

Characterizing Controllability and Observability for Systems with Locality, Communication, and Actuation Constraints

Lauren Conger
Caltech

lconger@caltech.edu

Yiheng Lin
Caltech

yihengl@caltech.edu

Adam Wierman
Caltech

adamw@caltech.edu

Eric Mazumdar
Caltech

mazumdar@caltech.edu

Abstract—This paper presents a closed-form notion of controllability and observability for systems with communication delays, actuation delays, and locality constraints. The formulation reduces to classical notions of controllability and observability in the unconstrained setting. As a consequence of our formulation, we show that the addition of locality and communication constraints may not affect the controllability and observability of the system, and we provide an efficient sufficient condition under which this phenomenon occurs. This contrasts with actuation and sensing delays, which cause a gradual loss of controllability and observability as the delays increase. We illustrate our results using linearized swing equations for the power grid, showing how actuation delay and locality constraints affect controllability.

I. INTRODUCTION

Sensor and actuator placement is a crucial problem in the design of real world systems in application areas ranging from power systems and traffic networks to environmental monitoring. Taking power systems as an example, devices such as generators and transformers can provide additional inputs to actuate the system, while other devices like meters report information about the system's state. A system designer may face the challenge of deploying a limited number of such devices to different locations in a power system with the goal of keeping the system stable or performing well in other control tasks. While sensor and actuator placement is an old problem that has been well studied [1], [2], the proliferation of low-cost sensors and the decentralization of control in large systems like power systems has led to new challenges. Many real-world systems are subject to delayed and information-constrained sensors that can significantly impact the quality of the control. However, we currently lack tools to rigorously characterize the performance loss under various sensor configurations. This leads us to study the problem of quantifying the quality of a set of controllers and sensors in systems subject to actuation and sensing delays as well as communication constraints.

While task-specific metrics like cost are commonly used to measure the quality of a set of sensors and actuators, there are many settings where the cost is unknown or time varying; here task-agnostic notions of controllability and observability

LC is supported by a NDSEG fellowship from the Air Force Office of Research and a PIMCO fellowship. YL is supported by Amazon AI4Science Fellowship and PIMCO Graduate Fellowship in Data Science. This work is supported by NSF Grants CNS-2146814, CPS-2136197, CNS-2106403, NGSDI-2105648, CNS-2240110.

of dynamical systems [2], [3] offer a suitable proxy. Controllability, defined as the determinant of the controllability Gramian [4], is the volume over which a control input with magnitude one can drive the state to the zero state. Likewise, observability, defined as the determinant of the observability Gramian, is the volume of the signal-to-noise ratio along any direction in the state space. These metrics are used to compare different sets of actuators and sensors for dynamical systems; a system with larger controllability, given fixed dynamics, is preferred to one with smaller controllability.

Although the notions of controllability and observability are well-defined for classic control problems, new challenges arise from the decentralized control of modern large-scale systems [5], [6], [7], [8], where the controllers may face different constraints on communication speed, actuation delays, and/or locality. Such constraints, which we refer to as *system level constraints*, cannot easily be captured in the classical formulations of controllability/observability based on the Gramian or control inputs [4]. Fortunately, a recently proposed framework called *system level synthesis (SLS)* [9] handles system level constraints efficiently, where the key technique is to formulate the original control problem in terms of closed-loop maps. Some recent works address controllability and observability in the specific setting of systems with delays [10], [11], [12]; SLS allows for a much broader class of constraints, including delays. Using SLS, a recent work [13] shows sufficient conditions under which locality constraints on model predictive control (MPC) leads to *no loss in optimality* for any convex global cost function. However, even with SLS, it is unclear whether one can formally extend the notion of controllability and observability to include system level constraints faced by large-scale networked systems, and whether such measures can be computed efficiently to quantify the loss of performance under different combinations of system level constraints.

Contributions. We propose a formal definition of controllability and observability subject to system level constraints (Section III). Our definitions use the SLS framework and we draw connections to classical distributed control, differential privacy, and speed-accuracy tradeoffs in neurons. We provide a provable closed-form solution for our definitions (Algorithm 1, Proposition 6), which strictly generalize the classic notions of controllability and observability (Theorem 4).

Importantly, we establish rank conditions for system level constraints under which there is no performance loss compared with the unconstrained system (Theorem 9). This rig-

orously characterizes the empirical observation that locality and communication constraints do not necessarily decrease controllability and observability. In contrast, actuation delay does not satisfy this condition, consistent with numerical results showing a gradual decrease in controllability as actuation delay increases. We perform numerical experiments on a synthetic example and a real-world model of linearized and discretized swing equations for the power grid (Section V).

Notation. Let $N_x \in \mathbb{Z}_+$ and $N_u \in \mathbb{Z}_+$ be the state and input dimensions respectively. Let $x_0 := x(0)$ and $x_T := x(T)$. A single index such as $Z[i]$ indicates the i^{th} row or block of rows of a matrix, and a double index such as $A[i, j]$ indicates the i^{th} row and j^{th} column of a matrix. Let $A[:, j]$ or $A[i, :]$ denote all rows or columns of A respectively, and $\text{blkdiag}([A, B, C])$ denote a matrix with submatrices A , B , and C on the block diagonals. Brackets around a scalar indicates the set of indices $[a] := [1, \dots, a]$. A vector over a variable, such as $\vec{\phi}_x$, is the vectorization of matrix, where the columns are stacked into a vector. A hat symbol over a matrix indicates the vectorization of matrix multiplication, that is, \hat{A} is defined such that if $Y = AX$ for two matrices X and Y , then $\vec{y} = \hat{A}\vec{x}$. We define $\text{supp } A$ to be an indicator matrix such that $(\text{supp } A)[i, j] = 0$ if $A[i, j]$ is constrained to be zero. A^\dagger is the Moore-Penrose pseudoinverse of A .

