1 \in \mathbb{R}^{n_x \times n_x}$ is an output matrix. As the system (10) is written in error coordinates, convergence to the saddle-point optimizer (x^*, ν^*) from (4) is equivalent to convergence of $(x(t), \nu(t))$ to the origin. A natural way to measure this convergence is to use the cost matrix Q from the optimization problem (1) as a weighting matrix, and to study the performance output $\|z(t)\|_2^2 = x(t)^\top Q x(t)$, which is obtained by choosing $C_1 = Q^{\frac{1}{2}}$. This performance output prioritizes the minimization of the objective function over the violation of the equality constraints during transients.

Theorem 3.1 (Saddle-Point Performance): Consider the input-output saddle-point dynamics (10) with $Q = Q^\top \succ \mathbb{0}$ diagonal, and let $C_1 = Q^{\frac{1}{2}}$ so that $\|z(t)\|_2^2 = x^\top(t)Qx(t)$. Then the squared \mathcal{H}_2 norm of the saddle-point system (10) is

$$\|G\|_{\mathcal{H}_2}^2 = \frac{t_c^2}{2} \text{Tr}(\tau_x^{-1}) + \frac{t_b^2}{2} \text{Tr}(W_b^\top \tau_\nu^{-1} W_b). \quad (11)$$

Proof of Theorem 3.1: We will directly construct the unique positive-definite observability Gramian; since the system is internally stable (Lemma 2.1), this also indirectly establishes observability [28, Exercise 4.8.1]. Assuming for the moment a block-diagonal observability Gramian $X = \text{blkdiag}(X_1, X_2)$, the Lyapunov equation (9) yields the two equations

$$X_1 \tau_x^{-1} Q + Q \tau_x^{-1} X_1 - C_1^2 = \mathbb{0}, \quad (12a)$$

$$X_2 \tau_\nu^{-1} S - S \tau_x^{-1} X_1 = \mathbb{0}, \quad (12b)$$

with the third independent equation trivially being $\mathbb{0} = \mathbb{0}$. By inspection, the solution to (12a) is diagonal and given by

$$X_1 = \frac{1}{2} \tau_x Q^{-1} C_1^2 = \frac{1}{2} \tau_x,$$

since Q is diagonal and $C_1 = Q^{\frac{1}{2}}$. Clearly X_1 is positive definite and symmetric. Since S has full row-rank, and X_2 can be uniquely recovered from (12b) as

$$X_2 = S \tau_x^{-1} X_1 S^\top \tau_\nu = \frac{1}{2} S S^\top \tau_\nu = \frac{1}{2} \tau_\nu.$$

It follows that X_2 is positive definite, and therefore $X = \frac{1}{2} \text{blkdiag}(\tau_x, \tau_\nu)$ is the unique positive definite solution to (9). Since X is block diagonal, we find from (8) and (10) that

$$\|G\|_{\mathcal{H}_2}^2 = t_c^2 \text{Tr}(\tau_x^{-1} X_1 \tau_x^{-1}) + t_b^2 \text{Tr}(W_b^\top \tau_\nu^{-1} X_2 \tau_\nu^{-1} W_b),$$

²In the remainder of the section we assume that we have made the change of state variables $\Delta x = x - x^*$, $\Delta \nu = \nu - \nu^*$, and with an abuse of notation we drop the Δ 's and simply refer to the error coordinates as x and ν .

from which the result follows. \square

We make two key observations about the result (11). First, (11) is independent of both the cost matrix Q and the constraint matrix S ; neither matrix has any influence on the value of the system norm. Second, the expression in (11) scales inversely with the time constants τ_x and τ_ν , which indicates an inherent trade-off between convergence speed and input-output performance. As a special case of Theorem 3.1, suppose that τ_x, τ_ν are multiples of the identity matrix, i.e., $\tau_x = \bar{\tau}_x I_{n_x}$, $\tau_\nu = \bar{\tau}_\nu I_{n_r}$, and that $W_b = I_{n_r}$, meaning there is one disturbance for each constraint. Then (11) reduces to

$$\|G\|_{\mathcal{H}_2}^2 = \frac{t_c^2}{2\bar{\tau}_x} n_x + \frac{t_b^2}{2\bar{\tau}_\nu} n_r, \quad (13)$$

meaning the squared \mathcal{H}_2 norm scales linearly in the number of disturbances to the primal dynamics and the number of disturbances to the dual dynamics. While this scaling is quite reasonable, the lack of tuneable controller gains other than the time constants means that convergence speed and input-output performance are always conflicting objectives.

A. Performance of Regularized Saddle-Point Methods

A common variation of the Lagrangian optimization framework includes a quadratic penalty [29]–[31] on the dual variable ν in the Lagrangian (2). The so-called *regularized* Lagrangian assumes the form

$$L_{\text{reg}}(x, \nu) = \frac{1}{2} x^\top Q x + c^\top x + \nu^\top (Sx - W_b b) - \frac{\epsilon}{2} \|\nu\|_2^2, \quad (14)$$

where $\epsilon > 0$ is small. The regularization term adds concavity to the Lagrangian, and has been shown to increase the convergence rate of optimization algorithms. However, this regularization alters the equilibrium of the closed-loop system, which moves from the value in (4) to the new value

$$\begin{bmatrix} x_{\text{reg}}^* \\ \nu_{\text{reg}}^* \end{bmatrix} = \begin{bmatrix} -Q^{-1}(S^\top \nu_{\text{reg}}^* + c) \\ -(SQ^{-1}S^\top + \epsilon I_{n_r})^{-1}(W_b b + SQ^{-1}c) \end{bmatrix}. \quad (15)$$

The penalty coefficient ϵ is chosen to strike a balance between the convergence rate improvement and the deviation of $(x_{\text{reg}}^*, \nu_{\text{reg}}^*)$ from (x^*, ν^*) . The Lagrangian (15) also admits a continuous-time saddle-point algorithm with *regularized saddle-point dynamics*

$$\begin{aligned} \tau_x \dot{x} &= -Qx - S^\top \nu - c \\ \tau_\nu \dot{\nu} &= Sx - W_b b - \epsilon \nu. \end{aligned} \quad (16)$$

Quite strikingly, we shall observe that a small regularization which has a minor effect on the equilibrium, achieves a tremendous improvement in performance.

As we did with the standard saddle-point dynamics, we can shift the equilibrium point $(x_{\text{reg}}^*, \nu_{\text{reg}}^*)$ of (16) to the origin and introduce disturbances to the parameters b and c , leading to the input-output model

$$\begin{aligned} \begin{bmatrix} \tau_x \dot{x} \\ \tau_\nu \dot{\nu} \end{bmatrix} &= \begin{bmatrix} -Q & -S^\top \\ S & -\epsilon I_{n_r} \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix} - \begin{bmatrix} t_c I_{n_x} & 0 \\ 0 & t_b W_b \end{bmatrix} \begin{bmatrix} \eta_c \\ \eta_b \end{bmatrix} \\ z &= \begin{bmatrix} Q^{\frac{1}{2}} & 0 \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix}, \end{aligned} \quad (17)$$

where we consider time constant matrices τ_x, τ_ν , and disturbances η_b, η_c as in (10).

Theorem 3.2 (Regularized Saddle-Point Performance): Consider the input-output regularized saddle-point dynamics (17) denoted by G_{reg} with $Q = Q^\top \succ 0$ a diagonal matrix. Then the squared \mathcal{H}_2 norm $\|G_{\text{reg}}\|_{\mathcal{H}_2}^2$ of the system (17) is upper-bounded by (11) with strict inequality.