II. MOTIVATION & BACKGROUND

We begin by introducing the problem of computing the controllability and observability of a dynamical system and illustrating the challenges that come from considering constraints in the classical formulation of the problem. We then show how system level synthesis (SLS) offers a framework for computing the controllability and observability of a system while incorporating system-level constraints.

A. Classical Controllability and Observability

We consider linear systems with dynamics given by

$$x(t+1) = Ax(t) + Bu(t), \quad y(t) = Cx(t), \quad (1)$$

where $x(t) \in \mathbb{R}^{N_x}$, $u(t) \in \mathbb{R}^{N_u}$, and $y(t) \in \mathbb{R}^{N_y}$. A system is said to be *controllable* in time T if it can be driven from any state x_0 to any terminal state x_T , which holds when the controllability Gramian W_c is full-rank. The controllability Gramian provides additional information about how easily the system is controlled in any direction, and the controllability volume is denoted V_c ,

$$W_c = \sum_{i=0}^{T-1} A^i B B^\top (A^\top)^i, \quad V_c = \det W_c. \quad (2)$$

The magnitude of each eigenvalue of W_c is the distance in the direction of the eigenvector that one unit of control provides (after accounting for the trajectory resulting from the initial condition alone). The determinant of W_c is then the state space volume that is reachable with one unit of control input. If the system is not controllable in a certain direction, the volume is zero. However, this computational method assumes that any sequence of control input can be selected; this assumption is not true when information among

states is restricted, actuation is delayed, or controllers are dependent on only local information.

Similarly, for an unconstrained system, one can define the observability Gramian and observability respectively, as

$$W_o = \sum_{i=0}^{T-1} A^\top C^\top C A^i, \quad V_o = \det W_o. \quad (3)$$

The magnitude of each eigenvalue of the Gramian W_o is the signal-to-noise ratio for a state in the direction of the corresponding eigenvector. The determinant of W_o corresponds to the ‘observability volume’ of the system. Similarly to controllability, this method assumes instantaneous and global observations for estimation. This does not apply to large-scale systems where state estimation depends only on local observation, and the observations are not immediate.

This paper proposes a new method for computing controllability and observability that accounts for delayed actuation and observation, local control and information sharing. Critical to the proposed approach is the system level synthesis (SLS) parameterization of the linear system. In what follows, we consider the state estimation problem as the dual of the control problem; in what follows, one can substitute

$$A \leftrightarrow A^\top, \quad B \leftrightarrow C^\top.$$

to obtain the corresponding observability framework.

B. System Level Synthesis (SLS)

The SLS parametrization of the state dynamics in (1) is

$$\phi_x(t+1) = A\phi_x(t) + B\phi_u(t), \quad (4)$$

where $x(t) = \phi_x(t)x_0$ and $u(t) = \phi_u(t)x_0$, $\phi_x(t) \in \mathbb{R}^{N_x \times N_x}$, $\phi_u(t) \in \mathbb{R}^{N_u \times N_x}$, and $\phi_x(0) = I$. Let the closed-loop maps ϕ_x , ϕ_u , and ϕ be defined as

$$\phi_x = \begin{bmatrix} \phi_x(0) \\ \vdots \\ \phi_x(T-1) \end{bmatrix}, \quad \phi_u = \begin{bmatrix} \phi_u(0) \\ \vdots \\ \phi_u(T-1) \end{bmatrix}, \quad \phi = \begin{bmatrix} \phi_x \\ \phi_u \end{bmatrix}.$$

A prior result from [9, Theorem 2] guarantees that the subspace defined by ϕ parameterizes all possible system responses; optimizing over the state feedback controller $u(t) = K(t)x(t)$ is equivalent to optimizing over ϕ .

This parameterization allows us to explicitly write constraints on ϕ_x and ϕ_u which encode the communication, actuation, and sensing constraints. To enforce that the i^{th} entry of x_0 does not affect the j^{th} entry of the state or input at time t , we set $\phi_x(t)[j, i] = 0$ or $\phi_u(t)[j, i] = 0$ respectively. We write the system level constraints as

$$S_x \vec{\phi}_x = 0, \quad S_u \vec{\phi}_u = 0, \quad (5)$$

where S_x (S_u resp.) has a single nonzero entry in each row, with the column corresponding to the entry of ϕ_x (ϕ_u resp.) that is constrained to be zero. We define the set of closed-loop maps that satisfy (5) as \mathcal{S} , such that $\phi \in \mathcal{S}$. Equipped with the ability to specify system level constraints, in the next section we present the definitions of controllability under these constraints.

III. CONSTRAINED SYSTEMS

Our focus in this paper is systems that are subject to actuation delay, communication delay and restrictions, and locality constraints. As discussed in the previous section, these constraints can be expressed as support constraints on the closed-loop maps ϕ_x and ϕ_u .

We formally define the different types of constraints considered in this work as follows. We assume the system is formed by a network of nodes, and each entry of the system state $x(t)$ corresponds to the local state of a node. The network is characterized by a directed graph whose adjacency matrix $\mathcal{G}[i, j]$ is defined such that $\mathcal{G}[i, j] = 1$ if $A[i, j] \neq 0$; $\mathcal{G}[i, j] = 0$ otherwise. We denote the graph distance between two nodes i and j as $\text{dist}_{\mathcal{G}}(i, j)$.

- **Actuation delay:** The system is subject to a τ -step actuation delay if the actuator can react to state information at least τ steps after, i.e., $\phi_u(t) = 0$, for $t < \tau$.
- **Communication speed:** The communication speed is v if the propagation of communication signals is v times the dynamics, i.e., $\text{supp } \phi_x(t) = \text{supp } \mathcal{G}^{\lfloor vt \rfloor}$ and $\text{supp } \phi_u(t) = \text{supp } (|B^\top| \text{supp } \phi_x(t))$, so that $x(t)[j]$ cannot depend on $x_0[i]$ if $\text{dist}_{\mathcal{G}}(i, j) > vt$. Likewise, $u(t)[i]$ cannot depend on $x_0[j]$ if $B[k, i] = 0$ for all k such that $\text{dist}_{\mathcal{G}}(j, k) \leq vt$.
- **Locality:** The system is subject to a locality constraint with radius r if the state of a node can only be affected by the states of its r -hop neighbors, i.e., $\phi_x(t)[j, i] = 0$ if $\text{dist}_{\mathcal{G}}(i, j) > r$, and a controller uses information from the nearest r -hops away from the states on which it acts, specified as $\text{supp } \phi_u = \text{supp } (|B^\top| \text{supp } \phi_x)$.

We formulate the definition of controllability and observability for constrained systems using SLS operators. We consider the magnitude of any x_0 such that one unit of control can drive the state to the origin, again scaled by the effect of the dynamics without any control. Using this notion of controllability from (2), we define the controllability with constraints in Definition 1.

Definition 1: The controllability Y_c of (1) with constraints (5) is given by

$$Y_c = \text{volume} \left\{ A^T x_0 : h_c(x_0) \leq 1 \right\}, \text{ where} \\ h_c(x_0) := \min_{\phi_u} x_0^\top \left(\sum_{i=0}^{T-1} \phi_u(i)^\top \phi_u(i) \right) x_0 \\ \text{s.t. } S_c \vec{\phi}_u + C_c = 0. \quad (6)$$

The function h_c is the squared norm of the input signal, and $A^\top x_0$ is the scaled set of initial conditions. The constraint for the terminal state condition, $x_T = 0$, as well as support constraints for ϕ_x and ϕ_u , are specified by the form of S_c and C_c . Recall that we use the hat symbol over a matrix to indicate the vectorization of matrix multiplication. The constraint matrices S_c and C_c are given by

$$S_c = \begin{bmatrix} S_u \\ D_1 \\ D_2 \end{bmatrix}, \quad C_c = \begin{bmatrix} 0 \\ v_1 \\ \hat{A}^T \vec{I} \end{bmatrix}, \quad v_1 = \begin{bmatrix} S_x[1] \vec{I} \\ \vdots \\ S_x[T-1] \hat{A}^{T-1} \vec{I} \end{bmatrix},$$

$$D_1 = \begin{bmatrix} S_x[1] \hat{B} & & & \\ S_x[2] \hat{A} \hat{B} & S_x[2] \hat{B} & & \\ \vdots & & \ddots & \\ S_x[T-1] \hat{A}^{T-2} \hat{B} & \dots & & \hat{B} \end{bmatrix}, \\ D_2 = [\hat{A}^{T-1} \hat{B} \quad \dots \quad \hat{B}].$$

The observability is likewise stated in Definition 2.

Definition 2: The observability Y_o of (1) with constraints (5) is given by

$$Y_o = \text{volume} \left\{ A^{T^\top} x_0 : h_o(x_0) \leq 1 \right\}, \text{ where} \\ h_o(x_0) := \min_{\phi_u} x_0^\top \left(\sum_{i=0}^{T-1} \phi_u(i)^\top \phi_u(i) \right) x_0 \\ \text{s.t. } S_o \vec{\phi}_u + C_o = 0, \quad (7)$$

The matrices S_o and C_o again utilize the terminal constraint $x_T = 0$, with S_o and C_o defined analogously to S_c and C_c .

Example 1 (Distributed Control): In large-scale systems with fast propagation of the dynamics, such as the power grid, the speed at which sensors send information can be on the same order of magnitude as the dynamics [14] and understanding controllability and observability with delays is an important problem [15]. This means that information about a distant state may not be available for a number of time steps, and controllers must be designed accounting for this lack of information. Additionally, if a large disturbance hits a system, ensuring that the disturbance will not propagate further than r neighbors ensures that distant states will not be affected by a disturbance about which they only have delayed information [9]. The sensor and actuator co-design problem with constraints [16] is important for addressing these new challenges; comparing the controllability and observability of sensors and actuators for large-scale systems with these delays and constraints helps designers understand how much faster or more densely-placed hardware improves performance.

Example 2 (Sensorimotor control): In neurophysiology, a nerve is made up of several heterogeneous communication channels formed by different types of axons. To achieve the control objective, the design of the communication channels faces a strict speed-accuracy trade-off [17] over channels with lower delay and less information or a channel with longer delay and more information. Our constrained observability problem (7) provides a method to compare the observability of different sets of axons that satisfy a fixed metabolic cost.

Example 3 (Privacy as Observability): The problem of preserving privacy in decentralized control of large-scale networked systems has received considerable attention recently, including privacy through the perspective of observability [18], [19], [20], [21], [22]. Using the interpretation of observability as a signal-to-noise ratio of an observer, less observability potentially corresponds to less noise needed for obtaining some fixed level of privacy. In particular, our method provides tools for analyzing this signal-to-noise ratio with sensing delay and locality of sensors.

IV. RESULTS

This section presents our main three technical results. Theorem 4 states that the controllability volume given by (6) is equivalent to classical controllability (2) when the constraint set is empty. Algorithm 1 offers a provably correct closed-form solution for controllability, and Theorem 9 provides conditions under which adding constraints to the closed-loop map ϕ does not change the set of feasible states and inputs. As a consequence, the controllability does not change with the addition of the constraints (Corollary 10).

A. Connection to classical methods

Our first result shows that, in the setting without support constraints, the generalized definition of controllability (Definition 1) reduces to the classical definition (2). This shows that the formulation is consistent with classical controllability.