Proof of Theorem 3.2: We rewrite (17) in the standard state-space form $G_{\text{reg}} := (A_{\text{reg}}, B, C, 0)$, where

$$\begin{aligned} \begin{bmatrix} \dot{x} \\ \dot{\nu} \end{bmatrix} &= \underbrace{\begin{bmatrix} -\tau_x^{-1}Q & -\tau_x^{-1}S^\top \\ \tau_\nu^{-1}S & -\tau_\nu^{-1}\epsilon I_{n_r} \end{bmatrix}}_{A_{\text{reg}}} \begin{bmatrix} x \\ \nu \end{bmatrix} \\ &\quad - \underbrace{\begin{bmatrix} \tau_x^{-1}t_c & 0 \\ 0 & \tau_\nu^{-1}W_b t_b \end{bmatrix}}_B \begin{bmatrix} \eta_c \\ \eta_b \end{bmatrix}, \quad z = \underbrace{\begin{bmatrix} Q^{\frac{1}{2}} & 0 \end{bmatrix}}_C \begin{bmatrix} x \\ \nu \end{bmatrix}. \end{aligned}$$

One may verify that A_{reg} is Hurwitz and that G_{reg} is observable. Consider the observability Gramian from (12), i.e., $X = \frac{1}{2} \text{blkdiag}(\tau_x, \tau_\nu)$. An easy computation shows that

$$X A_{\text{reg}} + A_{\text{reg}}^\top X + C^\top C = \begin{bmatrix} 0 & 0 \\ 0 & -\epsilon I_{n_r} \end{bmatrix} \preceq 0. \quad (19)$$

We conclude that X is a *generalized* observability Gramian for the regularized system G_{reg} . If X_ϵ is the true observability Gramian for G_{reg} , then $X_\epsilon \neq X$ and $X_\epsilon \preceq X$ [28, Chapter 4.7], and we conclude that

$$\|G_{\text{reg}}\|_{\mathcal{H}_2}^2 = \text{Tr}(B^\top X_\epsilon B) \leq \text{Tr}(B^\top X B) = \|G\|_{\mathcal{H}_2}^2. \quad (20)$$

It remains only to show that the above inequality holds strictly. Proceeding by contradiction, assume that $\text{Tr}(B^\top X_\epsilon B) = \text{Tr}(B^\top X B)$, which implies that $\text{Tr}(B^\top (X - X_\epsilon) B) = 0$. Since $X - X_\epsilon \succeq 0$, we may write $X - X_\epsilon = F^\top F$ for some matrix F , and

$$0 = \text{Tr}(B^\top (X - X_\epsilon) B) = \text{Tr}(B^\top F^\top F B).$$

Since B has full row rank, this implies that F must be zero, and thus $X = X_\epsilon$ which is a contradiction. \square

We conclude that regularization strictly improves the \mathcal{H}_2 performance of the saddle-point algorithm, but at the expense of shifting the equilibrium point away from the original optimum. In general it seems difficult to explicitly quantify the improvement, and as will be observed in the proof of Proposition 3.3, this amounts to computing the \mathcal{H}_2 -norm of (17) with the modified output matrix $\begin{bmatrix} 0 & \epsilon I_{n_r} \end{bmatrix}$. We can obtain a closed-form solution, however, for the case of a single constraint with uniform parameters.

Proposition 3.3 (Regularized Saddle-Point Performance with One Constraint): Consider the case with one constraint ($n_r = 1$) and one constraint disturbance ($n_b = 1$) with uniform problem parameters $Q = qI_{n_x}$, $\tau_x = \bar{\tau}_x I_{n_x}$, $\tau_\nu = \bar{\tau}_\nu I_{n_r}$ for scalars $q, \bar{\tau}_x, \bar{\tau}_\nu > 0$ and $W_b = 1$. Then, we have

$$\|G\|_{\mathcal{H}_2}^2 - \|G_{\text{reg}}\|_{\mathcal{H}_2}^2 = \alpha_\epsilon t_c^2 + \gamma_\epsilon t_b^2, \quad (21)$$

where $s = \|S\|_2$ and

$$\alpha_\epsilon = \frac{\epsilon s^2}{2(\epsilon q + s^2)(\epsilon \bar{\tau}_x + q \bar{\tau}_\nu)},$$

$$\gamma_\epsilon = \frac{\epsilon(\bar{\tau}_x q \epsilon + q^2 \bar{\tau}_\nu + \bar{\tau}_x s^2)}{2\bar{\tau}_\nu(\epsilon q + s^2)(\epsilon \bar{\tau}_x + q \bar{\tau}_\nu)}.$$

Proof of Proposition 3.3: Let $\Delta := X - X_\epsilon$, where $X = \frac{1}{2}\text{blkdiag}(\tau_x, \tau_\nu)$ is the observability Gramian of (10) with $C_1 = Q^{\frac{1}{2}}$, and X_ϵ is the observability Gramian of G_{reg} . As noted in the proof of Theorem 3.2, $X \succeq X_\epsilon$, and thus the matrix Δ is positive semidefinite. Clearly, Δ satisfies

$$A_{\text{reg}}^\top(X - \Delta) + (X - \Delta)A_{\text{reg}} + C^\top C = 0.$$

By (19), this reduces to

$$A_{\text{reg}}^\top \Delta + A_{\text{reg}} \Delta + \begin{bmatrix} 0 & 0 \\ 0 & \epsilon I_{n_r} \end{bmatrix} = 0. \quad (22)$$

Then it is easy to see that

$$\|G_{\text{reg}}\|_{\mathcal{H}_2}^2 = \|G\|_{\mathcal{H}_2}^2 - \|G_\epsilon\|_{\mathcal{H}_2}^2, \quad (23)$$

where $\|G\|_{\mathcal{H}_2}^2$ is as in (11), and $G_\epsilon(s) := C_\epsilon(sI - A_{\text{reg}})^{-1}B$ with $C_\epsilon = \begin{bmatrix} 0 & \epsilon I_{n_r} \end{bmatrix}$. Therefore, the improvement in the \mathcal{H}_2 -norm performance is equal to the squared \mathcal{H}_2 -norm of the axillary system given by G_ϵ . Next, we calculate the \mathcal{H}_2 -norm of G_ϵ , which requires computing the observability Gramian from (22). Consider the matrix

$$\bar{\Delta} = \begin{bmatrix} \alpha S^\top S & \beta S^\top \\ \beta S & \gamma \end{bmatrix} \quad (24)$$

where α, β, γ are constant and positive. Straightforward calculation shows that by choosing

$$\begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} = \begin{bmatrix} \frac{s^2}{\bar{\tau}_x} & \frac{\epsilon}{\bar{\tau}_\nu} + \frac{q}{\bar{\tau}_x} & -\frac{1}{\bar{\tau}_\nu} \\ \frac{q}{\bar{\tau}_x} & -\frac{1}{\bar{\tau}_\nu} & 0 \\ 0 & \frac{\bar{\tau}_\nu s^2}{\bar{\tau}_x \epsilon} & 1 \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ 0 \\ \frac{\bar{\tau}_\nu}{2} \end{bmatrix}, \quad (25)$$

the matrix $\bar{\Delta}$ is a solution to the Lyapunov equation (22). Given the fact that A_{reg} is Hurwitz, this solution is unique and the matrix $\Delta = \bar{\Delta}$ is the observability Gramian of the system given by G_ϵ . The proof is completed by calculating the inverse in (25), using (8), and noting $\text{Tr}(S^\top S) = s^2$. \square

Proposition 3.3 quantifies the performance improvement resulting from the regularization in the special case of a single constraint and uniform parameters. For sufficiently large ϵ , this improvement is approximated by $\frac{t_b^2}{2\bar{\tau}_\nu}$, which coincides with the second term in the right-hand side of (13). This means that, as expected, the constraints do not contribute to the \mathcal{H}_2 -norm in case the penalty term in the regularized Lagrangian (14) tends to infinity. On the other hand, for $\epsilon \rightarrow 0^+$, the \mathcal{H}_2 norm of the regularized dynamics (17) clearly converges to the \mathcal{H}_2 norm of (10). In general, the improvement in the input-output performance obtained due to regularization is not a monotonic function of the regularization parameter ϵ . To illustrate this, Figure 1 plots the right-hand side of (21) as a function of ϵ . It is noteworthy that for both the plots, even a modest ϵ improves the performance. Depending on the specific values

of the parameters, the maximum performance gain may be achieved at a finite value of ϵ (Figure 1(a)) or as $\epsilon \rightarrow \infty$ (Figure 1(b)).

B. Performance of Augmented Saddle-Point Methods

Another option for improving the \mathcal{H}_2 performance of saddle-point methods is to return to the Lagrangian function (2) and instead consider the *augmented Lagrangian* [32]

$$L_{\text{aug}}(x, \nu) \triangleq L(x, \nu) + \frac{\rho}{2} \|Sx - W_b b\|_2^2, \quad (26)$$

where we have incorporated the squared constraint $Sx - W_b b = 0_{n_r}$ into the Lagrangian with a gain $\rho \geq 0$. One way to interpret this is that the term $\frac{\rho}{2} \|Sx - W_b b\|_2^2$ adds additional convexity to the Lagrangian in the x variable. It follows that (x, ν) is a saddle point of $L_{\text{aug}}(x, \nu)$ if and only if it is a saddle point of $L(x, \nu)$, and hence the optimizer is unchanged. Applying the saddle-point method to the augmented Lagrangian $L_{\text{aug}}(x, \nu)$, we obtain the *augmented saddle-point dynamics*

$$\begin{aligned} \tau_x \dot{x} &= -(Q + \rho S^\top S)x - S^\top \nu - c + \rho S^\top W_b b \\ \tau_\nu \dot{\nu} &= Sx - W_b b. \end{aligned} \quad (27)$$

One may verify that as before, the unique stable equilibrium point of (27) is given by (4). We again consider disturbances η_b and η_c , and make the substitution $b \mapsto b + t_b \eta_b$ and $c \mapsto c + t_c \eta_c$. After translating the equilibrium point to the origin, we obtain the MIMO system

$$\begin{aligned} \begin{bmatrix} \tau_x \dot{x} \\ \tau_\nu \dot{\nu} \end{bmatrix} &= \begin{bmatrix} -(Q + \rho S^\top S) & -S^\top \\ S & 0 \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix} \\ &\quad - \begin{bmatrix} t_c I_{n_x} & -\rho t_b S^\top W_b \\ 0 & t_b W_b \end{bmatrix} \begin{bmatrix} \eta_c \\ \eta_b \end{bmatrix} \\ z &= \begin{bmatrix} Q^{\frac{1}{2}} & 0 \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix}. \end{aligned} \quad (28)$$

The additional term $-\rho S^\top S$ in the dynamics (28) complicates the solution of the Lyapunov equation, and we require additional assumptions to obtain an explicit formula. We consider the parametrically uniform case where $Q = qI_{n_x}$, $\tau_x = \bar{\tau}_x I_{n_x}$, and $\tau_\nu = \bar{\tau}_\nu I_{n_r}$ for scalars $q, \bar{\tau}_x, \bar{\tau}_\nu > 0$, and assume one disturbance per constraint ($W_b = I_{n_r}$). The next result appeared in [23] without proof.

Theorem 3.4 (Augmented Saddle-Point Performance): Consider the input-output augmented saddle-point dynamics (28), denoted by G_{aug} under the above assumptions, with performance output $\|z(t)\|_2^2 = x^\top(t)Qx(t) = q\|x(t)\|_2^2$. Then the squared \mathcal{H}_2 norm of the augmented saddle-point system (28) is

$$\begin{aligned} \|G_{\text{aug}}\|_{\mathcal{H}_2}^2 &= \frac{t_c^2}{2\bar{\tau}_x} (n_x - n_r) + \left(\frac{t_b^2}{2\bar{\tau}_\nu} + \frac{t_c^2}{2\bar{\tau}_x} \right) \sum_{i=1}^{n_r} \frac{q}{q + \rho \sigma_i^2} \\ &\quad + \frac{t_b^2}{2\bar{\tau}_x} \sum_{i=1}^{n_r} \frac{q \rho^2 \sigma_i^2}{q + \rho \sigma_i^2}, \end{aligned} \quad (29)$$

where σ_i is the i th non-zero singular value of S .




Fig. 1: System norm of regularized dynamics as a function of ϵ , for parameters $\bar{\tau}_x = \bar{\tau}_\nu = t_c = t_b = 1$, $Q = qI_5$, and $S = [0.82 \ 0.90 \ 0.13 \ 0.91 \ 0.63]$.

Proof of Theorem 3.4: Under the given assumptions, the system (28) simplifies to

$$\begin{bmatrix} \dot{x} \\ \dot{\nu} \end{bmatrix} = \begin{bmatrix} -\frac{1}{\bar{\tau}_x}(qI_{n_x} + \rho S^\top S) & -\frac{1}{\bar{\tau}_x}S^\top \\ \frac{1}{\bar{\tau}_\nu}S & 0 \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix} - \begin{bmatrix} \frac{t_c}{\bar{\tau}_x}I_{n_x} & -\frac{t_b\rho}{\bar{\tau}_x}S^\top \\ 0 & \frac{t_b}{\bar{\tau}_\nu}I_{n_r} \end{bmatrix} \begin{bmatrix} \eta_c \\ \eta_b \end{bmatrix}, \quad z = \begin{bmatrix} q^{\frac{1}{2}}I_{n_x} & 0 \end{bmatrix} \begin{bmatrix} x \\ \nu \end{bmatrix}.$$

Let $S = U\Sigma V^\top$ be the singular value decomposition of S , where $U \in \mathbb{R}^{n_r \times n_r}$ and $V \in \mathbb{R}^{n_x \times n_x}$ are both orthogonal matrices, $\Sigma = \begin{bmatrix} \bar{\Sigma} & 0_{n_r \times (n_x - n_r)} \end{bmatrix}$ and $\bar{\Sigma} \in \mathbb{R}^{n_r \times n_r}$ is the diagonal matrix of non-zero singular values. Consider now the invertible change of variables $\tilde{x} = V^\top x$, $\tilde{\nu} = U^\top \nu$. In these new coordinates, the dynamics become

$$\begin{bmatrix} \dot{\tilde{x}} \\ \dot{\tilde{\nu}} \end{bmatrix} = \begin{bmatrix} -\frac{1}{\bar{\tau}_x}(qI_{n_x} + \rho\Sigma^\top\Sigma) & -\frac{1}{\bar{\tau}_x}\Sigma^\top \\ \frac{1}{\bar{\tau}_\nu}\Sigma & 0 \end{bmatrix} \begin{bmatrix} \tilde{x} \\ \tilde{\nu} \end{bmatrix} - \begin{bmatrix} \frac{t_c}{\bar{\tau}_x}V^\top & -\frac{t_b\rho}{\bar{\tau}_x}V^\top S^\top \\ 0 & \frac{t_b}{\bar{\tau}_\nu}U^\top \end{bmatrix} \begin{bmatrix} \eta_c \\ \eta_b \end{bmatrix}, \quad y = \begin{bmatrix} q^{\frac{1}{2}}V & 0 \end{bmatrix} \begin{bmatrix} \tilde{x} \\ \tilde{\nu} \end{bmatrix}.$$

We now show the observability of the pair (C, A) ; note that this is equivalent to observability of $(C^\top C, A)$. First note that $\ker(C^\top C)$ is spanned by $\begin{bmatrix} 0_{n_x}^\top & \nu^\top \end{bmatrix}^\top$. Now, suppose that $\begin{bmatrix} 0_{n_x}^\top & \nu^\top \end{bmatrix}^\top$ is an eigenvector of A with eigenvalue λ :

$$\begin{bmatrix} -\frac{1}{\bar{\tau}_x}(qI_{n_x} + \rho\Sigma^\top\Sigma) & -\frac{1}{\bar{\tau}_x}\Sigma^\top \\ \frac{1}{\bar{\tau}_\nu}\Sigma & 0 \end{bmatrix} \begin{bmatrix} 0_{n_x} \\ \nu \end{bmatrix} = \lambda \begin{bmatrix} 0_{n_x} \\ \nu \end{bmatrix}.$$

Since by stability $\text{Re}(\lambda) < 0$, the above relation only holds for $\nu = 0_{n_r}$, which shows observability by the eigenvector test. Assuming a block-diagonal observability Gramian $X = \text{blkdiag}(X_1, X_2)$, the Lyapunov equation (9) yields the two

independent equations

$$X_1 \frac{1}{\bar{\tau}_x}(qI_{n_x} + \rho\Sigma^\top\Sigma) + \frac{1}{\bar{\tau}_x}(qI_{n_x} + \rho\Sigma^\top\Sigma)X_1 = qI_{n_x}, \quad (32a)$$

$$X_2 \frac{1}{\bar{\tau}_\nu}\Sigma - \Sigma \frac{1}{\bar{\tau}_x}X_1 = 0, \quad (32b)$$

where we have used the fact that $V^\top V = I_{n_x}$. Noting that $\Sigma^\top\Sigma = \text{blkdiag}(\bar{\Sigma}^\top, 0_{(n_x - n_r) \times (n_x - n_r)})$, we find by inspection that the solution to (32a) is diagonal and given by

$$X_{1ii} = \frac{1}{2} \frac{q}{q + \rho\sigma_i^2} \bar{\tau}_x, \quad i \in \{1, \dots, n_r\},$$

$$X_{1ii} = \frac{1}{2} \bar{\tau}_x, \quad i \in \{n_r + 1, \dots, n_x\}.$$

Observe that X_1 is positive definite and symmetric. The matrix equation (32b) admits a solution X_2 if and only if $\ker(\Sigma) \subseteq \ker(\Sigma X_1)$, which holds in this case since X_1 is diagonal. The lower block X_2 can therefore be uniquely recovered from (32b) as

$$X_2 = \bar{\tau}_x^{-1}(\Sigma X_1 \Sigma^\top)(\Sigma \Sigma^\top)^{-1} \bar{\tau}_\nu.$$