Theorem 4: When $S_u = 0$ and $S_x = 0$, $Y_c^2 = V_c$ for all $A \in \mathbb{R}^{N_x \times N_x}$, $B \in \mathbb{R}^{N_x \times N_u}$.

Proof: When the constraint set is empty, that is, when S_x and S_u are zero, the remaining constraint in S_c and C_c is the terminal constraint $x_T = 0$. Using $u = \phi_u x_0$, we can rewrite the objective of (6) as $\|u\|^2$. The constraint $S_c \vec{\phi}_u + C_c = 0$ can be written as $A^T x_0 + Du = 0$. Here, D takes the form $D = [A^{T-1} B \dots B]$ and $u = [u(0), \dots, u(T-1)]$. The optimal u is given by

$$u = -D^\dagger A^T x_0.$$

Plugging this into $h_c(x_0)$ results in

$$\begin{aligned} h_c(x_0) &= x_0^\top (A^{T^\top} D^{\dagger^\top} D^\dagger A^T) x_0 \\ &= \|D^\dagger A^T x_0\|_2^2. \end{aligned}$$

To compute the volume of all x_0 such that $h_c(x_0) \leq 1$, note that the volume in terms of a new variable $y = A^T x_0$ is the volume of an ellipse defined by D^\dagger :

$$\begin{aligned} \text{volume} \left\{ A^T x_0 \mid h_c(x_0) \leq 1 \right\} &= \text{volume} \left\{ y \mid \|D^\dagger y\|^2 \leq 1 \right\} \\ &= \sqrt{\det D D^\top}, \end{aligned}$$

and using the definition of D results in $Y_c^2 = \sum_{i=0}^{T-1} A^i B B^\top (A^\top)^i = V_c$. ■

B. Computing Constrained Controllability

Our second contribution is a novel algorithm (Algorithm 1) to calculate controllability with constraints in closed form. Algorithm 1 starts by eliminating the linear constraints in (6) to reformulate it as an unconstrained optimization problem. We first decompose $S_c \vec{\phi}_u + C_c = 0$ as n independent sets of linear constraints $S_i \Phi_i + C_i = 0$, for $i \in [N_x]$ in Step 1. Here, $\Phi_i \in \mathbb{R}^{TN_u}$ contains the entries in $\vec{\phi}_u$ that correspond to $\phi_u(t)[i, :]$ for $t < T$. Such decomposition can always be implemented due to the special structure of the support constraints induced by actuation delay, communication speed, and locality constraints. We state this result in Lemma 5.

Lemma 5: For any k , the constraint $(S_c \vec{\phi}_u + C_c)[k] = 0$ cannot depend on entries from both Φ_i and Φ_j ($i \neq j$).

Algorithm 1 Constrained Controllability

Require: Horizon T , Constraint matrices S_c and C_c .

- 1: Decompose $S_c \vec{\phi}_u + C_c = 0$ as $S_i \Phi_i + C_i = 0$, $i \in [N_x]$.
▷ Φ_i contains the entries in $\vec{\phi}_u$ that correspond to $x_0[i]$.
- 2: Compute the general solution $\Phi_i = \omega_i + V_i \lambda_i$, $i \in [N_x]$.
▷ $\lambda_i \in \mathbb{R}^{r_i}$ and $V_i \in \mathbb{R}^{(TN_u) \times r_i}$ has full column rank.
- 3: Define $M = [V_1, V_2, \dots, V_{N_x}] \in \mathbb{R}^{(TN_u) \times (\sum_{i=1}^{N_x} r_i)}$.
- 4: Define $N = [\omega_1, \omega_2, \dots, \omega_{N_x}] \in \mathbb{R}^{(TN_u) \times N_x}$.
- 5: Return the volume

$$\det(A^T) / \sqrt{\det((-MM^\dagger N + N)^\top (-MM^\dagger N + N))}.$$

Proof: Under the definitions of constraints on actuation delay, communication speed, and locality in Section III, the constraint $(S_c \vec{\phi}_u + C_c)[k] = 0$ corresponds to either $\vec{\phi}_u[\tau] = 0$ or $\vec{\phi}_x[\tau] = 0$ for some index $\tau \in [TN_u N_x]$. In the first case, it satisfies the statement of Lemma 5.

Otherwise, $(S_c \vec{\phi}_u + C_c)[k] = 0$ must correspond to $\vec{\phi}_x[\tau] = 0$ for some index $\tau \in [TN_u N_x]$. Suppose it is converted from $\phi_x(t)[j, i] = c$, where $c \in \{0, 1\}$. By (4), we see that $\phi_x(\tau+1)[:, i] = A\phi_x(\tau)[:, i] + B\phi_u(\tau)[:, i]$ for all $\tau < t$. Therefore, $\phi_x(t)[j, i]$ can be expressed as a linear combination of $\{\phi_u(\tau)[:, i] \mid 0 \leq \tau < t\}$. Since all entries in $\phi_u(\tau)[:, i]$ are in Φ_i , $(S_c \vec{\phi}_u + C_c)[k] = 0$ only involves entries from Φ_i . ■

After the decomposition, Step 2 solves the general solution $\Phi_i = \omega_i + V_i \cdot \lambda_i$ of each constraint set $S_i \Phi_i + C_i = 0$, so we can change the variables from $\vec{\phi}_u$ to $\{\lambda_i\}_{i \in [N_x]}$ in the optimization problem (6). With the matrices M and N defined in Steps 3-4, we see $h_c(x_0)$ can be expressed as

$$\min_{\lambda_1, \dots, \lambda_{N_x}} \|N x_0 + M(x_0[1]\lambda_1, \dots, x_0[N_x]\lambda_{N_x})^\top\|_2^2. \quad (8)$$

Finally, we conclude the discussion with a proposition that shows the output of Algorithm 1 is equal to the constrained controllability Y_c defined in (6).