A straightforward calculation shows that this is equivalent to the component formula

$$X_{2ii} = \frac{1}{2} \frac{q}{q + \rho\sigma_i^2} \bar{\tau}_\nu, \quad i \in \{1, \dots, n_r\}.$$

It follows that X_2 is diagonal and positive definite, and therefore $X = \text{blkdiag}(X_1, X_2)$ is the unique positive definite solution. A calculation using (8) now shows that

$$\begin{aligned} \|G_{\text{aug}}\|_{\mathcal{H}_2}^2 &= \frac{t_c^2}{\bar{\tau}_x^2} \text{Tr}(V X_1 V^\top) + \frac{t_b^2}{\bar{\tau}_\nu^2} \text{Tr}(U X_2 U^\top) \\ &\quad + \frac{t_b^2 \rho^2}{\bar{\tau}_x^2} \text{Tr}(S V X_1 V^\top S^\top) \end{aligned}$$

from which the result (29) follows by applying the cyclic property of the trace and simplifying the result. \square

Under the assumed restrictions on parameters, Theorem 3.4 generalizes Theorem 3.1, since when $\rho = 0$ the expression (29) reduces to (13). Consider now the dependence of the expression (29) on the augmentation gain ρ . First, in the case when $t_b = 0$ (meaning the vector b is not subject to disturbances), then as $\rho \rightarrow \infty$ the expression (29) reduces to only the first term: in this case, augmentation unambiguously improves input-output performance. In particular, note that a more favourable scaling than in (11) is achieved when n_r is comparable to n_x ; see the resource allocation problem in Section V.³ On the other hand, if $t_b \neq 0$, then as ρ becomes large, the second term in the expression vanishes, while the third term grows without bound. Therefore, a large augmentation gain will lead to poor input-output performance. This behaviour is explained by examining (27): the vector b enters the primal dynamics multiplied by ρ , and hence any noise in b is amplified as ρ grows. To remedy this deficiency in the augmented approach, the next section exploits a dual formulation of the optimization problem (1).

IV. DUAL AND DISTRIBUTED DUAL METHODS FOR IMPROVED SADDLE-POINT ALGORITHM PERFORMANCE

This section develops an approach to overcome the performance issues of augmented Lagrangian methods observed in Theorem 3.4 when disturbances enter the constraints. To focus in on these problematic disturbances, in this section we ignore possible disturbances to the vector c and set $t_c = 0$. Section IV-A contains a quick examination of dual ascent, before proceeding to a distributed dual formulation in Section IV-B.

A. Centralized Dual Ascent

To begin, we return to the Lagrangian (2) of the optimization problem (1), and compute

$$x^*(\nu) = \operatorname{argmin}_{x \in \mathbb{R}^n} L(x, \nu) = -Q^{-1}(c + S^T \nu).$$

It follows quickly that the dual function $\Phi(\nu)$ is given by

$$\begin{aligned} \Phi(\nu) &= \min_{x \in \mathbb{R}^n} L(x, \nu) \\ &= -\frac{1}{2} \nu^T S Q^{-1} S^T \nu - \nu^T (S Q^{-1} c + W_b b) - \frac{1}{2} c^T Q^{-1} c. \end{aligned} \quad (33)$$

With the primal variable eliminated, a possible approach is to simply maximize the dual function $\Phi(\nu)$ via gradient ascent. Introducing disturbance inputs $b \mapsto b + t_b \eta_b$ and performance outputs similar to before, and shifting the unique equilibrium point to the origin, one quickly obtains the input-output dual ascent dynamics

$$\begin{aligned} \tau_\nu \dot{\nu} &= -S Q^{-1} S^T \nu - t_b W_b \eta_b \\ z &= -Q^{-\frac{1}{2}} S^T \nu, \end{aligned} \quad (34)$$

³As an observation, we note that even when S^T is a sparse matrix, $S^T S$ typically will not be, and hence the augmented dynamics (27) may not be immediately implementable as a distributed algorithm. A notable exception is when S is the transposed incidence matrix of a sparse graph, which gives $S^T S$ as the corresponding sparse graph Laplacian; this will occur in Section V.

where $\tau_\nu \succ 0$. The performance of (34) is characterized by the following result.

Proposition 4.1 (Dual Ascent Performance): The \mathcal{H}_2 norm of the input-output dual ascent dynamics (34) is given by

$$\|G\|_{\mathcal{H}_2}^2 = \frac{t_b^2}{2} \operatorname{Tr}(W_b^T \tau_\nu^{-1} W_b)$$

Proof of Proposition 4.1: The Lyapunov equation (9) for this problem takes the form

$$-X \frac{1}{\tau_\nu} S Q^{-1} S^T - S Q^{-1} S^T \frac{1}{\tau_\nu} X + S Q^{-1} S^T = 0$$

from which we find the unique solution $X = \frac{1}{2} \tau_\nu \succ 0$. With $B = [t_b \tau_\nu^{-1} W_b]$, the result follows by applying (8). \square

Comparing the result of Proposition 4.1 to the unaugmented saddle-point result of Theorem 3.1, we observe that the terms proportional to t_b^2 are identical. Therefore, when considering algorithm performance with disturbances entering the constraints, the primal-dual and pure dual-ascent algorithms achieve identical performance.

B. A Distributed Dual Augmented Lagrangian Method

Building off the dual function (33), we now derive a modified augmented Lagrangian algorithm, which can overcome the performance issues posed by disturbances affecting the vector b . The particulars of the derivation below are tailored towards distributed solutions, which will be discussed further in Section V in the context of distributed resource allocation. With this application in mind, we will focus in on the case where $n_b = n_x$. We partition each of the following matrices according to their columns as

$$S = [s_1 \ s_2 \ \cdots \ s_{n_x}], \quad W_b = [w_1 \ w_2 \ \cdots \ w_{n_x}].$$

With this partitioning, one may quickly see that for $Q = \operatorname{diag}(q_1, \dots, q_{n_x})$, the dual function (33) may be written as

$$\Phi(\nu) = \sum_{i=1}^{n_x} \underbrace{\left[-\frac{1}{2q_i} \nu^T s_i s_i^T \nu - \nu^T \left(\frac{c_i}{q_i} s_i + w_i b_i \right) - \frac{c_i^2}{2q_i} \right]}_{:=\tilde{\Phi}_i(\nu)}.$$

The dual function appears to separate into a sum, except for the common multiplier ν . To complete the separation, for each $i \in \{1, \dots, n_x\}$ we introduce a local copy $\nu^i \in \mathbb{R}^{n_r}$ of the vector of Lagrange multipliers $\nu \in \mathbb{R}^{n_r}$, and require that $\nu^i = \nu^j$ for all $i, j \in \{1, \dots, n_x\}$. To enforce these so-called agreement constraints, let $E \in \mathbb{R}^{n_x \times |\mathcal{E}|}$ be the oriented node-edge incidence matrix [33, Chapter 8] of a weakly connected acyclic⁴ graph $\mathcal{G} = (\{1, \dots, n_x\}, \mathcal{E})$, where \mathcal{E} is the set of oriented edges. The dual problem $\max_{\nu \in \mathbb{R}^{n_r}} \Phi(\nu)$ is then

⁴The acyclic assumption implies that $\operatorname{rank}(E^T) = |\mathcal{E}|$, in line with our assumption from Section II-A that the constraint matrix has full row rank. This assumption can be relaxed at the expense of more complex stability/performance proofs.

equivalent to the constrained problem⁵

$$\begin{aligned} & \underset{\boldsymbol{\nu} \in \mathbb{R}^{n_x n_r}}{\text{minimize}} && J_{\text{dual}}(\boldsymbol{\nu}) := - \sum_{i=1}^{n_x} \tilde{\Phi}_i(\nu^i) \\ & \text{subject to} && (E^\top \otimes I_{n_r})\boldsymbol{\nu} = \mathbf{0}_{|\mathcal{E}|*n_r}, \end{aligned} \quad (35)$$

where $\boldsymbol{\nu} = (\nu^1, \dots, \nu^{n_x})^\top \in \mathbb{R}^{n_x n_r}$. Since the graph is acyclic, E^\top has full row rank, and therefore satisfies our assumption concerning the constraint matrix (Section II). The key observation now is that the parameters b do not enter into the equality constraints of the optimization problem (35); this permits the full application of augmented Lagrangian techniques for improving the \mathcal{H}_2 performance of the saddle point algorithm. Building the augmented Lagrangian (26) for the problem (35) we have

$$L_{\text{dual}}^{\text{aug}} = - \sum_{i=1}^{n_x} \tilde{\Phi}_i(\nu^i) + \boldsymbol{\mu}^\top (E^\top \otimes I_{n_r})\boldsymbol{\nu} + \frac{\rho}{2} \boldsymbol{\nu}^\top (\mathbf{L} \otimes I_{n_r})\boldsymbol{\nu},$$

where $\boldsymbol{\mu} = (\mu^1, \dots, \mu^{|\mathcal{E}|})^\top \in \mathbb{R}^{|\mathcal{E}|*n_r}$ is a stacked vector of Lagrange multipliers $\mu^\ell \in \mathbb{R}^{n_r}$ for $\ell \in \{1, \dots, |\mathcal{E}|\}$, and $\mathbf{L} = \mathbf{L}^\top = \mathbf{E}\mathbf{E}^\top \in \mathbb{R}^{n_x \times n_x}$ is the Laplacian matrix of the undirected graph \mathcal{G}_u , obtained by ignoring the orientation of the edges in \mathcal{G} ; we let $\mathcal{N}(i)$ denote the neighbours of vertex i in the graph \mathcal{G}_u . Applying the saddle-point method, the dynamics may be written block-component-wise as

$$\begin{aligned} \bar{\tau}_\nu \dot{\boldsymbol{\nu}} &= - \frac{s_i s_i^\top}{q_i} \boldsymbol{\nu}^i - \left(\frac{c_i}{q_i} s_i + w_i b_i \right) \\ &\quad - \sum_{j:(i,j) \in \mathcal{E}} \mu^{ij} + \sum_{j:(j,i) \in \mathcal{E}} \mu^{ji} \\ &\quad - \rho \sum_{j \in \mathcal{N}(i)} (\nu^i - \nu^j), \quad i \in \{1, \dots, n_x\}, \\ \bar{\tau}_\mu \dot{\boldsymbol{\mu}}^{ij} &= \nu^i - \nu^j, \quad (i, j) \in \mathcal{E}, \end{aligned} \quad (36)$$

The algorithm (36) is distributed, in that the i th update equation requires only the local parameters s_i, w_i, b_i, c_i, q_i along with communicated state variables ν^j, μ^{ij} which come from adjacent nodes and edges in the graph \mathcal{G} . We refer to this dynamical system as the augmented dual distributed saddle-point (ADD-SP) dynamics. Following (34), we equip this system with disturbance inputs and performance outputs

$$b_i \rightarrow b_i + t_b \eta_i, \quad z_i = -q_i^{-\frac{1}{2}} s_i^\top \nu^i, \quad i \in \{1, \dots, n_x\}. \quad (37)$$

While a closed-form expression for the \mathcal{H}_2 norm of the this system is difficult to compute, we can state the following comparative result.

Corollary 4.2 (Augmented Dual Distributed Saddle-Point Performance): Consider the ADD-SP dynamics (36), denoted by G_{ADD} , with disturbance inputs $\eta(t)$ and performance output $z(t)$ as in (37), under the same assumptions as Theorem 3.4.

⁵Another equivalent formulation can be obtained by using the Laplacian matrix $\mathbf{E}\mathbf{E}^\top = \mathbf{L} = \mathbf{L}^\top \in \mathbb{R}^{n_x \times n_x}$ of the graph, and using instead the constraint $(\mathbf{L} \otimes I_{n_r})\boldsymbol{\nu} = \mathbf{0}_{n_x n_r}$. This formulation is sometimes preferable for multi-agent implementations, the analysis of which requires only small modifications from the present analysis. We focus instead of formulations involving the incidence matrix.

Then the squared \mathcal{H}_2 norm of the system (36)-(37) satisfies the upper bound

$$\|G_{\text{ADD}}\|_{\mathcal{H}_2}^2 \leq \frac{t_b^2}{2\bar{\tau}_\nu} \text{Tr}(\mathcal{W}_b^\top \mathcal{W}_b), \quad (38)$$

where $\mathcal{W}_b = \text{blkdiag}(w_1, w_2, \dots, w_{n_x}) \in \mathbb{R}^{n_x n_r \times n_x}$. Moreover, (38) is satisfied with equality if and only if $\rho = 0$.

Proof of Corollary 4.2: To begin, we note that

$$\sum_{i=1}^{n_x} \tilde{\Phi}_i(\nu^i) = -\frac{1}{2} \boldsymbol{\nu}^\top \mathcal{S} Q^{-1} \mathcal{S}^\top \boldsymbol{\nu} - \boldsymbol{\nu}^\top \mathcal{S} Q^{-1} c - \boldsymbol{\nu}^\top \mathcal{W}_b b$$

where $\mathcal{S} = \text{blkdiag}(s_1, s_2, \dots, s_{n_x}) \in \mathbb{R}^{n_x n_r \times n_x}$. In vector notation after shifting the equilibrium to the origin, the system (36)-(37) can be written as

$$\begin{aligned} \bar{\tau}_\nu \dot{\boldsymbol{\nu}} &= -\mathcal{S} Q^{-1} \mathcal{S}^\top \boldsymbol{\nu} - t_b \mathcal{W}_b \eta_b - (E \otimes I_{n_r})\boldsymbol{\mu} - \rho (\mathbf{L} \otimes I_{n_r})\boldsymbol{\nu} \\ \bar{\tau}_\mu \dot{\boldsymbol{\mu}} &= (E^\top \otimes I_{n_r})\boldsymbol{\nu} \\ z &= -Q^{-\frac{1}{2}} \mathcal{S}^\top \boldsymbol{\nu} \end{aligned}$$

where $\boldsymbol{\eta} = (\eta_1, \dots, \eta_{n_x})^\top \in \mathbb{R}^{n_x}$. In state-space this translates to

$$\begin{aligned} \begin{bmatrix} \dot{\boldsymbol{\nu}} \\ \dot{\boldsymbol{\mu}} \end{bmatrix} &= \underbrace{\begin{bmatrix} -\frac{1}{\bar{\tau}_\nu} \mathcal{S} Q^{-1} \mathcal{S}^\top - \frac{1}{\bar{\tau}_\nu} \rho (\mathbf{L} \otimes I_{n_r}) & -\frac{1}{\bar{\tau}_\nu} (E \otimes I_{n_r}) \\ \frac{1}{\bar{\tau}_\mu} (E^\top \otimes I_{n_r}) & \mathbf{0} \end{bmatrix}}_{A_{\text{ADD}}} \begin{bmatrix} \boldsymbol{\nu} \\ \boldsymbol{\mu} \end{bmatrix} \\ &\quad - \underbrace{\begin{bmatrix} \frac{t_b}{\bar{\tau}_\nu} \mathcal{W}_b \\ \mathbf{0} \end{bmatrix}}_{B_{\text{ADD}}} \eta_b, \quad z = \underbrace{\begin{bmatrix} -Q^{-\frac{1}{2}} \mathcal{S}^\top & \mathbf{0} \end{bmatrix}}_{C_{\text{ADD}}} \begin{bmatrix} \boldsymbol{\nu} \\ \boldsymbol{\mu} \end{bmatrix}. \end{aligned} \quad (39)$$

As in our previous results, one may verify that A_{ADD} is Hurwitz and that $(A_{\text{ADD}}, C_{\text{ADD}})$ is observable. Consider the observability Gramian candidate $X_{\text{ADD}} = \frac{1}{2} \text{blkdiag}(\bar{\tau}_\nu I_{n_x n_r}, \bar{\tau}_\mu I_{|\mathcal{E}|*n_r})$. Straightforward algebra shows that

$$X_{\text{ADD}} A_{\text{ADD}} + A_{\text{ADD}}^\top X_{\text{ADD}} + Q_{\text{ADD}} = \begin{bmatrix} -\rho (\mathbf{L} \otimes I_{n_r}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \preceq \mathbf{0}, \quad (40)$$

where $Q_{\text{ADD}} = C_{\text{ADD}}^\top C_{\text{ADD}}$. We conclude that X_{ADD} is a generalized observability Gramian for the system G_{ADD} . Furthermore, if X_t is the true observability Gramian for G_{ADD} , then as in Theorem 3.2, we have $X_t \preceq X_{\text{ADD}}$, and we conclude that

$$\|G_{\text{ADD}}\|_{\mathcal{H}_2}^2 \leq \text{Tr}(B_{\text{ADD}}^\top X_{\text{ADD}} B_{\text{ADD}}) = \frac{t_b^2}{2\bar{\tau}_\nu} \text{Tr}(\mathcal{W}_b^\top \mathcal{W}_b). \quad (41)$$