Proposition 6: The constrained controllability Y_c satisfies $Y_c = \text{volume} \left\{ A^T x_0 : \|(-MM^\dagger N + N)x_0\|_2^2 \leq 1 \right\}$, which is the output of Algorithm 1.

Proof: By (8), we see $h_c(x_0) = \min_y \|N x_0 + M y\|_2^2$ holds for all x_0 that satisfies $x_0[i] \neq 0$ for all $i \in [N_x]$, because we can do the change of variables

$$y := (x_0[1]\lambda_1, \dots, x_0[N_x]\lambda_{N_x})^\top \in \mathbb{R}^{\sum_{i=1}^{N_x} r_i}.$$

Since this equation holds almost everywhere, we obtain that

$$\begin{aligned} Y_c &= \text{volume} \left\{ A^T x_0 : \min_y \|N x_0 + M y\|_2^2 \leq 1 \right\} \\ &= \text{volume} \left\{ A^T x_0 : \|(-MM^\dagger N + N)x_0\|_2^2 \leq 1 \right\}, \end{aligned}$$

which is equal to the return value of Algorithm 1. ■

It is worth noting that the time complexity of Algorithm 1 can grow polynomially in time horizon T and system dimensions (N_x, N_u) . In Section V-B, we show that the rank condition is faster than the volume computation, with the gap increasing as the state size increases.

C. Maintaining Performance Under Constraints

An important question for system design is whether the performance of the unconstrained system is preserved after imposing additional constraints on actuation delay, communication speed, and locality. While a straightforward answer is to compute the constrained controllability Y_c analytically with Algorithm 1, it can be inefficient for large-scale systems with high-dimensional state spaces. To address this challenge, we provide sufficient conditions under which

- 1) adding constraints does not decrease controllability;
- 2) constrained systems achieve the optimal unconstrained convex cost.

More importantly, they can be verified efficiently. Our result is inspired by a recent work [13] on comparing localized/global model predictive control (MPC), with the critical difference that we require $x_T = 0$ in (6). While [13] requires that one must check the conditions hold for each x_0 , we prove that if our conditions are satisfied, the performance is maintained for x_0 almost everywhere (a.e.).

First, we define sets $\mathcal{S}(x_0)$ and $\mathcal{T}(x_0)$ as the set of all feasible states and control inputs under the dynamics (1) and initial condition x_0 , without and with constraints. Let

$$Z_{AB0} := \begin{bmatrix} I & 0 \\ -A & -B \\ \ddots & \ddots \\ -A & -B \\ -A & -B \end{bmatrix},$$

such that $Z_{AB0} \in \mathbb{R}^{N_x T \times (N_x + N_u)T}$. Let $Z_p := Z_{AB0}^\dagger \begin{bmatrix} I \\ 0 \end{bmatrix}$, $Z_h := I - Z_p^\dagger Z_{AB0}$. Let $Z_h^{blk} = \text{blkdiag}(Z_h, \dots, Z_h)$, repeated N_x times. The indices R are defined as the indices of $\vec{\phi}$ where the entry is constrained to be zero. Denote the matrix $F = Z_h^{blk}[R]$; due to the structure of Z_h^{blk} , F has the form $F = \text{blkdiag}(\{Z_h[R_i]\}_{i \in [N_x]})$, where each R_i is the subset of indices in R that corresponds to the i^{th} block of Z_h^{blk} . Define $x_0 := [x_0[1]I \ x_0[2]I \ \dots \ x_0[N_x]I]$. We use the results of the following lemmas to prove the main results. Finally, let $H := [H_1, \dots, H_{N_x}]$ where $H_i = I - F[R_i]^\dagger F[R_i]$.

Lemma 7: The dimension of $\mathcal{S}(x_0)$ is given by $\dim \mathcal{S}(x_0) = \text{rank}(Z_h)$.

Lemma 8: The dimension of $\mathcal{T}(x_0)$ is given by $\dim \mathcal{T}(x_0) = \text{rank}(Z_h x_0 (I - F^\dagger F))$.

See the appendix for proofs, which follow closely from [13].

Theorem 9 (Optimality of Constrained Systems): For any constraint set S_x, S_u , if there exists a feasible solution $\phi \in \mathcal{S}$ satisfying (4), $x_T = 0$, and $\text{rank}(Z_h) = \text{rank}(Z_h H)$, then $\mathcal{S}(x_0) = \mathcal{T}(x_0)$ for all x_0 almost everywhere.

Proof: Note that F is a block-diagonal matrix; the product $Z_h x_0 (I - F^\dagger F)$ is therefore

$$Z_h x_0 (I - F^\dagger F) = [Z_h x_0[1] (I - Z_h[R_1]^\dagger Z_h[R_1]) \ \dots \ Z_h x_0[N_x] (I - Z_h[R_{N_x}]^\dagger Z_h[R_{N_x}])].$$

Let $a = \text{rank}(Z_h H)$. The matrix $Z_h H$ can be row-reduced into a matrix \hat{H} , where the first a rows of \hat{H} are linearly

independent and the last $N_x - a$ rows are linearly dependent on the first a rows. Select $\lambda \in \mathbb{R}^a$ such that

$$\sum_{j=1}^a \lambda_j [\hat{H}_1[j], \dots, \hat{H}_{N_x}[j]] = 0.$$

Then because the first a rows of \hat{H} are linearly independent, $\lambda_j = 0$ for $j \in [1, \dots, a]$. For any x_0 , define \tilde{x}_0 as x_0 re-weighted by the row-reduction used to define \hat{H} from H . If \tilde{x}_0 satisfies $\tilde{x}_0[i] \neq 0 \ \forall i \in [N_x]$, then, for any $\tilde{\lambda} \in \mathbb{R}^a$ such that