When $\rho = 0$, X_{ADD} is the exact observability Gramian and hence

$$\|G_{\text{ADD}}\|_{\mathcal{H}_2}^2 = \frac{t_b^2}{2\bar{\tau}_\nu} \text{Tr}(\mathcal{W}_b^\top \mathcal{W}_b) \quad (42)$$

To complete the proof, it remains to show that (42) in fact implies that $\rho = 0$. Suppose that (42) holds, and define $\Delta := X_{\text{ADD}} - X_t \succeq \mathbf{0}$. Then, following a similar argument as in the proof of Theorem 3.2, one may show that $\Delta B_{\text{ADD}} = \mathbf{0}$. From (40) and the fact that X_t is the actual observability Gramian of the system, we may subtract equations to obtain

$$\Delta A_{\text{ADD}} + A_{\text{ADD}}^\top \Delta = \begin{bmatrix} -\rho (\mathbf{L} \otimes I_{n_r}) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

Now, by multiplying each side of the equality above from the left by B_{ADD}^T and from the right by B_{ADD} , we find that

$$\begin{aligned} 0 &= \rho \mathcal{W}_b^T (L \otimes I_{n_r}) \mathcal{W}_b \\ &= \rho \mathcal{W}_b^T (E \otimes I_{n_r}) (E^T \otimes I_{n_r}) \mathcal{W}_b = \rho M^T M \end{aligned}$$

where $M = (E^T \otimes I_{n_r}) \mathcal{W}_b$. Since the graph \mathcal{G} is connected and W_b is square and of full rank, it always holds that $M \neq 0$, and therefore we conclude that $\rho = 0$. \square

The point of interest from Corollary 4.2 is that the bound on the \mathcal{H}_2 performance is *independent* of ρ . In particular then, as ρ becomes large the norm does not grow without bound, which resolves the issue observed in the result of Theorem 3.4. When applied to the resource allocation problem in Section V, we will in fact be able to make the stronger statement that the norm is a strictly decreasing function of ρ , and augmentation can therefore be used successfully to improve saddle-point algorithm performance.

V. APPLICATION TO RESOURCE ALLOCATION PROBLEMS

We now apply the results from the previous subsections to a particular class of problems. As a special case of the problem (1), consider the *resource allocation problem*

$$\begin{aligned} &\underset{x \in \mathbb{R}^n}{\text{minimize}} && \sum_{i=1}^n \frac{1}{2} q_i x_i^2 + c_i x_i \\ &\text{subject to} && \sum_{i=1}^n x_i = \sum_{i=1}^n d_i, \end{aligned} \quad (43)$$

where $q_i > 0$, $c_i \in \mathbb{R}$, and $d_i \in \mathbb{R}$. Comparing (43) to (1), we have $Q = \text{diag}(q_1, \dots, q_n)$, $S = \mathbb{1}_n^T$, $d = b$, and $W_b = \mathbb{1}_n^T$. The interpretation of (43) is that a resource must be obtained from one of n suppliers in amount x_i , subject to a total demand satisfaction constraint. The objective function of (43) can be interpreted as the sum of the utilities $-c_i x_i$ minus the sum of the costs $q_i x_i^2/2$. In a multi-agent context, each variable x_i is assigned to an agent, the parameters q_i, c_i, d_i are available locally to each agent, and the agents must collectively solve the problem (43) through local exchange of information.

We will consider (43) along with several equivalent reformulations, and apply our results from the previous sections to assess the input-output performance of the resulting saddle-point algorithms. External disturbances η_d will be integrated into the algorithms as $d_i \rightarrow d_i + \eta_i$, where η_i models the disturbances in demand d_i . For simplicity, all time-constant matrices are assumed to be multiples of the identity matrix. To most clearly indicate which algorithms require communication of which variables, in this section algorithms are *not* written in deviation coordinates with respect to the optimizer. In all cases, the performance output z is chosen such that $\|z(t)\|_2^2 = (x(t) - x^*)^T Q (x(t) - x^*)$, where x^* is the global primal optimizer.

To begin, the augmented Lagrangian of (43) is given in vector notation by

$$L(x, \nu) = \frac{1}{2} x^T Q x + c^T x + \nu \mathbb{1}_n^T (x - d) + \frac{\rho}{2} \|\mathbb{1}_n^T (x - d)\|_2^2. \quad (44)$$

where $\nu \in \mathbb{R}$. Applying the saddle-point method to the Lagrangian $L(x, \nu)$ and attaching the same disturbances $\eta \in \mathbb{R}^n$ and performance outputs $z \in \mathbb{R}^n$ as before, we obtain the *centralized saddle-point dynamics*

$$\text{RA}_{\text{cent}}(\rho) : \begin{cases} \bar{\tau}_x \dot{x} = -Qx - c - \nu \mathbb{1}_n - \rho \mathbb{1}_n \mathbb{1}_n^T (x - d - \eta) \\ \bar{\tau}_\nu \dot{\nu} = \mathbb{1}_n^T (x - d - \eta) \\ z = Q^{\frac{1}{2}} (x - x^*) \end{cases} \quad (45)$$

When $\rho = 0$, the algorithm (45) is of a gather-and-broadcast type [34], where all states x_i and disturbances d_i are collected and processed by a central agent with state ν . When $\rho > 0$, the additional term $\mathbb{1}_n \mathbb{1}_n^T$ in the algorithm requires all-to-all communication of the local imbalances $x_i - d_i$.

We now consider a reformulation that results in a distributed optimization algorithm. Let $\mathcal{G} = (\{1, \dots, n\}, \mathcal{E})$ denote a weakly connected acyclic graph over the agent set $\{1, \dots, n\}$, and let $E \in \mathbb{R}^{n \times |\mathcal{E}|}$ denote the oriented node-edge incidence matrix of \mathcal{G} . The constraint $\mathbb{1}_n^T x = \mathbb{1}_n^T d$ in the resource allocation problem (43) is equivalent to the existence⁶ of a vector $\delta \in \mathbb{R}^{|\mathcal{E}|}$ such that $E\delta = x - d$. The resource allocation problem (43) can therefore be equivalently rewritten as

$$\begin{aligned} &\underset{x \in \mathbb{R}^n, \delta \in \mathbb{R}^{|\mathcal{E}|}}{\text{minimize}} && \sum_{i=1}^n \frac{1}{2} q_i x_i^2 + c_i x_i \\ &\text{subject to} && E\delta = x - d, \end{aligned} \quad (46)$$

with associated augmented Lagrangian

$$L'(x, \delta, \nu) = \frac{1}{2} x^T Q x + c^T x + \nu^T (E\delta - x + d) + \frac{\rho}{2} \|E\delta - x + d\|_2^2,$$

where $\nu \in \mathbb{R}^n$. This reformulation can be interpreted as a version of (1) with an expanded primal variable (x, δ) and an expanded dual variable $\nu = (\nu_1, \dots, \nu_n)^T$. By applying the saddle-point method to the problem (46), we obtain the *distributed saddle-point dynamics*

$$\text{RA}_{\text{dist}}(\rho) : \begin{cases} \bar{\tau}_x \dot{x} = -Qx - c + \rho(E\delta - x + d + \eta) + \nu \\ \bar{\tau}_\delta \dot{\delta} = -E^T \nu - \rho E^T (E\delta - x + d + \eta) \\ \bar{\tau}_\nu \dot{\nu} = E\delta - x + d + \eta \\ z = Q^{\frac{1}{2}} (x - x^*) \end{cases} \quad (47)$$

When $\rho = 0$, the algorithm (47) is distributed with the topology of the graph \mathcal{G} , with states (x_i, ν_i) associated with each node and a state δ_i associated with each edge. When $\rho > 0$, the algorithm contains the so-called edge Laplacian matrix $E^T E$ [35], which under our acyclic assumption is positive definite.

Our third formulation is the dual ascent algorithm (33). Substitution of the appropriate matrices into (33) leads to the *centralized dual ascent dynamics*

$$\text{RA}_{\text{cent}}^{\text{dual}} : \begin{cases} \bar{\tau}_\nu \dot{\nu} = -(\mathbb{1}_n^T Q^{-1} \mathbb{1}_n) \nu - \mathbb{1}_n^T (Q^{-1} c + d + \eta) \\ z = Q^{\frac{1}{2}} (x - x^*) = -Q^{-\frac{1}{2}} (c + \nu \mathbb{1}_n) - Q^{\frac{1}{2}} x^*, \end{cases} \quad (48)$$

where $\nu \in \mathbb{R}$. Algorithm (48) is again centralized, with a single central agent with state ν performing all computations and broadcasting $x_i = -q_i^{-1}(c_i + \nu)$ back to each agent. For

⁶This follows since $\ker(E^T) = \text{span}(\mathbb{1}_n)$, and hence $\text{Im}(E)$ is the subspace orthogonal to the vector $\mathbb{1}_n$.

our fourth and final formulation, we apply the ADD-SP method developed in Section IV-B. For the problem (43), one quickly deduces that $\mathcal{S} = \mathcal{W}_b = I_n$, and the algorithm (36) reduces to

$$\text{RA}_{\text{dist}}^{\text{dual}}(\rho) : \begin{cases} \bar{\tau}_\nu \dot{\nu} = -Q^{-1}\nu - (d + \eta) - Q^{-1}c - E\mu - \rho L\nu \\ \bar{\tau}_\mu \dot{\mu} = E^\top \nu \\ z = Q^{\frac{1}{2}}(x - x^*) = -Q^{-\frac{1}{2}}(\nu + c) + Q^{\frac{1}{2}}x^*, \end{cases} \quad (49)$$

When $\rho = 0$, this algorithm is distributed with the graph \mathcal{G} associated with the incidence matrix E , with states ν_i associated with nodes and states μ_{ij} associated with edges. When $\rho > 0$, the algorithm additionally contains the undirected Laplacian matrix $L = EE^\top$ of \mathcal{G} , and thus remains distributed.

For each of the four formulations above, we compute the \mathcal{H}_2 norm from the disturbance input η to the performance output z . For $\text{RA}_{\text{cent}}(\rho)$, $\text{RA}_{\text{cent}}^{\text{dual}}(\rho)$, and $\text{RA}_{\text{dist}}^{\text{dual}}(\rho)$ this follows immediately from Theorem 3.4, Proposition 4.1, and Corollary 4.2, respectively. The algorithm $\text{RA}_{\text{dist}}(\rho)$ requires a modification of the proof of Theorem 3.4, since the objective function is no longer strongly convex in the primal variables (x, δ) ; we omit the details.

TABLE I: Comparison of squared \mathcal{H}_2 norm expressions

| System | $\rho = 0$ | $\rho \rightarrow \infty$ |
|-----------------------------------------------|-----------------------|---------------------------|
| $\text{RA}_{\text{cent}}(\rho)$ | $n/(2\bar{\tau}_\nu)$ | $+\infty$ |
| $\text{RA}_{\text{dist}}(\rho)$ | $n/(2\bar{\tau}_\nu)$ | $+\infty$ |
| $\text{RA}_{\text{cent}}^{\text{dual}}(\rho)$ | $n/(2\bar{\tau}_\nu)$ | — |
| $\text{RA}_{\text{dist}}^{\text{dual}}(\rho)$ | $n/(2\bar{\tau}_\nu)$ | $1/(2\bar{\tau}_\nu)$ |

The first column of Table I shows the \mathcal{H}_2 system norms for the four formulations when $\rho = 0$, i.e., the unaugmented versions of the various saddle-point algorithms. Despite substantial differences between the algorithms in terms of information structure and number of states, all four have the same input-output performance in the \mathcal{H}_2 norm. This implies that a distributed implementation will perform no worse than a centralized implementation.

While these four formulations all possess identical system norms under the basic primal-dual algorithm, augmentation differentiates these methods from one another, and substantial differences between the algorithms begin to appear as ρ is increased. The limiting results are tabulated in the second column of Table I. The input-output performance of the first two formulations becomes arbitrarily bad as the augmentation gain ρ increases, while the performance of the ADD-SP algorithm improves substantially, becoming independent of the system size in the limit $\rho \rightarrow \infty$.

We illustrate the results in Table I via time-domain simulations in Figures 2 and 3 for the system in (46) with $n = 2$ and an underlying line graph with $E = [1 \ -1]^\top$. With unit variance white noise as inputs, the unaugmented implementation in Figure 2 for the four different algorithms results in identical steady-state output variance, numerically computed as the squared \mathcal{H}_2 norm in (7).

In Figure 3, a sufficiently large augmentation factor ρ is introduced to penalise the constraint violations. It is observed

that with the augmentation, the steady-state variance of the outputs for the centralized and distributed implementations in $\text{RA}_{\text{cent}}(\rho)$, $\text{RA}_{\text{dist}}(\rho)$ worsens, while that of the distributed dual implementation from $\text{RA}_{\text{dist}}^{\text{dual}}(\rho)$ improves.

Figure 4(a) illustrates how the choice of communication graph topology influences the performance of the algorithm $\text{RA}_{\text{dist}}^{\text{dual}}(\rho)$. We consider $n = 4$ agents, and implement the algorithm with line, ring, and complete communication graphs. While Corollary 4.2 and the results of Table I hold only for acyclic graphs, Figure 4(a) shows that in all cases the algorithm's performance improves as ρ increases. For a given value of ρ , graphs with higher connectivity show a greater improvement. This behaviour is explained by noting that the algorithm (49) has the same form as the augmented saddle-point dynamics (27), and an analysis similar to that performed for Theorem 3.4 can in fact be performed for (49). For uniform cost function parameters and an acyclic graph, this leads to the expression

$$\|G_{\text{RA}_{\text{dist}}^{\text{dual}}(\rho)}\|_{\mathcal{H}_2}^2 = \frac{1}{2\bar{\tau}_\nu} \left(1 + \sum_{i=2}^n \frac{q}{q + \rho\lambda_i} \right) \quad (50)$$

where $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n$ are the eigenvalues of the Laplacian matrix. As graph connectivity increases, so does λ_2 , and performance therefore improves.

Finally, Figure 4(b) plots the system norm as a function of ρ for the three augmented algorithms, for a test case with $n = 4$ agents. The norm is not a monotonic function of the augmentation factor ρ for the implementations in $\text{RA}_{\text{cent}}(\rho)$ and $\text{RA}_{\text{dist}}(\rho)$, but is monotonic for $\text{RA}_{\text{dist}}^{\text{dual}}(\rho)$ applied to resource allocation problems, in agreement with the result for the parametrically uniform case in equation (50).

VI. CONCLUSIONS

We have studied the input-output performance of continuous-time saddle-point methods for solving linearly constrained convex quadratic programs, providing an explicit formula for the \mathcal{H}_2 norm under a relevant input-output configuration. We then studied the effects of Lagrangian regularization and augmentation on this norm, and derived a distributed dual version of the augmented algorithm which overcomes some of the limitations of naive augmentation. We then applied the results to compare several implementations of the saddle-point method to resource allocation problems.

Open directions for future research include input-output performance metrics for problems involving inequality constraints, for distributed implementations where communication delays occur between agents, and for other classes of distributed optimization algorithms. Another interesting question is how to further improve the \mathcal{H}_2 performance of saddle-point methods by designing auxiliary feedback controllers; augmentation is but one approach.

REFERENCES

- [1] T. Kose, "Solutions of saddle value problems by differential equations," *Econometrica*, vol. 24, no. 1, pp. 59–70, 1956.
- [2] K. Arrow, L. Hurwicz, and H. Uzawa, *Studies in linear and non-linear programming*. Stanford University Press, 2006.




Fig. 2: Steady-state variance for the four different implementations for the unaugmented case, for parameters $n = 2$, $Q = \text{diag}(4, 25)$, $\tau_x = \tau_\delta = \tau_\nu = \tau_\mu = 1$, $E = [1 \ -1]^\top$; the remaining parameters do not influence the results.




Fig. 3: Steady-state variance for the three different implementations for the augmented case, for parameters $n = 2$, $\rho = 100$, $Q = \text{diag}(4, 25)$, $\tau_x = \tau_\delta = \tau_\nu = \tau_\mu = 1$, $E = [1 \ -1]^\top$; the remaining parameters do not influence the results.