$$\sum_{j=0}^a \tilde{\lambda}_j [\tilde{x}_0[1] \hat{H}_1[j], \dots, \tilde{x}_0[N_x] \hat{H}_{N_x}[j]] = 0,$$

we have that $\tilde{\lambda}_j \tilde{x}_0[i] = 0$ for all $i \in [N_x]$ and $j \in [a]$. Because $\tilde{x}_0[i] \neq 0$ for all $i \in [N_x]$, so $\tilde{\lambda}_j = 0$ for all $j \in [a]$. This results in the desired rank condition $\text{rank}(Z_h x_0 H) = \text{rank}(Z_h)$ for all x_0 such that $\tilde{x}_0[i] \neq 0$ for all $i \in [N_x]$. Note that we do not analyze the linearly dependent rows of \hat{H} because they will remain linearly dependent when rescaled by \tilde{x}_0 . By Lemmas 7 and 8, $\dim \mathcal{S}(x_0) = \dim \mathcal{T}(x_0)$ for x_0 almost everywhere. Since $\mathcal{T}(x_0) \subseteq \mathcal{S}(x_0)$, the same dimensionality implies that the sets are equal. ■

Finally, we present a corollary of Theorem 9 that gives criteria for optimality conditions.

Corollary 10: If there exists a feasible solution $\phi \in \mathcal{S}$ and $\text{rank}(Z_h) = \text{rank}(Z_h H)$, then for any convex cost function $f(x, u)$, the unconstrained cost is achieved under the constraints $\phi \in \mathcal{S}$ for x_0 almost everywhere, with $V_c = Y_c^2$.

Proof: By Theorem 9, the theorem assumptions imply that $\mathcal{S}(x_0) = \mathcal{T}(x_0)$ for x_0 almost everywhere. Therefore any optimization problem over the constrained set obtains the same value as the unconstrained problem for x_0 almost everywhere.

Now select $f(x, u) = \|u\|_2^2$. The above result implies that the optimization $h_c(x_0)$ in (6) is over the same set with and without constraints, so the solution is the same for x_0 almost everywhere. Because the volume of an ellipse does not change over a measure-zero set, the volume over all x_0 is also the same. The result $V_c = Y_c^2$ then holds by Theorem 4. ■

From this result, we conclude that for commonly-used cost functions such as quadratic cost, the addition of constraints does not necessarily hurt the performance of the controller. This contrasts with some recent results on control [5], [6] and reinforcement learning [23], [24], where localized policies are less optimal than global policies. While this may seem surprising, it is worth noticing that we allow the controller to decide the (constrained) mapping $\vec{\phi}_u$ based on the initial state x_0 , which is easier than committing to a policy before x_0 is sampled or picked by an adversary. Further, the SLS formulation allows us to consider a more general policy class $u(t) = K(t)x(t)$ than the static policy $u(t) = Kx(t)$ in [5].

We show in the following numerical results how locality constraints satisfy this rank requirement and the controllability correspondingly does not change within this region. In contrast, actuation delay does not satisfy this condition

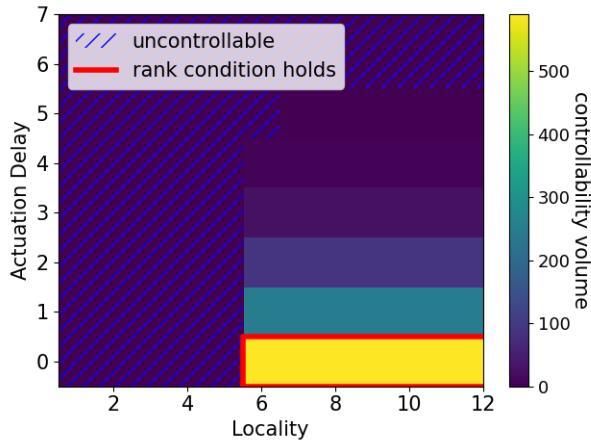


Fig. 1. When actuation delay is zero, the addition of locality constraints on the system does not change the controllability volume for locality greater than five. The controllability decreases with increasing actuation delay.

and we observe a decrease in the controllability. In addition to explaining the contrasting behavior of different types of constraints, computing the rank condition is faster and simpler to implement compared with computing the volume. We show a comparison in computation times in the numerical results.

Remark 1: While the formulation of $h_c(x_0)$ might suggest one needs global knowledge of x_0 to solve ϕ_u , one can show when locality constraints are enforced, the value of each entry of ϕ_u depends only on the local entries of x_0 . For example: let A be a tri-diagonal matrix (i.e., $A[i, j] = 0$ if $|i - j| > 1$), $B = I$, $T = 1$ and the locality constraint include only the immediate neighbors. Then $\text{supp } \phi_u(t)$ is also tri-diagonal. Let ϕ_u denote $\phi_u(0)$ in this setting with only one time step. The objective $\|\phi_u x_0\|^2$ is

$$\begin{aligned} \|\phi_u x_0\|^2 &= (\phi_u[0, 0]x_0[0] + \phi_u[0, 1]x_0[1])^2 + \\ &\sum_{j=1}^{N_x-2} (\phi_u[j, j-1]x_0[j-1] + \phi_u[j, j]x_0[j] + \phi_u[j, j+1]x_0[j+1])^2 \\ &+ (\phi_u[N_x-1, N_x-2]x_0[N_x-2] + \phi_u[N_x-1, N_x-1]x_0[N_x-1])^2. \end{aligned}$$

Note that each $\phi_u[j]$ depends only on $x_0[j-1]$ to $x_0[j+1]$. As this objective is *row-separable* and the dynamical constraints are *column separable*, ϕ_u can be computed locally using the alternating direction method of multipliers (ADMM), where at each iteration, a node only needs information from its direct neighbor. See [9, Section 5.5] for details.