- [3] D. Feijer and F. Paganini, "Stability of primal-dual gradient dynamics and applications to network optimization," *Automatica*, vol. 46, no. 12, pp. 1974–1981, 2010.
- [4] J. Wang and N. Elia, "A control perspective for centralized and distributed convex optimization," in *IEEE Conf. on Decision and Control and European Control Conference*, Orlando, FL, USA, Dec. 2011, pp. 3800–3805.
- [5] A. Cherukuri, E. Mallada, and J. Cortés, "Asymptotic convergence of constrained primal-dual dynamics," *Systems & Control Letters*, vol. 87, pp. 10 – 15, 2016.
- [6] S. H. Low and D. E. Lapsley, "Optimization flow control I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–74, 1999.
- [7] F. Dörfler, J. W. Simpson-Porco, and F. Bullo, "Breaking the hierarchy: Distributed control & economic optimality in microgrids," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 3, pp. 241–253, 2016.
- [8] N. Li, L. Chen, C. Zhao, and S. H. Low, "Connecting automatic generation control and economic dispatch from an optimization view," in *American Control Conference*, Portland, OR, USA, Jun. 2014, pp. 735–740.
- [9] A. Cherukuri and J. Cortés, "Initialization-free distributed coordination for economic dispatch under varying loads and generator commitment," *Automatica*, vol. 74, pp. 183 – 193, 2016.
- [10] E. Dall'Anese, S. V. Dhople, and G. B. Giannakis, "Regulation of dynamical systems to optimal solutions of semidefinite programs: Algorithms and applications to ac optimal power flow," in *American Control Conference*, Chicago, IL, USA, July 2015, pp. 2087–2092.
- [11] T. Stegink, C. D. Persis, and A. van der Schaft, "A unifying energy-based approach to stability of power grids with market dynamics," *IEEE Transactions on Automatic Control*, vol. 62, no. 6, pp. 2612–2622, Jun. 2017.
- [12] E. Mallada, C. Zhao, and S. Low, "Optimal load-side control for frequency regulation in smart grids," *IEEE Transactions on Automatic Control*,




Fig. 4: Squared \mathcal{H}_2 norm as a function of ρ , for parameters $n = 4$, $\tau_x = \tau_\delta = \tau_\nu = \tau_\mu = 1$ and unweighted graphs.

- vol. 62, no. 12, pp. 6294–6309, 2017.
- [13] J. T. Wen and M. Arcak, “A unifying passivity framework for network flow control,” *IEEE Transactions on Automatic Control*, vol. 49, no. 2, pp. 162–174, 2004.
- [14] J. Wang and N. Elia, “Control approach to distributed optimization,” in *Allerton Conf. on Communications, Control and Computing*, Monticello, IL, USA, 2010, pp. 557–561.
- [15] G. Droge, H. Kawashima, and M. B. Egerstedt, “Continuous-time proportional-integral distributed optimisation for networked systems,” *Journal of Control and Decision*, vol. 1, no. 3, pp. 191–213, 2014.
- [16] G. Droge and M. Egerstedt, “Proportional integral distributed optimization for dynamic network topologies,” in *American Control Conference*, Portland, OR, USA, June 2014, pp. 3621–3626.
- [17] T. Hatanaka, N. Chopra, T. Ishizaki, and N. Li, “Passivity-based distributed optimization with communication delays using pi consensus algorithm,” 2017, <https://arxiv.org/abs/1609.04666>.
- [18] A. Cherukuri, E. Mallada, S. Low, and J. Cortés, “The role of strong convexity-concavity in the convergence and robustness of the saddle-point dynamics,” in *Allerton Conf. on Communications, Control and Computing*, Monticello, IL, USA, Sep. 2016, pp. 504–510.
- [19] A. Cherukuri, E. Mallada, S. Low, and J. Cortés, “The role of convexity on saddle-point dynamics: Lyapunov function and robustness,” *IEEE Transactions on Automatic Control*, vol. PP, no. 99, pp. 1–1, 2017.
- [20] J. W. Simpson-Porco, “Input/output analysis of primal-dual gradient algorithms,” in *Allerton Conf. on Communications, Control and Computing*, Monticello, IL, USA, Sep. 2016, pp. 219–224.
- [21] L. Lessard, B. Recht, and A. Packard, “Analysis and design of optimization algorithms via integral quadratic constraints,” *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 57–95, 2016.
- [22] B. Hu and L. Lessard, “Control interpretations for first-order optimization methods,” in *American Control Conference*, Seattle, WA, USA, May 2017, pp. 3114–3119.
- [23] J. W. Simpson-Porco, B. K. Poola, N. Monshizadeh, and F. Dörfler, “Quadratic performance of primal-dual methods with application to secondary frequency control of power systems,” in *IEEE Conf. on Decision and Control*, Las Vegas, NV, USA, Dec. 2016, pp. 1840–1845.
- [24] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [25] M. Benzi, G. H. Golub, and J. Liesen, “Numerical solution of saddle point problems,” *Acta Numerica*, vol. 14, pp. 1–137, 2005.
- [26] A. Nedic and A. Ozdaglar, “Subgradient methods for saddle-point problems,” *Journal of Optimization Theory and Applications*, vol. 142, no. 1, pp. 205–228, 2009.
- [27] K. Zhou and J. C. Doyle, *Essentials of Robust Control*. New Jersey: Prentice Hall, 1998.
- [28] G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory*, ser. Texts in Applied Mathematics. Springer, 2000, no. 36.
- [29] M. B. Khuzani and N. Li, “Distributed regularized primal-dual method: Convergence analysis and trade-offs,” *arXiv preprint arXiv:1609.08262*, 2016.
- [30] A. Simonetto and G. Leus, “Double smoothing for time-varying distributed multiuser optimization,” in *IEEE Global Conference on Signal and Information Processing*, Atlanta, GA, USA, Dec. 2014, pp. 852–856.
- [31] E. Dall’Anese and A. Simonetto, “Optimal power flow pursuit,” *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 942–952, 2018.
- [32] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, 2008.
- [33] F. Bullo, *Lectures on Network Systems*. Version 0.85, May 2016, with contributions by J. Cortés, F. Dörfler, and S. Martínez. [Online]. Available: <http://motion.me.ucsb.edu/book-lns>
- [34] F. Dörfler and S. Grammatico, “Gather-and-broadcast frequency control in power systems,” *Automatica*, vol. 79, pp. 296 – 305, 2017.
- [35] D. Zelazo and M. Mesbahi, “Edge agreement: Graph-theoretic performance bounds and passivity analysis,” *IEEE Transactions on Automatic Control*, vol. 56, no. 3, pp. 544–555, 2011.




John W. Simpson-Porco (S’11–M’16) received the B.Sc. degree in engineering physics from Queen’s University, Kingston, ON, Canada in 2010, and the Ph.D. degree in mechanical engineering from the University of California at Santa Barbara, Santa Barbara, CA, USA in 2015.

He is currently an Assistant Professor of Electrical and Computer Engineering at the University of Waterloo, Waterloo, ON, Canada. He was previously a visiting scientist with the Automatic Control Laboratory at ETH Zürich, Zürich, Switzerland. His research

focuses on the control and optimization of multi-agent systems and networks, with applications in modernized power grids.

Prof. Simpson-Porco is a recipient of the 2012–2014 IFAC Automatica Prize and the Center for Control, Dynamical Systems and Computation Best Thesis Award and Outstanding Scholar Fellowship.



Bala Kameshwar Poolla (S'15) received his B.Tech. degree in Electrical Engineering and M.Tech. degree in Control Systems Engineering from the Indian Institute of Technology Kharagpur, India in 2013. He is currently a Ph.D. student at the Automatic Control Laboratory at ETH Zürich. His research interests include applications of control theory to power grids, renewable energy systems, and energy markets. He was a finalist for the ACC 2016 Best Student Paper Award.



Nima Monshizadeh obtained his B.Sc. degree in electrical engineering from the University of Tehran, the M.Sc. degree in control engineering from K.N. Toosi University of Technology, Tehran, Iran, and the Ph.D. degree with honor (cum laude) from the Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, The Netherlands, in December 2013. Afterwards, he was appointed as a Post-Doctoral Researcher at the Engineering and Technology Institute of the University of Groningen. He was a Research Associate with the

Control group of the University of Cambridge, before taking a tenure track position at the University of Groningen in January 2018. His research interests include power networks, model reduction, optimization and control of complex networks.



Florian Dörfler (S'09–M'13) is an Assistant Professor at the Automatic Control Laboratory at ETH Zürich. He received his Ph.D. degree in Mechanical Engineering from the University of California at Santa Barbara in 2013, and a Diplom degree in Engineering Cybernetics from the University of Stuttgart in 2008. From 2013 to 2014 he was an Assistant Professor at the University of California Los Angeles. His primary research interests are centered around distributed control, complex networks, and cyber-physical systems currently with applications

in energy systems and smart grids. His students were winners or finalists for Best Student Paper awards at the European Control Conference (2013), the American Control Conference (2016), and the PES PowerTech Conference 2017. His articles received the 2010 ACC Student Best Paper Award, the 2011 O. Hugo Schuck Best Paper Award, and the 2012-2014 Automatica Best Paper Award, and the 2016 IEEE Circuits and Systems Guillemin-Cauer Best Paper Award. He is a recipient of the 2009 Regents Special International Fellowship, the 2011 Peter J. Frenkel Foundation Fellowship, and the 2015 UCSB ME Best PhD award.