V. NUMERICAL RESULTS

A. Distributed Control of Linearized Swing Equations

To illustrate our results concretely, we consider the setting of the linearized swing equations, as used in [13], which models power generation at multiple substations connected together in a power grid. This example allows us to illustrate how controllability changes based on actuation delay and locality constraints. The model is given by

$$x(t+1)[i] = \sum_{j:G[i,j]=1} A_{ij}x(t)[j] + B_{iu}(t)[i]$$

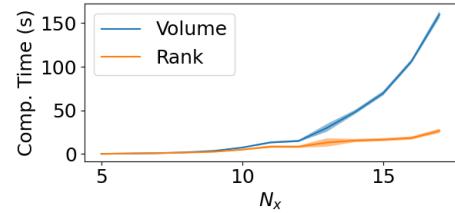


Fig. 2. As the state dimension grows, the amount by which the rank condition computation time is faster than the controllability volume time increases.

$$A_{ii} = \begin{bmatrix} 1 & \Delta t \\ -\frac{k_i}{m_i} \Delta t & 1 - \frac{d_i}{m_i} \Delta t \end{bmatrix}, \quad A_{ij} = \begin{bmatrix} 0 & 0 \\ \frac{k_{ij}}{m_i} \Delta t & 0 \end{bmatrix}, \quad B_i = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

The state $x(t)[i] = [\theta_i(t); \omega_i(t)]$ is the phase angle and frequency deviations at node i . The parameters k_{ij} , m_i^{-1} , and d_i are coupling, inertia, and damping parameters, with $k_i = \sum_{j:G[i,j]=1} k_{ij}$. In the numerical implementation, the discretization step is $\Delta t = 0.1$, $T = 12$. There are nine two-state subsystems with one actuator per subsystem, resulting in $N_x = 18$ and $N_u = 9$. The subsystems are arranged in a 3×3 grid, with possible neighbors of a particular subsystem being the other subsystems to the left, right, top or bottom. The coefficients were sampled uniformly $k_{ij} \sim [0.5, 1]$, $m_i^{-1} \sim [0, 2]$, $d_i \sim [1, 1.5]$. In Figure 1, we observe that the settings where rank condition and feasibility hold coincides with the maximum controllability; importantly, we see no decrease in controllability with locality constraints until the system is no longer controllable. The controllability decreases with increasing actuation delay, and as expected, the rank condition does not hold for any non-zero actuation delay. Large delay renders the system not controllable.

B. Computation of Rank Condition Versus Volume

The rank condition, which rigorously explains why communication speed and locality constraints may not result in decreased controllability, is easier to implement and faster to compute for large systems. We recorded the total time required to compute the controllability volume for all possible combinations of communication and locality constraints for a system where the dynamics matrix A is given by a tri-diagonal stochastic matrix, B is filled with one and zero entries with actuation density 0.75 relative to the state dimension, the time horizon is set to $T = 10$, and 40 trials were recorded for each combination of constraints. As seen in Figure 2, the time is comparable for N_x ranging from 5 – 10, but the volume computation becomes increasingly slower with larger N_x , with up to 6× speedup using the rank condition instead of the total volume. Users can compute the controllability volume for the unconstrained system and then compute the rank condition and feasibility of the locality- and communication-constrained system to determine the controllability, which will either be equal to the unconstrained system or possibly zero if the condition is not satisfied.

VI. CONCLUDING REMARKS

This work proposes notions of constrained controllability Y_c and constrained observability Y_o using SLS, and shows they reduce to classic definitions for unconstrained systems. We provide closed-form solutions for computing Y_c, Y_o and an efficient criteria to check whether a certain set of constraints can cause performance loss. In future work, we aim to apply this framework to solve real-world sensor and actuator problems, including air and water quality sensor placement for environmental monitoring, and investigate methods for computing the controllability in a distributed fashion for large-scale systems, which would greatly improve the range of settings in which this method can be used.

APPENDIX

A. Proof of Lemma 7

To prove Lemma 7, we first show that the set $\mathcal{S}(x_0)$ can be written as

$$\mathcal{S}(x_0) = \{\mathbf{s} : \mathbf{s} = Z_p x_0 + Z_h x_0 \lambda, \lambda \in \mathbb{R}^{(N_u + N_x)N_x T}\},$$

where $\mathbf{s} = [\mathbf{x}; \mathbf{u}]$. Using the dynamics and terminal constraint $x_T = 0$, the closed-loop map ϕ can be written as

$$Z_{AB0}\phi = \begin{bmatrix} I \\ 0 \end{bmatrix}. \quad (9)$$

This results in $\phi = Z_{AB0}^\dagger \begin{bmatrix} I \\ 0 \end{bmatrix} + (I - Z_{AB0}^\dagger Z_{AB0})\Lambda$, where Λ is a free variable. Using the definitions of Z_p and Z_h results in $\phi = Z_p + Z_h \Lambda$. Then the state and input signals \mathbf{x} and \mathbf{u} result are given by $\mathbf{s} = \phi x_0$, and we use the matrix x_0 to switch the multiplication of Λx_0 to $x_0 \lambda$. The dimension of $\mathcal{S}(x_0)$ is given by the degrees of freedom of \mathbf{s} , which is $\text{rank}(Z_h x_0) = \text{rank}(Z_h)$ because x_0 is full row-rank.

B. Proof of Lemma 8

This proof proceeds similarly to that of Lemma 7, except now the constraint on ϕ is defined in terms of $\vec{\phi}$. From (9) and the additional requirements that $S_u \vec{\phi}_u = 0$ and $S_x \vec{\phi}_x = 0$,

$$\begin{bmatrix} S_x \\ S_u \end{bmatrix} \left(\hat{Z}_p + \hat{Z}_h \vec{\Lambda} \right) = 0 \Rightarrow \hat{Z}_p[R] + \hat{Z}_h[R] \vec{\Lambda},$$

resulting in Λ being parameterized as

$$\vec{\Lambda} = -\hat{Z}_h[R]^\dagger \hat{Z}_p[R] + (I - Z_h[R]^\dagger Z_h[R])\mu,$$

for a free variable μ . Using the definition of F and noting that $\hat{Z}_h[R] = Z_h^{blk}[R]$, the dimension of $\mathcal{T}(x_0)$ is given by $\text{rank}(Z_h x_0(I - F^\dagger F))$.

REFERENCES

- [1] R. E. Kalman, Y. C. Ho, and K. C. Narendra, "Controllability of Linear Dynamical Systems," *Contributions to Differential Equations*, vol. 1, no. 1, 1961.
- [2] E. G. Gilbert, "Controllability and Observability in Multivariable Control Systems," *Journal of the Society for Industrial and Applied Mathematics Series A Control*, vol. 1, pp. 128–151, Jan. 1963.
- [3] D. Georges, "The use of observability and controllability gramians or functions for optimal sensor and actuator location in finite-dimensional systems," in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 4, pp. 3319–3324 vol.4, Dec. 1995. ISSN: 0191-2216.
- [4] T. Kailath, *Linear systems*. Englewood Cliffs, N.J. : Prentice-Hall, 1980.
- [5] S. Shin, Y. Lin, G. Qu, A. Wierman, and M. Anitescu, "Near-optimal distributed linear-quadratic regulator for networked systems," *SIAM Journal on Control and Optimization*, vol. 61, no. 3, pp. 1113–1135, 2023.
- [6] Y. Lin, J. Gan, G. Qu, Y. Kanoria, and A. Wierman, "Decentralized online convex optimization in networked systems," in *International Conference on Machine Learning*, pp. 13356–13393, PMLR, 2022.
- [7] Y. Li, Y. Tang, R. Zhang, and N. Li, "Distributed reinforcement learning for decentralized linear quadratic control: A derivative-free policy optimization approach," *IEEE Transactions on Automatic Control*, vol. 67, no. 12, pp. 6429–6444, 2021.
- [8] J. Yu, D. Ho, and A. Wierman, "Online Adversarial Stabilization of Unknown Networked Systems," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 7, pp. 1–43, Feb. 2023.
- [9] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni, "System level synthesis," *Annual Reviews in Control*, vol. 47, pp. 364–393, 2019.
- [10] Z. Wang and Q. Yuan, "A method for analyzing the state controllability of linear discrete time-varying time-delay systems," in *2020 IEEE 9th Data Driven Control and Learning Systems Conference (DDCLS)*, pp. 1176–1181, Nov. 2020.
- [11] B. Zhou, Z. Wang, Y. Zhai, and H. Yuan, "Data-Driven Analysis Methods for Controllability and Observability of A Class of Discrete LTI Systems with Delays," in *2018 IEEE 7th Data Driven Control and Learning Systems Conference (DDCLS)*, pp. 380–384, May 2018.
- [12] L. Jinna, Z. Qingling, and L. Yuan, "Controllability and observability of networked control systems with time-varying delays," *Journal of Systems Engineering and Electronics*, vol. 20, pp. 800–806, Aug. 2009.
- [13] J. S. Li and C. Amo Alonso, "Global Performance Guarantees for Localized Model Predictive Control," *IEEE Open Journal of Control Systems*, vol. 2, pp. 325–336, 2023.
- [14] J. Taft and P. De Martini, "Sensing and Measurement Architecture for Grid Modernization," tech. rep., Pacific Northwest National Laboratory, 2016.
- [15] C. O. Nwankpa, J. C. Jimenez, and S. Jayasuriya, "Observability of network-delayed multi-converter power systems," in *2015 IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 930–933, May 2015. ISSN: 2158-1525.
- [16] N. Matni and V. Chandrasekaran, "Regularization for design," in *53rd IEEE Conference on Decision and Control*, pp. 1111–1118, Dec. 2014. ISSN: 0191-2216.
- [17] Y. Nakahira, Q. Liu, N. Bernat, T. Sejnowski, and J. Doyle, "Theoretical foundations for layered architectures and speed-accuracy trade-offs in sensorimotor control," in *2019 American Control Conference (ACC)*, pp. 809–814, IEEE, 2019.
- [18] Y. Zhang, R. Cheng, and Y. Xia, "Observability Blocking for Functional Privacy of Linear Dynamic Networks," in *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 7469–7474, Dec. 2023. ISSN: 2576-2370.
- [19] Y. Kawano and M. Cao, "Design of Privacy-Preserving Dynamic Controllers," *IEEE Transactions on Automatic Control*, vol. 65, pp. 3863–3878, Sept. 2020.
- [20] Y. Kawano and M. Cao, "Revisit Input Observability: A New Approach to Attack Detection and Privacy Preservation," in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 7095–7100, Dec. 2018. ISSN: 2576-2370.
- [21] K. H. Degue and J. Le Ny, "Differentially Private Interval Observer Design with Bounded Input Perturbation," in *2020 American Control Conference (ACC)*, pp. 1465–1470, July 2020. ISSN: 2378-5861.
- [22] J. Le Ny, "Privacy-preserving nonlinear observer design using contraction analysis," in *2015 54th IEEE Conference on Decision and Control (CDC)*, pp. 4499–4504, Dec. 2015.
- [23] Y. Lin, G. Qu, L. Huang, and A. Wierman, "Multi-agent reinforcement learning in stochastic networked systems," *Advances in neural information processing systems*, vol. 34, pp. 7825–7837, 2021.
- [24] Y. Zhang, G. Qu, P. Xu, Y. Lin, Z. Chen, and A. Wierman, "Global convergence of localized policy iteration in networked multi-agent reinforcement learning," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 7, no. 1, pp. 1–51, 2023.